AN UPPER BOUND ON THE DEHN FUNCTION OF $Out(A_\Gamma)$

JAKUB F. TUCKER

Abstract. To obtain an upper bound on the Dehn function of the outer automorphism group $Out^0(A_\Gamma)$ for a right-angled Artin group A_Γ with defining graph Γ, we use the subnormal series defined by Day and Wade in [5] to decompose $Out^0(A_\Gamma)$. This yields a decomposition tree where each vertex G has two descendants N and Q, satisfying a short exact sequence

$$1 \to N \to G \to Q \to 1$$

We prove an upper bound for the Dehn function of the group G in relation to the Dehn functions of the groups N and Q. The Dehn functions of the leaves of the decomposition tree are known, with these we can bound above the Dehn function of their root, and by extension that of the group $Out^0(A_\Gamma)$.

1. Introduction

A right-angled Artin group (RAAG) is a finitely presented group determined solely by commutator relations. Given a graph Γ with no loops or multiple edges, one can define the RAAG A_Γ by the following presentation, where $V(\Gamma)$ is the set of vertices, and $E(\Gamma)$ is the set of edges of Γ:

$$A_\Gamma = \left\langle V(\Gamma) \;\middle|\; [u,v] \text{ with } (u,v) \in E(\Gamma) \right\rangle$$

The two extreme examples of RAAGs are free groups F_n - coming from graphs with no edges and n vertices, and free abelian groups - coming from complete graphs. Since a general RAAG lies somewhere between these two extremes, one might naturally think that the outer automorphism groups $Out(A_\Gamma)$ might lie somewhere between $Out(\mathbb{Z}^n) = GL(n,\mathbb{Z})$ and $Out(F_n)$ in some way. There are many properties shared by the extreme examples $Out(F_n)$ and $GL(n,\mathbb{Z})$, such as certain finiteness properties proven by Charney and Vogtmann in [1] and [2], utilising such tools as the restriction and projection homomorphisms. More recently, Day and Wade built upon this foundation in [5] where relative automorphism groups were applied to decompose these groups in a convenient way. Their decomposition was devised to compute exactly the virtual cohomological dimension of a group $Out^0\left(A_\Gamma; \mathcal{G}, \mathcal{H}^t\right)$, as explicitly computed with Sale in [4], but this decomposition proves useful in other cases.

Given an exact sequence of finitely presented groups

$$1 \to N \to G \to Q \to 1$$

where N and Q have superadditive Dehn functions, one can bound δ_G by

$$\delta_G(n) \preccurlyeq \delta_N\left(e^{\delta_Q(n)}\right)$$

Using this bound, one can apply Day and Wade's decomposition to the group $Out^0(A_\Gamma)$ and use the short exact sequences produced by the decomposition tree as in Proposition 2.4 to yield an upper bound on the Dehn function of $Out(A_\Gamma)$.

The bound achieved in this paper, while unlikely to be optimal, depends only on Day and Wade's decomposition. Given a decomposition series

$$1 = N_0 \leq N_1 \leq \cdots \leq N_k = Out^0\left(A_\Gamma; \mathcal{G}, \mathcal{H}^t\right)$$

as in Theorem 2.3, without explicit computation which type of group each quotient is, this result bounds the Dehn function above by

$$\delta_{Out(A_\Gamma)}(n) \preccurlyeq \left. e^{e^{\cdot^{\cdot^{\cdot^{e^n}}}}} \right\} k \leq 2^{|V(\Gamma)|+1}$$

Where the height of the exponential tower the number of terms in the decomposition series, which is less than $2^{|V(\Gamma)|+1} - 1$, for a much weaker bound dependent only on the number of vertices in the graph Γ. See Section 5 for the proof and a sharper bound in 5.1.

Structure of the paper. We begin by describing relevant definitions and results regarding RAAGs, before moving on to Dehn functions. In Section 4, the upper bound on the Dehn function of the short exact sequence is proven, then applied to RORGs in the following Section 5. Examples of explicit computation, as well as cases where the Dehn function bound from Section 4 is optimal, can be found in Section 6.

2. Background on right-angled Artin groups

There are several important definitions that will be essential to discuss the properties of RAAGs and their outer automorphism groups:

- The *link* $lk(v)$ of a vertex $v \in V(\Gamma)$ is the full subgraph of vertices adjacent to v in Γ.
- The *star* $st(v)$ of a vertex $v \in V(\Gamma)$ is the join $v * lk(v)$.
- The *link* $lk(\Delta)$ of a collection of vertices $\Delta \subseteq V(\Gamma)$ is the intersection the subgraphs $\cap_{v\in\Delta} lk(v)$
- The *star* $st(\Delta)$ of a collection of vertices $\Delta \subseteq V(\Gamma)$ is the join $\Delta * lk(\Delta)$.
- The partial order relation $\leq$ on $V(\Gamma)$ is defined by $v \leq w$ if $lk(v) \subseteq st(w)$. It is with respect to this relation that we discuss equivalence classes of vertices in Γ.
- A subgroup of the form $A_\Delta \leq A_\Gamma$ where Δ is a full subgraph of Γ is called a special subgroup.

2.1. The (pure) outer automorphism group.

In this article, we work in particular with the *outer* automorphism group of a RAAG A_Γ. The outer automorphism group $Out(A_\Gamma)$ is defined as the quotient of the full automorphism group $Aut(A_\Gamma)$ by the subgroup of inner automorphisms $Inn(A_\Gamma)$, which is generated by conjugations by a generator v of A_Γ of all generators in $V(\Gamma)$. We describe automorphisms in $Out(A_\Gamma)$ by a representative in $Aut(A_\Gamma)$.

In [8], Laurence describes a set of automorphisms:

- **Graph automorphisms** - automorphisms of A_Γ that arise from permuting vertices in Γ,
- **Transvections** - maps of the form $v \mapsto vw$ for generators $v \leq w$,
- **Partial conjugations** - conjugation by a generator v of connected components of $\Gamma - st(v)$, and
- **Inversions** of a single generator $v \mapsto v^{-1}$,

and proves that the set of all such automorphims is a finite generating set for $Out(A_\Gamma)$. It is also known that the groups $Out(A_\Gamma)$ is finitely presentable, as proven by Day in [3]. We don't need to understand the precise nature of the finite set of relations for $Out(A_\Gamma)$ for this article.

It is very often useful to consider the *pure* outer automorphism group $Out^0(A_\Gamma)$, which is the finite-index subgroup of $Out(A_\Gamma)$ obtained by omitting the graph automorphisms from the set of generators. Since $Out^0(A_\Gamma)$ has finite index in $Out(A_\Gamma)$, the two groups have equivalent Dehn functions. For this reason, the remainder of this article concerns $Out^0(A_\Gamma)$ in particular, rather than the full outer automorphism group. The group $Aut^0(A_\Gamma)$ is the preimage of $Out^0(A_\Gamma)$ under the quotient map by $Inn(A_\Gamma)$

In [1], Charney and Vogtmann define two key homomorphims from $Out^0(A_\Gamma)$ to the pure outer automorphism groups of certain special subgroups, those defined by the stars or links of *maximal equivalence classes.* In Proposition 3.2 of [1], it is shown that for a maximal equivalence class $[v] \subseteq \Gamma$, for any pure outer automorphism $\phi \in Out^0(A_\Gamma)$, there is a representative f_v in $Aut^0(A_\Gamma)$ that preserves both the special subgroups $A_{[v]}$ and $A_{st[v]}$ (and therefore $A_{lk[v]}$ too). With this property, one can define the *restriction homomorphism*

$$R_{[v]} : Out^0(A_\Gamma) \to Out^0(A_{st[v]})$$

induced by the restriction of each f_v to $A_{st[v]}$. Composing this restriction homomorphism with the map $Out^0(A_\Gamma) \to Out^0(A_{\Gamma-[v]})$, induced by mapping each generator $w \in [v]$ to the identity, yields the *projection homomorphism* $P_{[v]}$. The precise nature of the images and kernels of these homomorphisms are described in [5], with the introduction of relative outer automorphism groups by Day and Wade.

2.2. **Relative outer automorphism groups.** Given a graph Γ, and collections $\mathcal{G}, \mathcal{H}$ of special subgroups of A_Γ, the *relative outer automorphism group* (RORG) $Out^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ is the group of pure outer automorphisms which preserve each special subgroup $A_\Delta \in \mathcal{G}$ and act trivially on each special subgroup $A_\Delta \in \mathcal{H}$. That is, that $\Phi \in Out^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ if there exists a representative $\phi \in Aut^0(A_\Gamma)$ such that $\phi(G) = G$ for all $G \in \mathcal{G}$ and $\phi = (id)_H$ for each $H \in \mathcal{H}$.

Importantly, Fouxe-Rabinovitch groups are relative automorphism groups, where each group in the free-factor decomposition is preserved.

There are several more important definitions that are essential to understanding the results in [4] and [5]:

- We say that a collection of special subgroups $\mathcal{G}$ is *saturated* with respect to the pair $(\mathcal{G}, \mathcal{H})$ if $\mathcal{G}$ contains every special subgroup invariant under $Out^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$.
- We call two vertices v, w in Γ $\mathcal{G}$-adjacent if v is adjacent to w in the usual sense, or if there exists $A_\Delta \in \mathcal{G}$ with $v, w \in \Delta$.
- A finite sequence of vertices is a $\mathcal{G}$-path if each vertex is $\mathcal{G}$-adjacent to the next.
- A graph Γ $\mathcal{G}$-connected if there exists a $\mathcal{G}$-path between any two vertices in Γ.
- We call a subgraph Δ $(\mathcal{G}, \mathcal{H})$-*star-separated* by a vertex v if Δ intersects more than one $(\mathcal{G}^v \cup \mathcal{H})$-component of $\Gamma - st(v)$.
- We denote by $\mathcal{G}^v$ the set of special subgroups in $\mathcal{G}$ that do not contain v.

- We define a partial order $\leq_{(\mathcal{G},\mathcal{H})}$ on the vertices of Γ by $v \leq_{(\mathcal{G},\mathcal{H})} w \iff lk(v) \subseteq st(w)$ and $v \notin \mathcal{G}^w \cup \mathcal{H}$, and say that a subgraph Δ is *upwards closed under* $\leq_{(\mathcal{G},\mathcal{H})}$ if $v \in \Delta, v \leq_{(\mathcal{G},\mathcal{H})} \Rightarrow w \in \Delta$.

One can find the *saturation* of a pair $(\mathcal{G},\mathcal{H})$ using

Proposition 2.1. ([4], Proposition 2.8) *Let A_Δ be a special subgroup of A_Γ,*

- *A_Δ is invariant under $Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t)$ if and only if Δ is upwards closed under $\leq_{(\mathcal{G},\mathcal{H})}$ and Δ is not $(\mathcal{G},\mathcal{H})$-star-separated by a vertex $v \in \Gamma - \Delta$.*
- *$Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t)$ acts trivially on A_Δ if and only if $v \leq_{(\mathcal{G},\mathcal{H})} w \Rightarrow v = w$ for all $v \in \Delta$ and Δ is not $(\mathcal{G},\mathcal{H})$-star-separated by any vertex of Γ, and every element of Δ is contained in an element of $\mathcal{H}$.*

The restriction and projection homomorphisms can be restricted to relative automorphism groups, and the restriction homomorphism satisfies the following result:

Theorem 2.2. ([5], Theorem E) *Let $\mathcal{G},\mathcal{H}$ be collections of special subgroups of A_Γ, with $A_\Delta \in \mathcal{G}$. Suppose that $\mathcal{G}$ is saturated with respect to the pair $(\mathcal{G},\mathcal{H})$. Then the restriction homomorphism*

$$R_\Delta : Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t) \to Out(A_\Delta)$$

satisfies the short exact sequence

$$1 \to Out^0(A_\Gamma;\mathcal{G},(\mathcal{H}\cup\{A_\Delta\})^t) \to Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t) \overset{R_\Delta}{\to} Out^0(A_\Delta;\mathcal{G}_\Delta,\mathcal{H}^t_\Delta) \to 1$$

Where $\mathcal{G}_\Delta$ is defined as $\{A_{\Delta\cap\Theta} \big| A_\Theta \in \mathcal{G}\}$ and $\mathcal{H}_\Delta$ is defined similarly.

This restriction map is used to decompose $Out^0\left(A_\Gamma;\mathcal{G},\mathcal{H}^t\right)$ unless all the restriction maps have trivial image. If this happens, there are five subcases:

(1) Γ is disconnected and $\mathcal{G}$-disconnected,
(2) Γ is disconnected and $\mathcal{G}$-connected,
(3) Γ is connected, and the centre $Z(A_\Gamma)$ is trivial,
(4) Γ is connected, and the centre $Z(A_\Gamma)$ is a proper non-trivial subgroup, or
(5) Γ is complete and $A_\Gamma = \mathbb{Z}^n$ for some n.

Note that if every restriction homomorphism is trivial, then $Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t)$ acts trivially on every $A_\Delta \in \mathcal{G}$, so we may assume $\mathcal{G} = \mathcal{H}$ is saturated with respect to $(\mathcal{G},\mathcal{G})$ and write $Out^0(A_\Gamma;\mathcal{G}^t)$ in place of $Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t)$ Each of these cases has a corresponding result in ([5], Section 5.1) essential to the decomposition.

(1) There is a free-factor decomposition

$$A_\Gamma = A_{\Delta_1} * A_{\Delta_2} * \cdots * A_{\Delta_k} * F_m$$

In which case $Out^0(A_\Gamma;\mathcal{G}^t) = Out^0(A_\Gamma,\{A_{\Delta_i}\}_i^t)$ is a Fouxe-Rabinovitch group.

(2) $Out^0(A_\Gamma;\mathcal{G}^t)$ is a finite-rank free abelian group.
(3) $Out^0(A_\Gamma;\mathcal{G}^t)$ is a finite-rank free abelian group.
(4) Let $\Delta = \Gamma - z(\Gamma)$, where $z(\Gamma)$ is the full subgraph consisting of vertices which are adjacent to every other vertex, so that $A_{z(\Gamma)} = Z(A_\Gamma)$, then the projection homomorphism P_Δ satisfies the short exact sequence

$$1 \to K_{P_\Delta} \to Out^0(A_\Gamma;\mathcal{G}^t) \overset{P_\Delta}{\to} Out^0(A_\Delta;\mathcal{G}^t_\Delta) \to 1$$

with kernel K_{P_Δ} free abelian generated by leaf transvections in $Out^0(A_\Gamma;\mathcal{G}^t)$ - in this case the leaf transvections are those of the from $u \mapsto uw$ with $w \in z(\Gamma)$ and $u \notin z(\Gamma)$.

(5) $A_\Gamma = \mathbb{Z}^n$ and there is some $1 \leq m \leq n$ for which the exact sequence

$$1 \to A \to Out^0(A_\Gamma;\mathcal{G}^t) \to GL(m,\mathbb{Z}) \to 1$$

is satisfied where A is a finitely generated free abelian group of matrices with rank $m(n-m)$, note we can have $m = n$ and A trivial.

2.3. **Decomposing RORGs.** The key result for our Dehn function bound is the following theorem by Day and Wade.

Theorem 2.3. ([5],Theorem 5.9) *Let A_Γ be a right-angled Artin group, and $\mathcal{G},\mathcal{H}$ any collections of special subgroups of A_Γ. Then there is a subnormal series*

$$1 = N_0 \leq N_1 \leq \cdots \leq N_k = Out^0\left(A_\Gamma;\mathcal{G},\mathcal{H}^t\right)$$

where each quotient N_{i+1}/N_i is isomorphic to one of

(D1) *a finitely generated free abelian group,*
(D2) *$GL(n,\mathbb{Z})$ for some n, or*
(D3) *a Fouxe-Rabinovitch group $Out^0(A_\Gamma,\{A_{\Delta_i}\}_i^t)$ for a free factor decomposition $A_\Gamma = A_{\Delta_1} * A_{\Delta_2} * \cdots * A_{\Delta_k} * F_m$.*

With this result, we can use short exact sequences of the form

$$1 \to N_i \to N_{i+1} \to N_{i+1}/N_i \to 1$$

and a bound on the Dehn function of N_{i+1} in terms of the Dehn functions of N_i and the quotient to bound $\delta_{Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t)}$. In particular, we will use this result for $Out^0(A_\Gamma)$, which is a finite-index subgroup of $Out(A_\Gamma)$, to obtain an upper bound on the Dehn function of $Out(A_\Gamma)$.

While this result suffices to achieve a bound, it is not yet clear how one computes such a subnormal series. Proposition 4.2 of [4] describes the process of obtaining a *decomposition tree* for $Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t)$, with our initial group at the root, leaves labelled by a group of one of the quotient types (D1), (D2), or (D3). Each internal vertex of this tree is labelled by a RORG G, and has two descendants N and Q which satisfy the short exact sequence

$$1 \to N \to G \to Q \to 1$$

Proposition 2.4. ([4], Proposition 4.2) *There exists an algorithm that produces a decomposition tree for $Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t)$*

Proof. The process for obtaining a tree is iterative, given a vertex v in the tree, labelled by a group $Out^0(A_{\Gamma_v};\mathcal{G}_v,\mathcal{H}_v^t)$, one either recognises the group as one of the quotient types (D1), (D2), or (D3) or extends the tree by adding two new descendants v_1 and v_2 of v which are either RORGs of lower complexity (in the sense of [5], Theorem 5.9) or a group of one of the three quotient types. Since the complexity of RORGs decreases as we get further from the root, this process will terminate.

Given $Out^0(A_{\Gamma_v};\mathcal{G}_v,\mathcal{H}_v^t)$, first we extend $\mathcal{G}_v$ to its saturation $\mathcal{G}_v'$ with respect to the pair $(\mathcal{G},\mathcal{H})$ by determining the invariant subgroups.

If there exists a special subgroup A_Δ of A_{Γ_v} such that the restriction map R_Δ is non-trivial, we get 2 descendants from the kernel and image of the short exact sequence

$$\begin{aligned}1 \to Out^0(A_{\Gamma_v}; \mathcal{G}_v, (\mathcal{H}_v \cup \{A_\Delta\})^t) \to Out^0(A_{\Gamma_v}; \mathcal{G}_v, \mathcal{H}_v^t)\\ \xrightarrow{R_\Delta} Out^0(A_\Delta; (\mathcal{G}_v)_\Delta, (\mathcal{H}_v)^t_\Delta) \to 1\end{aligned}$$

Otherwise, we consider the five cases as above:

- In case 1, our group is of type (D3) and is a leaf.
- In cases 2 and 3, our groups is of type (D1) and is a leaf,
- In case 4, set $\Delta = \Gamma_v - z(\Gamma_v)$ and we use the short exact sequence

  $$1 \to K_{P_\Delta} \to Out^0(A_{\Gamma_v}; \mathcal{G}_v^t) \xrightarrow{P_\Delta} Out^0(A_\Delta; (\mathcal{G}_v)^t_\Delta) \to 1$$

  To obtain two new descendants, the kernel is of type (D1) so is a leaf.
- In case 5, we use the short exact sequence

  $$1 \to A \to Out^0(A_{\Gamma_v}; \mathcal{G}_v^t) \to GL(m, \mathbb{Z}) \to 1$$

  to obtain two new descendants, the kernel is of type (D1) and the image is of type (D2) so both are leaves

□

Note that each time a non-trivial restriction map is used to create two new descendants, a choice of subgraph $\Delta \subset \Gamma$ is made. It is unclear whether the quotients depend on the choice of decomposition tree.

3. Background on Dehn functions

Given a group G with finite presentation $\langle A \mid R \rangle$, the *Dehn function* $\delta_{\langle A|R\rangle} : \mathbb{N} \to \mathbb{N}$ of G with respect to this presentation is a bound on the *area* of a nullhomotopic word w, *i.e.* $w =_G 1$ in terms of its length. The notion of area can be viewed in multiple ways: algebraically, as the minimal number of applications of relators $r \in \langle\langle R \rangle\rangle$ required to reduce w to the empty word, or geometrically as the minimal number of 2-cells in a Van Kampen diagram for w - a minimal isoperimetric function. These two definitions are *equivalent* in the sense as described below and throughout this article I will regularly switch between these two perspectives when useful.

3.1. Ordering and Equivalence.

There is an ordering $\preccurlyeq$ for functions $[0, \infty) \to [0, \infty)$ that is commonly used for Dehn functions. Given two maps $a, b : [0, \infty) \to [0, \infty)$, we say $a \preccurlyeq b$ if there exists a constant $C \geq 1$ such that $a(n) \leq Cb(Cn + C) + Cn + C$ for all $n \in [0, \infty)$. For application to Dehn functions, this definition is restricted to $n \in \mathbb{N}$. We say that two functions $a, b : [0, \infty) \to [0, \infty)$ are $\simeq$-equivalent if $a \preccurlyeq b$ and $b \preccurlyeq a$. Whenever equivalence of Dehn functions is mentioned, it is referring to $\simeq$-equivalence.

A priori, the Dehn function depends on the specific presentation of a group G. A helpful result is that the Dehn function with respect to any two finite presentations of the same group are $\simeq$-equivalent. With this result, we consider Dehn functions up to $\simeq$-equivalence, and often compare them to simple functions such as polynomials or exponentials.

3.2. **Known Dehn functions.** For the purposes of this article, we are required to know some Dehn functions of specific groups. In particular, groups of type (D1), (D2), or (D3) as described in Theorem 2.3. Fortunately, these results are known.

(D1) Finitely generated free abelian groups G have quadratic Dehn function $\delta_G \simeq n^2$,

(D2) $GL(n,\mathbb{Z})$ has linear Dehn function for $n = 1,2$, Dehn function $\preccurlyeq e^n$ for $n = 3,4$, and $\preccurlyeq n^2$ for $n > 4$,

(D3) Fouxe-Rabinovitch groups $Out^0(A_{\Delta_1} * \cdots * A_{\Delta_k} * F_m)$ are believed to have exponential Dehn function, based on [9], though I have not yet been able to access this paper. In particular this case covers the groups $Out(F_n)$ which are known to satisfy an exponential isoperimetric inequlity as in [6].

Observe that all three types satisfy an exponential upper bound, so we know that $\delta_G(n) \preccurlyeq e^n$ for any group G of type (D1), (D2), or (D3). This is useful for a more general result, where we are not required to compute the subnormal series as in Theorem 2.3.

3.3. **Passing to finite-index subgroups.** Day and Wade's subnormal series in Theorem 2.3 will give us an upper bound on the Dehn function of $Out^0(A_\Gamma)$, but we can generalise this to a bound on the Dehn function of $Out(A_\Gamma)$.

The Cayley graph $Cay^1(G,A)$ of a group G with generating set A is the labelled and directed graph defined by assigning each element $g \in G$ a vertex, and for every $g \in G, a \in A$, there is a directed edge from g to ga labelled a. The metric on the Cayley graph is defined by the word metric on the vertices of $Cay^1(G,A)$, and extended to the edges by assigning them to each have length 1.

Let $(X,d_X),(Y,d_Y)$ be metric spaces, let $\lambda \geq 1, \varepsilon \geq 0$ be constants. Then a map $f : X \to Y$ is a (λ,ε)-*quasi-isometry* if for every two points $x,x' \in X$, we have

$$\frac{1}{\lambda} d_X(x,x') - \varepsilon \leq d_Y(x,x') \leq \lambda d_X(x,x') + \varepsilon$$

Two metric spaces $(X,d_X),(Y,d_Y)$ are called *quasi-isometric* if there exists such as (λ,ε)-quasi-isometry f between them, and there is a map $g : Y \to X$ such that there exists a constant $C \geq 0$ where $d_Y(f \circ g(y), y) \leq C$ for all $y \in Y$, and such a function g is a called a *quasi-inverse* for f.

It is known that the Dehn function is a *quasi-isometry* invariant. This means that two quasi-isometric groups have equivalent Dehn functions, where two groups are called quasi-isometric if their Cayley graphs are quasi-isometric as metric spaces.

Given a finite-index subgroup H in G, with the generating set of G being a generating set for H along with a finite transversal, then one can check that the inclusion map $H \hookrightarrow G$ induces a quasi-isometry between the Cayley graphs of G and H with respect to these generator, and that the map ϕ - that sends vertices g in $G \setminus H$ to the unique vertex $h \in H$ where $g = hg_i$ for a transversal letter g_i, sends vertices in H to themselves and points on edges to the image of either of their endpoints - defines a quasi-inverse to the inclusion map.

Since it is known that the group $Out^0(A_\Gamma)$ has finite index in the full outer automorphism group $Out(A_\Gamma)$, it follows that these two groups have equivalent Dehn functions.

4. Bounding Dehn functions in short exact sequences

In this section we will be dealing with short exact sequences of the form

$$1 \to N \to G \to Q \to 1$$

We will prove a bound on $\delta_G : \mathbb{N} \to \mathbb{N}$ as a function of δ_Q and δ_N, assuming that these two Dehn functions are known.

4.1. The split short exact sequence case.

Suppose we are dealing with the split short exact sequence

$$1 \to N \xrightarrow{f} G \underset{h}{\overset{g}{\rightleftarrows}} Q \to 1$$

This means that we can write $G \cong N \rtimes Q$, and say that G is generated by $\{n_1, \dots, n_k, q_1, \dots, q_r\}$ where $\{n_1, \dots, n_k\}$ is a generating set for N and $\{q_1, \dots, q_r\}$ is a generating set for Q. As for the relations in G, by the semi-direct product we have all the relations from N and all the relations from Q, as well as *conjugation-by-Q* relations of the form $q_i n_j q_i^{-1} = v_{i,j}(n_1, \dots, n_k)$, where $v_{i,j}$ is a word in N and $i \in \{1, \dots, r\}, j \in \{1, \dots, k\}$, coming from the normality of N in G. There are rk such $v_{i,j}$ - finitely many. Therefore we can set λ to be the maximal length of these $v_{i,j}$.

Suppose w is an arbitrary word of length t in G, with $w =_G 1$. If $\lambda > 1$, we can apply conjugation-by-Q relations no more than λ^{t-1} times to move any q_i to the right side. Each application of a conjugation-by-Q-relator adds at most λ to the number of N-letters in the word, and preserves the number of Q-letters - so we are left with a word of the form $u(n_1, \dots, n_k)v(q_1, \dots, q_r)$ of length $\leq \lambda^t$ in total where the N-part has length $\leq \lambda^t$ and the Q-part has length $\leq t$. If $\lambda = 1$, then every conjugation of an N-generator by a Q-generator yields an N-generator, then at most t applications of defining relators suffices to move all the N-generators past the Q-generators and we are left with a word still of length t.

We can consider the image of this word uv under the map g. Since the short exact sequence is split, and thanks to our choice of generating set, the map g is the projection of the Q-generators, so the image of this word uv under $g : G \to Q$ is precisely the subword v. Since g is a homomorphism, and $uv =_G 1$, we must have that v represents the identity in Q, and therefore in G. It follows that u must also represent the identity in G.

We can reduce both subwords u, v to the empty word in known time: u can be reduced with $\max\{\delta_N(\lambda^t), \delta_N(t)\}$ applications of defining relators (since if $\lambda > 1$, $\delta_N(\lambda^t)$ is larger, and otherwise $\delta_N(t)$ is larger - in both cases these are the Dehn function of N evaluated at the length of the subword u) and v can be reduced with $\delta_Q(t)$. It follows that we have an upper bound on the Dehn function of G in:

$$\delta_G(t) \preccurlyeq \max\left\{\delta_N(\lambda^t), \delta_N(t), \delta_Q(t)\right\}$$

We can try to sharpen this bound by reducing the exponential λ^t. In many cases, the length yielded from the application of a conjugation-by-Q-relator is less than the maximum λ. To account for this, we can define a *growth function*.

We need to account for the longest word in N that can be obtained by successive conjugations of a generator of N by different elements of Q. Let Q have a generating set $T = \{q_1, \dots, q_r\}$ and N have $S = \{n_1, \dots, n_k\}$. To find the worst case, the longest resulting word in N, we need to look at conjugation of each generator $s \in S$ of N by each possible sequence of length n of generators $q_i \in T$. This set of ordered n-tuples can be denoted T^n, and it follows that for the growth $g_Q : \mathbb{N} \to \mathbb{N}$ of an N-generator under conjugation by a word in Q, we have

$$g_Q(n) = \max_{x \in T^n} \left\{ \max_{s \in S} \left\{ \left| x^{-1} s x \right| \right\} \right\}$$

We can't immediately replace λ^t with $g_Q(t)$ in our upper bound. Previously, when we used the bound λ^t, we knew that the maximal length N-part would be obtained when the starting word had N-part of length 1, as with the case missed before, we need to be careful assuming this is the worst case now: what if at some n we have that $g_Q(n+1) < 2g_Q(n)$? Then there would be some length of word where the worst possible starting word would have 2 letters in N rather than 1. To account for such a case we cannot simply replace the λ^t upper bound with $g_Q(n)$ - we need to look at what is maximal of $g_Q(n), 2g_Q(n-1), 3g_Q(n-2)$, *etc.* To avoid so many computations, we can consider a simpler alternative

Consider any sequence $x \in T^{n-1}$. We can append any T-letter, say t_1, to the right hand side to get a sequence $\hat{x} \in T^n$. Consider the conjugation $\hat{x}^{-1} s \hat{x}$ for any $s \in S$. We have

$$\hat{x}^{-1} s \hat{x} = t_1^{-1} x^{-1} s x t_1$$

Since $|t_1^{-1} s t_1| \geq 1$ for all $s \in S$, we have that

$$\left| \hat{x}^{-1} s \hat{x} \right| = \left| t_1^{-1} x^{-1} s x t_1 \right| \geq \left| x^{-1} s x \right|$$

Taking x to be maximal in T^{n-1}, in the sense that the length of $x^{-1} s x$ is maximal, gives a way of generating a longer word from $\hat{x} \in T^n$, and therefore $g_Q(n) \geq g_Q(n-1)$. Repeating this process inductively, we have that $g_Q(n) \geq g_Q(k)$ for all $k \leq n$, and therefore that $g_Q(n) \geq g_Q(n+1-i)$. It follows immediately that $n g_Q(n) \geq i g_Q(n+1-i)$ for all $i \in \{1, \dots, n\}$ and so we can say that

$$\delta_G(t) \preccurlyeq \max \{ \delta_N(t g_Q(t)), \delta_Q(t) \}$$

Note that the original upper bound with λ^t is useful in cases where we are unable to compute the smaller growth functions g_Q, since $\lambda^t \simeq e^t$, and by lemma 4.2, we can interchange these. In particular, note that - by its construction - $g_Q(t) \leq \lambda^t$, and so $g_Q(t) \preccurlyeq e^t$.

4.2. **A useful presentation.** Given a short exact sequence

$$1 \to N \to G \to Q \to 1$$

and finite presentations $\langle Y | S \rangle$, $\langle \overline{X} | \overline{R} \rangle$ for N and Q respectively, we want to understand the Dehn function of G in terms of the Dehn functions of N and Q with respect to these presentations. In this case, we can think of N as a normal subgroup of G and we can think of Q interchangeably with the quotient groups G/N, that is, the group of cosets of N in G.

In order to even define the Dehn function, we must first have a finite presentation for the middle group G. In a post on math stack exchange [7], D. Holt gave the construction for such a presentation below.

Firstly, we want a choice of coset representatives in G that correspond to the generating set $\overline{X}$ of G/N, so for each $\overline{x} \in \overline{X}$ we choose an $x \in G$ such that $xN = \overline{x}$, and set $X := \left\{ x \mid \overline{x} \in \overline{X} \right\}$. I will call these generators the *lifts* of those in $\overline{X}$.

Observe that since the cosets partition G, every element $g \in G$ can be written as $g = u_g(X)v_g(Y)$ for some words u, v in X and Y respectively - it follows that X and Y generate G.

For any word $\overline{w}$ in $\overline{X}$, we can define a corresponding word w in X with $wN = \overline{w}$ by substituting each $\overline{x}$ or $\overline{x}^{-1}$ with x or x^{-1} respectively. In particular, for each relation $\overline{r} \in \overline{R}$ we have a word r, and since $\overline{r} =_{G/N} 1$, we know that $rN = N$, or $r \in N$. It follows that for every relation $\overline{r} \in \overline{R}$, there is a corresponding relation rw_r^{-1} for some word $w_r \in N$. Set $R \coloneqq \left\{ rw_r^{-1} \mid \overline{r} \in \overline{R} \right\}$. I will call these relators the *lifts* of those in $\overline{R}$.

Since N is a normal subgroup, there are conjugation relations for each $y \in Y, x \in X$ of the form $x^{-1}yxw_{xy}^{-1}$. Call the set of such relations T. These are the same as our *conjugation-by-Q-relators* that we used in the split case. It turns out that

$$\langle X \cup Y \mid R \cup S \cup T \rangle$$

Is a finite presentation for G. A proof can be found in [7].

When using this presentation, for a word w in G, I will call the subword consisting of the generators in X or Y the "X-part" or "Y-part" of the word respectively.

4.3. **A simple non-splitting example.** Before giving a general upper bound, it is insightful to first deal with a relatively simple example. In hopes that understanding this case will give techniques that can be generalised to any short exact sequence, We will consider the Heisenberg group

$$G \cong \left\langle a, b, c \;\middle|\; c = [a, b], [c, a], [c, b] \right\rangle$$

In relation to the construction above, this is obtained from the short exact sequence

$$1 \to \mathbb{Z} \to G \to \mathbb{Z}^2 \to 1$$

Where $\mathbb{Z} = \left\langle c \;\middle|\; \right\rangle$ and $\mathbb{Z}^2 = \left\langle \overline{a}, \overline{b} \;\middle|\; [\overline{a}, \overline{b}] \right\rangle$. This is a simple example of a *non-splitting* short exact sequence of groups.

It is known that the Dehn function of this G is $\delta_G(n) \simeq n^3$, so we expect to get an upper bound greater than or equal to this polynomial from the short exact sequence.

Initially, we will proceed similarly to the split short exact sequence case: given an arbitrary nullhomotopic word w in G of length $\leq n$, we first want to move all of the letters in the N subgroup to one side. Since $N = \langle c \rangle$ lies in the centre of G by its construction, this is easily done by applying at most n^2 commutator relations and does not increase the length of the word. After this is done, we have a word $w_1 =_G w$ of the same length as w with all copies of c at the left hand side.

Mapping w_1 into $G/N \cong \mathbb{Z}^2$, we see that since $w_1 =_G 1$, $w_1 \mapsto w_Q =_{G/N} 1$. All the quotient map here does is set $c = 1$, so it follows that the $\langle a, b \rangle$-part of w_1, w_Q, can be written as a word in the kernel N, *i.e.* as a power of c. The question is how many applications of defining relators does this take?

Consider a Van Kampen diagram in G/N. This has no more that $\delta_{G/N}(n)$ 2-cells. Consider the corresponding sequence of applications of defining relators in G/N to achieve this Van Kampen diagram. Consider the lift of the first applied relator: in G/N it will replace some part of $a^{-1}b^{-1}ab$ with the inverse of the other, but the lift in G will do so *and* introduce a copy of $c^{\pm 1}$. In order to apply the same sequence of defining relators in G as in G/N, we need to move all copies of c out of the way after each applied relator. Moving each copy of c takes no more than n^2 applications of defining relators and this whole process must be done up to $\delta_{G/N}(n)$ times.

After this, we are left with a word in N of length no more than $n + \delta_{G/N}(n)$, so this can be reduced to the identity in $\delta_N\left(n + \delta_{G/N}(n)\right)$-time.

Adding up all these, we have an upper bound of

$$n^2 + n^2 \times \delta_{G/N}(n) + \delta_{G/N}(n) + \delta_N\left(n + \delta_{G/N}(n)\right)$$

And if we plug in our known Dehn functions $\delta_N(n) \simeq n$ and $\delta_{G/N}(n) \simeq n^2$, we have a process of reducing an arbitrary nullhomotopic word in G of length up to n to the identity in no more than $n^2 + n^4 + n^2 + n + n^2 = n^4 + 3n^2 + n$ steps, so we can say with this construction that

$$\delta_G(n) \preccurlyeq n^4$$

Evidently, this is not an optimal upper bound (since δ_G is bounded above by n^3), but it is a valid one. Many of the steps in this construction can be generalised, but special care needs to be taken when rearranging the letters of a more general word, since the conjugation-by-Q-relators may not be as nice as the commutators they are in this example. To work with this, recall from "A simpler upper bound for the Dehn function of a split short exact sequence" the growth function g_Q. See section 6.1 for examples where this method yields an optimal upper bound.

4.4. **The general case.** The main difference between this Dehn function bound and the one for the split short exact sequence case is the Y-subwords that arise from lifting relations in Q (those in $\overline{R}$) to relations in G (those of the form $w_r r^{-1}$ for $\overline{r} \in \overline{R}$).

Proposition 4.1. *Let*

$$1 \to N \to G \to Q \to 1$$

be a short exact sequence of groups, where $N = \langle Y \mid S \rangle$, $Q = \langle \overline{X} \mid \overline{R} \rangle$ *and* $G = \langle X \cup Y \mid R \cup S \cup T \rangle$. *Then the Dehn function of G is bounded above by*

$$\delta_G(n) \preccurlyeq \delta_N\left((L_N + 1)\delta_Q(n) g_Q\left(\delta_Q(n) L_{\overline{R}}\right)\right)$$

Where L_N is the maximum length over the words w_r for relators $w_r r^{-1}$ in R and $L_{\overline{R}}$ is the maximum length over the relators $\overline{r} \in \overline{R}$.

Proof. We can begin similarly to the split case: let w be an arbitrary nullhomotopic word in G. The first step is to separate the Y-part and the X-part of the word - we can do this using our conjugation-by-Q-relators T. We can move all the X-letters past the Y-letters using no more than $n g_Q(n)$ commutators, yielding a new word w_1 with Y-part of length no more than $n g_Q(n)$ and X-part of length no more than n. In the figure below, w_1 is taken to be the inside loop once all the commutator relations have been applied.

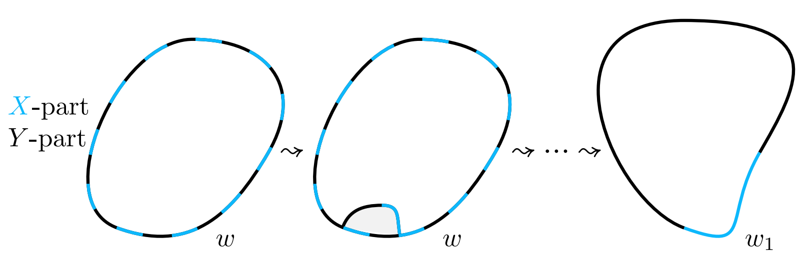

FIGURE 1. Rearranging w

Next, we know that since $w_1 =_G 1$, its image $\overline{w_Q}$ under the quotient map is equal to $1 \in Q$. The whole word w_1 lies in the kernel N of the quotient map, so there must be some way of rewriting the X-part w_Q of w_1 in N. Consider a Van Kampen diagram for w_Q in Q, with area no more than $\delta_Q(n)$, and consider the corresponding sequence of relations applied to obtain such a diagram. Each relation $\overline{r} \in \overline{R}$ applied in this sequence has a lift to a relation in $r \in R$, where the X-part of r is exactly $\overline{r}$ and there is an additional Y-part w_r^{-1} in the lifted relator.

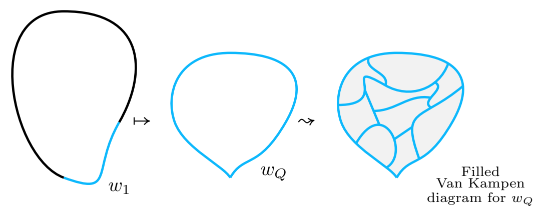

FIGURE 2. The image under the quotient map

We want to be able to lift this whole sequence into G, to rewrite the X-part w_Q of w_1 as a word in the letters of Y in $\delta_Q(n)$ steps. The appended Y-letters in the R-relators pose a problem: each time we lift an $\overline{R}$-relator to an R-relator, we need to account for the Y-subword before we can apply the next R-relator in the sequence. Since this Y-subword w_r^{-1} might be preventing us from adding the next 2-cell of the filling in the sequence - the Y-parts of the lifts of the boundaries of two adjacent cells in the filling of w_Q may not agree.

There are some useful constants to define to help here. Let $L_{\overline{R}}$ be the maximal length among relators in $\overline{R}$, and let L_N be the maximal length of the Y-part of the relators in R.

We want to move the any new Y-subwords created by applying these relations over to the side with the rest of the Y-part of w_1; we can do this using the conjugation-by-Q-relators in T. After each application of a relator in R, we can use the conjugation-by-Q-relators in T to create a corridor of T-relators with sides labeled by the same word in X, since the conjugation relators preserve the X-part of a word at the cost of multiplying the length of the Y-part. This allows us to then apply the next R-relator and proceed until the whole sequence of relators has been applied, leaving us with a closed loop w_N with boundary labelled only by Y.

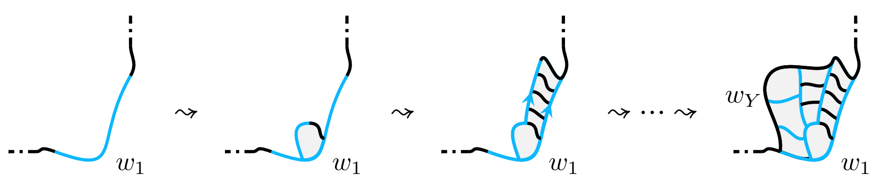

FIGURE 3. Translating the X-part into Y

To account for the area of each of these corridors, we consider a worst case scenario where the length of the corridor is the same as the length of w_Q, the X-part of w_1. One might have to move as many as $n + L_{\overline{R}}$ X-letters past L_N Y-letters, which would require as many as $L_N g_Q(n + L_{\overline{R}})$ applications of T-relators - this process also adds up to $L_N g_Q(n + L_{\overline{R}})$ Y-letters to the length of w_N. The following step might force moving $n + 2L_{\overline{R}}$ X-letters past L_N Y-letters, adding a further $L_N g_Q\left(n + 2L_{\overline{R}}\right)$ Y-letters to w_N and so on.

Taking a generous upper bound, we can say that each of the $\delta_Q(n)$ steps required to complete the translation process requires at most an additional $L_N g_Q\left(n+\delta_Q(n)L_{\overline{R}}\right)$ applications of T-relators, and can contribute as many Y-letters to w_N. Consider that after the k-th step, the remaining length of the Y-part can never exceed $(\delta_Q(n)-k)L_{\overline{R}}$, else it wouldn't be possible to reduce it to the empty word in time, and also at this step it can't have been made any longer than $n+kL_{\overline{R}}$. It follows that we can lower this bound by exchanging $n+\delta_Q(n)L_{\overline{R}}$ for $n+\frac{1}{2}\delta_Q(n)L_{\overline{R}}$, reducing the N-area of w_N.

Once this process is complete, we are left with a partially filled Van Kampen diagram, the remaining part to fill is w_N - a loop consisting only of Y-letters of length no more than $ng_Q(n)+\delta_Q(n)L_N g_Q\left(n+\frac{1}{2}\delta_Q(n)L_{\overline{R}}\right)$. Since this word is contained in the subgroup N, we can use the known Dehn function δ_N to tell us that we can fill this Van Kampen diagram with no more than $\delta_N\left(ng_Q(n)+\delta_Q(n)L_N g_Q\left(n+\frac{1}{2}\delta_Q(n)L_{\overline{R}}\right)\right)$ 2-cells. Note that this is an upper bound on a filling of w_N using only relations in S, which can be no better than a filling using all relations in $R\cup S\cup T$

So now we have fully reduced the word, we can add up the area to get a total upper bound

$$\begin{aligned}\delta_G(n) \leq\ & ng_Q(n)\\ &+\delta_Q(n)L_N g_Q\left(n+\frac{1}{2}\delta_Q(n)L_{\overline{R}}\right)\\ &+\delta_N\left(ng_Q(n)+\delta_Q(n)L_N g_Q\left(n+\frac{1}{2}\delta_Q(n)L_{\overline{R}}\right)\right)\end{aligned}$$

Since δ_N,δ_Q are at least linear, we can say

$$\delta_G(n)\preccurlyeq\delta_N\left((L_N+1)\delta_Q(n)g_Q\left(\delta_Q(n)L_{\overline{R}}\right)\right)$$

Note that I am implicitly assuming $L_{\overline{R}}\geq 2$ for this nicer-looking bound, but if $L_{\overline{R}}<2$, Q must be a free group: since all its relators are either $y=1$ for some generator, or Q has no relators; and so any nullhomotopic word w would have its X-part freely reduce once reshuffled to w_1 as above, in which case we would have $\delta_G(n)\leq ng_Q(n)+\delta_N(ng_Q(n))\preccurlyeq\delta_N(ng_Q(n))$. □

We often see the $\preccurlyeq$ relation ignore constants, such as our L_N and $L_{\overline{R}}$, but it is not immediately clear that we can omit these in the statement, since we have not proved any results about its behaviour under composition. With the following lemma, we will be able to simplify our upper bound, while being forced to include the condition that δ_Q is *superadditive.*

Remark. While it is not proven that every Dehn function is equivalent to a superadditive function, there are no known counterexamples. Particularly in our use case, the Dehn functions we deal with are all quadratic, exponential or various compositions of such functions - all of which are superadditive.

Lemma 4.2. *Let $a,b,a',b':\mathbb{N}\to\mathbb{N}$ be functions that are strictly increasing, at least linear and that b and b' are superadditive. Then*

$$a\circ b\simeq a'\circ b'$$

Proof. If $a\simeq a'$, we know that $a(n)\leq Ca'(Cn+C)+Cn+C$ for some constant $C\geq 1$, likewise $b(n)\leq Db'(Dn+D)+Dn+D$ for some $D\geq 1$. We need a bound

$a \circ b \leq a' \circ Eb'(En+E) + En + E$ for a constant $E \geq 1$.

$$
\begin{aligned}
a \circ b(n) &\leq a(Db'(Dn+D) + Dn + D) \\
&\leq Ca'(CDb'(Dn+D) + CDn + CD + C) + Cn + C \\
&\leq Ca'(b'(CD^2n + CD^2) + CDn + CD + C) + Cn + C \qquad \text{(superadditivity)} \\
&\leq Ca'\left(b'(CD^2n + CD^2 + CDn + CD + C)\right) + Cn + C \\
&\leq Ca' \circ b'\left((CD^2 + CD)n + (CD^2 + CD + C)\right) + Cn + C \\
&\leq Ea' \circ b'\left(En + E\right) + En + E
\end{aligned}
$$

Where $E \coloneqq CD^2 + CD + C$. The reverse inequality follows the same way, so $a \circ b \simeq a' \circ b'$. $\square$

Since Dehn functions δ_N, δ_Q are increasing and at least linear, we can say that if $\delta_Q : \mathbb{N} \to \mathbb{N}$ is superadditive, we can ignore the constants. If we take the upper bound $g_Q(n) \preccurlyeq e^n$, then we can use that $f(n)e^{f(n)} \preccurlyeq e^{f(n)}$ to obtain the following result.

Theorem 4.3. *Let*

$$1 \to N \to G \to Q \to 1$$

be a short exact sequence of groups. Suppose that the Dehn functions of N and Q are superadditive. Then the Dehn function of G is bounded above by

$$\delta_G(n) \preccurlyeq \delta_N\left(e^{\delta_Q(n)}\right)$$

5. Bounding the Dehn functions of RORGS

Applying Theorem 4.3 can be simple: we already know how to obtain a decomposition tree from a RORG thanks to [4], as outlined in 2.3, and we can use the known Dehn functions of the leaves to repeatedly apply our new bound to obtain a bound on the Dehn function of the group at the root of the tree.

Since it is not known whether the choices made in constructing a decomposition tree might result in different quotients, or quotients in a different order, to find the optimal upper bound by this method, one would have to check every possible decomposition tree and take the one that yields a minimal bound on the Dehn function. While such a computation is necessary to find the optimal bound, we can achieve a worse but simpler bound if we can bound the number of leaves of the decomposition tree (or equivalently the number of terms in the subnormal series).

Observe that the depth of a decomposition tree for $Out^0\left(A_\Gamma; \mathcal{G}, \mathcal{H}^t\right)$ can never be greater than the number of vertices plus one, since for each inner vertex, its left descendant acts trivially on at least one more vertex of the graph and its right descendant acts on a subgraph which has at least one fewer vertex than their root. Once there are no non-trivial restrictions, which must occur by the $|V(\Gamma)|$-th level, the vertex is either a leaf, or a projection homomorphism can be applied to obtain two descendants which are leaves. This upper bound on the depth of the tree as $|V(\Gamma)| + 1$ gives an upper bound on the number of leaves as $2^{|V(\Gamma)|+1}$. This naïve upper bound can be improved by having some greater understanding of the sets in $\mathcal{G} - \mathcal{H}$ and the order of decomposition.

Given a saturated pair $(\mathcal{G}, \mathcal{H})$, since every restriction homomorphism is applied to a subgraph defining a special subgroup in $\mathcal{G} - \mathcal{H}$, the depth of the decomposition tree can be similarly bounded by $|\mathcal{G}-\mathcal{H}|+1$. Observe that at each level of the tree, this number is reduced by at least one: for the left descendant, one special subgroup in $\mathcal{G}$ is added to $\mathcal{H}$, reducing $|\mathcal{G}-\mathcal{H}|$; for the right descendant, $\mathcal{G}_\Delta$ contains at least one fewer special subgroup than $\mathcal{G}$ (since A_Δ is removed), and every special subgroup removed from $\mathcal{H}$ when restricted to $\mathcal{H}_\Delta$ is also removed from $\mathcal{G}$ under the restriction, so $|\mathcal{G}_\Delta-\mathcal{H}_\Delta|<|\mathcal{G}-\mathcal{H}|$.

5.1. **Bounding the number of leaves on the decomposition tree for the outer automorphism group of a RAAG.** If we consider the root of the tree to be $Out^0(A_\Gamma)$ instead of a general RORG, we can describe an explicit decomposition.

Note first that the set $\mathcal{H}$ is empty for $Out^0(A_\Gamma)$, since every vertex is acted upon non-trivially by its respective inversion homomorphism, so we only need to understand the saturated $\mathcal{G}$. Also note that $Out^0(A_\Gamma)$ is defined with $\mathcal{G}$ and $\mathcal{H}$ empty, so the definitions of $\leq_{(\mathcal{G},\mathcal{H})}$, $\mathcal{G}$-connectedness and $(\mathcal{G},\mathcal{H})$-star-separation coincide with $\leq$, usual graph connectedness and usual star-separation respectively.

Recall that a *clique* is a complete subgraph, meaning that each of its vertices are adjacent to every other. Proposition 2.1 tells us that a subgraph defines a special subgroup in $\mathcal{G}$ if and only if it is upwards-closed under $\leq_{(\mathcal{G},\mathcal{H})}$ and not $(\mathcal{G},\mathcal{H})$-star-separated by an external vertex.

Cliques almost achieve these properties, since they cannot be star-separated by an external vertex, and there is only one case where they might not be upwards-closed. Suppose v is a vertex in a k-clique Δ, and there is some other vertex w, not in Δ, with $v\leq w$. Then we must have that $lk(v)\subseteq st(w)$, in particular, w is adjacent to all $k-1$ vertices in $\Delta-v$. If w is adjacent to v as well, then this was a $(k+1)$-clique all along - so we will suppose this is not the case, therefore $(\Delta-v)*w$ is a k-clique in Γ as well, so we see that if a k-clique Δ is not upwards-closed, then there is some other k-clique that intersects Δ at $k-1$ of its vertices. To ensure upwards-closure, I will define an *amalgamated clique*.

Definition 5.1.1. An amalgamated k-clique is the union of a maximal set of k-cliques, which have pairwise intersection a $(k-1)$-clique.

These amalgamated cliques satisfy some nice properties, which ensure that they define special subgroups in $\mathcal{G}$.

Lemma 5.1. *Amalgamated cliques are upwards closed*

Proof. Suppose that an amalgamated k-clique $\hat{\Delta}$ is not upwards closed, then there is a vertex $v\in\hat{\Delta}$ and a vertex $w\in\Gamma-\hat{\Delta}$ with $v\leq w$. Then w must be adjacent to the other $k-1$ vertices in a clique that v is contained in, so the rest of that clique and w form a k-clique, and by the definition of amalgamated clique, w was in $\hat{\Delta}$ all along. □

Lemma 5.2. *Amalgamated cliques are not star-separable*

Proof. Suppose that a vertex $u\in\Gamma-\hat{\Delta}$ star-separates an amalgamated k-clique $\hat{\Delta}$, then $\hat{\Delta}-st(u)$ is non-empty and not connected. Let v, w be any two vertices in $\hat{\Delta}-st(u)$. If v and w are in a k-clique, then they are adjacent. Suppose that v, w are in 2 distinct k-cliques Δ_v,Δ_w, then $\Delta_v\cap\Delta_w$ is a $(k-1)$-clique consisting of vertices adjacent to both v and w, so must be contained in the star of u, however, then $u\cup(\Delta_v\cap\Delta_w)$ is a k-clique, and intersects Δ_v at $k-1$ vertices, so u must have been in $\hat{\Delta}$ all along. □

It follows from these two results, and Proposition 2.1 that the set of amalgamated cliques is a subset of $\mathcal{G}$ for the group $Out^0(A_\Gamma)$. Now it is with these special subgroups that I would like to describe a specific process for obtaining a decomposition tree.

Firstly, one has to split up Γ into amalgamated cliques in some way, we will construct a set of amalgamated cliques $\mathcal{C}$. Suppose the maximal size of a clique in Γ is d, then find the set of all distinct amalgamated d-cliques in Γ, add these to $\mathcal{C}$. Note that some of the k-cliques in an amalgamated k-clique are allowed to be subcliques of larger cliques, it is only required that there is at least one k-clique in each amalgamated clique that is not strictly a subclique of another. After this, look for all the $(d-1)$-cliques that are not subcliques of any larger clique, and add all the distinct amalgamated $(d-1)$-cliques that can be obtained from these to $\mathcal{C}$. Repeat this process with cliques of all sizes. Since every vertex is contained in a clique of size at least 1, every vertex is contained in some amalgamated clique. Thus the set of subgraphs $\mathcal{C}$ covers Γ. Let m be the size of the set $\mathcal{C}$ and let $\hat{d}$ be the size of the largest amalgamated clique.

Since every subgraph in $\mathcal{C}$ defines a special subgroup in $\mathcal{G}$, one can apply the restriction homomorphism to these groups. If one first applies the restriction homomorphism to the d-amalgamated clique $\hat{\Delta}$, we see that the right descendant is a relative outer automorphism group of $A_{\hat{\Delta}}$, meaning that this branch of the tree has depth no more than $|\hat{\Delta}|+1$ by our naïve bound. If one repeats this process with all subgraphs in $\mathcal{C}$, we will have no more than $m+1$ left descendants from the root, the first m of which have a corresponding right descendant a branch with depth no more than the size of the corresponding amalgamated clique plus one.

Note that at each stage, the restriction homomorphism is non-trivial, since by its construction the set $\mathcal{C}$ does not allow for any k-amalgamated clique to be a subgraph of an ℓ-amalgamated clique for $k \leq \ell$, and so each $A_{\hat{\Delta}}$ is not acted upon trivially by any vertex of the decomposition tree until the left descendant of the short exact sequence for $R_{\hat{\Delta}}$ is reached. However, if there is an amalgamated clique $\hat{\Delta}$ with $\hat{\Delta} = \Gamma$, no restriction homomorphism can be applied. If this is the case, then one must apply the naïve upper bound as above, which will coincide with the bound below if there is only one amalgamated clique.

Combining these results gives a bound on the number of leaves of a decomposition tree for $Out^0(A_\Gamma)$ as

$$k = \#\text{leaves} \leq \sum_{\hat{\Delta}\in\mathcal{C}} 2^{|\hat{\Delta}|+1} \leq m2^{\hat{d}+1}$$

With this bound, we can also obtain an upper limit on the height of the exponential tree for $\delta_{Out^0(A_\Gamma)}$, since [4] tells us each leaf corresponds to a quotient in the decomposition series. We have the short exact sequences

$$1 \to N_i \to N_{i+1} \to N_{i+1}/N_i \to 1$$

which tells us

$$\delta_{N_{i+1}}(n) \preccurlyeq \delta_{N_i}\left(e^{\delta_{N+i+1/N_i(n)}}\right) \preccurlyeq \delta_{N_i}\left(e^{e^n}\right)$$

One can extrapolate this result, checking that $\delta_{N_1} \preccurlyeq e^{e^n}$, to the whole sequence to see that for a decomposition series of length k, the height of the exponential tower is $2k$, or that

$$\delta_{Out^0(A_\Gamma)}(n) \preccurlyeq \left. e^{e^{\cdot^{\cdot^{\cdot^{e^n}}}}} \right\} \leq \sum_{\hat{\Delta}\in\mathcal{C}} 2^{|\hat{\Delta}|+2} \leq m 2^{\hat{d}+2}$$

For a general RORG $Out^0\left(A_\Gamma;\mathcal{G},\mathcal{H}^t\right)$, one can take the intersection of $\mathcal{C}$ with $\mathcal{G}-\mathcal{H}$ and apply the same algorithm for generating the decomposition tree, to see that $\delta_{Out^0(A_\Gamma;\mathcal{G},\mathcal{H}^t)}$ satisfies the same upper bound.

Note that it has not been shown which of these three upper bounds is the smallest, and that it is most likely for and arbitrary RORG that none of the three give optimal decomposition trees - we have not considered that some of the leaves may have quadratic Dehn function, or that the depth of each branch will usually be less than $2^{|\hat{\Delta}|}+1$. However, this bound is much simpler and quicker to compute than taking the minimum obtained across all decomposition trees. Also observe that the number of amalgamated cliques in $\mathcal{C}$ is bounded above by the number of cliques in a decomposition of Γ into cliques of maximal size, but each amalgamated clique may be larger than these cliques.

6. Examples and computations

6.1. **Examples of an optimal short exact sequence upper bound.** Below I will discuss some cases where this bound *is* optimal.

Consider the short exact sequence

$$1 \to F_m \to F_m \times F_n \to F_n \to 1$$

where F_i is the free group of rank i and the maps are the inclusion and projection of each part of the direct product. It is known that $\delta_{G\times H} \simeq \max\left\{\delta_G, \delta_H, n^2\right\}$; in this case, since the Dehn function of a free group is linear, we can say $\delta_{F_m\times F_n} \simeq n^2$. To apply our new bound, we need to understand the growth function g_{F_n}, but since we have a direct product, this is simple: the generators of F_m commute with the generators of F_n, so $g_{F_n}(n) = n$.

Therefore our bound reads $\delta_{F_m\times F_n} \preccurlyeq \delta_{F_m}\left(\delta_{F_n}(n)^2\right) \simeq \delta_{F_m}(n^2) \simeq n^2$, the optimal quadratic upper bound.

For a more interesting example, consider the group

$$G = \left\langle a_1,\dots,a_n,s,t \;\middle|\; s^{-1}a_i s = \phi(a_i), t^{-1}a_i t = \phi(a_i)\right\rangle$$

Where ϕ is some automorphism of F_n, the subgroup generated by the a_i. This group G satisfies the short exact sequence

$$1 \to F_n = \langle a_1,\dots,a_n\rangle \to G \to F_2 = \langle s,t\rangle \to 1$$

as well as the Dehn function $\delta_G(n) = g_\phi(n)$ where g_ϕ the growth of ϕ is defined by

$$g_\phi(n) = \max_{x\in\{a_1\dots a_n\}} \{|\phi^n(x)|\}$$

Observe that since both generators of F_2 yield the automorphism ϕ when conjugating generators of F_n, this growth function g_ϕ coincides with the growth function g_{F_2} in this example.

Applying our upper bound, one can see

$$\begin{aligned}\delta_G &\preccurlyeq \delta_{F_n}\left(\delta_{F_2}(n) g_{F_2}\left(\delta_{F_2}(n)\right)\right) \\ &\simeq \delta_{F_n}\left(n g_{F_2}(n)\right) \\ &\simeq n g_\phi(n)\end{aligned}$$

If g_ϕ is polynomial $g_\phi \simeq n^d$, then we have achieved $n^d \preccurlyeq \delta_G(n) \preccurlyeq n^{d+1}$ - so our upper bound is nearly optimal. If g_ϕ grows faster than polynomial, *i.e.* exponentially, we have that $g_\phi(n) \preccurlyeq \delta_G(n) \preccurlyeq n g_\phi(n) \preccurlyeq g_\phi(n)$ - so our upper bound is again optimal.

6.2. **Computing an upper bound of the Dehn function for a simple RAAG.** Consider the graph Γ:

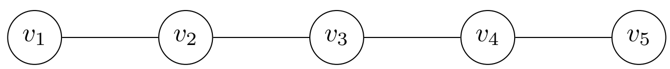

We want to compute the decomposition tree for $Out^0(A_\Gamma)$, following with Proposition 2.4, we need to first extend $\mathcal{G}$ to its saturation.

One can check that the ordering relations in Γ are $v_1 \leq v_2$, $v_1, \leq v_3$, $v_5 \leq v_3$, and $v_5, \leq v_4$. See that if the subgraph defining a special subgroup in $\mathcal{G}$ contains v_1, it must also contain v_2 and v_3 to be upwards-closed, and the same holds for v_5, v_4 and v_3 respectively.

For $Out^0(A_\Gamma)$, the definitions of $\leq_{(\mathcal{G},\mathcal{H})}$ and $(\mathcal{G},\mathcal{H})$ coincide with their standard, non-$(\mathcal{G},\mathcal{H})$, definitions. Observe that $\mathcal{H}$ is empty since each vertex is acted upon non-trivially by their respective inversion automorphism. The third condition of Proposition 2.8 from [4] can never be satisfied. It follows from these properties that $\mathcal{G}$ contains all special subgroups defined by graphs with only vertices in the set $\{v_2, v_3, v_4\}$ (graphs defined by the power set $\mathcal{P}(\{v_2, v_3, v_4\})$), since any combination of these elements will satisfy $v \leq_{(\mathcal{G},\mathcal{H})} w \Rightarrow v = w$ as all vertices in this set are incomparable. Such sets are also upwards-closed, not star-separated by any outside vertex in Γ and $Out^0(A_\Gamma)$ acts trivially on each of these three vertices, with the exception of the inversion automorphisms, so $\{v_i\} \in \mathcal{G}$ for $i \in \{2,3,4\}$ - these subsets satisfy both conditions to be in $\mathcal{G}$.

Observe that the special subgroup defined by the subgraph $\Delta =$

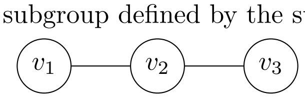

is upwards-closed and not $(\mathcal{G},\mathcal{H})$-star-separated by either v_4 or v_5, since $\Delta - st(v_5)$ and $\Delta - st(v_4)$ are both connected (and thus $(\mathcal{G},\mathcal{H})$-connected) graphs. So we can see that $A_\Delta \in \mathcal{G}$.

It then follows that $\mathcal{G} = \mathcal{P}(\{v_1, v_2, v_3\}) \cup \{A_\Delta, A_\chi, A_{\Delta'}, A_{\chi'}\}$ where $\chi = \{v_3, v_4, v_5\}$, $\Delta' = \{v_1, v_2, v_3, v_4\}$,$\chi' = \{v_2, v_3, v_4, v_5\}$, since the four special subgroups specified are defined by all the upwards-closed proper subgraphs containing either v_1 or v_5, one can easily check that $Out^0(A_\Gamma)$ preserves each of these special subgroups.

Furthermore, since the transvection $v_1 \mapsto v_1 v_3$ acts on Δ non-trivially, $\Delta \notin \mathcal{H}$, we apply the *restriction homomorphism* R_Δ to $Out^0(A_\Gamma)$ to find the first 2 descendants in our decomposition tree.

Applying the short exact sequence for R_Δ, we obtain the descendants A and B:

$$1 \to \underbrace{Out^0(A_\Gamma; \mathcal{G}, \{A_\Delta\}^t)}_{=A} \to \underbrace{Out^0(A_\Gamma)}_{=G} \to \underbrace{Out^0(A_\Delta; \mathcal{G}_\Delta, \mathcal{H}_\Delta)}_{=B} \to 1$$

First we consider the right descendant B. Recall the definition of $\mathcal{G}_\Delta$ and $\mathcal{H}_\Delta$: $\mathcal{J}_\Delta = \{A_{\Delta\cap\Theta} | A_\Theta \in \mathcal{J}\}$. Observe that $\mathcal{H}_\Delta = \varnothing$ and $\mathcal{G}_\Delta = \mathcal{P}(\{v_2, v_3\})$. We can thus apply the restriction homomorphism R_Λ for $\Lambda = \{v_2, v_3\}$:

$$1 \to \underbrace{Out^0(A_\Delta; \mathcal{G}_\Delta, \{A_\Lambda\}^t)}_{=C} \to \underbrace{Out^0(A_\Delta; \mathcal{G}_\Delta, (\mathcal{H}_\Delta)^t)}_{=B} \to \underbrace{Out^0(A_\Lambda; (\mathcal{G}_\Delta)_\Lambda, (\mathcal{H}_\Delta)_\Lambda)}_{=D} \to 1$$

Again we will consider the right descendant D. $(\mathcal{G}_\Delta)_\Lambda = \{\{v_2\}, \{v_3\}\}$, $(\mathcal{H}_\Delta)_\Lambda = \varnothing$, so we can apply the restriction homomorphism R_{v_2}:

$$1 \to \underbrace{Out^0(A_\Lambda; (\mathcal{G}_\Delta)_\Lambda, \{A_{v_2}\}^t)}_{=F} \to \underbrace{Out^0(A_\Lambda; (\mathcal{G}_\Delta)_\Lambda, (\mathcal{H}_\Delta)_\Lambda)}_{=D} \to \underbrace{Out^0(A_{v_2}; ((\mathcal{G}_\Delta)_\Lambda)_{v_2}, (((\mathcal{H}_\Delta)_\Lambda)_{v_2})^t)}_{=H} \to 1$$

Looking at the right descendant H, this group's only generator is the inversion of v_2, this group is isomorphic to $GL(1, \mathbb{Z}) \cong \mathbb{Z}/2\mathbb{Z}$ is finite and therefore has linear Dehn function.

Considering the left descendant F, the only generator of this group is the inversion of v_3 and so as above this group is $GL(1, \mathbb{Z})$ satisfies a linear isoperimetric inequality.

At the previous step, looking at the left descendant C, there are no non-trivial restriction maps, since the upwards closure of v_1 is all of Δ, and the other two vertices are acted upon trivially. Δ is a complete graph with proper non-trivial centre generated by v_2, so we apply the projection homomorphism on $\rho \coloneqq \{v_1, v_3\}$:

$$1 \to \underbrace{K_{P_\rho}}_{=I} \to \underbrace{Out^0(A_\Delta; \mathcal{G}_\Delta, \{A_\Lambda\}^t)}_{=C} \to \underbrace{Out^0(A_\rho; (\mathcal{G}_\Delta)_\rho^t)}_{=J} \to 1$$

The left descendant I is generated by leaf transvections in Δ, and therefore is $\mathbb{Z}$, generated by the transvection $v_1 \mapsto v_1 v_2$. The right descendant J is generated by both the transvection $v_1 \mapsto v_1 v_3$ and the two inversions. $(\mathcal{G}_\Delta)_\rho = \{v_3\}$ and J allows no non-trivial restrictions. It is disconnected and $(\mathcal{G}_\Delta)_\rho$-disconnected, meaning that it is the Fouxe-Rabinovitch group $Out^0(A_\rho; \{v_3\}^t)$.

Looking now at the left descendant of the inital step A, we can apply the restriction homomorphism R_χ to get 2 new descendants K and L:

$$1 \to \underbrace{Out^0 (A_\Gamma; \mathcal{G}, \{A_\Delta, A_\chi\})^t}_{=K} \to Out^0 \left(A_\Gamma; \mathcal{G}, \{A_\Delta\}^t\right) \to \underbrace{Out^0 \left(A_\chi; \mathcal{G}_\chi, (\{A_\Delta\})_\chi^t\right)}_{=L} \to 1$$

Since K acts trivially on A_Δ, A_χ, it must trivially on all of Γ, so is the trivial group. Exactness then tells us that A and L are isomorphic.

For L, see that $(\{A_\Delta\})_\chi = \{\{v_3\}\}$, and that $\mathcal{G}_\chi = \mathcal{P}(\{v_3, v_4\})$. We can apply the restriction homomorphism R_{v_4}:

$$1 \to \underbrace{Out^0 \left(A_\chi; \mathcal{G}_\chi, (\{A_\Delta\}_\chi \cup \{A_{v_4}\})^t\right)}_{=M} \to \underbrace{Out^0 \left(A_\chi; \mathcal{G}_\chi, (\{A_\Delta\})_\chi^t\right)}_{=L} \to \underbrace{Out^0 \left(A_{v_4}; (\mathcal{G}_\chi)_{v_4}, (\{A_\Delta\}_\chi)_{v_4}^t\right)}_{=N} \to 1$$

Since M acts trivially on A_Δ and A_{v_4}, it must act trivially on every vertex but v_5, and so we can apply the restriction homomorphism R_{v_5}:

$$\begin{aligned} 1 &\to \underbrace{Out^0\left(A_\chi; \mathcal{G}_\chi, \left(\{A_\Delta\}_\chi \cup \{A_{v_4}\} \cup \{A_{v_5}\}\right)^t\right)}_{=O} \\ &\to \underbrace{Out^0\left(A_\chi; \mathcal{G}_\chi, \left(\{A_\Delta\}_\chi \cup \{A_{v_4}\}\right)^t\right)}_{=M} \\ &\to \underbrace{Out^0\left(A_{v_5}; ((\mathcal{G}_\chi)_{v_4})_{v_5}, \left((\{A_\Delta\}_\chi)_{v_4}\right)^t_{v_5}\right)}_{=P} \to 1 \end{aligned}$$

Observe that O acts trivially on all of χ and so is the trivial group. See that the only generator of P is the inversion of v_5, so $P \cong GL(1,\mathbb{Z})$. Observe that N acts trivially on every vertex, so is the trivial group.

These computations yield a decomposition tree for $Out^0(A_\Gamma)$:

$$\begin{tikzcd}
 & & & G = Out^0(A_\Gamma) & & \\
 & A = Out^0\left(A_\Gamma; \mathcal{G}, (\mathcal{H} \cup \{A_\Delta\})^t\right) \arrow[urr] \arrow[dd] & & & B = Out^0(A_\Delta; \mathcal{G}_\Delta) \arrow[dd] & \\
K = 1 \arrow[ur] & & & C = Out^0(A_\Delta; \mathcal{G}_\Delta, \{A_\Lambda\}^t) \arrow[ur] \arrow[ddd] & & \\
 & L = Out^0\left(A_\chi; \mathcal{G}_\chi, (\{A_\Delta\})^t_\chi\right) \arrow[dr] & I = \mathbb{Z} \arrow[ur] & & D = Out^0(A_\Lambda; (\mathcal{G}_\Delta)_\Lambda, (\mathcal{H}_\Delta)_\Lambda) \arrow[dr] & \\
 & M \arrow[u] \arrow[d] & N = 1 & & F = GL(1,\mathbb{Z}) \arrow[u] & H = GL(1,\mathbb{Z}) \\
O = 1 \arrow[ur] & P = GL(1,\mathbb{Z}) & & J = \text{Fouxe-Rabinovitch} & &
\end{tikzcd}$$

With this tree, we can extract short exact sequences $1 \to X \to Y \to Z \to 1$, where Y is an internal vertex and X and Z are its two descendants. Using such sequences and our bound from Section 4, we can bound the Dehn function out $Out^0(A_\Gamma)$ by

$$\delta_{Out^0(A_\Gamma)} \preccurlyeq e^{e^{e^{e^{e^{e^{e^{e^{n}}}}}}}},$$

where the height of the exponential tower is 8. If we compare this to the bound given by studying the amalgamated cliques, one can observe that Γ contains three amalgamated 2-cliques, each of size three, which yields an exponential tower with height $3 \times 2^3 = 24$, evidently a far weaker bound. In fact, we can do much better using this decomposition tree - we can make many deductions rather than naïvely applying . In each short exact sequence, if a group X's two descendants are both finite, then Lagrange's theorem tells us that X must also be finite, and satisfy a linear isoperimetric inequality. Applying this, we see that M, L and A must all be finite on the left branch, and that D must be finite on the right branch, so C is a finite-index subgroup of B, meaning $\delta_B \simeq \delta_C$ and we can say

$$\delta_{Out^0(A_\Gamma)} \preccurlyeq e^{e^{e^{n}}}.$$

References

[1] R. Charney and K. Vogtmann. Finiteness properties of automorphism groups of right-angled Artin groups. *Bulletin of the London Mathematical Society*, 41(1):94–102, 2009.

[2] R. Charney and K. Vogtmann. Subgroups and quotient groups of automorphism groups of RAAGs, 2010.

[3] M. B. Day. Peak reduction and finite presentations for automorphism groups of right-angled Artin groups. *Geometry and Topology*, 13(2):817–855, Jan. 2009.

[4] M. B. Day, A. W. Sale, and R. D. Wade. Calculating the virtual cohomological dimension of the automorphism group of a RAAG. *Bulletin of the London Mathematical Society*, 53(1):259–273, 2021.

[5] M. B. Day and R. D. Wade. Relative automorphism groups of right-angled Artin groups. *Journal of Topology*, 12(3):759–798, 2019.

[6] A. Hatcher and K. Vogtmann. Isoperimetric inequalities for automorphism groups of free groups. *Pacific Journal of Mathematics*, 173(2):425 – 441, 1996.

[7] D. Holt. Finitely Presented is Preserved by Extension. Mathematics Stack Exchange, 2014. https://math.stackexchange.com/q/813692 (version: 2014-05-29).

[8] M. R. Laurence. A Generating Set for the Automorphism Group of a Graph Group. *Journal of the London Mathematical Society*, 52(2):318–334, 10 1995.

[9] A. Lowenstein. Deformation spines satisfy an exponential isoperimetric inequality, 2024.