6-18-2019

Dehn Functions of Bestvina-Brady Groups

Yu-Chan Chang
Louisiana State University and Agricultural and Mechanical College, yuchanchang74321@gmail.com

DEHN FUNCTIONS OF BESTVINA–BRADY GROUPS

A Dissertation

Submitted to the Graduate Faculty of the
Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

in

The Department of Mathematics

by
Yu-Chan Chang
B.S. in Math., National Central University, 2008
M.S., National Central University, 2010
M.S., Louisiana State University, 2015
August 2019

To my grandparents

Acknowledgments

I want to express my sincere gratitude and appreciation to my advisor, Pallavi Dani, for her time and encouragement. This dissertation would not have been possible without her infinite patience and guidance. I thank her for always being positive and supportive to me and for pushing me to be better.

I would like to thank Pramod Achar, Daniel Cohen, Patrick Gilmer, Daniel Sheehy, Stephen Shipman, and Martin Tzanov, for serving my general exam committee and/or my dissertation committee. I would like to thank Daniel Cohen for teaching me hyperplane arrangement. I would like to thank Patrick Gilmer for being my mentor for the Communicating Mathematics project. I would like to thank Shea Vela-Vick for teaching me many topics on geometry and topology, organizing the informal topology seminar, and giving me useful advice. I would like to thank Bogdan Oporowski for many discussions on graph theory. I would like to thank James Oxley for helpful conversations about teaching.

I would like to thank Max Forester at the University of Oklahoma, for his many helpful discussions. His ideas motivate part of this dissertation. Part of this dissertation was done during my visit to Cornell University in the spring semester in 2017; I would like to thank for their hospitality. This work was supported by NSF Grant No. DMS-1812061.

I would like to thank Hui-Hsiung Kuo, Ling Long, Richard Ng, and Fang-Ting Tu, for their help and support in many aspects. I would like to thank Bogdan Oporowski for everything during my six years at LSU. I would like to thank everyone in the Baton Rouge Table Tennis Club for making my life colorful. I would like to thank everyone whom I meet at LSU.

I would like to thank Yen-Mei Chen, Meng-Kai Hong, Mei-Lin Yau at National Central University (NCU) and Hung-Lin Chiu at National Tsing Hua University. I

could not possibly have had a chance to pursue my doctoral degree at LSU without their help. I would also like to thank Cheng-Hsiung Hsu and Hwa-Long Gau at NCU, for their guidance when I was an undergraduate student at NCU.

I would like to thank my friends Li-Ling Su, Yi-Jun Wang, Yu-Hao Li, and Chih-Jen Li, for their love and support.

Last but not least, I would like to thank my mother, father, and brother, for their endless love and support. I could not have achieved anything without them. I love you.

This dissertation is dedicated to my grandparents, who passed away in 2016 and 2017. I want to express my deepest gratitude to them for being part of my life for more than 30 years. I will always miss you.

Table of Contents

Abstract

In this dissertation, we prove that if the flag complex on a finite simplicial graph is a 2-dimensional triangulated disk, then the Dehn function of the associated Bestvina–Brady group depends on the maximal dimension of the simplices in the interior of the flag complex. We also give some examples where the flag complex on a finite simplicial graph is not 2-dimensional, and we establish a lower bound for the Dehn function of the associated Bestvina–Brady group.

Chapter 1. Introduction

The Dehn function of a group is analogous to the isoperimetric function in Riemannian geometry. One of the classical problems in Riemannian geometry is the isoperimetric problem, which may be stated as follows: given a closed curve C of a fixed length in a Riemannian manifold, what is the largest possible area of a region that has C as its boundary? An isoperimetric function describes a relationship between the length of a given closed curve and the area it encloses. In group theory, a word w that represents the identity in a finitely presented group G can be seen as a closed curve in the universal cover of the presentation complex of G. We can define the area of any region that has w as its boundary, denoted by $\text{Area}_G(w)$. The optimal function that bounds the area in terms of the length of w is called the Dehn function of G, and is denoted by δ_G. We refer to Chapter 2 for formal definitions.

In this paragraph, we explain some terminologies that will be used in the rest of the introduction. We refer to Definition 2.5 for more precise statements. We say that a function f is at most polynomial (linear, quadratic, cubic, quartic), exponential, or super-exponential, if f is bounded above by a such function. A function f is larger than a function g if g is bounded above by g. We say that a function f is polynomial (linear, quadratic, cubic, quartic), exponential, or super-exponential, if f is equivalent to a such function. We say that a group G has polynomial (linear, quadratic, cubic, quartic), exponential, or super-exponential Dehn function if δ_G is equivalent to a such function.

Dehn functions of finitely presented groups also measure, to a certain extent, how complicated the groups are. The word problem for a finitely presented group

G is an algorithmic problem to decide whether a given word in the generators of G represents the identity in G. The Dehn function δ_G provides an upper bound on the complexity of the word problem for G. Whether or not the word problem for a finitely presented group is solvable can be answered by studying the Dehn function of the given group: a finitely presented group G has a solvable word problem if and only if its Dehn function δ_G is recursive for every finite presentation of G.

Besides the solvability of the word problem, Dehn functions can also detect certain structures in groups. For example, a group is hyperbolic if and only if it has a linear Dehn function. On the other hand, Dehn functions cannot distinguish between many types of groups. For instance, Dehn functions of automatic groups, CAT(0) groups, mapping class groups, and right-angled Artin groups are bounded above by a quadratic function. In fact, right-angled Artin groups are CAT(0) groups [16], and mapping class groups are automatic groups [26].

The question of which functions are Dehn functions of finitely presented groups is quite well studied. One of the aspects of the question is investigating the following countable set of numbers, called the isoperimetric spectrum:

$$\mathrm{IP} = \big\{\rho \in [1,\infty) \;\big|\; \delta(n) = n^\rho \text{ is equivalent to a Dehn function}\big\}.$$

Gromov [24] stated that there is a gap in IP, which is $\mathrm{IP} \cap (1,2) = \emptyset$. This gap is called the Gromov gap. In other words, if a Dehn function of a finitely presented group is sub-quadratic, then it is linear. This fact was proved by Ol'shanskii [27], Bowdich [5], and Papasoglu [29]. Moreover, Brady and Bridson [6] proved that the Gromov gap is the only gap in IP. That is, they showed that the closure of IP is $\{1\} \cup [2,\infty)$. Not only functions of the form n^ρ can be Dehn functions of finitely presented groups, but so can exponential and super-exponential functions. For example, the Baumslag–Solitar group $\mathrm{BS}(1,2) = \{a, b \mid a^{-1}ba = b^2\}$ has an

exponential Dehn function ([31], Theorem 6.1); Platonov [30] showed that the Baumslag–Gersten group $\{a, b \mid a^{a^b} = a^2\}$ has a super-exponential Dehn function.

Given a finitely presented group G, the Dehn functions of subgroups of G can be larger than the Dehn function of G. As we will see later, CAT(0) groups and right-angled Artin groups have at most quadratic Dehn functions, but their finitely presented subgroups can have Dehn functions that are larger than quadratic. Because of the Gromov gap, Dehn functions of CAT(0) groups and right-angled Artin groups can only be either linear or quadratic. Brady and Forester [8] gave examples of CAT(0) groups that contain finitely presented subgroups whose Dehn functions are of the form n^ρ, for a dense set of $\rho \in [2, \infty)$. Recently, Brady and Soroko [9] proved that for each positive integer ρ, there is a right-angled Artin group that contains a finitely presented subgroup whose Dehn function is n^ρ. Larger functions can also appear as Dehn functions of finitely presented subgroups of a group. Bridson [12] proved that right-angled Artin groups and mapping class groups can have finitely presented subgroups whose Dehn functions are exponential.

Dehn functions are among quasi-isometry invariants for finitely presented groups: if two finitely presented groups are quasi-isometric, then their Dehn functions are equivalent in a certain precise sense; we refer to Section 2.4 for more details. Although Dehn functions have been studied for a long time, they are still difficult to compute. Perhaps it is because that there are only a few tools for establishing both the upper bound and the lower bound of a Dehn function; see [11] for a general survey. There are groups whose Dehn functions are still unknown. Thurston conjectured that $\mathrm{SL}_n(\mathbb{Z})$ has a quadratic Dehn function for $n \geq 4$ (see [22], the remark before Proposition 3.4). The Dehn function of $\mathrm{SL}_2(\mathbb{Z})$ is known to be quadratic; the Dehn function of $\mathrm{SL}_3(\mathbb{Z})$ is exponential [21]. Young [32] proved that $\mathrm{SL}_n(\mathbb{Z})$ is

quadratic for $n \geq 5$. Whether or not $\mathrm{SL}_4(\mathbb{Z})$ has a quadratic Dehn function is still unknown.

This dissertation is devoted to studying the Dehn functions of Bestvina–Brady groups, which are subgroups of right-angled Artin groups. Given a finite simplicial graph Γ, the right-angled Artin group A_Γ has a presentation whose generators are vertices of Γ, and whose relators are commutators $vw = wv$ whenever vertices v, w are connected by an edge in Γ. Define a homomorphism $\phi : A_\Gamma \to \mathbb{Z}$ by sending all the generators of A_Γ to 1. The Bestvina–Brady group is defined to be the kernel of ϕ, and is denoted by H_Γ. The right-angled Artin group A_Γ is the fundamental group of a compact, non-positively curved cube complex X_Γ, namely, the Salvetti complex. Consider a continuous map $f : X_\Gamma \to S^1$ such that $f_* = \phi$ and let $h : \widetilde{X}_\Gamma \to \mathbb{R}$ be the lifting of f. The Bestvina–Brady group H_Γ acts geometrically on the zero level set $Z_\Gamma = h^{-1}(0)$. The flag complex L on Γ is closely related to the finiteness properties of H_Γ. Specifically, H_Γ is finitely presented if and only if L is simply-connected [4].

Bestvina and Brady [4] introduced Bestvina–Brady groups to provide examples of groups which satisfy the finiteness property $\mathbf{FP}_n$ but not $\mathbf{FP}_{n+1}$. For instance, let F_2 be the free group of two generators, then the group $\ker(F_2 \times F_2 \to \mathbb{Z})$ is finitely generated ($\mathbf{FP}_1$) but not finitely presented ($\mathbf{FP}_2$) ([4], Example 6.3). Moreover, Bestvina–Brady groups are either counterexamples to the Eilenberg–Banea Conjecture or counterexamples to the Whitehead Conjecture ([4], Theorem 8.7). This implies that at least one of these two conjectures is false. Brady [7] provided an example showing that a torsion-free hyperbolic group G contains a finitely presented non-hyperbolic subgroup by considering the group $\ker(G \to \mathbb{Z})$ ([7], Theorem 1.1). In the sense of Dehn functions, Brady's example shows that a group with a linear Dehn function contains a finitely presented subgroup whose Dehn function is at

least quadratic. His example demonstrates again that subgroups can have larger Dehn functions than those of the ambient group.

Dison [20] proved that the Dehn functions of Bestvina–Brady groups are bounded above by a quartic function, while right-angled Artin groups have at most quadratic Dehn functions. There are examples showing that Bestvina–Brady groups can have linear, quadratic, cubic, or quartic Dehn functions. Instead of filling a closed curve with a disk, we can consider filling a k-sphere S^k by a $(k+1)$-ball D^{k+1}. The higher-order Dehn functions are generalizations of Dehn functions that describe the difficulty of the higher dimensional fillings. Abrams, Brady, Dani, Duchin, and Young [1] established the upper bound for the higher-order Dehn functions of Bestvina–Brady groups ([1], Theorem 5.1). Their result recovers Dison's quartic upper bound in [20].

In this dissertation, we study the relationship between the Dehn functions of Bestvina–Brady groups and the defining graphs of these groups. More precisely, we pose the following questions:

Main Questions.

(1) Are the Dehn functions of Bestvina–Brady groups always polynomial?

(2) Can we identify the Dehn function of a given Bestvina–Brady group directly from its defining graph?

The main method that will be used in this dissertation to compute Dehn functions is corridor schemes. For more details, we refer to Chapter 2 and Chapter 4. We say that a triangulated disk D has interior dimension d, denoted by $\dim_I(D) = d$, if the interior of D contains d-simplices but no i-simplices for $i > d$. Given a finite simplicial graph Γ, we consider the flag complex D on Γ. We answer the main

questions when D is a 2-dimensional disk with square boundary. Our main result is the following:

Theorem 4.1. *Let Γ be a finite simplicial graph such that the flag complex on Γ is a 2-dimensional triangulated disk D with square boundary. If $\dim_I(D) = d$ for $d \in \{0, 1, 2\}$, then $\delta_{H_\Gamma}(n) \cong n^{d+2}$.*

The case when $d = 0$ is a consequence of Theorem 3.10:

Theorem 3.10. *Let Γ be a finite simplicial graph such that the flag complex on Γ is a 2-dimensional triangulated disk D. If $\dim_I(D) = 0$, then $\delta_{H_\Gamma}(n) \simeq n^2$.*

Note that in Theorem 3.10, the boundary of the triangulated disk is not necessarily a square. The idea of the proof of Theorem 3.10 is as follows: the graph Γ is obtained by gluing fans and wheels in a specific way, see Lemma 3.2. The Bestvina–Brady groups on fans and wheels are isomorphic to some right-angled Artin groups (Proposition 3.3). Thus, they are fundamental groups of compact, non-positively curved spaces (Section 2.6). Proposition 3.8 states that a group obtained from the amalgamated product of two fundamental groups of compact, non-positively curved spaces over $\mathbb{Z}$ is also the fundamental group of a compact, non-positively curved space. This proposition implies that amalgamated product of two Bestvina–Brady groups whose defining graphs are fans and wheels over $\mathbb{Z}$ does not increase the Dehn function of the resulting Bestvina–Brady group and this implies the theorem; we refer to Chapter 3 for a full proof.

In general, Bestvina–Brady groups are not isomorphic to any right-angled Artin groups. Papadima and Suciu provided examples of finitely presented Bestvina–Brady groups that are not isomorphic to any Artin groups or arrangement groups ([28], Proposition 9.4 and 9.5). In particular, they constructed a Bestvina–Brady group that is not isomorphic to any right-angled Artin group ([28], Example 2.8).

For a special class of graphs, the conclusion of Theorem 3.10 can be recovered by a result of Carter and Forester [15]: for simply-connected cube complexes X_1, X_2, X_3 with some additional structures on the height function $h : X_1 \times X_2 \times X_3 \to \mathbb{R}$, if the Dehn functions of X_1, X_2, and X_3 are bounded above by n^α for $\alpha \geq 2$, then the Dehn function of the zero level set $h^{-1}(0)$ is also bounded above by n^α ([15], Theorem 4.2). As a corollary, if a finite simplicial graph is a join of three subgraphs $\Gamma = \Gamma_1 * \Gamma_2 * \Gamma_3$, then the Bestvina–Brady group H_Γ has a quadratic Dehn function ([15], Corollary 4.3). In Example 4.3, we give an explicit graph that is a join of three graphs and also satisfies the assumptions of Theorem 3.10.

We now sketch the proof of Theorem 4.1 for the case when $d = 1$ or $d = 2$; the detailed proof appears in Chapter 4. The lower bounds are given in Lemma 4.5, where the main tool used is the height-pushing map introduced in [1]; we refer to Theorem 4.4 for the formal definition of this map. The upper bound in the case $d = 2$ follows from Dison's universal quartic upper bound for Bestvina–Brady groups in [20]. The only case left to prove is the cubic upper bound for $d = 1$. We first show in Lemma 4.6 that when $d = 1$, the graph Γ is of the following form:

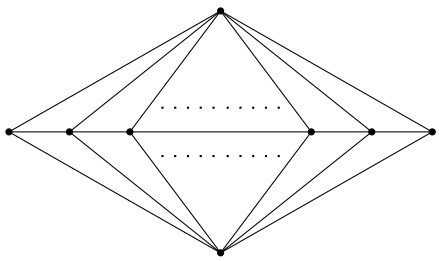

FIGURE 1.1. The graph for $d = 1$ in Theorem 4.1.

Let w be a freely reduced word of length at most n that represents the identity in H_Γ. A van Kampen diagram for w is a finite, compact, simply-connected, oriented, planar 2-complex with each oriented edge labeled by a generator of H_Γ so that the labeling on its boundary represents the word w; see Section 2.9 for more definitions.

Since w is freely reduced, by the van Kampen Lemma (Theorem 2.23), we can choose a reduced van Kampen diagram Δ for w such that $\text{Area}_G(w) = \text{Area}(\Delta)$. The van Kampen diagram Δ is cut up into specific objects, which we call stacks; see Definition 4.7. We show in Lemma 4.9 that the area of a single stack is at most cubic in terms of its perimeter. Then we show that the area of Δ is less than the area of a single stack T' whose perimeter is also at most n. Thus, by the van Kampen Lemma, we have

$$\text{Area}_{H_\Gamma}(w) = \text{Area}(\Delta) \leq \text{Area}(T') \leq n^3.$$

Hence, we obtain the cubic upper bound for δ_{H_Γ} in the case of $d = 1$. We refer to Lemma 4.9 and Lemma 4.14 for more details.

If a simplicial graph Γ contains K_4 subgraphs, then the flag complex D on Γ is not 2-dimensional. In this case, when a finite simplicial graph Γ contains some specific subgraphs, we establish the lower bound for δ_{H_Γ} in the following proposition; we refer to Chapter 5 for a detailed proof.

Proposition 5.15. *Let Γ be a finite simplicial graph such that the flag complex D on Γ is simply-connected. If Γ contains an induced subgraph Γ' and the flag complex on Γ' is a 2-dimensional triangulated subdisk D' of D that has square boundary and $\dim_I(D') = d$ for $d \in \{0, 1, 2\}$, then $\delta_{H_\Gamma}(n) \succeq n^{d+2}$.*

A brief sketch of the proof of Proposition 5.15 is as follows: the fact that Γ' is an induced subgraph of Γ gives a retraction from H_Γ to $H_{\Gamma'}$ (Proposition 5.10), and group retractions do not increase Dehn functions ([10], Lemma 2.2). Thus, δ_{H_Γ} is bounded below by $\delta_{H_{\Gamma'}}$. Since Theorem 4.1 tells us $\delta_{H_{\Gamma'}}$, the proposition follows.

Note that the flag complex on Γ in Proposition 5.15 need not be 2-dimensional, but the flag complex on the induced subgraph Γ' is 2-dimensional and whose boundary is a square.

This dissertation is organized as follows. Chapter 2 contains necessary background for the later chapters. In Chapter 3, we establish the quadratic upper bound of δ_{H_Γ} for graphs Γ whose flag complexes are 2-dimensional triangulated disks with interior dimension 0. In Chapter 4, we prove our main result Theorem 4.1. In Chapter 5, we compute Dehn functions δ_{H_Γ} for some graphs Γ whose flag complexes are not 2-dimensional.

Chapter 2. Preliminaries

2.1 Graphs

The main reference for this section is [19].

A *graph* $\Gamma = (V(\Gamma), E(\Gamma))$ is a pair which consists of two non-empty disjoint sets: $V(\Gamma)$, the set of *vertices* and $E(\Gamma) \subseteq V(\Gamma) \times V(\Gamma)$, the set of *edges*. If $V(\Gamma)$ is empty, then we call Γ an *empty graph* or a *null graph*. We say that two vertices v, w are *connected* by an edge, or *adjacent* if $\{v, w\} \subseteq E(\Gamma)$. Note that it is possible to have $v = w$. In this case, we call the edge (v, v) a *loop*. A graph is said to be *finite* if its vertex set and edge set are finite. The or *valency* of a vertex v is the number of edges that have v as one of their end points.

A *subgraph* $\Gamma' = (V(\Gamma'), E(\Gamma'))$ of a graph $\Gamma = (V(\Gamma), E(\Gamma))$ is a graph where $V(\Gamma') \subseteq V(\Gamma)$ and $E(\Gamma') \subseteq E(\Gamma)$. A subgraph $\Gamma' = (V(\Gamma'), E(\Gamma'))$ of Γ is called an *induced subgraph* if $E(\Gamma') = \big\{\{v, w\} \subseteq E(\Gamma) \;\big|\; v, w \in V(\Gamma')\big\}$.

A *path* P_n is a graph whose vertex set $V(P_n) = \{v_1, \cdots, v_{n+1}\}$ consists of consecutively adjacent vertices, that is, v_i and v_{i+1} are adjacent for $i = 1, \cdots, n$. If v_1 and v_{n+1} are also adjacent, then the graph is called a *cycle*, denoted by C_n. Define the *length* of a path and a cycle by the number of their edges. Thus, the lengths of P_n and C_n are n. Note that a cycle of length 1 is a loop. A *tree* is a graph with no cycles of any length. A graph Γ is said to be *connected* if for any pair of vertices $v, w \in V(\Gamma)$, there is a subgraph P_n of Γ whose vertex set is $\{v = v_1, v_2, \cdots, v_{n-1}, v_n = w\}$ for some n. A graph is said to be *complete* if any two vertices are adjacent. Denote the complete graph of n vertices by K_n. A *simplicial graph* is a graph with no loops, and such that adjacent vertices have exactly one edge connecting them.

Given two graphs $\Gamma_1 = (V(\Gamma_1), E(\Gamma_1))$ and $\Gamma_2 = (V(\Gamma_2), E(\Gamma_2))$. The *union* of Γ_1 and Γ_2 is $\Gamma_1 \cup \Gamma_2 = (V(\Gamma_1) \cup V(\Gamma_2), E(\Gamma_1) \cup G(\Gamma_2))$. The *join* of Γ_1 and Γ_2, denoted by $\Gamma_1 * \Gamma_2$, is defined to the graph union $\Gamma_1 \cup \Gamma_2$ and with every pair of vertices $(v, w) \in V(\Gamma_1) \times V(\Gamma_2)$ are adjacent. A *cone* on a graph Γ is a join of a vertex $v \notin V(\Gamma)$ and Γ. The *suspension* of Γ is the union of $v * \Gamma$ and $w * \Gamma$, where $v, w \notin V(\Gamma)$ are distinct vertices.

An edge e of a graph Γ that connects vertices v and w is called an *oriented edge* if there is an orientation assigned to e. We draw an arrow on the edge to indicate its orientation. For an edge, there are two orientations which can be assigned: from v to w or from w to v.

FIGURE 2.1. Two orientations of an oriented edge.

If the arrow of e points from v to w, then we call v the *initial vertex* of e and w the *terminal vertex* of e. An *oriented graph* Γ is a graph with an assigned orientation on each edge. Denote the set of oriented edges of Γ by $\overrightarrow{E}(\Gamma)$.

2.2 Cayley Graphs and Quasi-isometries

Let G be a group with a finite presentation $\mathcal{P} = \langle \mathcal{S} | \mathcal{R} \rangle$. The *Cayley graph* $\mathrm{Cay}(G, \mathcal{S})$ of G for a finite generating set $\mathcal{S}$ has the vertex set G, and there is an edge between g and h for $g, h \in G$ if $g^{-1}h \in G$.

We equip Cayley graphs with metrics to make them metric spaces. Let G be a group with a finite presentation $\mathcal{P} = \langle \mathcal{S} | \mathcal{R} \rangle$. The *word metric* $d_{\mathcal{S}}$ on $\mathrm{Cay}(G, \mathcal{S})$ with respect to a finite generating set $\mathcal{S}$ is defined as follows: for $g, h \in G$,

$$\begin{aligned} d_{\mathcal{S}}(g, h) &= \text{length of the shortest path between } g \text{ and } h \text{ in } \mathrm{Cay}(G, \mathcal{S}) \\ &= \text{the shortest word in the length of } \mathcal{S} \text{ which represents } g^{-1}h. \end{aligned}$$

Definition 2.1. *A map $f : X \to Y$ between metric spaces X and Y is a* quasi-isometric embedding *if there exist positive numbers k, c such that for all $a, b \in X$, we have*

$$\frac{1}{k} d_X(a, b) - c \leq d_Y(f(a), f(b)) \leq k d_X(a, b) + c.$$

In addition, f is a quasi-isometry *if $Y \subseteq N_c(f(X))$. If $c = 0$, then f is called* biLipschitz.

We say that two metric spaces X and Y are quasi-isometrically equivalent *(or simply* quasi-isometric*) if there is quasi-isometry $f : X \to Y$.*

Although Cayley graphs of a group are not unique, they are quasi-isometric to each other:

Proposition 2.2. ([25], Theorem 7.3) *If $\mathcal{S}$ and $\mathcal{S}'$ are finite generating sets for a group G, then Cay$(G, \mathcal{S})$ is quasi-isometric to Cay$(G, \mathcal{S}')$.*

With the help of the above proposition, we can view a group G as a metric space by taking its Cayley graph with any finite generating set, and this metric space is unique in the sense that all the Cayley graphs of G are quasi-isometric. It also makes sense to say that a group G is quasi-isometric to a metric space X, by which we mean the Cayley graph of G is quasi-isometric to X.

We say that two groups are quasi-isometric if their Cayley graphs are quasi-isometric. We say that a property P of a group G is a *quasi-isometry invariant* if whenever a group H is quasi-isometric to G, H also has the property P. That is, quasi-isometry invariants of groups are properties of groups that are preserved by quasi-isometries between groups.

2.3 Milnor-Švarc Lemma

Before stating the lemma, we need a few definitions. A metric space X is said to be *proper* if every closed ball is compact. An action of a group G on a metric space

X is called *cocompact* if there is a compact subset K of X such that $X \subseteq \bigcup_{g \in G} gK$. An action of a group G on a metric space X is called *properly discontinuous* if for each compact set $K \subseteq X$, the set $\{g \in G | gK \cap K \neq \phi\}$ is finite. A group action which is cocompact and properly discontinuous is called a *geometric action.*

Now we state the Milnor-Švarc Lemma. Its proof can be found in [25].

Lemma 2.3. ([25], Theorem 7.5) *(Milnor-Švarc) Let X be a proper geodesic metric space. If a group G acts geometrically by isometries on X, then*

(i) G is finitely generated

(ii) G is quasi-isometric to X

The following corollary is a consequence of the Milnor-Švarc Lemma.

Corollary 2.4. *If K is a compact simplicial complex, then $\pi_1(K)$ is quasi-isometric to $\widetilde{K}$, where $\widetilde{K}$ is the universal cover of K.*

2.4 Dehn Functions

Let G be a group with a finite presentation $\mathcal{P} = \langle \mathcal{S} | \mathcal{R} \rangle$. Let w be a word that represents the identity of G, denoted by $w \equiv_G 1$. The *area* of w, denoted by $\text{Area}(w)$, is defined as follows:

$$\text{Area}_G(w) = \min \left\{ N \;\middle|\; w = \prod_{i=1}^{N} x_i r_i^{\pm 1} x_i^{-1}, x_i \in F(\mathcal{S}), r_i \in \mathcal{R} \right\},$$

where $F(\mathcal{S})$ is the free group generated by S. The *Dehn function* $\delta_G : \mathbb{N} \to \mathbb{N}$ of a group G over the presentation $\mathcal{P} = \langle \mathcal{S} | \mathcal{R} \rangle$ is defined by

$$\delta_G(n) = \max \left\{ \text{Area}_G(w) \;\middle|\; w \equiv_G 1, |w| \leq n \right\},$$

where $|w|$ denotes the length of the word w.

Definition 2.5. *Let $f, g : [0, \infty) \to [0, \infty)$ be two functions. We say that f is* bounded above *by g, denoted by $f \preceq g$, if there is a number $C > 0$ such that*

$f(n) \leq Cg(Cn+C)+Cn+C$ *for all* $n > 0$. *We say that* f *and* g *are* $\simeq$-equivalent*, or simply* equivalent*, denoted by* $f \simeq g$*, if* $f \preceq g$ *and* $g \preceq f$.

We say that the Dehn function δ_G is *linear*, *quadratic*, *cubic* or *quartic* if for all $n \in \mathbb{N}$, $\delta_G(n) \simeq n$, $\delta_G(n) \simeq n^2$, $\delta_G(n) \simeq n^3$ or $\delta_G(n) \simeq n^4$, respectively.

Dehn functions are defined over a given finite presentation, but a finitely presented group can have many different finite presentations. It turns out that under the $\simeq$-equivalence, the Dehn function δ_G of a finitely presented group G does not depend on the presentations of G. The following proposition is standard, its proof can be found in [11].

Proposition 2.6. ([11], Proposition 1.3.3) *Let* $\mathcal{P}, \mathcal{P}'$ *be finite presentations for a group* G. *Then the Dehn function* δ_G *over* $\mathcal{P}$ *and* $\mathcal{P}'$ *are equivalent.*

Moreover, the Dehn function is a quasi-isometry invariant:

Theorem 2.7. ([3]) *If two finitely presented groups* G_1 *and* G_2 *are quasi-isometric, then* $\delta_{G_1} \simeq \delta_{G_2}$.

Let K be a finite simplicial CW complex and $\widetilde{K}$ its universal cover. Each word that represents the identity in $\pi_1(K)$ corresponds to a loop c in the 1-skeleton of $\widetilde{K}$, and c bounds a disk D in $\widetilde{K}$. The area of c is defined to be the minimal number of 2-cells in D. The *filling function* of $\widetilde{K}$ is

$$\mathrm{Fill}_{\widetilde{K}}(l) = \sup\left\{ \mathrm{Area}(c) \,\middle|\, c \text{ is a loop in } \widetilde{K}, \text{ length of } c \leq l \right\}.$$

The proof of the following fact is non-trivial.

Proposition 2.8. ([11], Theorem 5.0.1) *The functions* $\delta_{\pi_1(K)}$ *and* $Fill_{\widetilde{K}}$ *are equivalent.*

More generally, we have

Theorem 2.9. ([14], Theorem 1.1) *Let G be a finitely presented group and M a simply-connected Riemannian manifold. If G acts geometrically on M by isometries, then δ_G and $Fill_M$ are equivalent.*

2.5 CAT(0) Spaces and CAT(0) Groups

Let (X, d_X) be a geodesic metric space. A *geodesic triangle*, or simply a *triangle* in X is any three distinct points x, y, z in X joint by geodesic segments $[x, y]$, $[y, z]$, and $[z, x]$. Since we do not assume that X is a unique geodesic space, a geodesic triangle of any three distinct points is not unique.

Let Δ_{pqr} be a triangle in X whose vertices are p, q, r. A *comparison triangle* of Δ_{pqr} is a triangle $\overline{\Delta}_{\bar{p}\bar{q}\bar{r}}$ in the Euclidean space $\mathbb{R}^2$ such that $d_X(p, q) = d_{\mathbb{R}^2}(p, q)$, $d_X(q, r) = d_{\mathbb{R}^2}(q, r)$, $d_X(p, r) = d_{\mathbb{R}^2}(p, r)$, where $d_{\mathbb{R}^2}$ is the Euclidean metric on $\mathbb{R}^2$. A point $\bar{x} \in [\bar{p}, \bar{q}]$ (respectively $[\bar{p}, \bar{r}]$, $[\bar{q}, \bar{r}]$) is called a *comparison point* of a point $x \in [p, q]$ (respectively $[p, r]$, $[q, r]$) if $d_X(x, p) = d_{\mathbb{R}^2}(\bar{x}, \bar{p})$ (respectively $d_X(x, p) = d_{\mathbb{R}^2}(\bar{x}, \bar{p})$, $d_X(x, q) = d_{\mathbb{R}^2}(\bar{x}, \bar{q})$)

A geodesic metric space (X, d_X) is called CAT(0) if for any triangle Δ_{pqr}, any pair of points $x, y \in \Delta_{pqr}$ and their comparison points $\bar{x}, \bar{y} \in \overline{\Delta}_{\bar{p}\bar{q}\bar{r}}$, satisfying the following inequality:

$$d_X(x, y) \leq d_{\mathbb{R}^2}(\bar{x}, \bar{y}).$$

We list some properties of CAT(0) spaces.

Proposition 2.10. ([13], Chapter II.1, 1.4, 1.5, Chapter III.H, 2.4) *Let X be a* CAT(0) *space.*

(1) X is a unique geodesic space. That is, every pair of distinct points in X is connected by a unique geodesic.

(2) X is contractible.

(3) The filling function of X is at most quadratic.

A geodesic metric space X is called *locally* CAT(0) or *non-positively curved* if for every point $x \in X$, there is a neighborhood U_x of x with induced metric such that U_x is a CAT(0) space.

The following theorem stated by Gromov [24] is a generalization of the Cartan-Hadamard Theorem in the language of Riemannian geometry: the universal cover of a connected complete n-dimensional Riemannian manifold with non-positive sectional curvature is diffeomorphic to $\mathbb{R}^n$.

Theorem 2.11. ([13], Chapter II.4, 4.1) *If X is locally* CAT(0)*, then its universal cover $\widetilde{X}$ is* CAT(0)*.*

Now we define CAT(0) groups. A CAT(0) group is a group which acts geometrically on a CAT(0) space.

Let G be a CAT(0) group. By definition, G acts geometrically on a CAT(0) space X. By the Milnor-Švarc Lemma, G is quasi-isometric to X. It follows from Theorem 2.9 that the Dehn function δ_G is equivalent to the filling function of X, which is at most quadratic since X is CAT(0). We now have the following:

Proposition 2.12. *Dehn functions of* CAT(0) *groups are at most quadratic.*

2.6 Right-angled Artin Groups and Salvetti Complexes

Let Γ be a finite simplicial graph with vertex set $V(\Gamma)$. The *right-angled Artin group* A_Γ associated to Γ is defined as

$$A_\Gamma = \left\langle \mathrm{V}(\Gamma) \;\middle|\; [v_i, v_j] \text{ whenever } v_i \text{ and } v_j \text{ are connected by an edge of } \Gamma \right\rangle.$$

When Γ is a complete graph K_n on n vertices, $A_\Gamma = \mathbb{Z}^n$; when Γ is a set of n distinct points, $A_\Gamma = F_n$, the free group on n generators.

For each finite simplical graph Γ, its associated right-angled Artin group A_Γ is the fundamental group of a cubical complex X_Γ, called the **Salvetti complex**.

The construction is the following: start with a number of circles which are attached to a unique point, and label the circles by the generators $v_1, \cdots, v_n$ of A_Γ. These are 1-skeleton $X_\Gamma^{(1)}$ of X_Γ. For each edge connecting v_i and v_j, attach a square with boundary labeled by $v_i v_j v_i^{-1} v_j^{-1}$ to the $X_\Gamma^{(1)}$, we get the 2-skeleton $X_\Gamma^{(2)}$. For each triangle in Γ with vertices v_i, v_j, and v_k, attach a cube whose faces correspond to edges of the triangle to $X_\Gamma^{(2)}$, we get the 3-skeleton $X_\Gamma^{(3)}$. Similarly, for each complete subgraph K_d of Γ, attach a d-cube whose faces correspond to K_{d-1} subgraphs of K_d to $X_\Gamma^{(d-1)}$, we get the d-skeleton $X_\Gamma^{(d)}$. For example, if $\Gamma = K_2$, then $A_\Gamma = \mathbb{Z}^2$ and the Salvetti complex X_Γ is a torus and $\widetilde{X}_\Gamma$ is $\mathbb{R}^2$.

It is a fact that the Salvetti complex X_Γ is compact and non-positively curved ([17], Theorem 3.3.1). Therefore, by Theorem 2.11 we have

Theorem 2.13. ([16], Theorem 2.6) *The universal cover $\widetilde{X}_\Gamma$ of the Salvetti complex X_Γ is a* CAT(0) *cube complex. Hence, X_Γ is a* $\mathrm{K}(A_\Gamma, 1)$ *space.*

It follows from Theorem 2.13 that $\pi_1(X_\Gamma) = A_\Gamma$. Since fundamental groups act on universal covers geometrically, we get the following corollary

Corollary 2.14. *Right-angled Artin groups are* CAT(0) *groups.*

Since Dehn functions of CAT(0) groups are at most quadratic, we have

Corollary 2.15. *Dehn functions of right-angled Artin groups are at most quadratic.*

2.7 Bestvina–Brady Groups and Dick–Leary Presentation

Given a finite simplicial graph Γ, we define a group homomorphism $\phi : A_\Gamma \to \mathbb{Z}$ by sending all the generators to 1. The kernel of this homomorphism is called the *Bestvina–Brady group* defined by Γ, and is denoted by H_Γ.

There is a space which the Bestvina–Brady group acts on geometrically: Let X_Γ be the Salvetti complex of A_Γ and $f : X_\Gamma \to S^1$ a continuous map such that $f_* = \phi : A_\Gamma \to \mathbb{Z}$. The *height function* $h : \widetilde{X}_\Gamma \to \mathbb{R}$ is a lifting of f. The Bestvina–

Brady group H_Γ acts geometrically on the zero level set $Z_\Gamma = h^{-1}(0)$. We refer to [4] for more details.

The *flag complex* on a finite simplicial graph Γ is a simplicial complex L such that each complete subgraph K_n of Γ spans an $(n-1)$-simplex in L. Whether the Bestvina–Brady group H_Γ is finitely generated or finitely presented can be determined by the flag complex on Γ.

Theorem 2.16. ([4]) *Let Γ be a finite simplicial graph.*

(1) If the flag complex on Γ is connected, than H_Γ is finitely generated.

(2) If the flag complex on Γ is simply-connected, than H_Γ is finitely presented.

When the Bestvina–Brady group is finitely presented, Dicks and Leary [18] found a presentation for H_Γ.

Theorem 2.17. ([18], Corollary 3) *Let Γ be a finite simplicial graph with a given orientation. Suppose the flag complex on Γ is simply-connected. Then the corresponding Bestvina–Brady group H_Γ has the following finite presentation:*

$$H_\Gamma = \left\langle \overline{E}(\Gamma) \;\middle|\; ef = fe, ef = g \text{ whenever } e, f, g \text{ form an oriented triangle} \right\rangle,$$

where $\overline{E}(\Gamma)$ is the set of oriented edges of Γ, and the oriented triangle is as shown in Figure 2.2

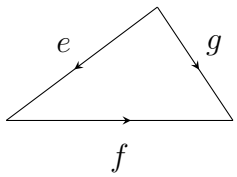

FIGURE 2.2. Realator of the Dicks–Leary presentation.

In fact, the generating set in the above theorem can be reduced:

Corollary 2.18. ([28], Corollary 2.3) *If the flag complex on a finite simplicial graph is simply-connected, then H_Γ has a presentation $H_\Gamma = F/R$, where F is the free group generated by the edges in a maximal tree of Γ, and R a finitely generated normal subgroup of the commutator group $[F, F]$.*

While Dehn functions of right-angled Artin groups are at most quadratic, Bestvina–Brady can have larger Dehn functions. Dison [20] proved that the Dehn functions of Bestvina–Brady groups are bounded above by a quartic function. In later chapters, we will give examples of Bestvina–Brady groups which have quadratic, cubic, and quartic Dehn functions.

Theorem 2.19. ([20]) *The Dehn functions of Bestvina–Brady groups are bounded above by a quartic function.*

2.8 Hyperbolic Spaces and Hyperbolic Groups

Let X be a geodesic metric space. A *triangle* in X is any three distinct points p, q, r joint by geodesic segments $[p, q], [q, r], [p, r]$. Note that since the geodesic is not unique, the triangle of any three distinct points is not unique. A triangle in X is called *δ-thin* if there is a $\delta > 0$ such that any side of the triangle is contained in the union of δ-neighborhoods of the other two sides. A geodesic metric space X is called *hyperbolic* if there is a $\delta > 0$ such that all triangles in X are δ-thin. A group G is said to be *hyperbolic* if its Cayley graph $Cay(G, S)$ is hyperbolic for some finite generating set S.

Hyperbolicity of metric spaces is a quasi-isometry invariant of spaces. That is,

Theorem 2.20. ([13], Chapter III.H, 1.9 Theorem) *Let X and Y be geodesic metric spaces and $f : X \to Y$ a quasi-isometry. If X is hyperbolic, then Y is also hyperbolic.*

Hyperbolic groups have many properties, see [24]. One of the properties is that: if G is a hyperbolic group, then G is finitely presented.

Theorem 2.21. ([24]) *Let G be a finitely presented group. The following statements are equivalent.*

(1) G is hyperbolic.

(2) The Duhn function if sub-quadratic

(3) The Dehn function of G is linear.

(4) G contains no $\mathbb{Z} \times \mathbb{Z}$ subgroups.

The following theorem follows from Theorem 2.20.

Theorem 2.22. *Let G be a hyperbolic group. If a group H is quasi-isometric to G, then H is also hyperbolic.*

2.9 van Kampen Diagram and van Kampen Lemma

Let G be a group with a finite presentation $\mathcal{P} = \langle \mathcal{S} | \mathcal{R} \rangle$. Let $\mathcal{D}$ be a finite, compact, simply-connected, oriented, planar 2-complex, where each oriented edge is labeled by a letter in $\mathcal{S}$. Each oriented path γ in the 1-skeleton of $\mathcal{D}$ is labeled by a word $s_1^{\epsilon_1} \cdots s_n^{\epsilon_n}$ in G, where $s_i \in \mathcal{S}$ is the labeling on the corresponding oriented edge; $\epsilon_i = 1$ when reading along the edge labeled by s_i agrees with the orientation of that edge, otherwise, $\epsilon_i = -1$. Suppose that for each 2-cell F of $\mathcal{D}$, the boundary of F is labeled by a cyclic permutation of an element in $\mathcal{R}$ or the inverse of an element in $\mathcal{R}$. Fix a vertex p on the boundary of $\mathbb{R}^2 \setminus \mathcal{D}$ and we call p a *base vertex*. The *boundary word* w of the diagram $\mathcal{D}$ is a labeling of the boundary of $\mathbb{R}^2 \setminus \mathcal{D}$, starting from the base vertex p and reading along the boundary counterclockwise. Then we say that $\mathcal{D}$ is a *van Kampen diagram* for the word w over the presentation $\mathcal{P}$. Denote the number of 2-cells of $\mathcal{D}$ by Area($\mathcal{D}$). We say that a van Kampen diagram D for

a word w is *reduced*, or *minimal*, if $\text{Area}(\mathcal{D}) \leq \text{Area}(\mathcal{D}')$, where $\mathcal{D}'$ is a van Kampen diagram for w.

Theorem 2.23 (van Kampen Lemma)**.** ([11], Theorem 4.2.2) *Let G be a group with a finite presentation $\mathcal{P} = \langle \mathcal{S} | \mathcal{R} \rangle$.*

(1) A word w represents the identity in G if and only if there is a van Kampen diagram $\mathcal{D}$ for w over $\mathcal{P}$.

(2) If a word w represents the identity in G, then

$$Area_G(w) = \min \left\{ Area(\mathcal{D}) \;\middle|\; \mathcal{D} \text{ is a van Kampen diagram for } w \text{ over } \mathcal{P} \right\}.$$

(3) A word $w \in G$ is freely reduced word and $w =_G 1$ if and only if there is a minimal van Kampen diagram $\mathcal{D}$ for w over $\mathcal{P}$.

Consider a van Kampen diagram $\mathcal{D}$ for a word w over a finite presentation $\mathcal{P} = \langle \mathcal{S} | \mathcal{R} \rangle$. Fix an edge e on the boundary of $\mathbb{R}^2 \setminus \mathcal{D}$. If e is part of the boundary of a 2-cell, then the boundary of this 2-cell is labeled by a word $r \in \mathcal{R}$, and the labeling on the edge e is part of the word r. In this 2-cell, fix an edge $e' \notin \partial(\mathbb{R}^2 \setminus \mathcal{D})$. Since $\mathcal{D}$ is simply-connected, e' must be a face of another 2-cell in $\mathcal{D}$. The labeling on the boundary of this 2-cell is also a word in $\mathcal{R}$. By continuing this process, we get a chain of 2-cells that starts from an edge on the boundary of $\mathbb{R}^2 \setminus \mathcal{D}$ and ends at another edge on the boundary of $\mathbb{R}^2 \setminus \mathcal{D}$. This chain of 2-cells is called a *corridor*.

Fix a 2-cell in the interior of $\mathcal{D}$. Using the similar argument in defining the corridor, we get a chain of 2-cells that starts from a 2-cell inside $\mathcal{D}$. If this chain ends at an edge of the boundary of the starting 2-cell, we call this chain of 2-cells a *ring*.

Topologically, a corridor is homeomorphic to $[0,1]\times[0,1]$ and a ring is homeomorphic to $[0,1]\times S^1$. Corridors and rings have orientations induced by the orientation of $\mathcal{D}$.

2.10 Flag Complexes and Interior Dimensions

Let D be a triangulated 2-disk. An *interior i-simplex* of D is an i-simplex whose faces do not intersect the boundary of D. We also call an interior 0-simplex an *interior vertex*, an interior 1-simplex an *interior edge*, and an interior 2-simplex an *interior triangle*. We say that a triangulated disk D has interior dimension d, denoted by $\dim_I(D) = d$, if the interior of D contains d-simplices but no i-simplices for $i > d$.

Example 2.24. *The flag complexes on the following graphs are triangulated disks with square boundaries, and they have interior dimensions* 0, 1, *and* 2, *respectively:*

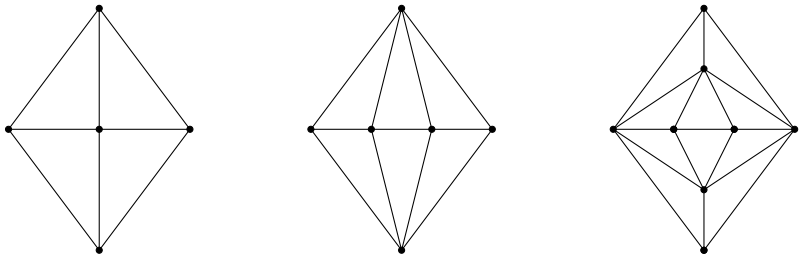

FIGURE 2.3. Disks of interior dimensions $0, 1, 2$, respectively, and with square boundaries.

Chapter 3. Disk with Interior Dimension 0

In this section, we prove that if Γ is a finite simplicial graph such that the flag complex on Γ is a 2-dimensional triangulated disk D with $\dim_I(D) = 0$, then δ_{H_Γ} is quadratic. Such Γ can be obtained by gluing *fans* and *wheels* in a certain way.

Definition 3.1. *A* fan *is the join of a vertex and a path P_n. A* wheel *is the join of a vertex and a cycle C_n, $n \geq 4$.*

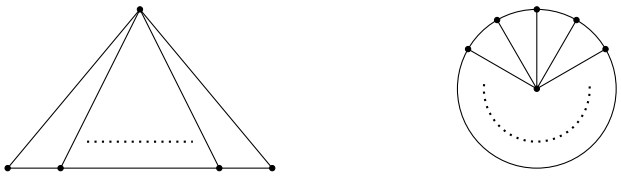

FIGURE 3.1. Fan and Wheel.

In other words, fans and wheels are cones on paths and cycles, respectively. Notice that a triangle is also a fan, a join of a vertex and a path of length 1. We also require that the minimal wheel have at least five vertices. If Γ contains a wheel with four vertices, then the flag complex on Γ will be a 3-dimensional complex.

Lemma 3.2. *Let Γ be a finite simplicial graph such that the flag complex on Γ is a 2-dimensional triangulated disk D. If $\dim_I(D) = 0$, then Γ is obtained by starting with a fan or a wheel and repeatedly gluing on such graphs, one at a time, along a 1-simplex on the boundary of the flag complex on the previous graph.*

Proof. Since $\dim_I(D) = 0$, we have two cases: either D contains no interior vertices or D contains interior vertices. Observe that in both cases, if D contains any 1-simplex not on ∂D such that all its faces are on ∂D, then cutting along that 1-simplex separates D into two disjoint triangulated disks. (Otherwise, D would be a triangulated annulus.)

Case I. Suppose that D contains no interior vertices. There are two kinds of 1-simplices of D: those are on ∂D and those are not on ∂D. We claim that Γ is obtained by gluing fans in the desired way. We prove this claim by induction on the number of the 1-simplices of D which are not on ∂D. Note that such 1-simplices have all their faces on ∂D. If all the 1-simplices of D are on ∂D, then there are exactly three 0-simplices on ∂D. Thus, D is a single 2-simplex and Γ is a K_3, which is a fan. If D contains one 1-simplex that is not on ∂D, cut along that 1-simplex, we get two disjoint triangulated disks with all the 1-simplices are on their boundaries. Thus, each of these two disks is a single 2-simplex and D is obtained by gluing two 2-simplices along one 1-simplex. That is, Γ is obtained by gluing two K_3 along a single edge. Suppose that the claim holds for D having k 1-simplices that are not on ∂D. If D contains $k+1$ 1-simplices that are not on ∂D, cut along any such 1-simplex, we get two disjoint triangulated disks and each of these two disks contains at most k 1-simplices that are not on its boundary. By the induction hypothesis, each of these two disks is obtained by gluing 2-simplices one by one, along a 1-simplex on the boundary of the previous disk. Thus, D is obtained by gluing these two disks along a 1-simplex on their boundaries. Hence, Γ is obtained by gluing fans together in the desired way.

Case II. Suppose that D contains interior vertices. Notice that any two pair of interior vertices can not be adjacent. Otherwise, $\dim_I(D) > 0$. We claim that Γ is obtained by gluing wheels and fans in the desired way. We argue by induction on the number of interior vertices of D. Suppose D has one interior vertex, then cutting along all the 1-simplices that have all the faces on ∂D gives disjoint triangulated subdisks, only one of them has an interior vertex. Denote this subdisk by D'. Note that each of the 1-simplices of D' that is not on $\partial D'$ has the interior vertex as one of its faces and the other face is on $\partial D'$. There is an induced subgraph Γ' of

Γ such that the flag complex on Γ' is D'. Clearly, Γ' must be a wheel. All other subdisks of D have no interior vertices, so they are 2-simplices by Case I. Thus, D is obtained by gluing D' and 2-simplices together in the desired way. Hence, Γ is obtained by gluing one wheel and fans together in the desired way. Suppose the claim holds for D having k interior vertices, that is, Γ is obtained by gluing k wheels and some number of fans together in the desired way. Now, suppose that D has $k+1$ interior vertices. Since none of the interior vertices of D is adjacent to another interior vertex, there is a 1-simplex not on ∂D and whose faces are on ∂D such that when cutting along that 1-simplex, we get two triangulated subdisks D'_1 and D'_2 of D, where one of them, say D'_1, has one interior vertex, and the other, D'_2, has k interior vertices. By the induction hypothesis, Γ has an induced subgraph Γ'_2 such that the flag complex on Γ'_2 is D'_2 and Γ'_2 is obtained by gluing wheels and fans in the desired way. By the previous discussion, Γ has an induced subgraph Γ'_1 such that Γ'_1 is a wheel and the flag complex on Γ'_1 is D'_1. Since D is obtained by gluing D'_1 and D'_2 together along a 1-simplex on their boundaries, Γ is obtained by gluing wheels and fans together in the desired way. □

In [28], the authors give an example (Example 2.5) that if a graph Γ is the cone on Γ', then $H_\Gamma \cong A_{\Gamma'}$. In the following proposition, we give the explicit isomorphism. Furthermore, this isomorphism gives us the upper bound of δ_{H_Γ} since the Dehn functions of right-angled Artin groups are bounded above by a quadratic function.

Proposition 3.3. *Let Γ be a finite simplicial graph. If Γ can be decomposed as a graph join $\Gamma = \{v\} * \Gamma'$ where v is a vertex of Γ, then δ_{H_Γ} is at most quadratic. In particular, if Γ' contains an edge, then δ_{H_Γ} is quadratic.*

Proof. Since $\Gamma = \{v\} * \Gamma'$, we have $A_\Gamma = \mathbb{Z} \times A_{\Gamma'}$. We claim that $H_\Gamma \cong A_{\Gamma'}$. Label the edges that have v as an endpoint by $e_1, \cdots, e_k$. Give an orientation on all the

edges of Γ such that v is the initial point of $e_1, \cdots, e_k$. Denote the terminal points of $e_1, \cdots, e_k$ by $w_1, \cdots, w_k$. Since $e_1, \cdots, e_k$ form a maximal tree of Γ, they form a generating set of H_Γ, see Corollary 2.18. Meanwhile, $w_1, \cdots, w_k$ form a generating set of $A_{\Gamma'}$. Define a map $\psi : H_\Gamma \to A_{\Gamma'}$ by $\psi(e_i) = w_i$, $i = 1, \cdots, k$. We claim that ψ is an isomorphism. Note that the relators of $A_{\Gamma'}$ and H_Γ are commutators. The generators e_i and e_j commute when they are two edges of the same triangle, that is, when their terminal points w_i and w_j are connected by an edge. This is equivalent to saying that w_i and w_j commute. We have

$$\psi([e_i, e_j]) = [w_i, w_j].$$

Since all the relators of H_Γ and $A_{\Gamma'}$ are commutators and ψ preserves all the commutators, ψ is a homomorphism. Define a map $\varphi : A_{\Gamma'} \to H_\Gamma$ by $\varphi(w_i) = e_i$, $i = 1, \cdots, k$. By the same reason for ψ being a homomorphism, φ is also a homomorphism. Obviously, we have $\varphi = \psi^{-1}$. Thus, $\psi : H_\Gamma \to A_{\Gamma'}$ is an isomorphism. This proves the claim. Since $H_\Gamma \cong A_{\Gamma'}$ and $\delta_{A_{\Gamma'}}$ is at most quadratic, δ_{H_Γ} is at most quadratic.

If Γ' contains an edge, then $H_\Gamma \cong A_{\Gamma'}$ contains $\mathbb{Z} \times \mathbb{Z}$ as a subgroup. Therefore, $H_{\Gamma'}$ cannot be hyperbolic and δ_{H_Γ} has to be quadratic. □

Corollary 3.4. *If Γ is a fan or a wheel, then H_Γ is a right-angled Artin group, whose Dehn function δ_{H_Γ} is quadratic.*

Proof. Let Γ be either a fan or a wheel. Then Γ is a join of one vertex and a path graph or a cycle graph, respectively. By Proposition 3.3, H_Γ is a non-hyperbolic right-angled Artin group. Thus, δ_{H_Γ} is quadratic. □

Proposition 3.5. *Let $\Gamma = \Gamma_1 \cup \Gamma_2$, where Γ_1, Γ_2 are finite simplicial graphs and their flag complexes are simply-connected. If $\Gamma_3 = \Gamma_1 \cap \Gamma_2$ is a connected induced subgraph of Γ, then $H_\Gamma = H_{\Gamma_1} *_{H_{\Gamma_3}} H_{\Gamma_2}$.*

Proof. Since Γ_1 and Γ_2 are finite simplicial graphs, Γ is also a finite simplicial graph. Note that the connectivity of $\Gamma_3 = \Gamma_1 \cap \Gamma_2$ implies that the flag complex on Γ is simply-connected as well. Moreover, since the flag complexes on $\Gamma, \Gamma_1, \Gamma_2$ are simply-connected and Γ_3 is connected, the corresponding Bestvina–Brady groups $H_\Gamma, H_{\Gamma_1}, H_{\Gamma_2}$ are finitely presented and H_{Γ_3} is finitely generated by Theorem 2.16. Let $H_{\Gamma_1} = \langle \mathcal{S}_1 \big| \mathcal{R}_1 \rangle$, $H_{\Gamma_2} = \langle \mathcal{S}_2 \big| \mathcal{R}_2 \rangle$ be the Dicks–Leary presentations (see Theorem 2.17), and $H_{\Gamma_3} = \langle \mathcal{S}_3 \big| \mathcal{R}_3 \rangle$, where $\mathcal{S}_i$ is the set of directed edges of Γ_i, $i = 1, 2, 3$. Let $i_1 : H_{\Gamma_3} \hookrightarrow H_{\Gamma_1}$, $i_2 : H_{\Gamma_3} \hookrightarrow H_{\Gamma_2}$ be natural inclusions, we have

$$H_{\Gamma_1} *_{H_{\Gamma_3}} H_{\Gamma_2} = \left\langle \mathcal{S}_1 \cup \mathcal{S}_2 \middle| \mathcal{R}_1 \cup \mathcal{R}_2 \cup \left\{ i_1(h) i_2^{-1}(h) \middle| h \in \mathcal{S}_3 \right\} \right\rangle.$$

Since $\Gamma_3 = \Gamma_1 \cap \Gamma_2$, we have $\mathcal{S}_3 = \mathcal{S}_1 \cap \mathcal{S}_2$ and i_1, i_2 are identities. That is, $i_1(h) i_2^{-1}(h) = h h^{-1} = 1$ for all h in $\mathcal{S}_3$. Thus, $H_{\Gamma_1} *_{H_{\Gamma_3}} H_{\Gamma_2}$ has the presentation

$$H_{\Gamma_1} *_{H_{\Gamma_3}} H_{\Gamma_2} = \left\langle \mathcal{S}_1 \cup \mathcal{S}_2 \middle| \mathcal{R}_1 \cup \mathcal{R}_2 \right\rangle,$$

which is exactly the Dicks–Leary presentation for H_Γ. □

In Proposition 3.5, when $\Gamma_1 \cap \Gamma_2$ is a vertex or an edge, we get a free product or a amalgamated product over $\mathbb{Z}$, respectively.

Corollary 3.6. *Let $\Gamma = \Gamma_1 \cup \Gamma_2$, where Γ_1, Γ_2 are finite simplicial graphs and their flag complexes are simply-connected. If $\Gamma_1 \cap \Gamma_2$ is a single vertex, then $H_\Gamma = H_{\Gamma_1} * H_{\Gamma_2}$.*

Corollary 3.7. *Let $\Gamma = \Gamma_1 \cup \Gamma_2$, where Γ_1, Γ_2 are finite simplicial graphs and their flag complexes are simply-connected. If $\Gamma_1 \cap \Gamma_2$ is a single edge, then $H_\Gamma = H_{\Gamma_1} *_{\mathbb{Z}} H_{\Gamma_2}$.*

If each of Γ_1 and Γ_2 is either a fan or a wheel, Lemma 3.4 tells us that H_{Γ_1} and H_{Γ_2} are right-angled Artin groups, hence, CAT(0) groups. For our purpose,

we would like to know whether the group $H_{\Gamma_1} *_{\mathbb{Z}} H_{\Gamma_2}$ is also CAT(0). The next proposition gives an affirmative answer.

Proposition 3.8. ([13], Chapter II.11, 11.17 Proposition) *If each of the groups H_1 and H_2 is a fundamental group of a compact metric space of non-positive curvature, then so is the amalgamated product $H_1 *_{\mathbb{Z}} H_2$.*

We immediately have the following lemma.

Lemma 3.9. *If H_i is the fundamental group of a non-positively curved compact metric space, $i = 1, \cdots k$. Then*

$$H = ((((H_1 *_{\mathbb{Z}} H_2) *_{\mathbb{Z}} H_3) *_{\mathbb{Z}} \cdots) *_{\mathbb{Z}} H_{k-1}) *_{\mathbb{Z}} H_k$$

is also the fundamental group of a non-positively curved compact metric space.

Proof. By Proposition 3.8, each group in the parentheses is the fundamental group of a non-positively curved compact metric space, and amalgamated H_k over $\mathbb{Z}$ is also the parentheses is the fundamental group of a non-positively curved compact metric space. □

We are ready to prove our main result in this section.

Theorem 3.10. *Let Γ be a finite simplicial graph such that the flag complex on Γ is a 2-dimensional triangulated disk D. If $\dim_I(D) = 0$, then $\delta_{H_\Gamma}(n) \simeq n^2$.*

Proof. By Lemma 3.2, Γ is obtained by gluing fans and wheels, say k pieces of them. That is, let Γ_i be either a fan or a wheel and D_i be the flag complex on Γ_i, $i = 1, \cdots, k$. Define the graph $\Gamma_{1,\cdots,i+1}$ by gluing the flag complex $D_{1,\cdots,i}$ of $\Gamma_{1,\cdots,i}$ and D_{i+1} along a single 1-simplex on their boundaries. Then we have $\Gamma = \Gamma_{1,\cdots,k}$ and Lemma 3.7 gives the following decomposition:

$$H_\Gamma = ((((H_{\Gamma_1} *_{\mathbb{Z}} H_{\Gamma_2}) *_{\mathbb{Z}} H_{\Gamma_3}) *_{\mathbb{Z}} \cdots) *_{\mathbb{Z}} H_{\Gamma_{k-1}}) *_{\mathbb{Z}} H_{\Gamma_k}.$$

By Lemma 3.4, each H_{Γ_i} is a right-angled Artin group and non-hyperbolic. Therefore, it is the fundamental group of a non-positively curved compact metric space X_i, that is, the Salvetti complex associated to H_{Γ_i}. Lemma 3.9 tells us that H_Γ is also the fundamental group of a non-positively curved compact metric space X. This implies that the universal cover $\widetilde{X}$ of X is a CAT(0) space. Since H_Γ acts geometrically on $\widetilde{X}$, we have $\delta_{\widetilde{X}} \simeq \delta_{H_\Gamma}$. Thus, the fact that $\delta_{\widetilde{X}}$ is bounded above by a quadratic function gives us the quadratic upper bound of δ_{H_Γ}. Combining the natural quadratic lower bound, we obtain $\delta_{H_\Gamma}(n) \simeq n^2$. □

Chapter 4. Disk with Square Boundary

In this section, we prove that when the flag complex on a finite simplicial graph Γ is a triangulated disk D with square boundary, the Dehn function δ_{H_Γ} depends on the interior dimension of D.

Theorem 4.1. *Let Γ be a finite simplicial graph such that the flag complex on Γ is a 2-dimensional triangulated disk D with square boundary. If $\dim_I(D) = d$ for $d \in \{0, 1, 2\}$, then $\delta_{H_\Gamma}(n) \cong n^{d+2}$.*

The case of $\dim_I(D) = 0$ is a consequence of Theorem 3.10. For the case of $d = 1$ and $d = 2$, the lower bound is obtained by using the *height-pushing maps* introduced in [1], more details will be given in Section 4.1. For the upper bound, since Dison [20] gave an universal quartic upper bound for Dehn functions of Bestvina–Brady groups, we only need to establish the cubic upper bound for $d = 1$. This will be dealt with in Section 4.2.

When $d = 0$, some special cases of Theorem 4.1 can be recovered by a result of Carter and Forester [15]:

Lemma 4.2. ([15], Corollary 4.3) *Let Γ be a finite simplicial graph. Suppose Γ is a join of three graphs $\Gamma = \Gamma_1 * \Gamma_2 * \Gamma_3$, then δ_{H_Γ} is quadratic.*

Example 4.3. *Let Γ be a wheel with five vertices. Label the vertices as follows:*

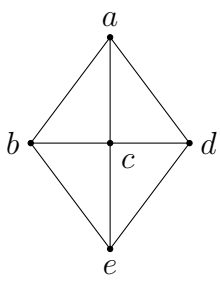

Let $\Gamma_1 = \{c\}$, $\Gamma_2 = \{b, d\}$, *and* $\Gamma_3 = \{a, e\}$, *then* $\Gamma = \Gamma_1 * \Gamma_2 * \Gamma_3$. *Therefore,* δ_{H_Γ} *is quadratic by Lemma 4.2.*

The rest of this chapter is devoted to the proof of Theorem 4.1.

4.1 Lower Bound

In [1], the authors introduced the *height-pushing maps* to obtain the lower bound on the Dehn functions of *orthoplex groups*. The same technique can be adapted here to obtain the lower bound. We only give necessary definitions here, more details and properties of height-pushing maps can be found in [1].

Let Γ be a finite simplicial graph. Let X_Γ be the Salvetti complex of A_Γ and $\widetilde{X}_\Gamma$ be the universal cover of X_Γ. Recall that the complex X_Γ is non-positively curved and $\widetilde{X}_\Gamma$ is a CAT(0) space. Let $Z_\Gamma = h^{-1}(0)$, where $h : \widetilde{X}_\Gamma \to \mathbb{R}$ is the height function. Recall that H_Γ acts geometrically on Z_Γ. Equip the space $\widetilde{X}_\Gamma$ with l^2-metric. Denote $B_r(x)$ to be the ball centered at x and whose radius is r, and $S_r(x)$ is the boundary of $B_r(x)$.

A subspace $F \subseteq \widetilde{X}_\Gamma$ is called a *flat* of dimension k if it is isometric to the Euclidean space $\mathbb{E}^k$.

Theorem 4.4. ([1], Theorem 4.2) *There is an* H*-equivariant retraction, called the* height-pushing map

$$\mathbf{P} : \widetilde{X}_\Gamma \setminus \cup_{v \notin Z_\Gamma} B_{1/4}(r) \to Z_\Gamma$$

such that when $\mathbf{P}$ *is restricted to* $h^{-1}([-t, t])$, $\mathbf{P}$ *is a* $(ct + c)$*-Lipschitz map, where* c *is a uniform constant which only depends on the defining graph* Γ.

Lemma 4.5. *Let* Γ *be a finite simplicial graph such that the flag complex on* Γ *is a* 2*-dimensional triangulated disk* D *with square boundary. If* $\dim_I(D) = d$, *then* $\delta_{H_\Gamma}(n) \succcurlyeq n^{d+2}$.

The proof of Lemma 4.5 is similar to the proof of Theorem 5.3 in [1] which does the case $d = 2$. Before proceeding the proof, we describe the main idea here. In order to establish the lower bound, we need to find a loop of length $\simeq n$ in $Z_\Gamma = h^{-1}(0)$ that has area at least n^{d+2}. We define a non-standard 2-dimensional flat $\overline{F}$ in $\widetilde{X}_\Gamma$ at non-negative height. At each height $r > 0$, the intersection $\overline{F} \cap h^{-1}(r)$ is homeomorphic to S^1. This intersection is a loop in $\widetilde{X}_\Gamma$ and it bounds a 2-disk $\overline{F} \cap h^{-1}([0, r])$ in $\overline{F}$. Translate the 2-disk $\overline{F} \cap h^{-1}([0, r])$ to $D_r := (a_0)^{-r}\overline{F} \cap h^{-1}([-r, 0])$ so that the loop $\overline{F} \cap h^{-1}(r)$ is translated to the loop $S_r := (a_0)^{-r}\overline{F} \cap h^{-1}(0)$ in Z_Γ. This loop S_r in Z_Γ has length $\simeq r$ and it bounds the 2-disk D_r in $(a_0)^{-r}\overline{F}$. That is, S_r has a natural filling D_r. Then we use the height-pushing map described in Theorem 4.4 to push the filling D_r of S_r to Z_Γ, and use that fact that the height-pushing map is a Lipschitz map and its Lipschitz constants grows linearly with respect to the height r. This gives us that the loop S_r has area at least r^{d+2} in Z_Γ. Thus, we get the desired lower bound.

Proof of Lemma 4.5. The case $d = 0$ follows from Theorem 3.10. We prove the case when $d = 1$ and the case $d = 2$ appears in [1]. Label the four vertices on the boundary as follows:

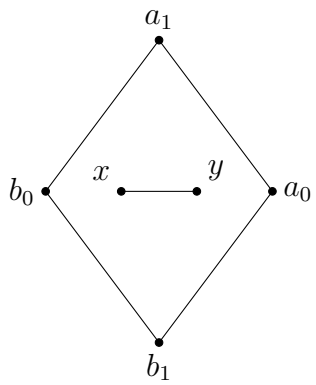

FIGURE 4.1. The square boundary of D and an interior 1-simplex σ.

Define a bi-infinite geodesic ray γ_i, $i = 0, 1$, in the 1-skeleton of $\widetilde{X}_\Gamma$ as follows:

$$\gamma_i(0) = e, \quad \gamma_i|_{\mathbb{R}^+} = a_i b_i a_i b_i \cdots, \quad \gamma_i|_{\mathbb{R}^-} = b_i a_i b_i a_i \cdots .$$

Define a map $F : \mathbb{Z} \times \mathbb{Z} \to A_\Gamma$ by

$$F(x) = \gamma_0(x_0)\gamma_1(x_1), \quad x = (x_0, x_1) \in \mathbb{Z} \times \mathbb{Z}.$$

The image of F consists of elements in the non-abelian group $\langle a_0, a_1, b_0, b_1 \rangle$. Let $\overline{F}$ be the non-standard 2-dimensional flat in $\widetilde{X}_\Gamma$ such that the set of its vertices is the image of F. Since

$$h(x) = h(F(x_0, x_1)) = |x_0| + |x_1|, \quad x = (x_0, x_1) \in \mathbb{Z} \times \mathbb{Z},$$

the flat $\overline{F}$ is at non-zero height and has a unique vertex $F(0, 0) = \gamma_0(0)\gamma_1(0) = e$ at height 0. We think of $\overline{F}$ as the boundary of a reversed infinite square pyramid with a unique vertex $F(0, 0)$ at height zero. For each $r > 0$, the intersection $\overline{F} \cap h^{-1}(r)$ is homeomorphic to a 1-sphere S^1. Note that this 1-sphere $\overline{F} \cap h^{-1}(r)$ bounds a 2-disk $\overline{F} \cap h^{-1}([0, r])$ in $\overline{F}$.

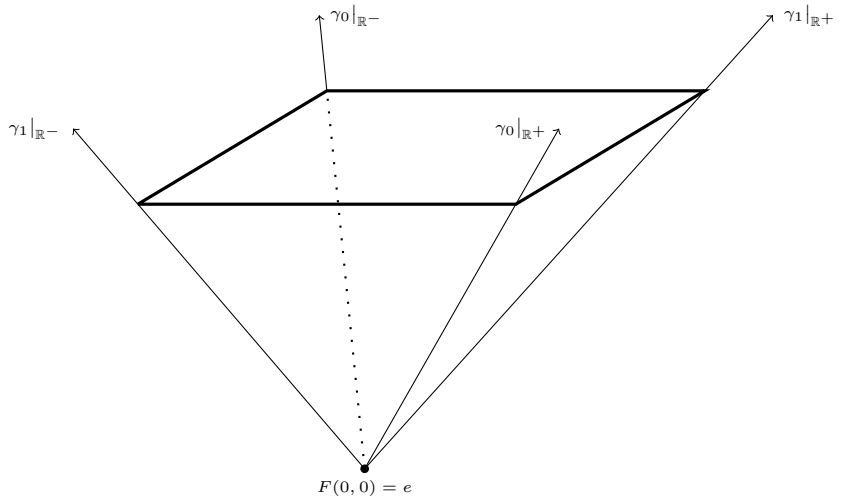

FIGURE 4.2. The non-standard 2-dimensional flat $\overline{F}$. The parallelogram drew in the thick lines is the intersection $\overline{F} \cap h^{-1}(r)$.

Using the group action A_Γ on $\widetilde{X}_\Gamma$, we translate the 2-disk $\overline{F} \cap h^{-1}([0,r])$ to a 2-disk $(a_0)^{-r}\overline{F} \cap h^{-1}([-r,0])$, denoted by D_r and its interior by $\mathring{D}_r$. Thus, the 1-sphere $\overline{F} \cap h^{-1}(r)$ at height r is translated to a 1-sphere S_r at height zero:

$$S_r := [(a_0)^{-r}\overline{F}] \cap h^{-1}(0) = [(a_0)^{-r}\overline{F}] \cap Z_\Gamma.$$

Notice that after performing this translation, the unique vertex $(a_0)^{-r}F(0,0)$ is at height $-r$. Now we have that S_r is a loop in Z_Γ and it bounds a 2-disk D_r. The next step is to use the height-pushing map to push this D_r to Z_Γ.

In order to use the height-pushing map, we need to do a surgery on D_r. At each vertex $v \in \mathring{D}_r$, replace $B_{1/4}(v)$ by the 2-dimensional flag complex D. We want to show that after applying the height-pushing map to all vertices of $\mathring{D}_r$, the scaled copies of interior 1-simplices do not intersect much in Z_Γ. If σ_1 and σ_2 are two different interior 1-simplices of D, then they are either disjoint or intersect at a 0-simplex. Let $\sigma = [x,y]$ be an interior 1-simplex of D, see Figure 4.1. We claim that after applying the height-pushing map $\mathbf{P}$ to D_r, the scaled copies of σ in Z_Γ are disjoint. Let v_1' and v_2' be vertices of scaled copies of σ in Z_Γ, based at vertices v_1 and v_2 in $\overline{F}$, respectively. Then we have

$$v_1' = (a_0)^{-r}v_1x^{s_1}y^{t_1} \quad \text{and} \quad v_2' = (a_0)^{-r}v_2x^{s_2}y^{t_2}$$

for some $s_1, t_2, s_2, t_2 \in \mathbb{Z}_{\geq 0}$. If $v_1' = v_2'$, then we have $v_1 = v_2$ since $\langle a_0, a_1, b_0, b_1\rangle \cap \langle x, y\rangle = \{0\}$. Thus, no vertices of different scaled copies of σ are the same. That is, scaled copies of σ in Z_Γ are disjoint.

Note that each interior 1-simplex is an intersection of two 2-simplices in the 2-dimensional flag complex D. For each interior 1-simplex σ in D based at a vertex v of D_r, it follows by Theorem 4.4 that the length of the scaled copied of σ in Z_Γ grows linearly in terms of r, that is, the length of the scaled copy depends on how far is σ pushed into Z_Γ by the height-pushing map. Thus, the further the σ is from

Z_Γ, the longer the length of its scaled copy. So, the σ in D based at the vertex $(a_0)^{-r+1}$ has the shortest length of scaled copy in Z_Γ, say $c(r-1)$, where c is a constant. Since there are $2r^2+2r+1$ vertices in $\mathring{D}_r$, we have

$$\text{Area}(S_r) \geq 2c(r-1)\cdot(2r^2+2r+1) \simeq r^3.$$

Thus, the filling function of Z_Γ is at least cubic. Hence, δ_{H_Γ} is at least cubic. This proves the case $d=1$. □

4.2 Upper Bound

Since there is a natural quartic upper bound, we only need to establish the upper bound for Γ satisfying $\dim_I(D)=1$. The following lemma shows that the shape of Γ is restricted in this case.

Lemma 4.6. *Let Γ be a finite simplicial graph such that the flag complex on Γ is a 2-dimensional triangulated disk D with square boundary. If $\dim_I(D)=1$, then Γ is the suspension of a path of length at least 3.*

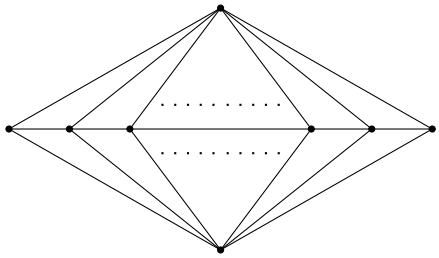

FIGURE 4.3. The suspension of a path of length at least 3.

Proof. Denote Γ' to be the graph consists of all the interior edges of D. Label the four vertices on ∂D by a,b,c,d counterclockwisely, starting from the right-most vertex in Figure 4.3. Since D is a triangulated disk, each interior edge of D is a common edge of two adjacent triangles of D. That is, for each interior edge of D there are exactly two vertices on ∂D such that this interior edge together with these two vertices on ∂D form two adjacent triangles of D. These two vertices on

∂D have to be non-adjacent. Otherwise, there would be a K_4 subgraph contained in Γ.

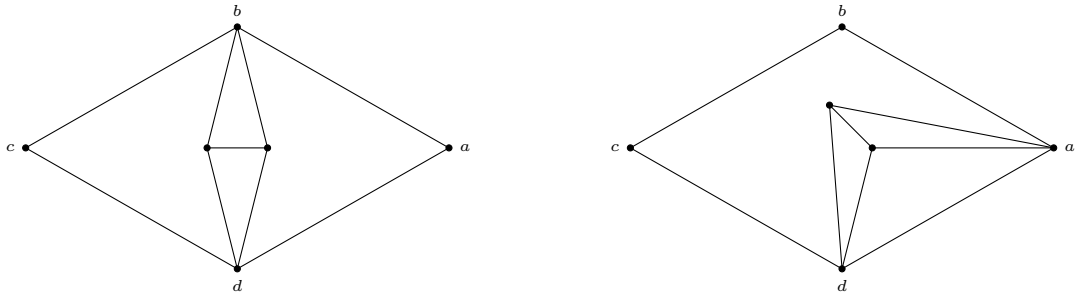

FIGURE 4.4. An interior edge together with vertices on ∂D form two adjacent triangles. These two vertices on ∂D cannot be adjacent, see the picture on the left. Otherwise, the interior edge with two adjacent vertices on ∂D form a K_4 subgraph of Γ, see the picture on the right.

Next, we prove two claims on Γ'. The first claim is that: if Γ' is a connected path, then all the pairs of adjacent vertices on Γ' are connected to vertices either a, c or b, d simultaneously. As we have seen in the previous paragraph, every pair of adjacent vertices on Γ' is connected to either vertices a, c or b, d. Pick any pair of adjacent vertices on Γ' and connect these two vertices to b, d, as shown in the following figure:

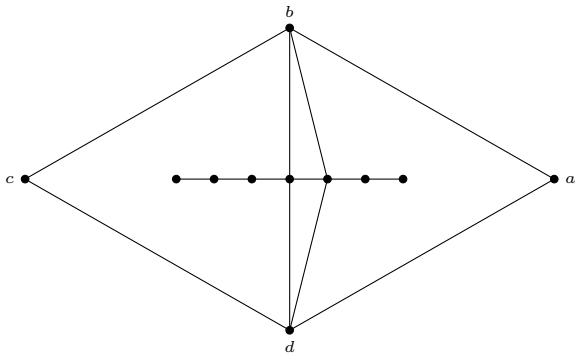

FIGURE 4.5.

It is obvious that other vertices on Γ' have to be connected to vertices b, d as well. This proves the first claim.

The second claim is that: if Γ' contains more than two edges, then Γ' has no vertices whose valency is ≥ 3. Suppose Γ' consists of three edges and they form a tripod, that is, Γ' has a vertex whose valency is 3. Each of the edges of Γ' together with two non-adjacent vertices on ∂D, say b, d, form two adjacent triangles. Pick any two edges of Γ' and connect the vertices to vertices b, d. Then the vertices of the third edge have to be connected to the same vertex, as shown in Figure 4.6. This contradicts the fact that Γ is a simplicial graph.

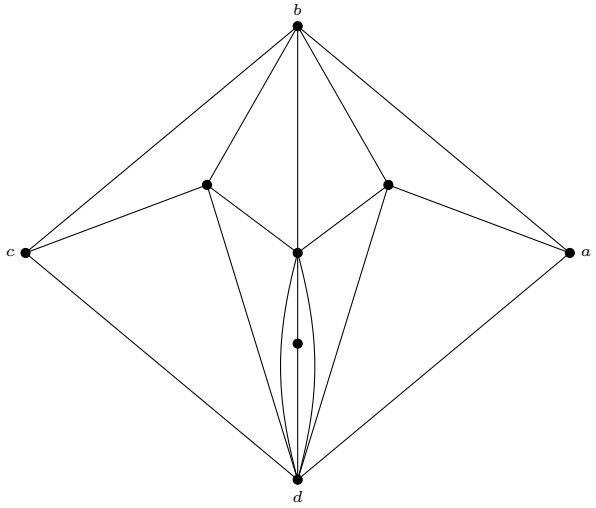

FIGURE 4.6. The case when Γ' is a tripod. There are two vertices which connect to the same vertex d.

Suppose Γ' contains more than three edges and has a vertex whose valency is greater than 3. Then Γ' contains a tripod and the above argument also shows that Γ would be a non-simplicial graph. This proves the second claim.

Now, we prove that Γ' is a connected path. Suppose Γ' is not connected, say Γ' has k connected components $\Gamma'_1, \cdots, \Gamma'_k$. Since none of vertices of Γ' has valency ≥ 3 by the second claim that we prove, each of $\Gamma'_1, \cdots, \Gamma'_k$ is a connected path. Also, all the pairs of adjacent vertices of $\Gamma'_1, \cdots, \Gamma'_k$ are connected either to vertices a, c or vertices b, d simultaneously, say vertices b, d:

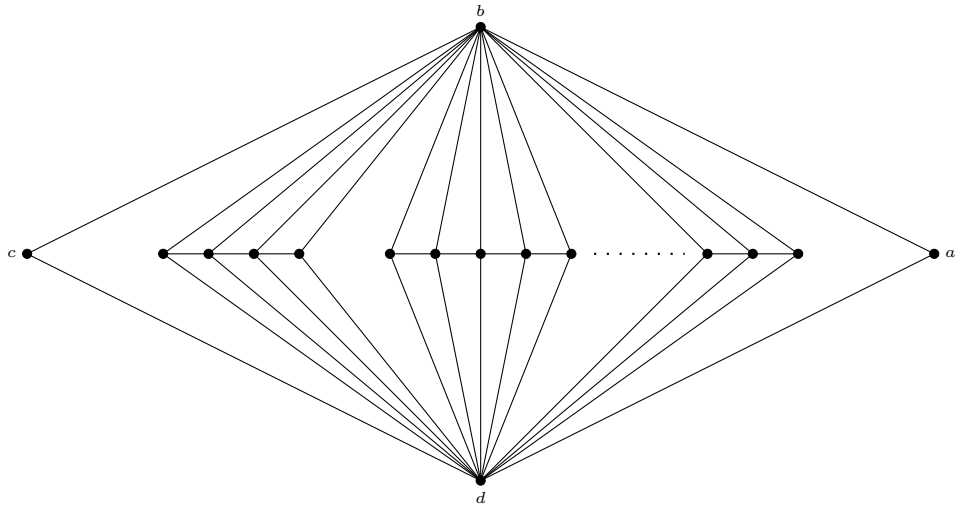

FIGURE 4.7.

In Figure 4.7, label from the left-most connected component to the right-most component of Γ' by $\Gamma'_1, \cdots, \Gamma'_k$. Denote Γ_i to be the join of Γ'_i and $\{b, d\}$. Since D is a triangulated disk, there must be an edge connecting b and d between Γ_i and Γ_{i+1}, $i = 1, \cdots k-1$:

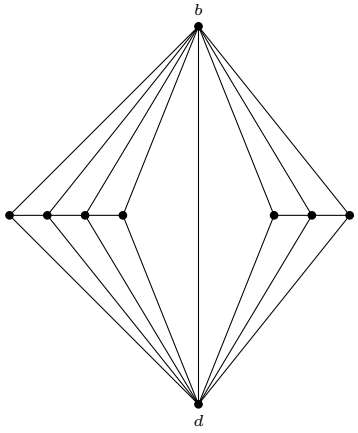

FIGURE 4.8. Between Γ_i and Γ_{i+1}, vertices b, d are connected by an edge.

In Figure 4.8 we see that connecting b, d by an edge creates K_4 subgraphs in Γ. This contradicts out assumption. Thus, the graph Γ' is a connected path.

Since Γ' is a connected path, it follows by the first claim that Figure 4.3 is the only possible scenario. This proves the lemma. □

To establish the cubic upper bound, we use *corridor schemes*. Here, we only give the definition in our setting, more details can be found in [8]. Let Γ be a finite simplicial graph such that the flag complex on Γ is a 2-dimensional triangulated disk D. A *corridor scheme* for Γ is a collection σ of labels of edges of Γ such that every triangle of D has either zero or two edges in σ. Given a van Kampen diagram Δ, a σ-*corridor* is a corridor in Δ that consists of triangles, each triangle has exactly two edges labeled by letters in σ and every pair of adjacent triangles intersect at an edge in σ. If α and β are two disjoint corridor schemes, then α-corridors and β-corridors never intersect.

Let Γ be the suspension of a path of length at least 3. Give an orientation to each edge of Γ so that we can write down the Dicks–Leary presentation for H_Γ. Label all the edges as Figure 4.9

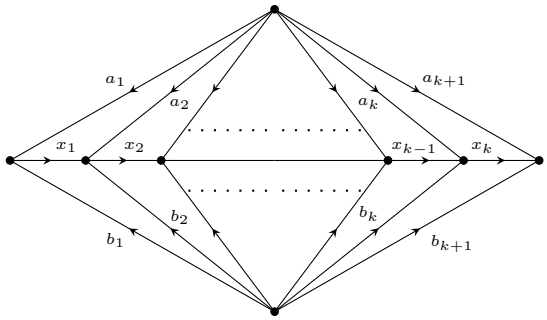

FIGURE 4.9. The suspension of a path of length at least 3 with orientation.

Let w be a word of length at most n that represents the identity in H_Γ and Δ a minimal van Kampen diagram for w. Choose two disjoint corridor schemes $\alpha = \{a_1, \cdots, a_{k+1}\}$ and $\beta = \{b_1, \cdots, b_{k+1}\}$ for Γ. Note that α-corridors and β-corridors never intersect. For each corridor in Δ, fix a vertex p of the corridor that is on $\partial\Delta$. When reading along the boundary of the corridor starting from p, we get

a boundary word $u'v''u''v'$ or its cyclic permutation, where u', u'' are both words consist of letters in α or β, and v', v'' are words in the free group $F(x_1, \cdots, x_k)$.

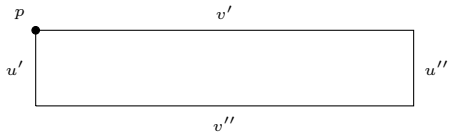

FIGURE 4.10. Boundary of a single corridor.

Definition 4.7. *Let C_i, $1 \leq i \leq h$, be a corridor in Δ with the boundary word $u_i'v_i''u_i''v_i'$ (or its cyclic permutation) where u_i', u_i'' are both words consist of letters in α or β and v_i', v_i'' are words in the free group $F(x_1, \cdots, x_k)$. If $u_1', \cdots, u_h'$ and $u_1'', \cdots, u_h''$ are two sets of consecutive letters on $\partial\Delta$ and $v_i'' = v_{i+1}'$ for $i = 1, \cdots h-1$, then the subdiagram T obtained by gluing corridors $C_1, \cdots, C_h$ along the words $v_1'' = v_2', \cdots, v_{h-1}'' = v_h'$ is called a* stack. *The shorter word of words v_1' and v_h'' is called the* top *of T; the longer word of the words v_1' and v_h'' is called the* bottom *of T; the number h is called the* height *of T. The words $u_1' \cdots u_h'$ and $u_1'' \cdots u_h''$ are called the* legs *of T.*

By Definition 4.7, a single corridor is a stack of height 1.

Roughly speaking, a stack is a pile of corridors where one corridor sits on top of another. The shape of a stack looks like an isosceles trapezoid, whose legs are parts of $\partial\Delta$. The height of a stack is the length of its legs. The bases of a stack, the top and the bottom, are not parts of $\partial\Delta$. The bases of a stack are words in the free group $F(x_1, \cdots, x_k)$. The top of a stack is the shorter base and the bottom of a stack is the longer base.

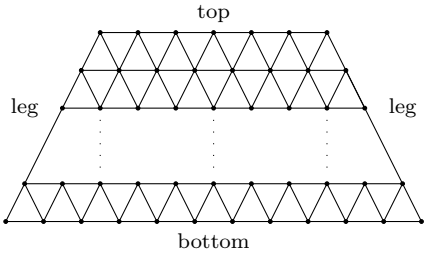

FIGURE 4.11. A simple example of a stack.

Definition 4.8. *Let C be either a α-corridor or a β-corridor with boundary word $u'v''u''v'$ (or its cyclic permutation), as shown in Figure 4.10. We say that a vertex q on the part of boundary of C labeled by the word v' (respectively v'') is a* j-vertex *or q has* type j *if there are $j+1$ edges connecting q to $j+1$ distinct consecutive vertices on the part of boundary of C labeled by the word v'' (respectively v').*

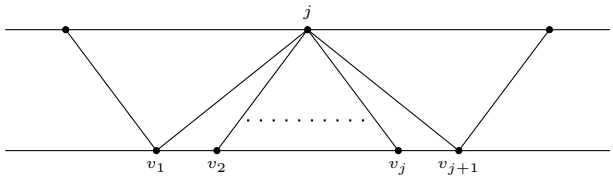

FIGURE 4.12. Picture of a j-vertex. The number j around the vertex indicates that the vertex is a j-vertex.

We say that a pair of adjacent vertices q_1 and q_2 in v' (respectively v'') generate *a vertex q_3 in v'' (respectively v') if there is a triangle shown as follows:*

FIGURE 4.13. A pair of adjacent vertices p_1 and p_2 generates w.

To get control of the area of w, we need to control the area of the corridors inside Δ. We now prove the main technical lemma.

Lemma 4.9. *Let Γ be the suspension of a path of a fixed length k, $k \geq 3$. Let $w \in H_\Gamma$ be a freely reduced word that represents the identity. Let T be a stack in a reduced van Kampen diagram Δ of w such that the top of T has length l and the height of T is h. Then the area of T satisfies*

$$Area(T) \leq C(lh^2 + h^3),$$

where C is a positive constant which does not depend on l and h.

Recall that we give an orientation and a labeling to each edge of Γ as shown in Figure 4.9, and we choose two disjoint corridor schemes $\alpha = \{a_1, \cdots, a_{k+1}\}$ and $\beta = \{b_1, \cdots, b_{k+1}\}$ for Γ. Denote the label of the top and the label of the bottom of T by t_0 and t_h, respectively. Denote the labels of the intermediate boundaries of the corridors of T by $t_1, \cdots, t_{h-1}$, see Figure 4.14.

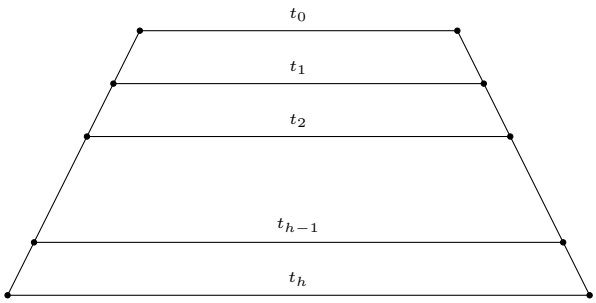

FIGURE 4.14. Stack with labeling of boundaries of corridors.

The proof of Lemma 4.9 relies on a series of lemmas, that analyze the the boundaries of corridors of a stack. In the proofs of the following lemmas, we use the fact that all the corridors and stacks are in a reduced van Kampen diagram Δ. In particular, for a fixed $k \geq 3$ in Lemma 4.9, each vertex on the boundary of a corridor in Δ is of type at most k.

Lemma 4.10. *Let C_i be either an α-corridor or a β-corridor in a stack T. Denote the top and the bottom of C_i by t_i and t_{i+1}, respectively. Then*

(1) The area of C_i is $|t_i| + |t_{i+1}|$.

(2) The length of t_{i+1} depends on the types of vertices on t_i. More precisely,

$$|t_{i+1}| = \sum_j j \cdot \big|\{j\text{-vertices on } t_i\}\big|.$$

Proof.

(1) Since each edge on t_i and t_{i+1} is part of a unique triangle, the result follows.

(2) Since each j-vertex on t_i contributes j edges on t_{i+1}, the statement follows immediately.

□

The following lemma shows how to recognize a 1-vertex in the boundary of a single corridor.

Lemma 4.11. *Let C be either a single α-corridor or a single β-corridor. The following are the possible combinations that will create a 1-vertex on the boundary of C, for a given m, $2 \leq m \leq k-1$:*

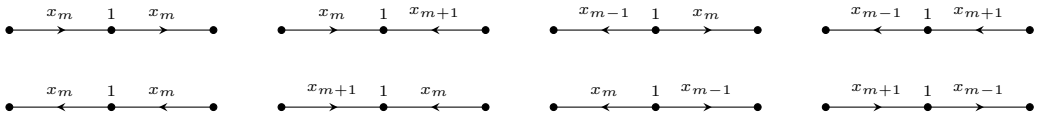

FIGURE 4.15. Possible combinations of 1-vertices.

Proof. We prove the case when C is a single α-corridor; the case for a single β-corridor is similar. To see the lemma, we only need to see how can we fill C. Take the left-most case in the first row of Figure 4.15:

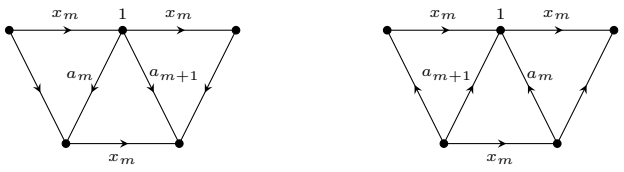

FIGURE 4.16.

The rest of the cases are similar. Thus, the lemma follows. □

Lemma 4.12. *Let C_i and C_{i+1} be two consecutive corridors in a stack T and let the arrows on the edges labeled by $u'_i, u''_i, u'_{i+1}, u''_{i+1}$ have the same orientation, as shown in the following picture. There are 2 cases. Assume that there are no 0-vertices on t_i. Then*

(1) All the vertices on t_{i+1} are 1-vertices, 2-vertices, or 3-vertices, except possibly the two vertices at the ends of t_{i+1}.

(2) All the vertices on t_{i+2} are either 1-vertices or 2-vertices, except possibly the two vertices at the ends of t_{i+1}.

(3) We have $|t_i| \leq |t_{i+1}| \leq |t_{i+2}|$.

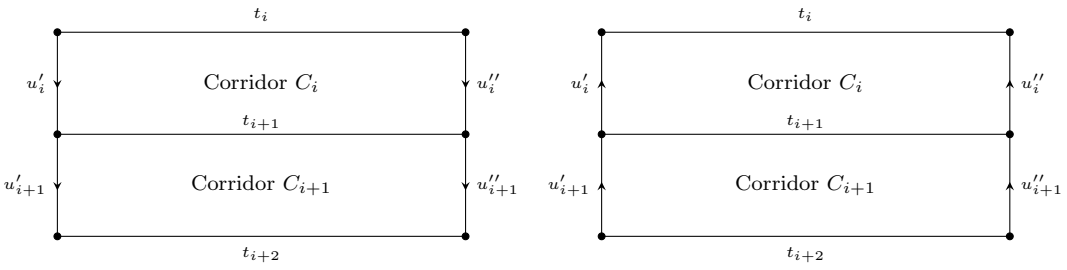

FIGURE 4.17. The figure shows the 2 cases in which u'_i, u''_i and u'_{i+1}, u''_{i+1} have the same orientation.

Proof. We prove the case when C_i, C_{i+1} are α-corridors and $u'_i, u''_i, u'_{i+1}, u''_{i+1}$ are pointing away from t_i. Other cases are similar.

(1) First, we claim that every pair of adjacent vertices on t_i generates either a 1-vertex or a 3-vertex on t_{i+1}. We prove this claim by showing all the eight possibilities of the combinations of labels on t_{i+1} that create a 1-vertex and a 3-vertex:

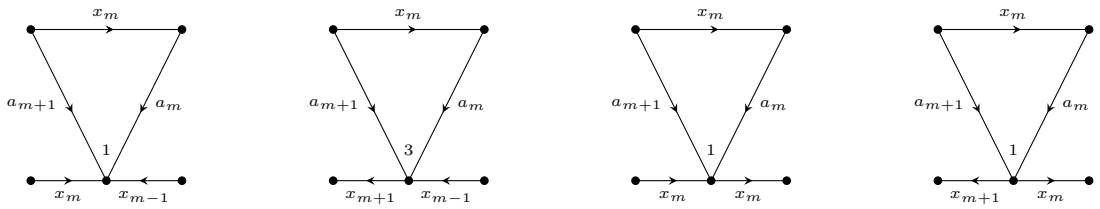

FIGURE 4.18. A pair of adjacent vertices on t_i generates either a 1-vertex or a 3-vertex on t_{i+1}.

The 1-vertices in Figure 4.18 are recognized by Lemma 4.11. The following picture shows that the vertex in Figure 4.18 is a 3-vertex:

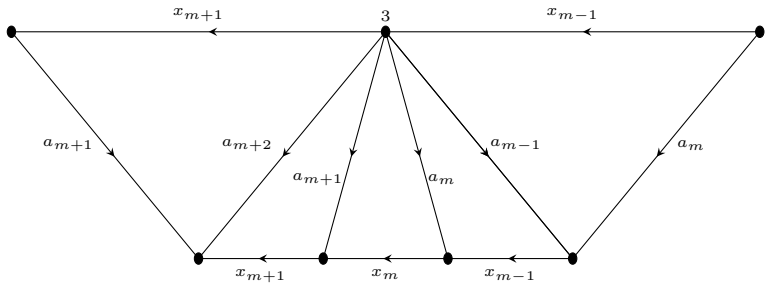

FIGURE 4.19.

The other four cases are obtained by flipping the pictures in Figure 4.18 such that the arrow on the edge labeled by x_m on t_i has the opposite orientation. This proves the claim.

Next, we show that the vertices on t_{i+1} that are not generated by pairs of adjacent vertices on t_i are 2-vertices. We claim that each j-vertex on t_i create $j-1$ vertices on t_{i+1} that are all 2-vertices. For a j-vertex on t_i, we have the following local pictures:

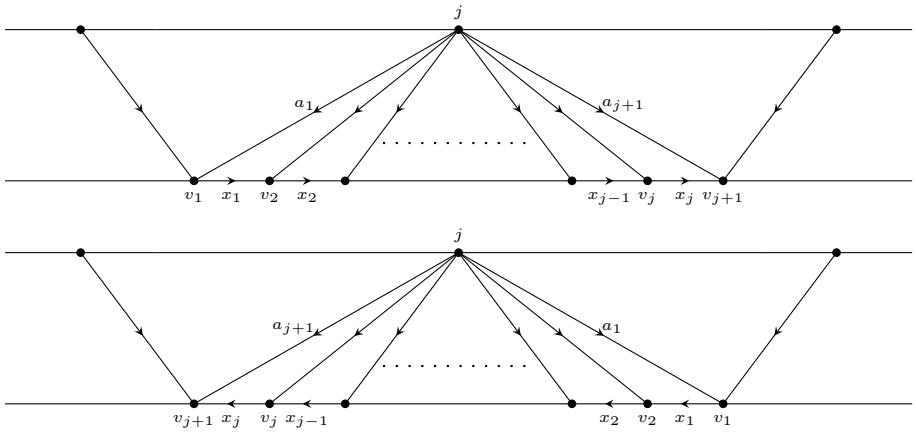

FIGURE 4.20. The figures show that each j-vertex on t_i create $j-1$ vertices on t_{i+1} that are all 2-vertices.

Note that each of v_1 and v_{j+1} is either either a 1-vertex or a 3-vertex as we showed in the previous paragraph. Also, note that for $m = 2, \cdots, j$, the vertex v_m is the terminal vertex of an edge labeled by x_{m-1} and the initial vertex of an edge labeled by x_m. These vertices $v_2, \cdots, v_j$ are 2-vertices, as shown in the following pictures:

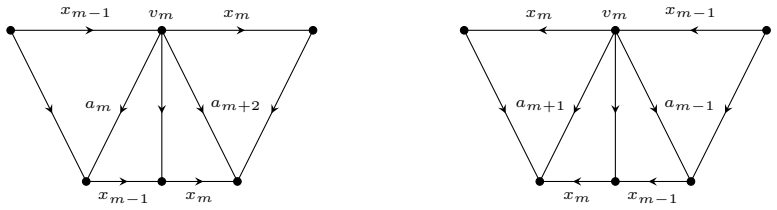

FIGURE 4.21. Vertex v_m on t_{i+1} is a 2-vertex.

This proves the claim.

(2) In (1) we have shown that all the vertices on t_{i+1} are 1-vertices, 2-vertices, or 3-vertices. Each of the vertices on t_{i+1} creates numbers of 0, 1, or 2 vertices on t_{i+2}, respectively, and they are all 2-vertices. The rest of the vertices on t_{i+2} are generated by pairs of adjacent vertices on t_{i+1}. We claim that these

vertices on t_{i+2} are 1-vertices. We prove the claim by showing all the possible combinations. In the following pictures, all the 1-vertices are recognized by Lemma 4.11.

The following pictures show that a pair of adjacent 1-vertices on t_{i+1} generates a 1-vertex on t_{i+2}:

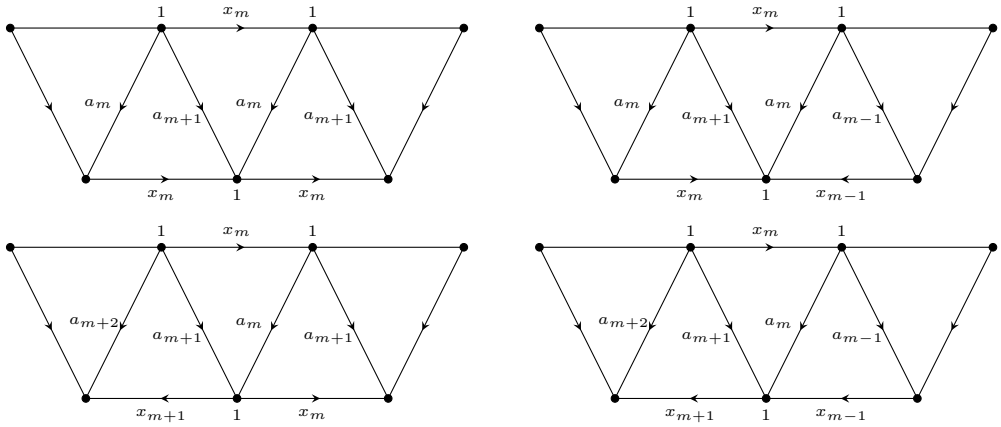

FIGURE 4.22. Two consecutive 1-vertices on t_i generate a 1-vertex on t_{i+1}.

Next, the following pictures show that a pair consisting of a 1-vertex adjacent to a 2-vertex on t_{i+1} generates a 1-vertex on t_{i+2}. Note that as we have seen in the proof of (1), there is only one combination that can give a 2-vertex on t_{i+1}, up to orientation, and there are two possibilities for a 1-vertex.

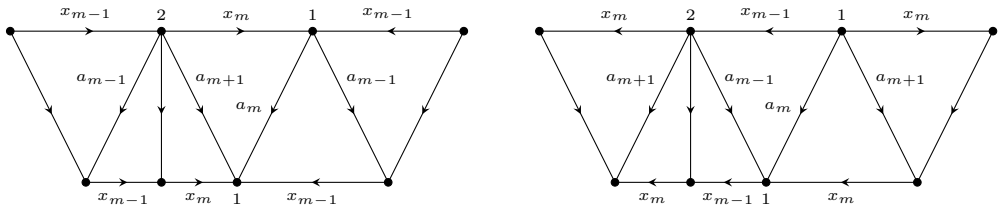

FIGURE 4.23. A pair of adjacent 1-vertex and 2-vertex on t_{i+1} generates a 1-vertex on t_{i+2}.

For adjacent 2-vertices on t_{i+1}, we have one possibility up to orientation:

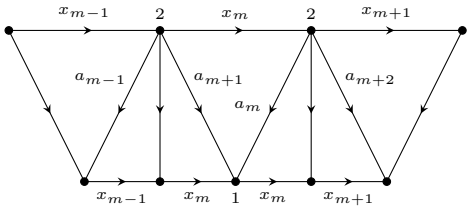

FIGURE 4.24. A pair of adjacent 2-vertices on t_{i+1} generates a 1-vertex on t_{i+2}.

For a pair consisting of a 1-vertex adjacent to a 3-vertex on t_{i+1}, we have

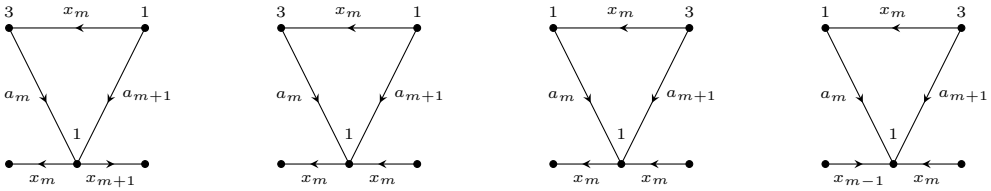

FIGURE 4.25. A pair consisting of a 1-vertex adjacent to a 3-vertex on t_{i+1} generates a 1-vertex on t_{i+2}.

For a pair consisting of a 2-vertex adjacent to a 3-vertex on t_{i+1}, since the number of combinations which will create 2-vertices and 3-vertices are limited, we have fewer possible cases here:

FIGURE 4.26. A pair of adjacent 2-vertex and 3-vertex on t_{i+1} generates a 1-vertex on t_{i+2}.

Finally, for a pair of adjacent 3-vertices, we have one possibility up to orientation

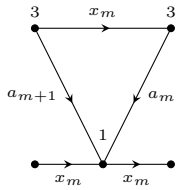

FIGURE 4.27. A pair of adjacent 3-vertex on t_{i+1} generates a 1-vertex on t_{i+2}.

This proves (2).

(3) Since there are no 0-vertices on t_i, by the assumption in the statement of the lemma, the number of vertices on t_{i+1} is no less than the number of vertices on t_i; the number of vertices on t_{i+2} is also no less than the number of vertices on t_{i+1}. Thus, we have

$$\begin{aligned} |t_i| &= (\text{number of vertices on } t_i) - 1 \\ &\leq (\text{number of vertices on } t_{i+1}) - 1 \\ &= |t_{i+1}| \end{aligned}$$

and

$$\begin{aligned} |t_{i+1}| &= (\text{number of vertices on } t_{i+1}) - 1 \\ &\leq (\text{number of vertices on } t_{i+2}) - 1 \\ &= |t_{i+2}|. \end{aligned}$$

Hence, $|t_i| \leq |t_{i+1}| \leq |t_{i+2}|$. A similar argument gives $|t_i| \leq |t'_{i+1}| \leq |t''_{i+2}|$.

□

Lemma 4.13. *Let C_i and C_{i+1} be two consecutive corridors in a stack T as in the Lemma 4.12. Let the arrows on the edges labeled by u'_i, u''_i and u'_{i+1}, u''_{i+1} have different orientations. There are two cases, as shown in Figure 4.28. Assume that*

there are no 0*-vertices on* t_i*. Then* $|t'_{i+1}| \leq |t_{i+2}|$*, where* t_{i+2} *is as shown in Figure 4.17.*

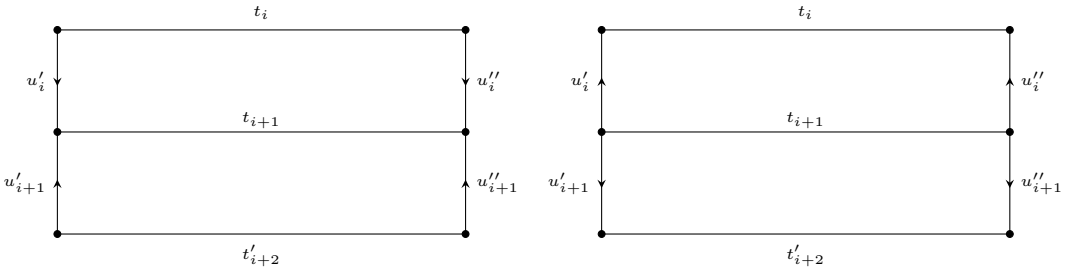

FIGURE 4.28. u'_i, u''_i and u'_{i+1}, u''_{i+1} have different orientations.

Proof. We prove the case when $u'_i, u''_i, u'_{i+1}, u''_{i+1}$ are pointing toward t_{i+i}. The other case is similar. We claim that in this case, C_i and C_{i+1} cannot be both α-corridors or both β-corridors. If C_i and C_{i+1} are both α-corridors (respectively β-corridors), then there are two triangles that share a unique common edge on t_{i+1}:

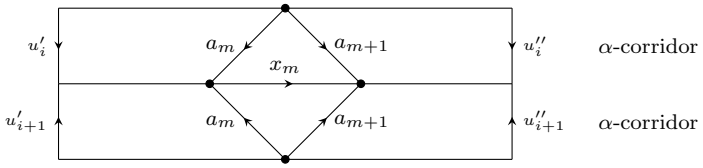

FIGURE 4.29. Two triangles from different α-corridors with opposite orientations.

This contradicts the fact that T is part of a minimal van Kampen diagram since we can obtain a smaller van Kampen diagram by canceling the two triangle as shown in Figure 4.29. This proves the claim.

The only possible case is when C_i and C_{i+1} are not both α-corridors or both β-corridors, say C_i is an α-corridor and C_{i+1} is a β-corridor. For each j-vertex on t_i together with its adjacent vertices, there are $j+1$ vertices on t_{i+1}:

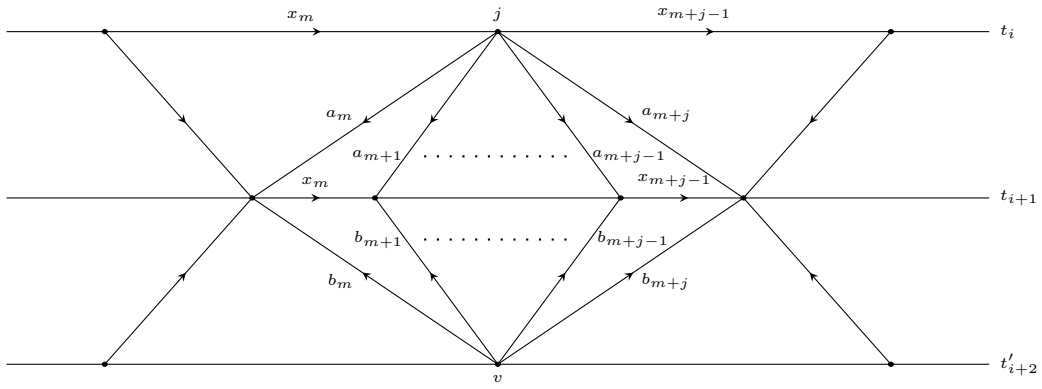

FIGURE 4.30. The j at the very top indicates the vertex is a j-vertex. Every pair of adjacent vertices on t_{i+1} generates the same vertex v on t'_{i+2}.

From Figure 4.30 we see that every pair of adjacent vertices of these $j+1$ vertices on t_{i+1} generates the same vertex on t'_{i+2}. That is, each j-vertex on t_i creates $j-1$ vertices on t_{i+1} and they are 0-vertices. Since there are no 0-vertices on t_{i+1} by Lemma 4.12, the number of vertices on t'_{i+2} is less than the numbers of vertices on t_{i+2} in Lemma 4.12. Thus, we have

$$\begin{aligned} |t'_{i+2}| &= (\text{number of vertices on } t'_{i+2}) - 1 \\ &\leq (\text{number of vertices on } t_{i+2}) - 1 \\ &\leq |t_{i+2}| \end{aligned}$$

This completes the proof. □

Now, we prove Lemma 4.9.

Proof of Lemma 4.9. Label the boundaries of corridors of the stack T as follows:

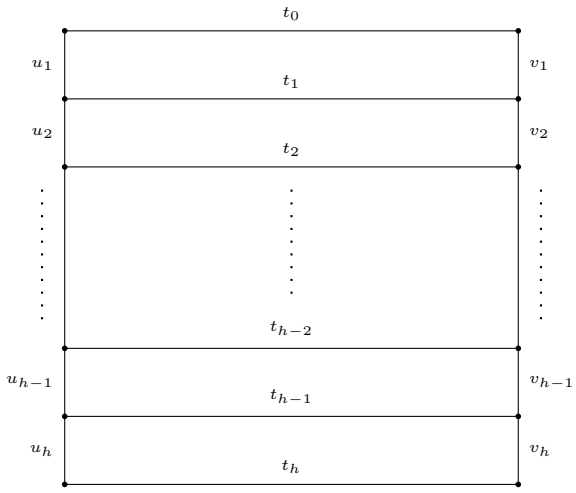

FIGURE 4.31. Stack T.

Recall that $t_0, \cdots, t_h$ are words in $F(x_1, \cdots, x_k)$ and u_i, v_i are both letters in the same corridor schemes $\alpha = \{a_1, \cdots, a_{k+1}\}$ or $\beta = \{b_1, \cdots, b_{k+1}\}$. Since the area of a stack T is the sum of the areas of the corridors that contained in T, by Lemma 4.10 we have:

$$\text{Area}(T) = |t_0| + 2|t_1| + \cdots + 2|t_{h-1}| + |t_h| \leq 2\left(|t_0| + \cdots + |t_h|\right).$$

Thus, to get the largest possible area of T, we need the longest possible length of the word t_i for each $i = 1, \cdots, h$. From Lemma 4.10 we know that the length of t_i depends on the number of vertices and the types of vertices on t_{i-1}. We may assume that all the vertices on t_0 are k-vertices, and this assumption gives the longest possible length of the word t_1. We may also assume that the arrows on the edges labeled by u_1, v_1 are pointing away from t_0. In order to have the longest possible length of the word t_2, we assume that the arrows on the edges labeled by u_2, v_2 are pointing away from t_1. Otherwise, the word t_2 would be shorter by Lemma 4.13. We also assume that the two vertices at the two ends of t_1 are k-vertices. Continue this process, we may assume that the arrows on the edges labeled by u_i, v_i are pointing away from t_{i-1}, and the vertices at the two ends of t_i are k-

vertices. These two assumptions give the longest possible length of the word t_i for $i = 1, \cdots, h$.

On t_0, since we assume that all the vertices are k-vertex, we have $|t_1| = k(l+1)$.

On t_1, there are two k-vertices at the two ends. Lemma 4.12 tells us that every pair of adjacent vertices on t_0 generates either a 1-vertex or a 3-vertex on t_1. But we may assume that every pair of adjacent vertices on t_0 generates a 3-vertex on t_1 so that we can get the largest possible $|t_2|$. So the number of 3-vertices on t_1 is l. Other vertices on t_1 are 2-vertices and there are $(k-1)(l+1)$ of them. Note that the total number of vertices on t_1 is $kl + k + 1$, which matches $|t_1| + 1$. Knowing the types of vertices on t_1 gives the length of t_2:

$$|t_2| = 2 \cdot k + l \cdot 3 + (k-1)(l+1) \cdot 2 = 2kl + l + 4k - 2$$

On t_2, there are two k-vertices at the two ends and each of these two k-vertices creates $(k-1)$ 2-vertices on t_3. By Lemma 4.12, every pair of adjacent vertices on t_1 generates a 1-vertex on t_2. So the number of 1-vertices on t_2 is $|t_1|$. Since each 2-vertex on t_1 creates a 2-vertex on t_2 and each 3-vertex on t_1 creates a 2-vertex on t_2, the number of 2-vertices on t_2 is

$$2 \cdot (k-1) + l \cdot 2 + (l+1) \cdot (k-1) = kl + l + 3k - 3.$$

Thus, the total number of vertices on t_2 is

$$2 + |t_1| + kl + l + 3k - 3 = 2kl + l + 4k - 1,$$

which matches $|t_2| + 1$. We have the length of t_3:

$$|t_3| = 2 \cdot k + (kl + l + 3k - 3) \cdot 2 + |t_1| \cdot 1 = 3kl + 2l + 9k - 6.$$

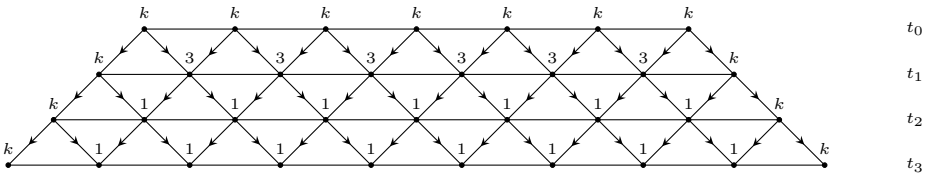

FIGURE 4.32. Vertices and their types on t_0, t_1, t_2, t_3. Starting from t_1, vertices lie between pairs of 1-vertices, pairs of 3-vertices, pairs of 1-vertex and k-vertex, pairs of 3-vertex and k-vertex, are all 2-vertices.

Now, for $i \geq 2$, every pair of adjacent vertices on t_{i-1} generates a 1-vertex on t_i by Lemma 4.12, so the number of 1-vertices on t_i is $|t_{i-1}|$. Every 2-vertex on t_{i-1} creates a 2-vertex on t_i and two k-vertices at the ends of t_{i-1} creates $k-1$ 2-vertices on t_i. So the number of 2-vertices on t_i is the number of 2-vertices on t_{i-1} plus $2(k-1)$:

$$(kl + l + k - 1) + (i-1) \cdot 2(k-1) = kl + l + (2i-1)k - 2i + 1.$$

Note that the two vertices at the ends of t_i are k-vertices by assumption. Having the information of the vertices on t_i we get

$$\begin{aligned} |t_{i+1}| &= |t_{i-1}| \cdot 1 + [kl + l + (2i-1)k - 2i + 1] \cdot 2 + 2 \cdot k \\ &= |t_{i-1}| + 2kl + 2l + 4ik - 4i + 2 \end{aligned}$$

for $i = 2, \cdots, h-1$. Let

$$d(i) = |t_{i+1}| - |t_{i-1}| = 2kl + 2l + 4ik - 4i + 2, \quad i = 2, \cdots, h-1,$$

then $\{d(i)\}$ is an arithmetic sequence whose difference is $4k-4$ and

$$d(i+2) - d(i) = 8k - 8.$$

When i is even, we have

$$\begin{aligned}
|t_{i+1}| &= |t_1| + d(2) + d(4) + \cdots + d(i) \\
&= |t_1| + \left[\frac{i}{2} \cdot \frac{d(2) + d(i)}{2}\right] \\
&= |t_1| + \frac{i}{4}\left[2d(2) + \left(\frac{i}{2} - 1\right)(8k - 8)\right] \\
&= |t_1| + \frac{i}{2}d(2) + i(i-2)(k-1).
\end{aligned}$$

When i is odd, we have

$$\begin{aligned}
|t_{i+1}| &= |t_2| + d(3) + d(5) + \cdots + d(i) \\
&= |t_2| + \left[\frac{i-1}{2} \cdot \frac{d(3) + d(i)}{2}\right] \\
&= |t_2| + \frac{(i-1)}{4}\left[2d(3) + \left(\frac{i-3}{2}\right)(8k - 8)\right] \\
&= |t_2| + \left(\frac{i-1}{2}\right)d(3) + (i-1)(i-3)(k-1).
\end{aligned}$$

When h is odd, we have

$$\text{Area}(T) \leq 2\sum_{i=1}^{h} |t_i| = 2\left(|t_0| + |t_1| + |t_2| + \sum_{\substack{i=2 \\ i \text{ is even}}}^{h-1} |t_{i+1}| + \sum_{\substack{i=3 \\ i \text{ is odd}}}^{h-2} |t_{i+1}|\right).$$

Furthermore,

$$\begin{aligned}
\sum_{\substack{i=2\\ i \text{ is even}}}^{h-1} |t_{i+1}| &= \sum_{\substack{i=2\\ i \text{ is even}}}^{h-1} \left[|t_1| + \frac{i}{2} d(2) + i(i-2)(k-1) \right] \\
&= \sum_{\substack{i=2\\ i \text{ is even}}}^{h-1} |t_1| + \left[\frac{d(2)}{2} - 2(k-1) \right] \sum_{\substack{i=2\\ i \text{ is even}}}^{h-1} i + (k-1) \sum_{\substack{i=2\\ i \text{ is even}}}^{h-1} i^2 \\
&= \left(\frac{h-1}{2} \right) |t_1| + \left[\frac{d(2)}{2} - 2(k-1) \right] \left(\frac{h^2-1}{4} \right) + (k-1) \frac{h(h^2-1)}{24} \\
&\leq h|t_1| + [d(2) - 2(k-1)]\, h^2 + (k-1)h^3 \\
&= hk(l+1) + (2kl + 2l + 6k - 4)h^2 + (k-1)h^3 \\
&\leq hk(l+1) + (2kl + 2l + 6k)h^2 + kh^3 \\
&\leq 2klh + 10klh^2 + kh^3 \\
&\leq 12klh^2 + kh^3.
\end{aligned}$$

and

$$\begin{aligned}
\sum_{\substack{i=3\\ i \text{ is odd}}}^{h-2} |t_{i+1}| &= \sum_{\substack{i=3\\ i \text{ is odd}}}^{h-2} \left[|t_2| + \left(\frac{i-1}{2} \right) d(3) + (i-1)(i-3)(k-1) \right] \\
&= \left(\frac{h-3}{2} \right) |t_2| + \frac{d(3)}{2} \sum_{\substack{i=3\\ i \text{ is odd}}}^{h-2} (i-1) + (k-1) \sum_{\substack{i=3\\ i \text{ is odd}}}^{h-2} (i^2 - 4i + 3) \\
&= \left(\frac{h-3}{2} \right) |t_2| + \frac{d(3)}{8}(h-3)(h-1) \\
&\qquad + (k-1) \left[\frac{1}{24}(h-3)(h-2)(h-1) - (h-3)(h+1) + \frac{3}{2}(h-3) \right] \\
&\leq h|t_2| + d(3)h^2 + kh^3 + kh^2 + 2h \\
&= h(2kl + l + 4k - 2) + (2kl + 2l + 12k - 10)h^2 + kh^3 + kh^2 + 2h \\
&= (2kl + l + 4k)h + (2kl + 2l + 13k - 10)h^2 + kh^3 \\
&\leq 7klh^2 + 17klh^2 + kh^3 \\
&= 24klh^2 + kh^3.
\end{aligned}$$

Thus,

$$\begin{aligned}\text{Area}(T) &\leq 2\left[|t_0| + |t_1| + |t_2| + (12klh^2 + kh^3) + (24klh^2 + kh^3)\right] \\ &= 2\left[l + k(l+1) + 2kl + l + 4k - 2 + 36klh^2 + 2kh^3\right] \\ &\leq (6kl + 4l + 10k) + 72klh^2 + 4kh^3 \\ &\leq 92k(lh^2 + h^3).\end{aligned}$$

When h is even, the computation is similar. Hence, $\text{Area}(T) \leq C(h^3 + lh^2)$, where C is a positive constant which does not depend on l and h. □

Lemma 4.14. *Let Γ be the suspension of a path of length at least 3 with labeling and a given orientation as shown in Figure 4.9. Then $\delta_{H_\Gamma}(n) \preceq n^3$.*

Proof. Let w be a freely reduced word of length at most n that represents the identity in H_Γ. Let Δ be a reduced van Kampen diagram for w such that $\text{Area}(w) = \text{Area}(\Delta)$. Choose corridor schemes $\alpha = \{a_1, \cdots, a_{k+1}\}$ and $\beta = \{b_1, \cdots, b_{k+1}\}$ for Γ. The van Kampen diagram Δ is cut up by α-corridors and β-corridors and some of the corridors form stacks whose heights are greater than 1. Recall that a single corridor is a stack of height 1. Thus, the diagram Δ is cut up by stacks.

There are three possible cases that the van Kampen diagram Δ could be. The first case is that every stack has a base that is part of $\partial\Delta$, as shown in Figure 4.33. The second case is that at least one of the stacks such that parts of its bases are on $\partial\Delta$, as shown in Figure 4.34. The third case is that at least one of the stacks whose bases are not part of $\partial\Delta$, as shown in Figure 4.35. Note that the following three pictures are only local diagrams of Δ.

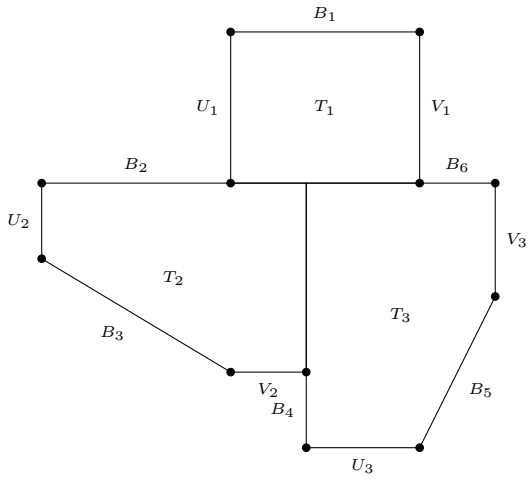

FIGURE 4.33. Case 1: Every stack has a base as a part of $\partial\Delta$.

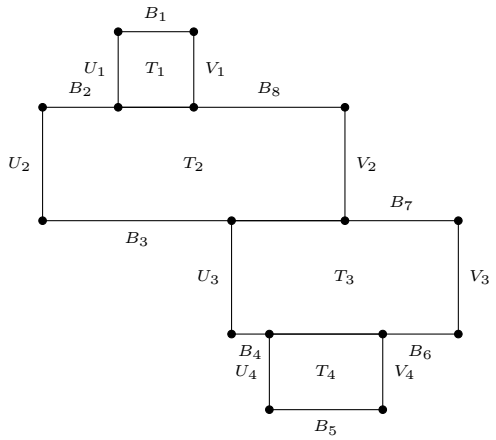

FIGURE 4.34. Case 2: At least one of the stacks such that part of its bases is part of $\partial\Delta$.

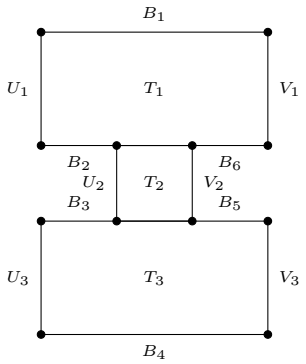

FIGURE 4.35. Case 3: At least one of the stacks whose bases are not part of $\partial\Delta$.

Denote the stacks by $T_1, \cdots, T_m$ that cut up Δ; denote the two legs of T_j by U_j and V_j, and the height of T_j by $|U_j| = |V_j| = h_j$, $j = 1, \cdots, m$. Note that U_j and V_j are words on $\partial\Delta$ that consist of letters in the corridor schemes α and β. Now, the boundary word w of Δ is a cyclic of the following word, and we also denote it by w;

$$w = A_1 B_1 A_2 B_2 \cdots A_r B_s,$$

where each of the words $A_1, \cdots, A_r$ consists of one or more words from $U_1, \cdots, U_m$ and $V_1, \cdots, V_m$; each of the words $B_1, \cdots, B_s$ is a word in the free group $F(x_1, \cdots, x_k)$. The length of each of the words $A_1, \cdots, A_r$ is either the height of a single stack or the sum of the heights of multiple stacks. Each of the words $B_1, \cdots, B_s$ is a base of a stack or part of a base of a stack. Denote the length of the words $B_1, \cdots, B_s$ by $l_1, \cdots, l_s$. Let

$$h = |A_1| + \cdots + |A_r| = 2(h_1 + \cdots + h_m) = 2\sum_{i=1}^{m} h_i$$

and

$$l = |B_1| + \cdots + |B_s| = \sum_{i=1}^{s} l_i.$$

Consider a stack T' whose top is the word $B_1 \cdots B_s$ of length l and whose legs are $U_1 \cdots U_m$ and $V_1 \cdots V_m$ and the height $h = |U_1 \cdots U_m| = |V_1 \cdots V_m|$, assuming that all the the arrows on the edges labeled by U_i, V_i are pointing away from the top. We claim that the area of Δ is less than the area of the stack T' in all three cases.

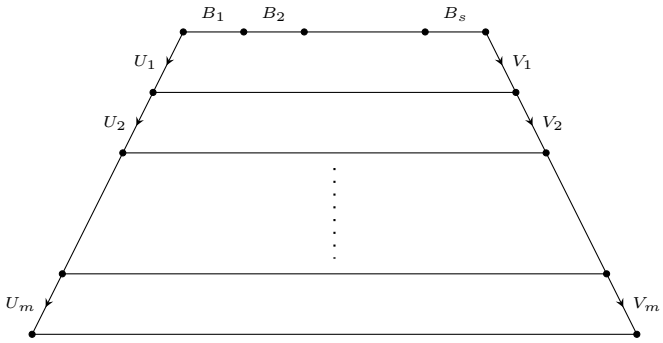

FIGURE 4.36. The stack T with labels on the top and the legs.

For the first case, recall that each of A_i is a leg of one or more stacks; each B_i is part of a base of a stack or a base of a stack. By assumption, each stack T_i has a base that is on $\partial\Delta$, say B_i is a base of the stack T_i. If B_i is the top of the stack T_i, then the top of the stack T_i' in Figure 4.37 is longer than B_i. Thus, $\text{Area}(T_i) \leq \text{Area}(T_i')$. If B_i is the bottom of the stack T_i, then consider a stack whose heights are U_i, V_i and whose top is B_i. The area of this stack is obviously greater than the area of T_i, but less than the area of T_i' in Figure 4.37. Thus, for each stack T_i, $j = 1, \cdots, m$, there is a stack T_i' in T' satisfying $\text{Area}(T_j) \leq \text{Area}(T_j')$. These substacks $T_1', \cdots, T_m'$ are disjoint inside T' because their legs are disjoint. Each substack T_j' is at different height inside T', as shown in Figure 4.37:

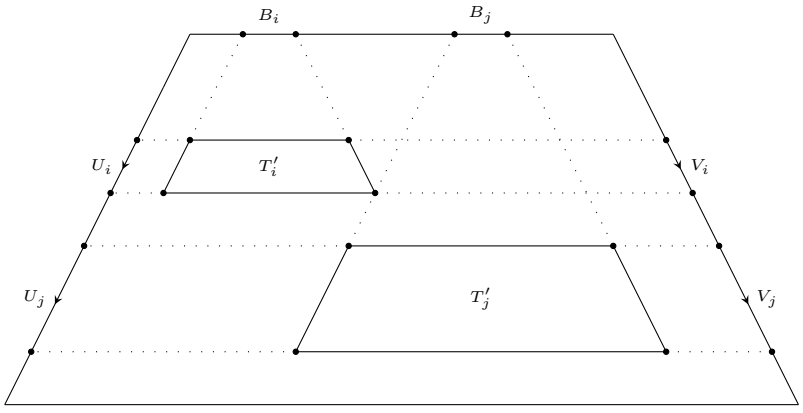

FIGURE 4.37. Disjoint substacks T_i' and T_j' in T'.

We have

$$\begin{aligned}
\text{Area}(w) &= \text{Area}(\Delta) \\
&= \text{Area}(T_1) + \cdots + \text{Area}(T_m) \\
&\leq \text{Area}(T_1') + \cdots + \text{Area}(T_m') \\
&\leq \text{Area}(T') \\
&\leq C(h^3 + lh^2) \\
&\leq 2Cn^3
\end{aligned}$$

for some positive constant C which does not depend on $|w| = n$. The second last inequality follows by Lemma 4.9 and the last inequality holds since $l \leq n$ and $h \leq n$. Thus, we prove the claim for the first case.

For the second case, suppose the van Kampen diagram Δ contains Figure 4.34 as a subdiagram; we denote the subdiagram by Δ'. Divide Δ' as shown in Figure 4.38:

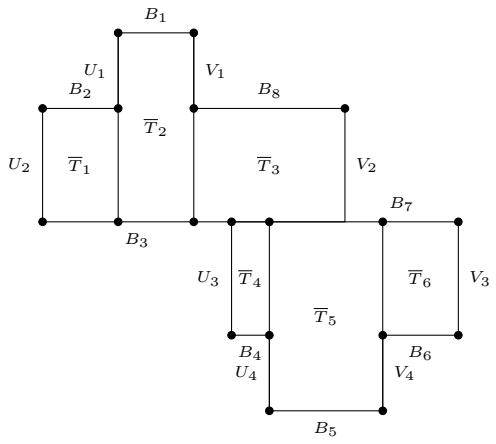

FIGURE 4.38. Divide the subdiagram Δ' in Δ.

The subdiagram Δ' is the same as the subdiagram Figure 4.34, but with a different of stacks. It consists of stacks $\overline{T}_1, \cdots, \overline{T}_6$ and each of these stacks has at least one base on $\partial\Delta$. Thus, it follows from the first that the claim holds for the second case.

For the third case, suppose the van Kampen diagram Δ contains Figure 4.35 as a subdiagram, and is denoted by Δ''. Similar to the second case, we divide Δ'' into substacks and each of the substacks has a base that is part of Δ, as shown in Figure 4.39:

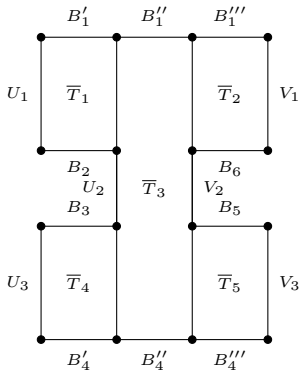

FIGURE 4.39. Divide the subdiagram Δ'' in Δ.

The words $B_1'B_1''B_1'''$ and $B_4'B_4''B_4'''$ in Figure 4.39 are B_1 and B_4, respectively. It follows again from the first case that the claim holds for the third case.

This completes the proof of the lemma. $\square$

Chapter 5. Graph Contains K_4 Subgraphs

In the previous sections, we considered finite simplicial graphs Γ that contain no K_4 as induced subgraphs. In this section, we give examples of graphs that contain k_4 as induced subgraphs. Notice that when a graph contains induced K_4 subgraphs, the flag complex on a such graph is not 2-dimensional. This indicates that the assumption that the flag complex on Γ is 2-dimensional is necessary.

5.1 Graph Product of Groups

Graph product of groups was defined by Green in her PhD thesis [23]. Many nice properties about graph products have been discovered by many people. For example, Alonso [2] [JM Alonso, DE Cohen, John Meier] etc.

Definition 5.1. *Let Γ be a finite simplicial graph with vertices $v_1, \cdots, v_m$. For each vertex v_i, assign a group G_i and call it the vertex group of v_i. Then the* graph product $G\Gamma$ *of the vertex groups of Γ is defined as*

$$G\Gamma = \left\langle G_1 * \cdots * G_n \;\middle|\; [G_i, G_j] \text{ whenever } v_i \text{ and } v_j \text{ are adjacent} \right\rangle$$

Example 5.2. *Let Γ be a finite simplicial graph with no edges. Then the graph product $G\Gamma$ is the free product of the vertex groups.*

Example 5.3. *Let Γ be a finite simplicial graph. If all the vertex groups are $\mathbb{Z}$, then the graph product $G\Gamma$ is the right-angled Artin group A_Γ.*

It is natural to ask: what is the Dehn function of the graph product if we know the Dehn functions of the vertex groups? Alonso [2] proved the following:

Theorem 5.4. ([2], Theorem 1) *Let Γ be a finite simplicial graph (not necessary connected) with vertices $v_1, \cdots, v_m$. Let G_i be the vertex group of v_i.*

(1) If δ_{G_i} is polynomial of degree at least 2 for all i, then $\delta_{G\Gamma} = \max\{\delta_{G_i}\}$.

(2) *If* δ_{G_i} *is linear for all* i*, then* $n \preceq \delta_{G\Gamma}(n) \preceq n^2$.

(3) *If* Γ *contains an edge which connects vertices* v_i *and* v_j *and* $\delta_{G_i}, \delta_{G_j}$ *are polynomials of degrees at least* 2*, then* $n^2 \preceq \delta_{G\Gamma}(n)$.

(4) *The Dehn function* $\delta_{G\Gamma}$ *is exponential if and only if* $\max\{\delta_{G_i}\}$ *is exponential.*

Theorem 5.4 can recover the Dehn functions of right-angled Artin groups.

Corollary 5.5. *Let* Γ *be a finite simplicial graph, not necessary connected. Then* δ_{A_Γ} *is at most quadratic. It is quadratic if* Γ *contains an edge.*

Proof. Recall that A_Γ is the graph product $G\Gamma$ where the vertex groups of Γ are all $\mathbb{Z}$. Since $\delta_{\mathbb{Z}}$ is linear, by Theorem 5.4 (2), δ_{A_Γ} is at most quadratic. If Γ contains an edge, then it follows by Theorem 5.4 (3) that δ_{A_Γ} is quadratic. □

Definition 5.6. *We say that a subgroup* H *is a* retract *of a group* G *if there is a homomorphism* $r : G \to H$ *such that* $r : H \to H$ *is the identity. We call the homomorphism* r *a* retraction.

A standard fact about group retract is that if H is a retract of a finitely presented group G, then H is also finitely presented. Moreover, we can write down a presentation of H which is a subpresentation of G. That is, if $\langle \mathcal{S}_G | \mathcal{R}_\mathcal{G} \rangle$ is a finite presentation of G, then we can give H a presentation $\langle \mathcal{S}_H | \mathcal{R}_\mathcal{H} \rangle$ such that $\mathcal{S}_H \subseteq \mathcal{S}_G$ and $\mathcal{R}_H \subseteq \mathcal{R}_G$. The following lemma can be found in [10]. We reproduce the proof for completeness.

Lemma 5.7. ([10], Lemma 2.2) *If* H *is a retract of a finitely presented group* G*, then* $\delta_H \preceq \delta_G$.

Proof. Let $r : G \to H$ be a retraction. Note that H is also finitely presented. Let $G = \langle \mathcal{S}_G | \mathcal{R}_\mathcal{G} \rangle$ and $H = \langle \mathcal{S}_H | \mathcal{R}_\mathcal{H} \rangle$ be finite presentations for G and H, respectively,

where $\mathcal{S}_H \subseteq \mathcal{S}_G$ and $\mathcal{R}_H \subseteq \mathcal{R}_G$. Let w be a word which represents the identity in H. Note that w also represents the identity in G. Write

$$w = \prod_{i=1}^{m} x_i r_i^{\pm 1} x_i^{-1},$$

where $x_i \in F(\mathcal{S}_G)$, $r_i \in F(\mathcal{R}_G)$, and m is the area of w over the presentation $\langle \mathcal{S}_G | \mathcal{R}_\mathcal{G} \rangle$. Applying the retraction r gives

$$w = r(w) = \prod_{i=1}^{m} r(x_i) r(r_i)^{\pm 1} r(x_i)^{-1}.$$

If $r_i \in \mathcal{R}_H$ for some i, then $r(x_i)r(r_i)^{\pm 1}r(x_i)^{-1} = 1$ since $r(r_i) = r_i = 1$. Eliminate the terms if $r(r_i) = 1$, we have

$$w = \prod_{i=1}^{m'} r(x_i) r(r_i)^{\pm 1} r(x_i)^{-1},$$

where m' is the area of w over the presentation $\langle \mathcal{S}_H | \mathcal{R}_\mathcal{H} \rangle$ and $m' \leq m$. Hence, $\delta_H \preceq \delta_G$. □

The proof of the following Lemma is more involved, its proof can be found in [2].

Lemma 5.8. ([2], Lemma 1) *If Γ' is an induced subgraph of Γ, then $G\Gamma'$ is a retract of $G\Gamma$.*

Corollary 5.9. *If Γ' is an induced subgraph of Γ, then $A_{\Gamma'}$ is a retract of A_Γ.*

Proof. Since right-angled Artin groups are graph products of groups $\mathbb{Z}$, the result follows by Lemma 5.8. □

Proposition 5.10. *Let Γ be a finite simplicial graph. If Γ' is a connected induced subgraph of Γ, then $H_{\Gamma'}$ is a retract of H_Γ.*

Proof. Since Γ' is an induced subgraph of Γ, $A_{\Gamma'}$ is a retract of A_Γ by Lemma 5.9. Let $r : A_\Gamma \to A_{\Gamma'}$ be a retraction and $r' = r|_{H_\Gamma} : H_\Gamma \to H_{\Gamma'}$. We argue that r' is

a retraction. Since Γ and Γ' are connected, H_Γ and $H_{\Gamma'}$ are finitely generated and their generating sets are sets of directed edges of Γ and Γ', respectively. It suffices to show that r' is the identity on a generating set of $H_{\Gamma'}$. Let e be a directed edge of Γ' with initial vertex v and terminal vertex w. The generator e of $H_{\Gamma'}$ can be expressed in terms of the generators of $A_{\Gamma'}$, $e = vw^{-1}$. Since $r : A_\Gamma \to A_{\Gamma'}$ is a retraction, we have

$$r'(e) = r(vw^{-1}) = vw^{-1} = e.$$

This shows that $r' : H_\Gamma \to H_{\Gamma'}$ is a retraction. □

5.2 Examples

Example 5.11. *Let Γ be a K_4, that is, a wheel with four vertices. Observe that the flag complex on Γ is a 3-cell. By Proposition 3.4, δ_{H_Γ} is quadratic.*

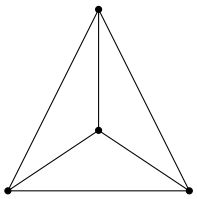

FIGURE 5.1. A wheel with four vertices.

Example 5.12. *Let Γ_1 and Γ_2 be fans of possibly different sizes (we draw two identical ones only for simplicity). Connect Γ_1 and Γ_2 by a path graph P_n for any $n \geq 1$ to get a new graph Γ' as shown in Figure 5.2:*

FIGURE 5.2. The graph Γ'.

Let Γ be the cone on Γ', that is, Γ is a join of Γ' and a vertex v:

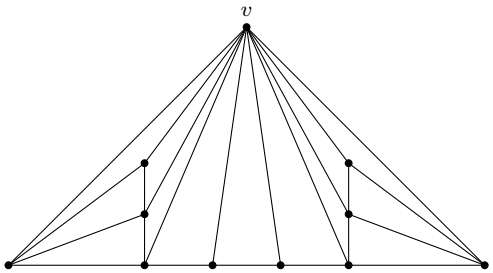

FIGURE 5.3. The graph Γ.

The graph Γ contains induced K_4 subgraphs and the flag complex on Γ is not a 2-dimensional complex. By Proposition 3.3, we have $\delta_{H_\Gamma}(n) \simeq \delta_{A_{\Gamma'}}(n) \simeq n^2$.

Remark 5.13. *In Example 5.12, we actually have a family of graphs such that the associated Bestvina–Brady groups have quadratic Dehn functions since the size of fans and the path graph P_n can be arbitrary.*

Example 5.14. *Take two copies of Γ in Example 5.12 and glue them along the path from the left-most vertex to the right-most vertex. Call the new graph $\overline{\Gamma}$:*

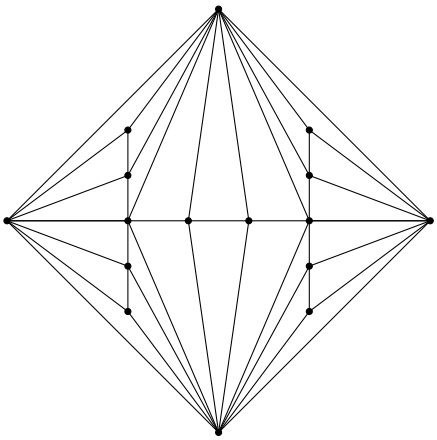

FIGURE 5.4. The graph $\overline{\Gamma}$.

Again, the flag complex on $\overline{\Gamma}$ is not a 2-dimensional complex. Observe that $\overline{\Gamma}$ contains the induced subgraph Γ'':

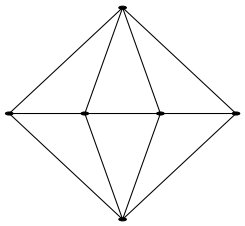

FIGURE 5.5. An induced subgraph of $\overline{\Gamma}$.

By Proposition 5.10, the group $H_{\Gamma''}$ is a retract of $H_{\overline{\Gamma}}$ and $\delta_{H_{\Gamma''}}$ is cubic by Theorem 4.1. Then it follows by Lemma 5.7 that $n^3 \simeq \delta_{H_{\Gamma''}}(n) \preceq \delta_{H_\Gamma}(n)$.

Proposition 5.15. *Let Γ be a finite simplicial graph such that the flag complex D on Γ is simply-connected. If Γ contains an induced subgraph Γ' and the flag complex on Γ' is a 2-dimensional triangulated subdisk D' of D that has square boundary and $\dim_I(D') = d$, $d \in \{0, 1, 2\}$, then $\delta_{H_\Gamma}(n) \succeq n^{d+2}$.*

Proof. Let Γ' be the subgraph of Γ such that the flag complex on Γ' is D'. Then Γ' is a connected induced subgraph of Γ and $\delta_{H_{\Gamma'}}(n) \simeq n^{d+2}$. By Proposition 5.10 and Lemma 5.7, $H_{\Gamma'}$ is a retract of H_Γ and $\delta_{H_{\Gamma'}} \preceq \delta_{H_\Gamma}$. Combining all these sentences, we have $\delta_{H_\Gamma}(n) \succeq n^{d+2}$. □

The flag complex in Proposition 5.15 has no restrictions on its dimension, so the graphs in Examples 5.11, 5.12 and 5.14 are included.

References

[1] Aaron Abrams, Noel Brady, Pallavi Dani, Moon Duchin, and Robert Young. Pushing fillings in right-angled Artin groups. *J. Lond. Math. Soc. (2)*, 87(3):663–688, 2013.

[2] J. M. Alonso. Dehn functions and finiteness properties of graph products. *J. Pure Appl. Algebra*, 107(1):9–17, 1996.

[3] Juan M. Alonso. Inégalités isopérimétriques et quasi-isométries. *C. R. Acad. Sci. Paris Sér. I Math.*, 311(12):761–764, 1990.

[4] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.*, 129(3):445–470, 1997.

[5] B. H. Bowditch. A short proof that a subquadratic isoperimetric inequality implies a linear one. *Michigan Math. J.*, 42(1):103–107, 1995.

[6] N. Brady and M. R. Bridson. There is only one gap in the isoperimetric spectrum. *Geom. Funct. Anal.*, 10(5):1053–1070, 2000.

[7] Noel Brady. Branched coverings of cubical complexes and subgroups of hyperbolic groups. *J. London Math. Soc. (2)*, 60(2):461–480, 1999.

[8] Noel Brady and Max Forester. Snowflake geometry in CAT(0) groups. *J. Topol.*, 10(4):883–920, 2017.

[9] Noel Brady and Ignat Soroko. Dehn functions of subgroups of right-angled Artin groups. *arXiv e-prints*, page arXiv:1709.04066, Sep 2017.

[10] Stephen G. Brick. On Dehn functions and products of groups. *Trans. Amer. Math. Soc.*, 335(1):369–384, 1993.

[11] Martin R. Bridson. The geometry of the word problem. In *Invitations to geometry and topology*, volume 7 of *Oxf. Grad. Texts Math.*, pages 29–91. Oxford Univ. Press, Oxford, 2002.

[12] Martin R. Bridson. On the subgroups of right-angled Artin groups and mapping class groups. *Math. Res. Lett.*, 20(2):203–212, 2013.

[13] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[14] José Burillo and Jennifer Taback. Equivalence of geometric and combinatorial Dehn functions. *New York J. Math.*, 8:169–179, 2002.

[15] William Carter and Max Forester. The Dehn functions of Stallings-Bieri groups. *Math. Ann.*, 368(1-2):671–683, 2017.

[16] Ruth Charney. An introduction to right-angled Artin groups. *Geom. Dedicata*, 125:141–158, 2007.

[17] Ruth Charney and Michael W. Davis. Finite $K(\pi, 1)$s for Artin groups. In *Prospects in topology (Princeton, NJ, 1994)*, volume 138 of *Ann. of Math. Stud.*, pages 110–124. Princeton Univ. Press, Princeton, NJ, 1995.

[18] Warren Dicks and Ian J. Leary. Presentations for subgroups of Artin groups. *Proc. Amer. Math. Soc.*, 127(2):343–348, 1999.

[19] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Berlin, fifth edition, 2018. Paperback edition of [MR3644391].

[20] Will Dison. An isoperimetric function for Bestvina-Brady groups. *Bull. Lond. Math. Soc.*, 40(3):384–394, 2008.

[21] David B. A. Epstein, James W. Cannon, Derek F. Holt, Silvio V. F. Levy, Michael S. Paterson, and William P. Thurston. *Word processing in groups*. Jones and Bartlett Publishers, Boston, MA, 1992.

[22] Steve M. Gersten. Isoperimetric and isodiametric functions of finite presentations. In *Geometric group theory, Vol. 1 (Sussex, 1991)*, volume 181 of *London Math. Soc. Lecture Note Ser.*, pages 79–96. Cambridge Univ. Press, Cambridge, 1993.

[23] Elisabeth Ruth Green. *Graph products of groups*. 1990. Thesis (Ph.D.)–University of Leeds.

[24] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[25] Dan Margalit and Anne Thomas. Quasi-isometries. In *Office hours with a geometric group theorist*, pages 125–145. Princeton Univ. Press, Princeton, NJ, 2017.

[26] Lee Mosher. Mapping class groups are automatic. *Ann. of Math. (2)*, 142(2):303–384, 1995.

[27] A. Yu. Ol' shanskii. Hyperbolicity of groups with subquadratic isoperimetric inequality. *Internat. J. Algebra Comput.*, 1(3):281–289, 1991.

[28] Stefan Papadima and Alexander Suciu. Algebraic invariants for Bestvina-Brady groups. *J. Lond. Math. Soc. (2)*, 76(2):273–292, 2007.

[29] Panagiotis Papasoglu. On the sub-quadratic isoperimetric inequality. In *Geometric group theory (Columbus, OH, 1992)*, volume 3 of *Ohio State Univ. Math. Res. Inst. Publ.*, pages 149–157. de Gruyter, Berlin, 1995.

[30] A. N. Platonov. An isoparametric function of the Baumslag-Gersten group. *Vestnik Moskov. Univ. Ser. I Mat. Mekh.*, (3):12–17, 70, 2004.

[31] Timothy Riley. Dehn functions. In *Office hours with a geometric group theorist*, pages 146–175. Princeton Univ. Press, Princeton, NJ, 2017.

[32] Robert Young. The Dehn function of $\mathrm{SL}(n;\mathbb{Z})$. *Ann. of Math. (2)*, 177(3):969–1027, 2013.

Vita

Yu-Chan Chang was born in Taipei, Taiwan, in 1985. He finished his undergraduate studies at National Central University University, Taiwan, in 2008. He earned a master of science degree in mathematics from National Central University and Louisiana State University, in 2010 and 2015, respectively. In August 2013 he came to Louisiana State University to pursue graduate studies in mathematics. He is currently a candidate for the degree of Doctor of Philosophy in mathematics, which will be awarded in August 2019.