Homological and homotopical Dehn functions are different

Aaron Abrams[a,1], **Noel Brady**[b], **Pallavi Dani**[c], and **Robert Young**[d]

[a]Mathematics Department, Washington and Lee University, Lexington, VA 24450; [b]Department of Mathematics, University of Oklahoma, Norman, OK 73019; [c]Department of Mathematics, Louisiana State University, Baton Rouge, LA 70803-4918; and [d]Department of Mathematics, University of Toronto, Toronto, ON, Canada, M5S 2E4

Edited by Robion C. Kirby, University of California, Berkeley, CA, and approved May 3, 2013 (received for review May 2, 2012)

The homological and homotopical Dehn functions are different ways of measuring the difficulty of filling a closed curve inside a group or a space. The homological Dehn function measures fillings of cycles by chains, whereas the homotopical Dehn function measures fillings of curves by disks. Because the two definitions involve different sorts of boundaries and fillings, there is no a priori relationship between the two functions; however, before this work, there were no known examples of finitely presented groups for which the two functions differ. This paper gives such examples, constructed by amalgamating a free-by-cyclic group with several Bestvina–Brady groups.

1. Introduction

The classical isoperimetric problem is to determine the maximum area that can be enclosed by a closed curve of fixed length in the plane. This problem has been generalized in many different ways. For example, in a metric space X, one can study the homotopical filling area of a curve γ, denoted $\delta_X(\gamma)$ and defined to be the infimal area of a disk whose boundary is γ. This leads to the idea of the *homotopical Dehn function* of X, defined as the smallest function δ_X such that any closed curve γ of length ℓ has filling area at most $\delta_X(\ell)$. A remarkable result of Gromov (1) states that if X is simply connected and if there is a group G that acts on X geometrically (i.e., cocompactly, properly discontinuously, and by isometries), then the growth rate of δ_X depends only on G; indeed, δ_X is related to the difficulty of determining whether a product of generators of G represents the identity. We can thus define the Dehn function of a group to be the (homotopical) Dehn function of any simply connected space that the group acts on geometrically; this is well-defined up to certain constants (details are provided in *Section 2*).

Another way to generalize the isoperimetric problem is to consider fillings of 1-cycles by 2-chains instead of fillings of curves by disks. If α is a 1-cycle in X, we can define its homological filling area, $\mathrm{FA}(\alpha)$, to be the infimal mass of a 2-chain in X with integer coefficients whose boundary is α. This leads to the *homological Dehn function* FA_X, defined as the smallest function such that any 1-chain α of mass at most ℓ has a homological filling of area at most $\mathrm{FA}_X(\ell)$. Like its homotopical counterpart, FA_X can be used to construct a group invariant: If $H_1(X)=0$ and if there is a group G that acts on X geometrically, then the growth rate of FA_X depends only on G and we can define $\mathrm{FA}_G=\mathrm{FA}_X$. Again, this is well-defined up to constants.

The exact relationship between these two filling functions has been an open question for some time. The homological Dehn function deals with a wider class of possible fillings (surfaces of arbitrary genus) and a wider class of possible boundaries (sums of arbitrarily many disjoint closed curves), so it is not a priori clear whether FA_H is always the same as δ_H when they are both defined. Some hints that they may differ come from a construction, due to Bestvina and Brady (2), of groups with unusual finiteness properties. Bestvina and Brady used a combinatorial version of Morse theory to construct a group that is FP_2 but not finitely presented. Such a group does not act geometrically on any simply connected space but does act geometrically on a space with trivial first homology, so its homological Dehn function is defined but its homotopical Dehn function is undefined.

In this paper, we will construct a family of finitely presented groups such that FA_H grows strictly more slowly than δ_H. Specifically, we will show the following.

Theorem 1.1. *For every $d\in\mathbb{N}\cup\{\infty\}$, there is a CAT(0) group G containing a finitely presented subgroup H such that $\mathrm{FA}_H(\ell)\preceq\ell^5$ and the homotopical Dehn function satisfies*

$$\ell^d\preceq\delta_H(\ell)\qquad \text{if } d\in\mathbb{N},$$

$$e^\ell\preceq\delta_H(\ell)\qquad \text{if } d=\infty.$$

Remark: Using the methods of Brady, Guralnik, and Lee (3), one can show that in the $d\in\mathbb{N}$ case, the group H constructed in the theorem satisfies $\delta_H(\ell)\preceq\ell^{d+3}$.

Our construction uses methods of Brady, Guralnik, and Lee (3) to create a hybrid of a Bestvina–Brady group with a group having a large Dehn function. The resulting group is finitely presented, so both δ_H and FA_H are defined, and we will show that the unusual finiteness properties coming from the Bestvina–Brady construction lead to a large gap between homological and homotopical filling functions.

Similar results are known for higher dimensional versions of δ and FA. One can define k-dimensional homotopical and homological Dehn functions by considering fillings of k-spheres or k-cycles by $(k+1)$-balls or $(k+1)$-chains; by historical accident, the corresponding homotopical and homological filling functions have come to be called δ^k_X and FV^{k+1}_X, respectively. The relationship between δ^k_X and FV^{k+1}_X is better understood when $k\geq 2$, because in this case, the Hurewicz theorem can be used to replace cycles and chains by spheres and balls.

If X is k-connected and β is a $(k+1)$-chain with $k\geq 2$, then the Hurewicz theorem can be used to show that β is the image of the fundamental class of a ball under a map $b: D^{k+1}\to X$ with $\mathrm{Vol}\ b=\mathrm{Mass}\ \beta$. Thus, if $a: S^k\to X$ is a map of a sphere and α is the image of the fundamental class of S^k under a, then $\delta^k_X(a)=\mathrm{FV}^{k+1}_X(\alpha)$, so $\delta^k_X\preceq\mathrm{FV}^{k+1}_X$ for $k\geq 2$ (see appendix 2 in ref. 4, 5).

Likewise, if X is k-connected and α is a k-cycle for $k\geq 3$, then the Hurewicz theorem can be used to show that α is the image of the fundamental class of a sphere under a map $a: S^k\to X$ such that $\mathrm{Vol}\ a=\mathrm{Mass}\ \alpha$ [see remark 2.6.(4) in ref. 6]. Consequently, because $\delta^k_X(a)=\mathrm{FV}^{k+1}_X(\alpha)$, we have $\delta^k_X\sim\mathrm{FV}^{k+1}_X$ for $k\geq 3$.

Thus, if $k\geq 3$ and if H is a group that acts geometrically on a k-connected complex, the above results imply that $\delta^k_H\sim\mathrm{FV}^{k+1}_H$. When $k=2$, Young (7) constructed examples of groups for which $\delta^2_H\precnsim\mathrm{FV}^3_H$. The examples in this paper are the earliest known examples of groups for which $\delta^1_H\nsim\mathrm{FV}^2_H$.

2. Preliminaries

2.1. Dehn Functions. For a full exposition of Dehn functions, we recommend the study by Bridson (8). We will briefly review the

---

Author contributions: A.A., N.B., P.D., and R.Y. performed research and wrote the paper.

[1]To whom correspondence should be addressed. E-mail: abramsa@wlu.edu.

www.pnas.org/cgi/doi/10.1073/pnas.1207377110

definitions that we will need. Let X be a simply connected Riemannian manifold or simplicial complex. If $\alpha: S^1 \to X$ is a Lipschitz map, define the *homotopical filling area* of α to be

$$\delta_X(\alpha) = \inf_{\substack{\beta: D^2 \to X \\ \beta|_{S^1} = \alpha}} \operatorname{Area} \beta,$$

where β ranges over Lipschitz maps $D^2 \to X$ which agree with α on ∂D^2. Because X is simply connected and any continuous map can be approximated by a Lipschitz map, such maps exist. We can define an invariant of X by letting

$$\delta_X(\ell) = \sup_{\ell(\alpha) \le \ell} \delta_X(\alpha).$$

We call this the *homotopical Dehn function* of X.

We define a relation on functions $\mathbb{R}^{\ge 0} \to \mathbb{R}^{\ge 0}$ by $f \preceq g$ if there is a $c > 0$ such that for all n,

$$f(n) \le cg(cn + c) + cn + c.$$

If $f \preceq g$ and $g \preceq f$, we write $f \sim g$. Thus, $\sim$ distinguishes all functions x^μ for $\mu \ge 1$, and all functions of the form λ^x are equivalent for $\lambda > 1$. Gromov (1) stated [and Bridson (8) proved] that if H acts geometrically on X (for instance, if $H = \pi_1(M)$ and $X = \widetilde{M}$ for some compact M), then $\delta_X(n)$ is determined up to $\sim$-equivalence by H. If H is finitely presented, then H acts geometrically on the universal cover $\widetilde{X}_H$ of a presentation complex, which is a 2-complex X_H with $\pi_1(X_H) = H$. Thus, $\delta_H := \delta_{\widetilde{X}_H}$ is well-defined up to $\sim$.

To define the homological invariant FA_H, suppose that X is a polyhedral complex with $H_1(X) = 0$. If α is a 1-cycle in X, we let

$$\mathrm{FA}_X(\alpha) = \inf_{\partial \beta = \alpha} \operatorname{Mass} \beta,$$

where Mass β is defined to be $\sum |b_i|$ if $\beta = \sum b_i \Delta_i$ is a sum of 2-simplices in X with integer coefficients. We can define an invariant of X by letting

$$\mathrm{FA}_X(\ell) = \sup_{\operatorname{Mass} \alpha \le \ell} \mathrm{FA}_X(\alpha).$$

We call this the *homological Dehn function* of X. Like the homotopical Dehn function, if H acts geometrically on X, then FA_X is determined up to $\sim$ by H, and if X_H is a presentation complex for a finitely presented group H, we define $\mathrm{FA}_H = \mathrm{FA}_{\widetilde{X}_H}$.

2.2. Right-Angled Artin Groups. If Λ is a simple graph (i.e., without loops or multiple edges), we can define a *right-angled Artin group* (RAAG) based on Λ. If $V(\Lambda)$ and $E(\Lambda)$ are the vertex set and edge set of Λ, respectively, we define

$$A_\Lambda = \langle V(\Lambda) \mid [i(e), t(e)] = 1 \text{ for all } e \in E(\Lambda) \rangle,$$

where $i, t: E(\Lambda) \to V(\Lambda)$ are the functions taking an edge to its start and end. We say that Λ is the *defining graph* of A_Λ. These RAAGs generalize free groups and free abelian groups; if Λ is a complete graph, there is an edge between every pair of vertices, so every pair of generators of A_Λ commutes and A_Λ is free abelian. On the other hand, if Λ has no edges, then A_Λ is a free group.

A full exposition of RAAGs can be found in the paper by Charney (9). One important fact that we will use is that for every Λ, there is a one-vertex locally CAT(0) cube complex X_Λ with $\pi_1(X_\Lambda) = A_\Lambda$; this is called the *Salvetti complex*. This complex can be built directly from the graph Λ: it has one vertex, one edge for every vertex of Λ, one square for each edge of Λ, and in general one n-cube for each n-vertex clique in Λ.

Bestvina and Brady (2) used RAAGs to construct subgroups of nonpositively curved groups with unusual finiteness properties. They defined a homomorphism $h_{A_\Lambda}: A_\Lambda \to \mathbb{Z}$ which sends each generator of A_Λ to 1, and, viewing A_Λ as the 0-skeleton of $\widetilde{X}_\Lambda$, they extended h_{A_Λ} to a map $h_{X_\Lambda}: \widetilde{X}_\Lambda \to \mathbb{R}$. This map is linear on each cube of $\widetilde{X}_\Lambda$, so the level set $L_{A_\Lambda} := h_{X_\Lambda}^{-1}(0)$ can be given the structure of a polyhedral complex. The subgroup $H_{A_\Lambda} := \ker h_{A_\Lambda}$ acts vertex-transitively on L_{A_Λ}, so if L_{A_Λ} is connected, the 1-skeleton of L_{A_Λ} is a Cayley graph for H_{A_Λ}. In this case, we can construct a generating set for H_{A_Λ} explicitly: each edge of L_{A_Λ} is a diagonal of a square of $\widetilde{X}_\Lambda$, so H_{A_Λ} has a generating set consisting of elements of the form ab^{-1}, where a and b are generators of A_Λ.

Recall that a complex is *flag* if every clique of n vertices spans an $(n-1)$-dimensional simplex. Bestvina and Brady (2) proved that the topology of Λ determines the topology of $h_{X_\Lambda}^{-1}(0)$.

Theorem 2.1 (Bestvina and Brady). *If Λ is the 1-skeleton of a flag complex Y, and $h_{A_\Lambda}, h_{X_\Lambda}$ are the maps defined above, then $H_{A_\Lambda} = \ker h_{A_\Lambda}$ acts on the complex $L_{A_\Lambda} = h_{X_\Lambda}^{-1}(0)$, which is homotopy equivalent to a wedge product of infinitely many copies of Y, indexed by the vertices in $\widetilde{X}_\Lambda \setminus L_{A_\Lambda}$. In fact, $h_{X_\Lambda}^{-1}(0)$ is a union of infinitely many scaled copies of Y.*

The main tool used to prove this theorem is a combinatorial version of Morse theory. If X is a complex, $x \in X$ is a vertex, and $h: X \to \mathbb{R}$ is a function that is linear on each cell and is not constant on any edge, one may define subcomplexes $\mathrm{Lk}_\uparrow(x)$ and $\mathrm{Lk}_\downarrow(x)$ of the link $\mathrm{Lk}(x)$ called the *ascending* and *descending links* of x. To define these, we identify the vertices of $\mathrm{Lk}(x)$ with the neighbors of x. The ascending link $\mathrm{Lk}_\uparrow(x)$ is the full subcomplex spanned by vertices y such that $h(y) > h(x)$; likewise, the descending link $\mathrm{Lk}_\downarrow(x)$ is the full subcomplex spanned by vertices y such that $h(y) < h(x)$. These ascending and descending links play a role similar to the ascending and descending manifolds in classical Morse theory.

If X has one vertex, then all vertices of $\widetilde{X}$ have the same link, so we will write $\mathrm{Lk}(\widetilde{X})$, $\mathrm{Lk}_\uparrow(\widetilde{X})$, and $\mathrm{Lk}_\downarrow(\widetilde{X})$. The link $\mathrm{Lk}(\widetilde{X}_\Lambda)$ has two vertices $s^\pm$ for each generator s of A_Λ; the ascending link $\mathrm{Lk}_\uparrow(\widetilde{X}_\Lambda)$ is spanned by the s^+'s, and the descending link $\mathrm{Lk}_\downarrow(\widetilde{X}_\Lambda)$ is spanned by the s^-'s. If Y is the flag complex with 1-skeleton Λ, then $\mathrm{Lk}_\uparrow(\widetilde{X}_\Lambda)$ and $\mathrm{Lk}_\downarrow(\widetilde{X}_\Lambda)$ are isomorphic to Y.

2.3. Labeled Oriented Graph Groups. We can also construct groups using *labeled oriented graphs* (LOGs). A LOG on a set S is a directed multigraph Γ with vertex set S and a labeling of the edges given by $l: E(\Gamma) \to S$; loops and multiple edges are allowed. We say that Γ presents the group:

$$B_\Gamma := \left\langle S \,\middle|\, i(e)^{l(e)} = t(e), e \in E(\Gamma) \right\rangle,$$

where the notation a^b represents the conjugation $b^{-1}ab$ and, again, $i(e), t(e) \in S$ are the end points of e. Because each relation has length 4, the presentation 2-complex X_Γ of B_Γ is a two-dimensional cube complex.

Note that although $i(e) = t(e)$ is possible, we may assume that $i(e) \ne l(e) \ne t(e)$ because, otherwise, we could contract such an edge without changing the group. This implies that $\mathrm{Lk}(X_\Gamma)$ contains no loops or edges of the form s^+s^-.

As with RAAGs, we can apply Morse theory to LOG groups. Let $h_{B_\Gamma}: B_\Gamma \to \mathbb{Z}$ be the homomorphism mapping each $s \in S$ to 1. This homomorphism can be extended linearly over each cell of $\widetilde{X}_\Gamma$ to get a map $h_{X_\Gamma}: \widetilde{X}_\Gamma \to \mathbb{R}$. Consider the level set $L_{B_\Gamma} = h_{X_\Gamma}^{-1}(0) \subset \widetilde{X}_\Gamma$. As in the RAAG case, the group $\ker h_{B_\Gamma}$ acts on L_{B_Γ} vertex-transitively, so if L_{B_Γ} is connected, then its 1-skeleton is a Cayley graph for $\ker h_{B_\Gamma}$. Edges in L_{B_Γ} are diagonals of squares in X_Γ, so each orbit of squares labeled $a^b = c$ contributes a generator that can be written as cb^{-1} or $b^{-1}a$ (Fig. 1).

As was the case with RAAGs, the link $\mathrm{Lk}(\widetilde{X}_\Gamma)$ has two vertices $s^\pm$ for each vertex s of Γ. The ascending link $\mathrm{Lk}_\uparrow(\widetilde{X}_\Gamma)$ is the full subcomplex of $\mathrm{Lk}(\widetilde{X}_\Gamma)$ spanned by the s^+'s, and the descending link $\mathrm{Lk}_\downarrow(\widetilde{X}_\Gamma)$ is spanned by the s^-'s. Brady (10) showed the following.

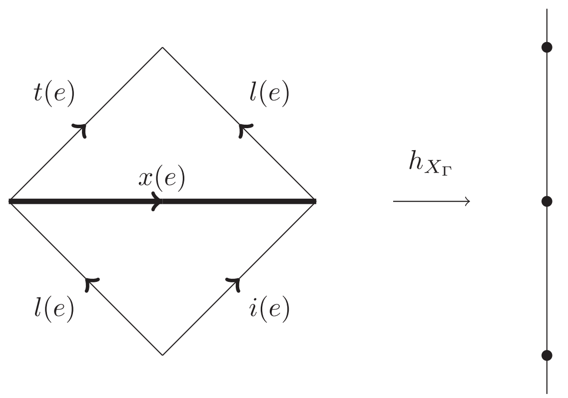

Fig. 1. The map h_{X_Γ} on a 2-cell in $\widetilde{X}_\Gamma$. The group element $x(e)$ is in ker h.

Theorem 2.2 (Brady). *Suppose B_Γ is a group presented by a LOG Γ such that*:

- *The ascending and descending links $\mathrm{Lk}_\uparrow(\widetilde{X}_\Gamma)$ and $\mathrm{Lk}_\downarrow(\widetilde{X}_\Gamma)$ are trees.*
- *The full link $\mathrm{Lk}(\widetilde{X}_\Gamma)$ has girth at least 4.*

Then, (1) *X_Γ is locally CAT(0), hence a $K(B_\Gamma, 1)$;* (2) *the level set L_{B_Γ} is a tree; and* (3) *B_Γ is isomorphic to the free-by-cyclic group $F_n \rtimes_\phi \mathbb{Z}$, where $F_n \cong \ker\ h_{B_\Gamma}$.*

Brady, Guralnik, and Lee (3) used these groups to construct Stallings-type examples of groups that are of type F_2 but not of type F_3 and that have Dehn functions with prescribed polynomial or exponential growth rates.

3. Main Theorem

To understand our construction, first consider the problem of constructing a space where the homological and homotopical filling functions differ. Suppose W is a simply connected space with a large Dehn function and α is a closed curve in W. To reduce the homological filling area but not the homotopical filling area of α, we could attach a 2-complex Z to α in which α is the boundary of a 2-chain but not the boundary of a disk. If $\pi_1(Z)/\langle\alpha\rangle = 0$, the resulting space is still simply connected. By attaching copies of Z to infinitely many closed curves, we can obtain a complex that has large δ but small FA.

Our construction will be based on a graph of groups with each vertex labeled by one of two groups, A and Q. The first group, A, will be a RAAG with a kernel H_A that is FP_2 but not finitely presented. This subgroup acts geometrically on a space that has trivial H_1 and nontrivial π_1, which will provide the Z's in the construction.

We define a *Thompson complex* to be a connected, finite, two-dimensional flag complex Y whose fundamental group is a simple group with an element of infinite order. (The name comes from the earliest known group with these properties, Thompson's group T.)

Let Y be a Thompson complex (for example, a triangulation of a presentation complex for Thompson's group). Note that because $\pi_1(Y)$ is simple, $H_1(Y) = 0$, and every $g \neq 1 \in \pi_1(Y)$ normally generates all of $\pi_1(Y)$. Let $g \in \pi_1(Y)$ be an element of infinite order. By gluing an annulus to Y, we may assume that there is a path of length 4 in the 1-skeleton of Y which represents g. We label the vertices of this path a, u, s, v, and we label the rest of the vertices of Y by $y_1, \ldots, y_d$. Since Y is flag, the subcomplex spanned by a, u, s, v must be a cycle of length 4. We will consider the RAAG A_Λ, where Λ is the 1-skeleton of Y.

As we will not need to refer to Λ explicitly, we drop it from the notation and set $A = A_\Lambda$. We denote the associated homomorphism by $h_A : A \to \mathbb{Z}$, its extension to a Morse function by $h_{X_A} : \widetilde{X}_A \to \mathbb{R}$, the level set $h_{X_A}^{-1}(0)$ by L_A, and so on. By results of Bestvina and Brady (2), the group $H_A = \ker\ h_A$ is FP_2 but not finitely presented.

The second group, Q, will be a product of a LOG group and a free group. Suppose we are given a LOG Γ' that satisfies the hypotheses of *Theorem 2.2*. We may form a new LOG Γ by adding an isolated vertex s to Γ', adding a loop connecting a vertex t of Γ' to itself, and labeling the new edge by s. This corresponds to adding a generator s and a relation $[s, t] = 1$ to B_Γ. We call a LOG Γ obtained this way a *special LOG* (SLOG), and the corresponding group a *SLOG group*. Note that Γ still satisfies the hypotheses of *Theorem 2.2*.

As with Λ, we will often omit Γ from the notation when it is easily understood. We will abbreviate B_Γ by B, X_Γ by X_B, h_{B_Γ} by h_B, and so on.

If B is a SLOG group, then by *Theorem 2.2*, it can be written as a free-by-cyclic group $B = F_n \rtimes_\phi \mathbb{Z}$. (The notation F_n indicates a rank n free group; if we want to emphasize a particular free basis $\{x_i\}$, we will write $F_n(x_1, \ldots, x_n)$.) We define Dist_B to be the distortion of F_n inside B; precisely:

$$\mathrm{Dist}_B(\ell) = \max_g \left\{ |g|_{F_n} : |g|_B \leq \ell \right\},$$

where $|g|$ is the word length of g in the subscripted group. Brady, Guralnik, and Lee (3) give constructions of SLOG groups $B = F_n \rtimes_\phi \mathbb{Z}$ with $\mathrm{Dist}_B \sim e^\ell$, and also with $\mathrm{Dist}_B \sim \ell^d$ for all integers $d \geq 3$. The Bieri–Stallings double of B, denoted $D = B *_{F_n} B$, has large Dehn function resulting from this distortion. Specifically, Bridson and Haefliger (theorem III.Γ.6.20 of ref. 11) showed the following.

Theorem 3.1 (Bridson and Haefliger). *If B and D are as above, then*

$$\mathrm{Dist}_B(\ell) \preceq \delta_D(\ell).$$

The group D will serve as W in our construction; because its Dehn function is large, it has many curves that are difficult to fill by disks. By an embedding trick appearing in a paper by Baumslag, Bridson, Miller, and Short (12), D can be viewed as a subgroup of the product $Q := B \times F_2$; in fact, we will see that D is the kernel of a map $h_Q : Q \to \mathbb{Z}$.

We will construct a finitely presented CAT(0) group G as a graph product of Q with several copies of A. The subgroup H will be the kernel of a map $h : G \to \mathbb{Z}$, and H will have the structure of a graph product of copies of D and H_A. We will show that attaching H_A to D does not affect δ but that the copies of Y that lie in L_A can be used to replace fillings by disks with more efficient fillings by chains.

Theorem 3.2. *Let A be a RAAG based on a Thompson complex as described above, and let $B = F_n \rtimes_\phi \mathbb{Z}$ be a SLOG group. Then there exists a finitely presented CAT(0) group G containing A and B such that the homomorphisms $h_A : A \to \mathbb{Z}$ and $h_B : B \to \mathbb{Z}$ extend to $h : G \to \mathbb{Z}$ and such that $H = \ker(h)$ is finitely presented and satisfies*:

$$\mathrm{FA}_H(\ell) \preceq \ell^5$$

$$\delta_H(\ell) \succeq \mathrm{Dist}_B(\ell).$$

Using the examples of SLOG groups constructed by Brady, Guralnik, and Lee (3), this implies Theorem 1.1.

4. Constructing *G* and *H*

In this section, we construct the groups G and H of *Theorem 3.2*. The construction is similar to the perturbed RAAGs in the paper by Brady, Guralnik, and Lee (3), but we glue several RAAGs (rather than just one) to a free-by-cyclic group.

Throughout this paper, if g is a group element, $\bar{g}$ will represent its inverse.

4.1. The SLOG Piece. Let $B = B_\Gamma = F_n \rtimes_\phi \mathbb{Z}$ be as in *Theorem 3.2*. The first step of the construction is to use B to construct a group $D \lhd Q \cong B \times F_2$ with large Dehn function. The group Q will contain several copies of the group $F_2 \times F_2$, and we will attach RAAGs A_i to Q along some of these groups. The result of this gluing will be G.

Since Γ is a SLOG, it contains an isolated vertex s, which is the label of a single loop in Γ. Call the vertex of that loop t. Call the rest of the vertices $\{a_1,\dots,a_{n-1}\}$. We have two presentations for B, namely, the SLOG presentation with generating set $\{s,t,a_i\}$ and a free-by-cyclic presentation. For the latter, we may take $\{x_i,t \mid 1\le i\le n\}$ as a generating set, where $x_i=a_i\bar{t}$ for $1\le i<n$ and $x_n=s\bar{t}$ (Fig. 1). Thus, $B=F_n(x_1,\dots,x_n)\rtimes_\phi\mathbb{Z}$.

Let D be the double $D=B*_{F_n}B$, where the F_n is generated by the x_i. *Theorem 3.1* implies that D has Dehn function at least as large as Dist_B. By results of Baumslag, Bridson, Miller, and Short (12), D is isomorphic to the subgroup

$$D\cong F_n(x_1,\dots,x_n)\rtimes_{(\phi,\phi)}F_2\left(t\bar{u},t\bar{v}\right)$$

of the group $Q=B\times F(u,v)$. Furthermore, if $h_Q:Q\to\mathbb{Z}$ is the group homomorphism taking the elements a_i,s,t,u,v to $1\in\mathbb{Z}$, then the kernel of h_Q is precisely D.

Because Γ is a SLOG, the group $B=B_\Gamma$ contains many copies of F_2. Recall that the presentation 2-complex X_B of B is a locally CAT(0) two-dimensional cube complex and thus a $K(B,1)$. For any i, consider the subgroup of B generated by a_i and s. It is easy to check that any two of the vertices $a_i^\pm$ and $s^\pm$ in the link of X_B are separated by a distance of at least 2 in $\mathrm{Lk}(X_B)$. Consequently, the Cayley graph of the subgroup generated by a_i and s is convexly embedded in $\widetilde{X}_B$. It is therefore a copy of F_2.

Define $X_Q=X_B\times R_2$, where R_2 is a wedge of two circles, so that X_Q is a $K(Q,1)$. This is locally CAT(0), and by the argument above, for all i, the subgroup generated by a_i,s,u,v is a convexly embedded copy of $F_2\times F_2$.

4.2. Attaching the RAAG Pieces. Let $A=A_\Lambda$ be a RAAG constructed from a Thompson complex Y as in *Section 3*. Thus, $g\in\pi_1(Y)$ has infinite order and is represented by a path *ausv* in the defining graph Λ for A. Let X_A be the Salvetti complex of A. Because a,u,s, and v span a square in Y, A has a convex subgroup isomorphic to $F_2\times F_2$, generated by a,u,s,v. Let $E:=F_2\times F_2$.

We form the group G by gluing copies of A to Q along copies of E. Specifically, consider a graph of groups with vertex groups $Q,A_1,\dots,A_{n-1}$, where $A_i=A$, and with each vertex A_i connected to Q by an edge. Each edge group will be isomorphic to E. As noted above, for each i, the elements $a_i,s,u,v\in Q$ generate a copy of E; denote this copy by E_i. We identify E_i with the copy of E in A_i by $a_i\leftrightarrow a$, $u\leftrightarrow u$, $s\leftrightarrow s$, $v\leftrightarrow v$. Let G be the fundamental group of this graph of groups. This is a group generated by

$$\left\{a_1,\dots,a_{n-1},s,t,u,v,y_j^i\right\}\quad i=1,\dots,n-1,\ j=1,\dots,d.$$

We can define subgroups Q, E_i, and A_i, where $E_i\cong E$, $A_i\cong A$, and

$$Q=B_\Gamma\times F_2(u,v)=\langle a_1,\dots,a_{n-1},s,t,u,v\rangle,$$

$$A_i=\langle a_i,s,u,v,y_1^i,\dots,y_d^i\rangle,$$

$$E_i=A_i\cap Q=F_2(a_i,s)\times F_2(u,v).$$

The homomorphisms $h_Q:Q\to\mathbb{Z}$ and $h_A:A\to\mathbb{Z}$ agree on the edge groups, so we can extend them to a function $h:G\to\mathbb{Z}$.

Let $H=\ker\ h$.

5. Finite Presentability

In this section, we will construct a space on which H acts and consider its topology. This will let us prove that H is finitely presented and will help us bound the Dehn functions of H.

We can realize the above construction of G geometrically to construct a $K(G,1)$ as follows. Let $X_E:=R_2\times R_2$, where R_2 is the wedge of two circles; this is a $K(E,1)$, and each edge group corresponds to copies of X_E in X_A and X_Q. Each of these copies of X_E is convex, so we can glue $n-1$ copies of X_A to X_Q along the X_E's to obtain a locally CAT(0) cube complex that we call X_G. This is a $K(G,1)$, and the 1-skeleton of its universal cover $\widetilde{X}_G$ is a Cayley graph of G. In particular, we have the following.

Lemma 5.1. *The group G is CAT(0).* □

Now, H is the kernel of the homomorphism $h:G\to\mathbb{Z}$. As before, the vertices of $\widetilde{X}_G$ are in correspondence with the elements of G, so by viewing h as a function on the vertices of $\widetilde{X}_G$, we may extend h linearly over cubes to obtain a Morse function $h:\widetilde{X}_G\to\mathbb{R}$. Let $L_G=h^{-1}(0)$.

Because h cuts cubes of L_G "diagonally," L_G is a polyhedral 2-complex whose cells are slices of cubes. The subgroup H acts freely on L_G, and since the vertices of L_G are in one-to-one correspondence with the elements of H, the action is cocompact and thus geometric. We now show the following.

Lemma 5.2. *L_G is simply connected and thus $H=\ker(h)\subset G$ is finitely presented.*

Proof: Because $\widetilde{X}_G$ is contractible and h is a Morse function on $\widetilde{X}_G$, theorem 4.1 in ref. 2 implies that it is enough to show that the ascending and descending links of any vertex in $\widetilde{X}_G$ are simply connected. Since X_G has only one vertex, it is enough to show this for that vertex.

Since the 1-skeleton of $\widetilde{X}_G$ is a Cayley graph of G, we can label the vertices of $\mathrm{Lk}(X_G)$ by $g^\pm$, where g ranges over the generating set S. The link $\mathrm{Lk}(X_G)$ of the vertex of X_G is obtained by gluing $\mathrm{Lk}(X_Q)$ and the various $\mathrm{Lk}(X_{A_i})$. For each $1\le i\le n-1$, the links $\mathrm{Lk}(X_Q)$ and $\mathrm{Lk}(X_{A_i})$ each contain a subcomplex with vertices $a_i^\pm$, $u^\pm$, $s^\pm$, and $v^\pm$, and gluing the links along these subcomplexes gives $\mathrm{Lk}(X_G)$.

Likewise, we can form $\mathrm{Lk}_\uparrow(X_G)$ by gluing $\mathrm{Lk}_\uparrow(X_{A_i})$, $1\le i\le n-1$, to $\mathrm{Lk}_\uparrow(X_Q)$ along subcomplexes S_i spanned by a_i^+,u^+,s^+, and v^+. We claim that $\mathrm{Lk}_\uparrow(X_G)$ is simply connected. Since $\mathrm{Lk}_\uparrow(X_{B_\Gamma})$ is a tree by hypothesis, $\mathrm{Lk}_\uparrow(X_Q)$ is the suspension of a tree (with suspension points u^+ and v^+), and thus, it is simply connected. Since A_i is a RAAG with defining complex Y, each $\mathrm{Lk}_\uparrow(X_{A_i})$ is isomorphic to Y, and each S_i is a square such that the normal closure of $\pi_1(S_i)$ in $\pi_1(\mathrm{Lk}_\uparrow(X_{A_i}))=\pi_1(Y)$ is all of $\pi_1(Y)$. By the Seifert–van Kampen theorem,

$$\begin{aligned}\pi_1(\mathrm{Lk}_\uparrow(X_G))=(\pi_1(\mathrm{Lk}_\uparrow(X_{A_1}))*\dots*\pi_1(\mathrm{Lk}_\uparrow(X_{A_{n-1}})))/\\ \langle\pi_1(S_1),\dots,\pi_1(S_{n-1})\rangle=0,\end{aligned}$$

so the ascending link is simply connected. The same argument with +'s changed to −'s shows that the descending link is also simply connected, so L_G is simply connected. Therefore, $H=\pi_1(L_G/H)$ and H is finitely presented. □

For an alternate description of H, recall that the group G is the fundamental group of a graph of groups, with one vertex labeled Q connected to $n-1$ vertices labeled A_i by edges labeled E_i. Since $H\subset G$, G induces a graph of groups structure on H. Indeed, G acts on a tree T whose vertices correspond to the cosets of Q and A_i, whose edges correspond to cosets of E_i, and whose quotient $G\backslash T$ is a star with $n-1$ edges. We can restrict the action of G on T to an action of H, and since any coset of Q, A_i, or E_i has nontrivial intersection with H, the orbit of any vertex or edge under H is the same as its orbit under G. Therefore, H acts on T with vertex stabilizers conjugate to $H_Q:=H\cap Q$ and $H_{A_i}:=H\cap A_i$, edge stabilizers conjugate to $H_{E_i}:=H\cap E_i$, and quotient $H\backslash T=G\backslash T$. This shows that H is the fundamental group of the graph of groups with a central vertex labeled H_Q connected to $n-1$ vertices labeled H_{A_i} by edges labeled H_{E_i}.

The level set $L_G:=h^{-1}(0)\subset\widetilde{X}_G$, however, is *not* the universal cover of a corresponding graph of spaces. To describe L_G, we define $L_Q:=L_G\cap\widetilde{X}_Q$, $L_{A_i}:=L_G\cap\widetilde{X}_{A_i}$, and $L_{E_i}:=L_G\cap\widetilde{X}_{E_i}$. These level sets have geometric actions by H_Q, H_{A_i}, and H_{E_i}, respectively. We can write the quotient L_G/H as L_Q/H_Q with the L_{A_i}/H_{A_i} attached along copies of L_{E_i}/H_{E_i}, but because L_{A_i} and L_{E_i} are not simply connected, $\pi_1(L_{A_i}/H_{A_i})\ne A_i$ and $\pi_1(L_{E_i}/H_{E_i})\ne E_i$; this is a graph of spaces for a different graph of groups.

The fact that L_{A_i} and L_{E_i} are not simply connected will be important in the rest of this paper, so we will go into some more detail. All the L_{A_i}'s and all the L_{E_i}'s are isometric, so when i is unimportant, we will denote them by L_A and L_E, respectively. To understand the topology of L_A and L_E, consider them as subsets of $\widetilde{X}_A$. According to Bestvina and Brady (2), L_A is a union of scaled copies of Y, indexed by vertices in $\widetilde{X}_A\setminus L_A$. Likewise, L_E is composed of scaled copies of a square, which we denote $\diamondsuit$, indexed by vertices in $\widetilde{X}_E\setminus L_E$. Translating L_E by elements of H_A gives infinitely many disjoint copies of L_E inside L_A.

By a result of Bestvina and Brady (theorem 8.6 in ref. 2), L_A is homotopy equivalent to an infinite wedge sum of copies of Y and L_E is homotopy equivalent to an infinite wedge sum of copies of $\diamondsuit$, so L_A and L_E have infinitely generated π_1. Generators of $\pi_1(L_A)$ can be filled in two ways. First, each generator can be freely homotoped into some copy of L_E. Each copy of L_E is contained in some L_Q, and because L_Q is simply connected, each generator of $\pi_1(L_A)$ is filled by a disk in one of the copies of L_Q. Second, although $\pi_1(L_A)$ is infinitely generated, $H_1(L_A)$ is trivial; thus, any curve in L_A can be filled by some 2-chain entirely inside L_A. Our goal in the rest of this paper is to use these two types of fillings to show that the homological and homotopical Dehn functions of L_G are different.

6. Upper Bound on the Homological Dehn Function

In this section, we prove the following.

Proposition 6.1. *With H as above, we have* $\mathrm{FA}_H(\ell)\preceq\ell^5$.

Proof: We show that any 1-cycle in L_G of mass at most ℓ can be filled by a 2-chain of mass $\preceq\ell^5$. Since ℓ^5 is a superadditive function, it is enough to prove this for loops in $L_G^{(1)}$.

L_G consists of copies of L_A and L_Q glued together along copies of L_E. We first show how to homologically fill loops that lie in a single copy of L_A or L_Q, and we then use these fillings to fill arbitrary loops. Note that each 1-cell of L_G is a diagonal of some square in $\widetilde{X}_G$, so each 1-cell corresponds to a product $x\bar{y}$, where x and y are (certain) generators of G.

Consider a loop α of length ℓ that lies in a copy of L_A. Recall that L_A is a level set of the Morse function $h_A:\widetilde{X}_A\to\mathbb{R}$. Because $\widetilde{X}_A$ is CAT(0), there exists a 2-chain β with boundary equal to α and mass $\preceq\ell^2$. Further, β lies in $h_A^{-1}[-c\ell,c\ell]$ for some $c>0$ (13; cf. proposition 2.2 in ref. 14). We will use β to produce a filling of α in L_A using a pushing map as in the paper by Abrams, Brady, Dani, Duchin, and Young (14).

Let Z be the space obtained by deleting open neighborhoods of the vertices of X_A outside L_A, with the induced cell structure; that is, $Z=X_A\setminus\bigcup_{v\notin L_A}B^{\circ}_{1/4}(v)$. According to Abrams, Brady, Dani, Duchin, and Young (theorem 4.2 in ref. 14), there is an H_A-equivariant locally Lipschitz retraction (pushing map) $\mathcal{Q}:Z\to L_A$ such that the Lipschitz constant grows linearly with distance from L_A. Furthermore, if $S_v=\partial B^{\circ}_{1/4}(v)$ is the boundary of one of the deleted neighborhoods, the image of S_v is a copy of Y with the metric scaled by a factor of $h_A(v)$. In particular, if γ is a 1-cycle in S_v of length $\ell(\gamma)$, then $\mathcal{Q}_\sharp(\gamma)$ is a 1-cycle in a scaled copy Y_v of Y. Since $H_1(Y)=0$ and Y is compact, the corresponding 1-cycle in Y has homological filling area $\preceq\ell(\gamma)$. Therefore, the original cycle, $\mathcal{Q}_\sharp(\gamma)$, has homological filling area $\preceq\ell(\gamma)h_A(v)^2$.

Consider the restriction β' of β to Z; this is a 2-chain in Z, and

$$\partial\beta'=\alpha+\sum_{v\in\operatorname{supp}\beta\setminus L_A}\gamma_v,$$

where γ_v is a 1-cycle in S_v. We can construct a filling β'' of α in L_A by combining the image $\mathcal{Q}_\sharp(\beta')$ with fillings of each of the $\mathcal{Q}_\sharp(\gamma_v)$'s. Since β' lies in $h_A^{-1}[-c\ell,c\ell]$, the restriction of $\mathcal{Q}$ to supp β' has $O(\ell)$ Lipschitz constant, and

$$\operatorname{Mass}(\beta'')\preceq\ell^2\operatorname{Mass}(\beta)+\ell^2\sum_v\ell(\gamma_v)\preceq\ell^4+\ell^2\sum_v\ell(\gamma_v).$$

Each 2-cell of β contributes at most four 1-cells to the γ_v's; thus, $\sum_v\ell(\gamma_v)\preceq\ell^2$, and we have $\operatorname{Mass}(\beta'')\preceq\ell^4$.

Next, we produce a quartic mass filling of any loop that lies entirely in a copy of L_Q. Such a loop α is labeled by generators of

$$D\cong F_n(x_1,\ldots,x_n)\rtimes_{(\phi,\phi)}F_2(t\bar{u},t\bar{v}),$$

where $x_i=a_i\bar{t}$ for $1\le i<n$ and $x_n=s\bar{t}$. In this section, we will use the notation $t_1=t\bar{u}$ and $t_2=t\bar{v}$. Let w denote the word labeling α, where

$$w=w_1\ldots w_\ell.$$

Here, the w_i are generators and w represents the identity.

Let $\vartheta:F(x_1,\ldots,x_n)\rtimes F_2(t_1,t_2)\to F(x_1,\ldots,x_n)\rtimes\langle t_1\rangle$ send t_2 to t_1 and each other generator to itself. The word $\vartheta(w)$ lies in a free-by-cyclic group which is CAT(0) and therefore has quadratic Dehn function. Thus, to fill our loop with quartic mass, it will be enough to reduce w to $\vartheta(w)$ in such a way that the reduction takes quartic mass.

We will achieve this reduction by first decomposing w into subwords as follows. Let $p:F(x_1,\ldots,x_n)\rtimes F_2(t_1,t_2)\to F_2(t_1,t_2)$ be the projection map. Define

$$w(i)=w_1\ldots w_i$$

so that $w(0)=w(\ell)=1$. Now, decompose w as follows:

$$w=\left[p(w(0))w_1p(w(1))^{-1}\right]\left[p(w(1))w_2p(w(2))^{-1}\right]\\ \ldots\left[p(w(\ell-1))w_\ell p(w(\ell))^{-1}\right].$$

We can reduce w to $\vartheta(w)$ by reducing each subword in this decomposition to its image under ϑ. If $w_j=t_1^{\pm1}$ or $t_2^{\pm1}$, then $p(w(j-1))w_jp(w(j))^{-1}$ is freely equal to the identity and no reduction is necessary. Otherwise, $p(w(j-1))=p(w(j))\in F(t_1,t_2)$. Thus, it suffices to reduce words of the form $gx_i\bar{g}$, with $g\in F(t_1,t_2)$, to $\vartheta(gx_i\bar{g})$ or, equivalently, to fill loops with labels of the form $v=gx_i\bar{g}\vartheta(g\bar{x}_i\bar{g})$.

Write $g=t_{r_1}^{d_1}\cdots t_{r_m}^{d_m}$, where r_i alternates between 1 and 2 and $d_i\in\mathbb{Z}\setminus\{0\}$. We proceed by induction on m.

If $m=1$ and $g=t_1^d$, there is nothing to do. If $g=t_2^d$, then $v=t_2^dx_it_2^{-d}t_1^d\bar{x}_it_1^{-d}$ can be written as the sum of four 1-cycles as shown in Fig. 2. Writing $s_1=s\bar{u}$ and $s_2=s\bar{v}$, the words labeling the 1-cycles are $t_2^ds_2^{-d}s_1^dt_1^{-d}$, $s_2^d(a_i\bar{v})^{-d}(a_i\bar{u})^ds_1^{-d}$, $(a_i\bar{v})^{d+1}s_2^{-(d+1)}s_1^{d+1}(a_i\bar{u})^{-(d+1)}$, and $s_2^{d+1}t_2^{-(d+1)}t_1^{(d+1)}s_1^{-(d+1)}$. The first and last are words representing the identity in the CAT(0) group $\langle t\rangle\times\langle s\rangle\times F(u,v)$, and thus can be filled with quadratic mass. The middle two are generators of $\pi_1(L_A)$ and can be filled in L_A with a scaled copy of Y with quadratic mass. These four fillings fit together to give a filling of v with quadratic mass.

If $m>1$, then $g=g_0t_{r_m}^{d_m}$. Let $r=r_m$ and $d=d_m$. Let $g'=g_0t_{3-r}^d$ and note that $\vartheta(g)=\vartheta(g')$. As in the $m=1$ case, we can reduce $t_r^dx_it_r^{-d}$ to $t_{3-r}^dx_it_{3-r}^{-d}$ using quadratic area. This immediately lets us reduce

$$gx_i\bar{g}\vartheta\left(g\bar{x}_i\bar{g}\right)$$

to

$$g'x_ig'^{-1}\vartheta\left(g'x_i^{-1}g'^{-1}\right).$$

Because we use m steps, and $m\le\ell(g)\le\ell$, it takes area $\preceq\ell^3$ (and linear genus) to reduce $gx_i\bar{g}$ to $\vartheta(gx_i\bar{g})$.

Since each of the ℓ subwords in the decomposition of w above can be reduced to its image under ϑ using mass $\preceq\ell^3$, the word w can be reduced to $\vartheta(w)$ with mass $\preceq\ell^4$.

Finally, we consider curves that travel through multiple copies of L_{A_i} and L_Q. We will need to make arguments based on the

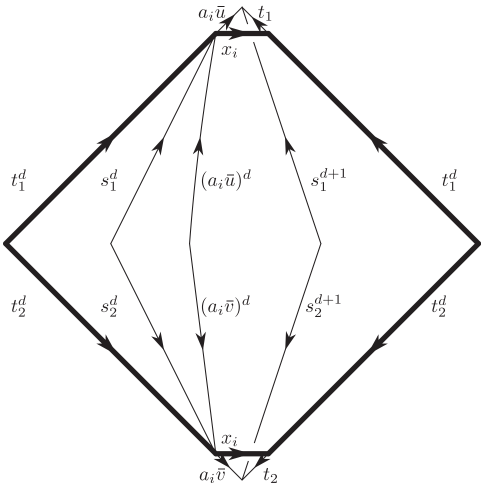

Fig. 2. A homological filling of v. The curve v (thick line) is a sum of four 1-cycles (thin lines).

graph product decomposition of H, so it will be helpful to have a slightly different complex L on which H acts. We construct L by "stretching" each copy of L_E in L_G into a product $L_E\times[0,1]$. Let Z be the complex obtained by gluing X_Q and $n-1$ copies of X_A to $n-1$ copies of $X_E\times[0,1]$ according to the graph product decomposition of G, and let $\widetilde{Z}$ be the universal cover of Z. Then, the homotopy equivalence $p:Z\to X_G$ that collapses each copy of $X_E\times[0,1]$ to a copy of X_E lifts to a homotopy equivalence $\widetilde{p}:\widetilde{Z}\to\widetilde{X}_G$. Then, $H=\ker(h)\subset G$ acts geometrically on the level set $L=(h\circ\widetilde{p})^{-1}(0)$, because H acts geometrically on L_G. The following lemma describes the structure of L.

Lemma 6.2. *The level set L intersects each vertex space of $\widetilde{Z}$ of the form $\widetilde{X}_Q$ in a copy of L_Q. Likewise, L intersects each vertex space $\widetilde{X}_A$ in a copy of L_A and each edge space $\widetilde{X}_E\times[0,1]$ in a copy of $L_E\times[0,1]$.* □

Each edge in $L^{(1)}$ either lies in a copy of L_Q, lies in a copy of L_{A_i} for some i, or crosses from $L_E\times\{0\}$ to $L_E\times\{1\}$; thus, we can classify the edges as Q-edges, A-edges, or E-edges. Consider a path of edges in $L^{(1)}$ of length ℓ and call it α.

By standard arguments (i.e., the normal form theorem for graphs of groups), α must have an "innermost piece," i.e., a subpath that enters a copy of L_{A_i} or L_Q through a copy of L_E and then leaves through the same L_E. We write this as $t\gamma t'$, where t and t' lie in the same copy of $L_E\times[0,1]$ and where γ is either a path of A-edges or a path of Q-edges.

Without loss of generality, suppose that the endpoints of γ lie in $L_E\times\{0\}$. Call them $(v,0)$ and $(w,0)$, and let γ' be a geodesic in $L_E\times\{0\}$ from $(v,0)$ to $(w,0)$. Then, p and γ' form a loop θ, and because they lie in the union of a copy of $L_E\times[0,1]$ and a copy of L_Q or L_{A_i}, there is a 2-chain filling θ whose mass is $\preceq(\ell+\ell(\gamma'))^4+\ell(\gamma')$. But L_E is undistorted in L, because L_E is undistorted in X_E and X_E is convex in X_G, so $\ell(\gamma')\preceq\ell$ and the filling area of θ is $\preceq\ell^4$.

Repeating this process for the loop $\alpha-\theta$ inductively, we obtain a filling of α. Each time we repeat the process, the number of E-edges in α decreases by 2, and we use a filling of mass $\preceq\ell^4$, so the total filling area of α is $\preceq\ell^5$. □

7. Lower Bound on Homotopical Dehn Function

Recall the situation we are in: The group G is a graph product of groups Q and (copies of) A along edge groups E, and $\widetilde{X}_G,\widetilde{X}_Q,\widetilde{X}_A,\widetilde{X}_E$ are contractible spaces on which G,Q,A,E, respectively, act geometrically. Meanwhile $H<G$ is a graph product of D (the double of a SLOG group B) and copies of H_A (the Bestvina–Brady group associated with the RAAG A). In this section, we prove the following.

Theorem 7.1. *With G (hence, H, D, Q, A, B) as above, we have*

$$\delta_H\succeq\mathrm{Dist}_B,$$

where Dist_B *is as in Section 3.*

To prove this, we will need the following refinement of *Theorem 3.1.*

Lemma 7.2. *For all $\ell>0$, there is a curve $\gamma:S^1\to L_Q$ of length $\sim\ell$ such that if $\Delta\in C_2(L_Q)$ is a filling of γ, then*

$$\text{Area supp }\Delta\succeq\mathrm{Dist}_B(\ell),$$

where supp Δ *is the support of* Δ.

We can take γ to be the curve $w_n^{-1}w_n'$ used by Bridson and Haefliger (proof of theorem III. Γ.6.20 in ref. 11); their proof shows that the image of any disk filling γ has to have area $\succeq\mathrm{Dist}_B(\ell)$, but the same bound applies to any chain filling γ as well.

Proof of Theorem 7.1: Let L be the complex constructed in the previous section, on which H acts; this is made up of copies of L_A and L_Q, joined along copies of $L_E\times[0,1]$. The translates of the set $L_E\times\{1/2\}$ separate L into infinitely many components, each of which is either a copy of L_A glued to copies of $L_E\times[0,1/2]$ or a copy of L_Q glued to copies of $L_E\times[1/2,1]$. Let L_A' and L_Q' be complexes isometric to each type of piece. We will refer to the union of the copies of $L_E\times\{1/2\}$ that lie in L_A' and L_Q' as $\partial L_A'$ and $\partial L_Q'$.

Fix a "root" copy L_0 of L_Q' in L. Let γ be a curve in L_0 as in *Lemma 7.2*, and let $[\gamma]\in C_1(L_0;\mathbb{R})$ be its fundamental class. Let τ:$D^2\to L$ be a filling of γ in L. We will show that there is a chain $\Delta\in C_2(L_0;\mathbb{R})$ that fills $[\gamma]$ and is supported on $L_0\cap\tau(D^2)$. Then by Lemma 7.2,

$$\text{Area }\tau\ \geq\ \text{Area supp }\Delta\ \succeq\ \mathrm{Dist}_B(\ell),$$

which proves the theorem.

First, we use the structure of L to break a disk in L into punctured disks whose images lie in copies of L_Q' and L_A'. Homotope τ so that it is transverse to each copy of $L_E\times\{1/2\}$. The preimages of the copies of $L_E\times\{1/2\}$ then divide D^2 into *pieces* $M_1,\ldots,M_n\subset D^2$, and for each i, $\tau|_{M_i}$ is a punctured disk in a copy of L_A' or L_Q'. Each M_i has a distinguished boundary component that is homotopic to ∂D^2 in $D^2\setminus M_i$, and we call that component the *outer boundary* of M_i, denoted $\partial_o M_i$; we call the other boundary curves *inner boundaries*. Each boundary curve of M_i either coincides with ∂D^2 or lies in one of the $L_E\times\{1/2\}$ preimages.

Next, we claim that if M is a punctured disk in L_A' with boundary in $\partial L_A'$, the topology of L_A places strong restrictions on M. More precisely, we show the following.

Lemma 7.3. *Suppose that $f:M\to L_A'$ is a punctured disk in L_A' with boundary in $\partial L_A'$. Suppose that M has a distinguished boundary component $\partial_o M$ and other boundary components $\partial_1 M\ldots\partial_m M$. If $[\partial_i M]\in C_1(\partial L_A';\mathbb{R})$ is the fundamental class of $\partial_i M$, then there are $a_i\in\mathbb{R}$ such that*

$$f_\sharp[\partial_o M]=\sum_{i=1}^m a_i f_\sharp[\partial_i M].$$

Furthermore, we may assume that $a_i\neq 0$ only if $f(\partial_i M)$ and $f(\partial_o M)$ lie in the same copy of $L_E\times\{1/2\}$.

We will prove this lemma at the end of the section after we use it to construct Δ. Let M_i be one of the pieces of D^2, and suppose that τ takes $\partial_o M_i$ to a curve in L_0. (For example, take M_i such that $\partial_o M_i=\partial D^2$.) Then $\tau(M_i)$ is either a punctured disk in L_0 or a punctured disk in a copy of L_A' that neighbors L_0. Let $\overline{M}_i\subset D^2$ be the disk bounded by $\partial_o M_i$. We claim that there is a chain Δ_i in $\tau(\overline{M}_i)\cap L_0$ such that $\partial\Delta_i=\tau_\sharp[\partial_o M_i]$.

We proceed by induction on the number n of pieces of D^2 contained in $\overline{M}_i$. If $n=1$, then $\overline{M}_i=M_i$. Then, as before, $\tau(M_i)$ is either a disk in L_0 or a disk in a copy of L_A' that neighbors L_0. In the first case, we can take $\Delta_i=\tau_\sharp[M_i]$. In the second case, the lemma implies that $\tau_\sharp[\partial M_i]=0$, so we can take $\Delta_i=0$.

Suppose that the claim is true for $n-1$ and let M_i be a piece of D^2 such that $\tau(\partial_o M_i) \subset L_0$ and $\overline{M}_i$ is comprised of n pieces of D^2. If $\tau(M_i)$ is a punctured disk in L_0, then τ takes each inner boundary of M_i to a curve in L_0 and each inner boundary bounds a disk in D^2 with at most $n-1$ pieces. By induction, each $\tau_\sharp[\partial_j M_i]$ bounds a chain Δ_{ij} in $\tau(\overline{M}_i) \cap L_0$. Consequently, we can get the required filling of $\tau_\sharp[\partial_o M_i]$ as

$$\Delta_i = \sum_j \Delta_{ij} + \tau_\sharp[M_i].$$

On the other hand, if $\tau(M_i)$ is a punctured disk in a copy of L_A', then *Lemma 7.3* implies that we can write

$$\tau_\sharp[\partial_o M_i] = \sum_j a_{ij} \tau_\sharp\left[\partial_j M_i\right],$$

where each $\partial_j M_i$ is an inner boundary component of M_i that lies in ∂L_0. By induction, each of these can be filled by a chain Δ_{ij} in $\tau(\overline{M}_i) \cap L_0$, and if

$$\Delta_i = \sum_j a_{ij} \Delta_{ij},$$

then Δ_i fills $\tau_\sharp[\partial_o M_i]$.

Therefore, $\tau(\overline{M}_i) \cap L_0$ supports a chain filling γ. ☐

Proof of Lemma 7.3: Let Y be the complex used in the construction of A and let $\diamondsuit \subset Y$ be as in *Section 5*. Choose a basepoint $* \in Y$ such that $* \in \diamondsuit$. Then, as discussed in Section 5, L_A' is homotopy equivalent to an infinite wedge sum of copies of Y,

$$Y_\infty = \bigvee_{\alpha \in S_A} Y_\alpha,$$

where S_A is the set of vertices in $\widetilde{X}_A \setminus L_A$. Similarly, $\partial L_A'$ is made up of disjoint copies of $L_E \times \{1/2\}$; call these $L_{E,0}, L_{E,1}, \ldots$, where $f(\partial_o M) \subset L_{E,0}$. Each of these is homotopy equivalent to a wedge sum indexed by a subset $S_{E,i} \subset S_A$,

$$\diamondsuit_{\infty,i} = \bigvee_{\alpha \in S_{E,i}} \diamondsuit_\alpha.$$

The $S_{E,i}$'s partition S_A into disjoint sets. We can consider each of the $\diamondsuit_{\infty,i}$'s as a subset of Y_∞, and we can define a homotopy equivalence $h: L_A' \to Y_\infty$ that restricts to a homotopy equivalence $L_{E,i} \to \diamondsuit_{\infty,i}$ on each $L_{E,i}$.

Let $* \in Y_\infty$ be the basepoint of the wedge sum. Let

$$f' : (M, \partial M) \to \left(Y_\infty, \bigcup_i \diamondsuit_{\infty,i}\right)$$

be a map that differs from $h \circ f$ by a small homotopy such that $f'^{-1}(*)$ is a graph in M. We can further require that if λ is a boundary component of M and $f(\lambda) \subset L_{E,i}$, then $f'(\lambda) \subset \diamondsuit_{\infty,i}$. Then $f'^{-1}(*)$ cuts M into punctured disks $P_1, \ldots, P_k$, and for each i, we can choose an $\alpha_i \in S_A$ such that $f'(P_i) \subset Y_{\alpha_i}$ and $f'(\partial P_i) \subset \diamondsuit_{\alpha_i}$.

Consider M as a subset of $\mathbb{R}^2$, embedded so that $\partial_o M$ is the outer boundary of the subset. For each i, let D_i be the disk bounded by $\partial_i M$. This embedding lets us choose an outer boundary component $\partial_o P_i$ for each P_i. Furthermore, each closed curve λ in M has an inside, and we can use this to put a partial ordering on the boundary curves of the P_i's. If λ and λ' are two boundary curves, we write $\lambda \prec \lambda'$ if the inside of λ is a subset of the inside of λ'.

If $\lambda_1, \ldots, \lambda_l$ are the inner boundary components of P_i, we claim that $f_\sharp'[\partial_o P_i] \in C_1(\diamondsuit_{\alpha_i})$ is a linear combination of the $f_\sharp'[\lambda_j]$'s. All these chains are in fact 1-cycles in $\diamondsuit_{\alpha_i}$, and because $Z_1(\diamondsuit_{\alpha_i}) \cong \mathbb{R}$, it is enough to show that if $f_\sharp'[\partial_o P_i] \neq 0$, then one of the $f_\sharp'[\lambda_j]$'s is nonzero too. However, if $f_\sharp'[\lambda_j] = 0$ for all j, then $f'(\lambda_j)$ is a null-homotopic curve for all j, and so $f'(\partial_o P_i)$ is the boundary of a disk in Y_{α_i}. Since the inclusion $\diamondsuit \subset Y$ is π_1-injective, this means that if $f_\sharp'[\lambda_j] = 0$ for all j, then $f_\sharp'[\partial_o M] = 0$ as well, which proves the claim.

Next, we claim that if λ is a boundary curve of some P_i, then $f_\sharp'[\lambda] \in C_1(\bigcup_i \diamondsuit_{\infty,i})$ is a linear combination of the $f_\sharp'[\partial_j M]$'s. We proceed by induction on the number of boundary curves λ' with $\lambda' \prec \lambda$. If this number is 0, then λ is one of the $\partial_j M$'s and there is nothing to prove. Otherwise, the inside of λ is a union of P_i's and D_i's, so $[\lambda]$ is a sum of $[\partial P_i]$'s and $[\partial_i M]$'s. By induction, if P_i is inside λ and $\lambda \neq \partial_o P_i$, then $f_\sharp'[\partial P_i]$ is a linear combination of the $f_\sharp'[\partial_j M]$'s, so $f_\sharp'[\lambda]$ is a linear combination of the $f_\sharp'[\partial_j M]$'s too.

Therefore, there are $a_i \in \mathbb{R}$ such that

$$f_\sharp'[\partial_o M] = \sum_i a_i f_\sharp'[\partial_i M],$$

where the equality is taken in $C_1(\bigcup_i \diamondsuit_{\infty,i})$, so

$$f_\sharp[\partial_o M] = \sum_i a_i f_\sharp[\partial_i M]$$

in $H_1(\partial L_A')$. Because $\partial L_A'$ is one-dimensional, the equality in fact holds in $C_1(\partial L_A')$ as well. Finally, $C_1(\partial L_A') = \oplus_i C_1(L_{E,i})$, and if we project the above equation to the $C_1(L_{E,0})$ factor, we get

$$f_\sharp[\partial_o M] = \sum_{\left\{i \mid f(\partial_i M) \subset L_{E,0}\right\}} a_i f_\sharp[\partial_i M],$$

as desired. ☐

ACKNOWLEDGMENTS. We thank Dan Guralnik and Sang Rae Lee for helpful discussions and the American Institute of Mathematics for supporting this research. We gratefully acknowledge additional funding from National Science Foundation Award DMS-0906962 (to N.B.), the Louisiana Board of Regents Support Fund Contract LEQSF(2011-14)-RD-A-06 (to P.D.), and the Natural Sciences and Engineering Research Council of Canada (to R.Y.).

1. Gromov M (1993) Asymptotic invariants of infinite groups. *Geometric Group Theory, volume 2* London Mathematical Society Lecture Note Series (Cambridge Univ Press, Cambridge, UK), Vol 182, pp 1–295.
2. Bestvina M, Brady N (1997) Morse theory and finiteness properties of groups. *Inventiones Mathematicae* 129(3):445–470.
3. Brady N, Guralnik D, Lee SR (2011) Dehn functions and finiteness properties of subgroups of perturbed right-angled artin groups. *arXiv 1102.5551*.
4. Gromov M (1983) Filling Riemannian manifolds. *Journal of Differential Geometry* 18(1):1–147.
5. White B (1984) Mappings that minimize area in their homotopy classes. *Journal of Differential Geometry* 20(2):433–446.
6. Brady N, Bridson MR, Forester M, Shankar K (2009) Snowflake groups, Perron-Frobenius eigenvalues and isoperimetric spectra. *Geometry & Topology* 13(1): 141–187.
7. Young R (2011) Homological and homotopical higher-order filling functions. *Groups, Geometry, and Dynamics* 5(3):683–690.
8. Bridson MR (2002) The geometry of the word problem. *Invitations to Geometry and Topology*, Oxford Graduate Texts in Mathematics, eds Bridson MR, and Salamon SM (Oxford Univ Press, Oxford), Vol 7, pp 29–91.
9. Charney R (2007) An introduction to right-angled Artin groups. *Geometriae Dedicata* 125:141–158.
10. Brady N, Riley T, Short H (2007) *The Geometry of the Word Problem for Finitely Generated Groups, Advanced Courses in Mathematics* (Birkhäuser, Basel).
11. Bridson MR, Haefliger A (1999) Metric spaces of non-positive curvature. *Grundlehren der Mathematischen Wissenschaften* [*Fundamental Principles of Mathematical Sciences*] (Springer, Berlin), Vol 319.
12. Baumslag G, Bridson MR, Miller CF III, Short H (1997) Finitely presented subgroups of automatic groups and their isoperimetric functions. *Journal of the London Mathematical Society* 56(2):292–304.
13. Wenger S (2008) A short proof of Gromov's filling inequality. *Proc Am Math Soc* 136(8): 2937–2941.
14. Abrams A, Brady N, Dani P, Duchin M, Young R (2013) Pushing fillings in right-angled Artin groups. *Journal of the London Mathematical Society* 87(3):663–688.