A BROWN THEOREM FOR DEHN FUNCTIONS OF GRAPHS OF GROUPS

CLAUDIO LLOSA ISENRICH AND JANNIS WEIS

Abstract. We prove an upper bound on the Dehn function of a group G acting cellularly, cocompactly, and without inversions on a simply connected CW complex X in terms of the Dehn functions of the vertex stabilizers, the Dehn function of X, and the distortion of the edge stabilizers, provided that X is either a tree or the stabilizer of each 2-cell has finite index in the stabilizer of every edge in its boundary. This provides a Dehn function analogue of Brown's Theorem for finiteness properties and an answer to a question of Zaremsky in these cases. We also prove analogues of our result for higher Dehn functions when X is a tree.

1. Introduction

The Dehn function $\delta_G(n)$ of a finitely presented group G provides a quantitative measure of the worst-case complexity of a brute-force approach to the word problem for words of length $\leq n$. It can be interpreted geometrically as an optimal isoperimetric function δ_Y for null-homotopic loops of length $\leq n$ in the universal cover $\widetilde{Y}$ of a closed manifold or compact CW complex Y with fundamental group $\pi_1(Y) \cong G$. This geometric interpretation readily generalizes to the higher-dimensional Dehn functions $\delta_Y^k(n)$, which are defined as optimal isoperimetric functions for null-homotopic k-spheres of volume $\leq n$ in $\widetilde{Y}$ under the additional assumption that the universal cover $\widetilde{Y}$ is k-connected. Under these assumptions, δ_Y^k depends only on the fundamental group $G = \pi_1(Y)$, up to asymptotic equivalence, and is a quasi-isometry invariant among groups of finiteness type F_{k+1} [AWP99]; here we say that a group is of type F_k if it admits a classifying space with finite k-skeleton. We thus often write δ_G^k instead of δ_Y^k, with the understanding that this denotes an equivalence class of functions.

Taking the point of view of finiteness properties and recalling that type F_n is itself a quasi-isometry invariant [Alo94], one can view k-dimensional Dehn functions δ_G^k as quasi-isometry invariants that further quantify the finiteness properties F_{k+1}. This naturally raises the question of the extent to which results on finiteness properties admit quantitative analogues in terms of Dehn functions. A classical result of this kind is Brown's Theorem.

Theorem A (Brown [Bro87])**.** *Let G be a group that acts cellularly and cocompactly on a $(k-1)$-connected CW complex X. Assume that the cell stabilizer G_σ of every i-cell σ is F_{k-i} for every $0 \leq i \leq k$. Then G is F_k.*

Date: August 7, 2026.

2020 *Mathematics Subject Classification.* 20F65; 20F69; 20J05; 20E08; 57M07.

Key words and phrases. Dehn functions, Higher filling functions, Brown's Theorem.

The purpose of this work is to explore the existence of analogues of Brown's criterion for Dehn functions. A version of the question whether such an analogue exists for the first Dehn function δ_G^1 was raised by Zaremsky [Zar26, Problem 1.17]. We give an affirmative answer if X is 1-dimensional. To state it, we recall that a relative 0-dimensional analogue of the Dehn function is the distortion dist_H^G of a subgroup $H \leq G$. We denote by $\overline{f}$ the superadditive closure of a function f.

Theorem B. *Let G be a group that acts cellularly and cocompactly without inversion on a tree T. Assume that all vertex stabilizers G_v are finitely presented and that all edge stabilizers G_e are finitely generated. Then*

$$\delta_G(n) \preccurlyeq n \cdot \max_{v \in V(T)} \delta_{G_v}\left(\overline{\operatorname{edist}_T}(n)\right),$$

where $\operatorname{edist}_T(n) = \max_{e \in E(T)} \operatorname{dist}_{G_e}^G(n)$.

The assumptions in Theorem B are equivalent to saying that G is the fundamental group of a finite graph of groups, where, up to conjugation, the G_v and G_e are its vertex and edge groups. In particular, the maxima in the statement exist. Similar considerations imply that the maxima in all subsequent statements exist. The form of the estimate in Theorem B is sharp in general. Indeed, the Baumslag–Solitar group $BS(1,2) = \langle a, b \mid bab^{-1} = a^2 \rangle$ decomposes as an HNN extension with exponentially distorted infinite cyclic edge and vertex stabilizers. Hence, Theorem B yields an exponential upper bound on $\delta_{BS(1,2)}$, which is known to be optimal [Ger92].

While Theorem B does not seem to appear in the literature for arbitrary graphs of groups, previous results cover several special cases. Brick studied the behavior of Dehn functions under amalgamated products and HNN extensions, while Brick and Corson obtained sharper estimates for amalgamated products along strongly undistorted subgroups [Bri93, BC00]. Brick and Corson also proved a general estimate for finite developable complexes of groups with finite edge groups in terms of the Dehn functions of the vertex groups and the Howie function of the complex [BC98]. Estimates for doubles $G *_H G$ and HNN extensions $G *_H$ of a finitely presented group G with respect to a finitely generated subgroup $H \leq G$ appear in [BH99, Thm. III.Γ.6.20]. Doubles and certain iterated ascending HNN extensions have played an important role in gaining new insights into Dehn functions, and they admit generalizations to higher dimensions and to homological Dehn functions [BBFS09, LM22].

We generalize our results to Dehn functions of groups acting on higher-dimensional simplicial complexes and to higher Dehn functions of groups acting on trees.

Theorem C. *Let G be a group that acts cellularly and cocompactly without inversion on a simply connected simplicial complex X. Assume that all vertex stabilizers G_v are finitely presented, that all edge stabilizers G_e are finitely generated, and that $[G_e : G_\sigma] < \infty$ for all 2-cells σ and edges $e \subset \sigma$. Then*

$$\delta_G(n) \preccurlyeq \delta_X(n) \cdot \max_{v \in X^{(0)}} \delta_{G_v}\left(\overline{\operatorname{edist}_X}(\delta_X(n))\right),$$

where $\operatorname{edist}_X(n) = \max_e \operatorname{dist}_{G_e}^G(n)$ *and e runs over all edges of X.*

For higher Dehn functions, we obtain the following result, where FV_H^{k+1} denotes the $(k+1)$-dimensional integral homological filling function of H; see Section 2 for a definition and note that, for historical reasons, there is an index shift.

Theorem D. *For $k \geq 2$, let G be a group that acts cellularly and cocompactly without inversion on a tree T. Assume that all vertex stabilizers G_v are F_{k+1} and that all edge stabilizers G_e are F_k. Then*

$$\mathrm{FV}_G^{k+1}(n) \preccurlyeq \overline{f}(\overline{g}(n)),$$

where $f(n) = \max_{v \in V(T)} \mathrm{FV}_{G_v}^{k+1}(n)$ and $g(n) = \max_{e \in E(T)} \mathrm{FV}_{G_e}^{k}(n)$.

Note that Barnard, Brady, and Dani [BBD12, Proposition 6.1] proved an analogue of Theorem D for the second-order homotopical Dehn function δ_G^2. Moreover, in very recent work, Sauer and the second author proved a version of Theorem D with a polynomial upper bound for complexes of groups under the assumption that the filling functions of the complex, the homological Dehn functions of all cell stabilizers, and the distortion of cell stabilizers are polynomially bounded [SW26].

It is known that δ_H^k and FV_H^{k+1} need not be equivalent for $k = 1, 2$, while for $k \geq 3$ we have $\delta_H^k \approx \mathrm{FV}_H^{k+1}$ [ABDY13, You11]. Thus, we deduce the following higher-dimensional version of [BBD12, Proposition 6.1] from Theorem D.

Corollary E. *For $k \geq 4$, let G be a group that acts cellularly and cocompactly without inversion on a tree T. Assume that all vertex stabilizers G_v are F_{k+1} and that all edge stabilizers G_e are F_k. Then*

$$\delta_G^k(n) \preccurlyeq \overline{f}(\overline{g}(n)),$$

where $f(n) = \max_{v \in V(T)} \delta_{G_v}^k(n)$ and $g(n) = \max_{e \in E(T)} \delta_{G_e}^{k-1}(n)$.

Note that, for $k = 2$, we have $\delta_H^2 \preccurlyeq \mathrm{FV}_H^3$. Thus, we cannot deduce Corollary E in the stated form for $k = 3$. However, we can still bound δ_G^3 analogously in terms of the Dehn functions $\delta_{G_v}^3$ and $\mathrm{FV}_{G_e}^3$ for $v \in V(T)$ and $e \in E(T)$.

Structure. In Section 2 we introduce all required background on distortion, Dehn functions, collapses, and complexes of groups. In Section 3 we prove Theorem B and Theorem C. In Section 4 we prove Theorem D. For completeness, we also include a proof of the analogue of Corollary E for $k = 2$, although this is [BBD12, Proposition 6.1]. Finally, in Section 5 we discuss the special case of group extensions as well as an application to amalgamated products of nilpotent groups over cyclic groups.

Acknowledgements. The authors thank Uri Bader, Jerónimo García-Mejía and Roman Sauer for helpful discussions.

2. Preliminaries

In this section, we introduce all relevant background and notions concerning distortion, Dehn functions, collapses, and classifying spaces for complexes of groups.

2.1. **Distortion, van Kampen diagrams and Dehn functions.** We will give a brief introduction to distortion, van Kampen diagrams and Dehn functions. For more detailed accounts we refer to [Bri02a, BRS07] in the case of 1-dimensional homotopical Dehn functions and to [AWP99, Bri02b, BBFS09, You11] for higher-dimensional homotopical and homological Dehn functions.

The notions of distortion and Dehn functions for groups are well-defined only up to certain equivalence relations on functions $f, g\colon \mathbb{N}_0 \to \mathbb{R}_{\geq 0}$. We write $f \lesssim g$ (resp. $f \preccurlyeq g$) if there is a constant $C > 0$ such that $f(n) \leq Cg(Cn + C) + C$ (resp. $f(n) \leq Cg(Cn + C) + Cn + C$). We denote by $\simeq$ (resp. $\approx$) the equivalence relation induced by $\lesssim$ (resp. $\preccurlyeq$). We say that a function $f\colon \mathbb{N}_0 \to \mathbb{R}_{\geq 0}$ is *superadditive* if, for every partition $n = n_1 + \cdots + n_k$ of $n \in \mathbb{N}_0$ into nonnegative integers, we have $f(n) \geq f(n_1) + \cdots + f(n_k)$. The *superadditive closure* of a function $f\colon \mathbb{N}_0 \to \mathbb{R}_{\geq 0}$ is the smallest superadditive function $\overline{f}\colon \mathbb{N}_0 \to \mathbb{R}_{\geq 0}$ with $f \leq \overline{f}$; it can be defined by $\overline{f}(n) \coloneqq \max_{n=n_1+\cdots+n_k} \sum_{i=1}^{k} f(n_i)$, where the maximum ranges over partitions with $n, n_i \geq 1$.

For a CW complex X, we denote by $X^{(n)}$ its n-skeleton and by $d_X : X^{(0)} \times X^{(0)} \to \mathbb{N}_0$ the combinatorial metric on its 1-skeleton, where every edge has length 1. A *combinatorial path* γ in the 1-skeleton $X^{(1)}$ is a path defined by a consecutive choice of edges. Its length $\ell_X(\gamma)$ is its number of edges.

For a finitely generated group $G = \langle S \rangle$, we denote by $d_{G,S}$ its word metric, which is the combinatorial metric on its Cayley graph, and by $\ell_{G,S}(g) \coloneqq d_{G,S}(g, 1)$ the word length of an element $g \in G$. When a generating set is understood, we usually write simply d_G and ℓ_G.

Definition 2.1. Let $G = \langle S \rangle$ and $H = \langle T \rangle$ be finitely generated groups with $H \leq G$. The *distortion* $\mathrm{dist}_H^G\colon \mathbb{N}_0 \to \mathbb{R}_{\geq 0}$ of H in G is

$$\mathrm{dist}_H^G(n) \coloneqq \max \{\ell_H(h) \mid h \in H \text{ with } \ell_G(h) \leq n\}.$$

The distortion is independent of the choice of finite generating set up to the $\simeq$-equivalence.

We will now define higher-dimensional Dehn functions, following the definitions in [Bri02b, Ril03]. For further references with equivalent definitions, see [AWP99, BBFS09, You11].

We start by recalling that a *combinatorial map* between CW complexes is a map that maps open n-cells homeomorphically onto open n-cells. A *combinatorial complex* K is defined iteratively as a CW complex such that the attaching maps $\phi_\sigma\colon \partial D^{n+1} \to K^{(n)}$ of $(n+1)$-cells are combinatorial with respect to some combinatorial structure on $S^n = \partial D^{n+1}$. A *singular combinatorial map* $f\colon L \to K$ between CW complexes is a continuous map such that, for every open n-cell e in L, either $f|_e$ is a homeomorphism onto an open n-cell of K or $f(e) \subseteq K^{(n-1)}$; in the latter case, we say that e is *singular*. A *singular combinatorial complex* K is a CW complex such that the attaching maps $\phi_\sigma\colon \partial D^{n+1} \to K^{(n)}$ are singular combinatorial maps with respect to some combinatorial structure on $S^n = \partial D^{n+1}$.

Definition 2.2. Let $k \geq 1$, let X be a singular combinatorial complex, and let W be a compact k-dimensional manifold. An *admissible map* $f\colon W \to X$ is a continuous map together with a combinatorial cell structure on W such that f is a singular combinatorial map. We define the volume $\mathrm{Vol}_X(f)$ to be the number of open k-cells of W on which the restriction of f is a homeomorphism onto an open k-cell of X.

Definition 2.3. Let X be a k-connected singular combinatorial complex and let $\alpha\colon S^k \to X$ be an admissible map. Its filling volume is

$$\mathrm{FVol}_X(\alpha) \coloneqq \min \left\{ \mathrm{Vol}_X(\beta) \mid \beta\colon D^{k+1} \to X \text{ admissible with } \beta|_{\partial D^{k+1}} = \alpha \right\}.$$

The *k-dimensional (homotopical) Dehn function* of X is

$$\delta_X^k(n) \coloneqq \max \left\{ \mathrm{FVol}_X(\alpha) \mid \mathrm{Vol}_X(\alpha) \leq n,\ \alpha : S^k \to X \text{ admissible} \right\}.$$

The *k-dimensional (homotopical) Dehn function* of a group G of type F_{k+1} is $\delta_G^k(n) \coloneqq \delta_X^k(n)$, where X is any k-connected singular combinatorial complex that admits a free, cocompact, cellular G-action.

Remark 2.4. The definition of δ_X^k in [Bri02b, Ril03] uses based spheres and a basepoint in X. Since we consider only complexes with a cocompact group action, our basepoint-independent definition is equivalent to their definition up to $\approx$-equivalence of functions.

For every group G of type F_{k+1}, $k \geq 1$, there is a k-connected singular combinatorial complex that admits a free, cocompact, cellular G-action. It can be constructed as the universal cover of the $(k+1)$-skeleton of a classifying space for G. However, it is not unique in general, and the definition of δ_G^k depends on its choice. Thus, formally, δ_G^k is defined only up to $\approx$-equivalence. More generally, its $\approx$-equivalence class is a quasi-isometry invariant among groups of type F_{k+1}.

In the special case $k = 1$, one can also work with combinatorial maps by considering compact planar combinatorial 2-complexes instead of disks. These are known as van Kampen diagrams. We recall the main facts and definitions that we will use and refer to [Bri02a, BRS07] for further details.

Definition 2.5. Let X be a combinatorial 2-complex. A *van Kampen diagram* (Δ, f, ι) for a null-homotopic combinatorial loop $\gamma\colon S^1 \to X^{(1)}$ is a compact contractible planar combinatorial 2-complex Δ together with a combinatorial map $f\colon \Delta \to X^{(2)}$ such that $f|_{\partial\Delta} \circ \iota = \gamma$, where $\iota\colon S^1 \to \partial\Delta$ is a combinatorial parametrization of $\partial\Delta$. We denote by $\mathrm{Area}(f) \coloneqq \mathrm{Vol}(f)$ the number of 2-cells of Δ.

Remark 2.6. If G is a group that acts freely, cellularly, and cocompactly on a combinatorial 2-complex X, then replacing admissible maps by van Kampen diagrams in the definition of the Dehn function of X yields a function that is $\approx$-equivalent to δ_G^1. We can thus restrict to van Kampen diagrams when working with the one-dimensional Dehn function $\delta_G \coloneqq \delta_G^1$.

Remark 2.7. For every finitely presented group $G \cong \mathcal{P} = \langle S \mid R\rangle$, the universal cover $\widetilde{X}_{\mathcal{P}}$ of its presentation complex is a combinatorial 2-complex. It can be equipped with a labeling of its edges by generators and its 2-cells by relations. For a van Kampen diagram $f\colon \Delta \to X^{(2)}$, this induces a labeling of its edges and 2-cells by generators and relations. Van Kampen's Lemma then provides a translation between van Kampen diagrams for combinatorial loops and free decompositions of null-homotopic words in S as products of conjugates of relations from R. This provides a translation between isoperimetric functions and the complexity of a brute-force approach to the word problem.

There is a homological version of k-dimensional Dehn functions for groups G of type F_{k+1}, defined in terms of cycles α and chains β with $\partial\beta = \alpha$ in the chain complex of the universal cover of a classifying space with finite $(k+1)$-skeleton. In fact, the weaker homological finiteness property FP_{k+1}, defined in terms of finitely generated resolutions by projective $\mathbb{Z}G$-modules, would suffice to define it. We do not discuss these finiteness properties further because we do not require them. As with homotopical Dehn functions, we define homological Dehn functions for more general complexes. We refer to [You11] for further details.

For a CW complex X, we equip its cellular chain complex with integral coefficients $C_i(X,\mathbb{Z})$ with the 1-norm $\|\alpha\| \coloneqq \sum_{i\in I} |a_i|$ for $\alpha = \sum_{i\in I} a_i\sigma_i \in C_i(X,\mathbb{Z})$.

Definition 2.8. Let X be a k-connected CW complex. We define the filling volume of a k-cycle $\alpha \in Z_k(X,\mathbb{Z})$ as

$$\mathrm{FVol}_X^{k+1}(\alpha) \coloneqq \inf\{\|\beta\| \mid \beta \in C_{k+1}(X,\mathbb{Z}),\ \partial\beta = \alpha\}.$$

The $(k+1)$-*dimensional homological Dehn function of* X (with integer coefficients) is

$$\mathrm{FV}_X^{k+1}(n) \coloneqq \sup\left\{\mathrm{FVol}_X^{k+1}(\alpha) \mid \alpha \in Z_k(X,\mathbb{Z}),\ \|\alpha\| \leq n\right\}.$$

For a group G of type F_{k+1}, let X be a k-connected CW complex with a free cellular cocompact G-action. The $(k+1)$-*dimensional homological Dehn function of* G (with integer coefficients) is $\mathrm{FV}_G^{k+1}(n) \coloneqq \mathrm{FV}_X^{k+1}(n)$.

Note that, for historical reasons, there is an index shift by $+1$ between homotopical and homological Dehn functions. As for homotopical Dehn functions, the $(k+1)$-dimensional homological Dehn function of a group G of type F_{k+1} is well-defined up to $\approx$-equivalence [You11, Lemma 1]. One can also define homological Dehn functions with other normed coefficient rings. Here we restrict to integral coefficients for simplicity.

2.2. **Collapsible complexes.** We recall the terminology of collapses and collapsible complexes, following Whitehead [Whi39].

Definition 2.9. Let K be a simplicial complex. A simplex $\tau \in K$ is called a *free face* of a simplex $\sigma \in K$ if τ is a proper face of σ and σ is the unique maximal simplex in K containing τ.

Removing all intermediate simplices between τ and σ from K, that is, removing the set $\{\tau' \mid \tau \subseteq \tau' \subseteq \sigma\}$ yields a subcomplex L. This operation is called a *collapse*, and we write $K \searrow L$. If τ has codimension 1 in σ, it is called an *elementary collapse*.

The complex K is called *collapsible* if there exists a sequence of elementary collapses collapsing K to a point.

Lemma 2.10 ([Whi39]). *Let K be a finite collapsible simplicial complex of dimension n. Then there exists a sequence of collapses*

$$K = K_0 \searrow K_1 \searrow \ldots \searrow K_r = \{*\},$$

such that, for every $1 \leq i \leq n-1$, every collapse of an $(i+1)$-cell precedes every collapse of an i-cell.

Proof. Collapsing an i-cell cannot create new free faces of $(i+1)$-cells; hence, if a collapse of an i-cell occurs after a collapse of an $(i+1)$-cell, they are independent and can also be performed in the opposite order. $\square$

Example 2.11. There is only one type of collapse of a 1-cell, namely the elementary collapse of the cell along one of its vertices. There are two types of collapses of a 2-cell: collapse along a free vertex and collapse along a free edge. The non-elementary collapse along a free vertex is equivalent to first collapsing an adjacent free edge and then collapsing the remaining edge through an elementary collapse of a 1-cell (see Fig. 1).

Lemma 2.12. *Every finite contractible planar simplicial 2-complex K is collapsible.*

Proof. If K is not a tree, then it has a 1-simplex e on its boundary that is a face of a 2-simplex σ. Performing an elementary collapse of σ along e reduces the number of 2-simplices of the complex by one. Once all 2-simplices have been collapsed, the remaining complex is 1-dimensional and still contractible, hence a tree, which is clearly collapsible. $\square$

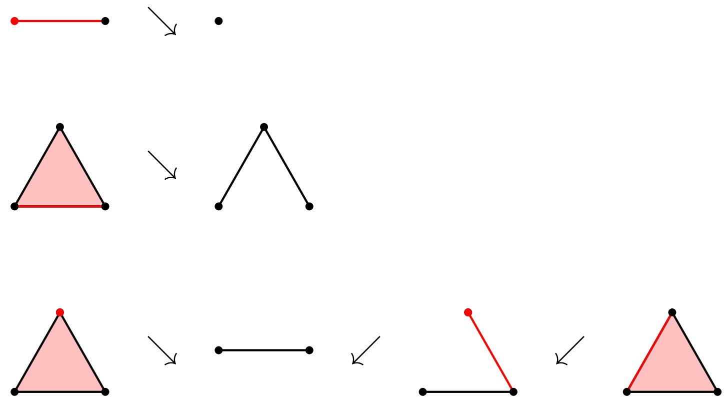

FIGURE 1. Types of collapses of a 1-cell and a 2-cell. The second type of 2-cell collapse is the composition of two elementary collapses.

2.3. **G-complexes and the Haefliger construction.** To prove our main results, we use the following construction of classifying spaces for complexes of groups due to Haefliger [Hae92]; see also [BH99, III.$\mathcal{C}$.2]. We refer to [DK18] for a more detailed account of the construction. We follow the conventions in [BH99, III.$\mathcal{C}$.2] and use similar notation.

Theorem 2.13 (Haefliger construction [Hae92]). *Let X be a cocompact G-simplicial complex. Then there exists a free regular G-CW complex $\widehat{X}$ together with a G-equivariant map $\pi\colon \widehat{X} \to X$ and a filtration $\widehat{X} = \operatorname{colim}_n \widehat{X}^{[n]}$ with $\widehat{X}^{[-1]} = \emptyset$ such that*

(1) *π restricts to a homotopy equivalence $\widehat{X}^{[n]} \to X^{(n)}$;*
(2) *$\widehat{X}^{[n]}$ fits into a pushout diagram of the form*

$$\begin{tikzcd}
\coprod_{\sigma\in I_n} G \times_{G_\sigma} EG_\sigma \times \partial\sigma \arrow[r] \arrow[d, hook] & \widehat{X}^{[n-1]} \arrow[d] \\
\coprod_{\sigma\in I_n} G \times_{G_\sigma} EG_\sigma \times \sigma \arrow[r] & \widehat{X}^{[n]}.
\end{tikzcd}$$

Construction 2.14. Let $Y = G\backslash X$ and $p\colon X \to Y$ be the projection map. For every simplex σ of Y, choose a simplex $\overline{\sigma}$ in X such that $p(\overline{\sigma}) = \sigma$. Let $\tau \subset \sigma$. Because G acts without inversion, there exists a unique face τ_σ of $\overline{\sigma}$ such that $p(\tau_\sigma) = \tau$. Then in general $\tau_\sigma \neq \overline{\tau}$. However, τ_σ and $\overline{\tau}$ lie in the same G-orbit, so we can choose an element $h_{\tau,\sigma} \in G$ such that $h_{\tau,\sigma}\tau_\sigma = \overline{\tau}$. For each $\sigma \in Y$, let G_σ be the stabilizer of $\overline{\sigma}$ and define

$$c_{\tau,\sigma}\colon G_\sigma \to G_\tau, \quad g \mapsto h_{\tau,\sigma} g h_{\tau,\sigma}^{-1}$$

for every face inclusion $\tau \subseteq \sigma$. The fiber $\pi^{-1}(\sigma)$ of a cell $\sigma \in X$ is the copy of $EG_{p(\sigma)} \times \overline{p(\sigma)}$ in

$$G \times_{G_{p(\sigma)}} EG_{p(\sigma)} \times \overline{p(\sigma)} = \coprod_{G/G_{p(\sigma)}} EG_{p(\sigma)} \times \overline{p(\sigma)}$$

indexed by the coset $gG_{p(\sigma)}$ satisfying $g \cdot \overline{p(\sigma)} = \sigma$. It is sometimes necessary to distinguish how these copies of $EG_{p(\sigma)}$ sit inside $\widehat{X}$. In this case we write $\widehat{X}_\sigma$ for the copy such that $\pi^{-1}(\sigma) = \{gG_\sigma\} \times \widehat{X}_\sigma \times \overline{p(\sigma)}$.

In particular, for an edge $e = \{u, v\} \in G\backslash X$, let u_e and v_e be the faces of $\overline{e}$ with $p(u_e) = u$ and $p(v_e) = v$, respectively. Since $h_{u,e}u_e = \overline{u}$ by construction, the vertex of the edge $g \cdot \overline{e}$ lying over u is $gu_e = gh_{u,e}^{-1}\overline{u}$. The attaching map of the edge fiber to the vertex fiber at u is therefore the G-equivariant map

$$\begin{aligned} \psi_{u,e}\colon G \times_{G_e} EG_e \times \{u_e\} &\longrightarrow G \times_{G_u} EG_u \\ [g, x] &\longmapsto [gh_{u,e}^{-1}, \psi_{u,e}(x)], \end{aligned}$$

where $\psi_{u,e}\colon EG_e \to EG_u$ is the map induced by $c_{u,e}$, and analogously for the attaching map to the fiber of v. More generally, if $\tau \subseteq \sigma$ and $\tau_\sigma \subseteq \overline{\sigma}$ projecting to τ, we consider the map

$$\begin{aligned} \Phi_{\tau,\sigma}\colon G \times_{G_\sigma} EG_\sigma \times \tau_\sigma &\longrightarrow G \times_{G_\tau} EG_\tau \times \overline{\tau} \\ [g, x, y] &\longmapsto [gh_{\tau,\sigma}^{-1}, \psi_{\tau,\sigma}(x), h_{\tau,\sigma}y], \end{aligned}$$

where again $\psi_{\tau,\sigma}\colon EG_\sigma \to EG_\tau$ is induced by $c_{\tau,\sigma}$. This is not the attaching map itself because, for iterated face inclusions $\rho \subseteq \tau \subseteq \sigma$, we cannot in general assume that $\psi_{\rho,\sigma} = \psi_{\rho,\tau} \circ \psi_{\tau,\sigma}$. However, they always differ by homotopies

$$\Psi_{\rho,\tau,\sigma}\colon EG_\sigma \times [0, 1] \to EG_\rho$$

depending on the monodromy element

$$g_{\rho,\tau,\sigma} := h_{\rho,\tau}h_{\tau,\sigma}h_{\rho,\sigma}^{-1}.$$

The explicit construction of the correct attaching maps is more involved; we refer to [DK18, Section 5.8.2] for details.

For our purposes, we require the CW structure on $\widehat{X}$ and the map π to be well behaved up to a certain dimension. This can be achieved by performing a more careful version of Haefliger's construction.

Remark 2.15. Choose finite generating sets S_e for the edge stabilizers and S_v for the vertex stabilizers in $G\backslash X$. After replacing S_v by $S_v \cup \bigcup_{v\subset e} c_{v,e}(S_e)$, we may assume that $c_{v,e}(S_e) \subseteq S_v$ for every incidence $v \subset e$. Choose the models EG_e and EG_v such that their 1-skeleta are the corresponding Cayley graphs. Then $\psi_{v,e}\colon EG_e^{(1)} \to EG_v^{(1)}$ sends each edge labeled by $s \in S_e$ to the edge labeled by $c_{v,e}(s)$, and is therefore injective. This injectivity will be used repeatedly in our proofs.

The setup in Remark 2.15 is sufficient for the proof of Theorem B in Section 3 after observing that one can choose the cells over edges of X to be product cells (see Remark 2.22). In higher dimensions we will require $\widehat{X}$ and π to be polyhedral up to a given dimension and more work is required to guarantee this. We start by introducing the terminology of polyhedral complexes and maps.

Definition 2.16 (Polyhedral complex)**.** A *polyhedral complex* is a combinatorial complex in which every n-cell is equipped with the structure of an n-dimensional compact convex polytope and the attaching maps are affine on open cells.

Definition 2.17 (Polyhedral map)**.** A map $f\colon X \to Y$ between polyhedral complexes is *polyhedral* if, for every cell σ of X, there is a cell τ of Y such that $f(\sigma^\circ) = \tau^\circ$ and the restriction $f|_\sigma\colon \sigma \to \tau$ is an affine surjection.

Definition 2.18 (G-polyhedral complex)**.** Let G be a group. A *G-polyhedral complex* is a G-space X that admits a filtration $\emptyset = X^{(-1)} \subseteq X^{(0)} \subseteq X^{(1)} \subseteq \cdots$ with $X = \operatorname{colim}_n X^{(n)}$, such that each $X^{(n)}$ fits into a polyhedral equivariant pushout diagram of the form

$$\begin{tikzcd} \coprod_{i\in I_n} G/G_i \times \partial P_i \arrow[r, "\alpha_n"] \arrow[d, hook] & X^{(n-1)} \arrow[d] \\ \coprod_{i\in I_n} G/G_i \times P_i \arrow[r, "e_n"] & X^{(n)}, \end{tikzcd}$$

where

(i) I_n is an index set;
(ii) each P_i is a compact convex n-polytope;
(iii) each $G_i \leq G$ is a subgroup called the *cell stabilizer* of the cell represented by i.

The complex X is *cocompact* if all index sets I_n are finite and only finitely many are non-empty. It is called *regular* if every attaching map e_n is a homeomorphism onto its image.

Remark 2.19. Equivalently, a G-polyhedral complex is a polyhedral complex X on which G acts by polyhedral automorphisms without inversions.

Definition 2.20. A *G-simplicial complex* is a G-polyhedral complex, where all cells P_i, $i \in I_n$ are simplices.

Proposition 2.21. *Let X be a cocompact G-simplicial 1-complex and let $n \geq 1$. Assume that, G_v is of type F_n for every vertex $v \in X$ and G_e is of type F_{n-1} for every edge e of X. Then $\widehat{X}$ can be constructed such that $\widehat{X}^{(n)}$ is a cocompact G-polyhedral complex and*

(1) *$\pi|_{\widehat{X}^{(n)}}$ is a polyhedral map;*
(2) *for every incidence $v \subset e \subset X$ of a vertex and an edge, the attaching map $\psi_{v,e}\colon EG_e^{(n-1)} \to EG_v^{(n-1)}$ is a simplicial embedding;*
(3) *for every edge $e \subset X$, each cell of $\widehat{X}^{(n)}$ whose interior projects onto e° is of the form $P \times e$, where P is a cell of EG_e of dimension at most $n-1$.*

Proof. Let $Y = G\backslash X$. Since X is cocompact, Y has only finitely many cells. For every p-cell σ of Y with $p \in \{0,1\}$, choose a model EG_σ whose $(n-p)$-skeleton is a cocompact free G_σ-simplicial complex. Such models can be obtained by iterative simplicial approximation of the cell attaching maps. For a vertex v of an edge e, one can similarly arrange that the map $\psi_{v,e}$ is simplicial on the relevant skeleta. Moreover, one can arrange that $\psi_{v,e}$ is injective by passing to mapping cylinders. Running the Haefliger construction with these maps as input data gives a complex $\widehat{X}$ satisfying (2); see [DK18, Section 5.8.2] for details. To obtain (1), (3), and cocompactness of $\widehat{X}^{(n)}$, choose the polyhedral structure on $G \times_{G_e} EG_e \times e$ to be induced by the product cells $P \times e$, where P is a cell of EG_e. $\square$

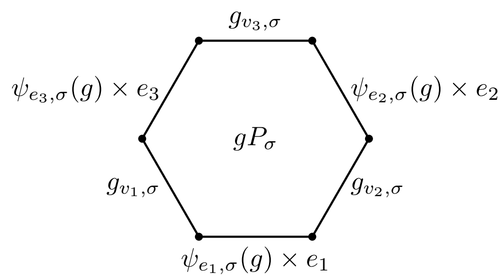

FIGURE 2. The hexagonal cell gP_σ with its boundary labels.

Remark 2.22. One can strengthen the construction from Remark 2.15 to also satisfy the conditions (2) and (3) of Proposition 2.21 for $n = 2$ using the same techniques. This is all we need for the proof of Theorem B.

For the proof of Theorem C we will have to work with a 2-dimensional complex X. In this case we replace the Haefliger construction by a more direct construction of a simply connected singular combinatorial complex $\widehat{X}$ together with a singular combinatorial map $\pi\colon \widehat{X} \to X$, which we will describe now.

Start with a complex Y obtained by applying Proposition 2.21 to the restriction of the G-action to $X^{(1)}$, that is, copies of EG_v for every vertex $v \in X$ and $EG_e \times e$ for every edge e, glued together using the maps $\psi_{v,e}$.

We now attach a G_σ-orbit of hexagonal 2-cells gP_σ for every 2-cell σ of X. To do so identify $EG_\sigma^{(0)}$ with G_σ and for every edge $e \subset \sigma$ fix a G_σ-equivariant inclusion $\psi_{e,\sigma}\colon EG_\sigma^{(0)} \to EG_e^{(0)}$ induced by $c_{e,\sigma}$. For the edges $e, e' \subset \sigma$ adjacent to a vertex $v \subset \sigma$, the compositions $\psi_{v,e} \circ \psi_{e,\sigma}$ and $\psi_{v,e'} \circ \psi_{e',\sigma}$ differ by an element $g_{v,\sigma} \in G_v$ only depending on the orbit of σ. Note that the elements $g_{v,\sigma}$ are products of monodromy elements $g_{v,e,\sigma}$. Since $X\backslash G$ is finite, we may assume that the finite generating set S_v contains all such elements $g_{v,\sigma}$. For $g \in G_\sigma$ we attach the hexagonal 2-cell gP_σ along the loop described by the edges corresponding to the $g_{v,\sigma}$ inside EG_v for the vertices $v \subset \sigma$ and the edges $\psi_{e,\sigma}(g) \times e$ inside $EG_e \times e$ for the edges $e \subset \sigma$ as depicted in Fig. 2.

By construction $\widehat{X}$ is a singular combinatorial complex and the singular combinatorial projection $\pi\colon \widehat{X} \to X$ is the obvious extension of the projection of Y to $X^{(1)}$ by the projections $gP_\sigma \to \sigma$. Finally, for simple connectivity of $\widehat{X}$ note that this is a direct consequence of Theorem 3.3 and Proposition 3.6 which form part of the proof of Theorem C and do not rely on it.

Definition 2.23. A G-polyhedral complex X is called *locally finite in degree n* if, for every n-cell σ of X and every $(n-1)$-face τ of σ, the index $[G_\tau : G_\sigma]$ is finite.

Example 2.24. Let $\mathcal{G} = (\Gamma, G_e, G_v)$ be a graph of groups with fundamental group $G = \pi_1(\mathcal{G})$ and corresponding Bass–Serre tree T. Then T is a G-polyhedral complex and is cocompact if and only if Γ is a finite graph, as $G\backslash T = \Gamma$.

For each directed edge $e = (u, v)$ of T, consider the edge $p(e) = (p(u), p(v))$. Then $[G_v : G_e] = [G_{p(v)} : \iota_{p(e)}(G_{p(e)})]$ and T is locally finite in degree 1 if and only if T is locally finite as a graph. Since T is 1-dimensional, there are no higher-dimensional cells, so T is locally finite in degree n for every $n \geq 2$.

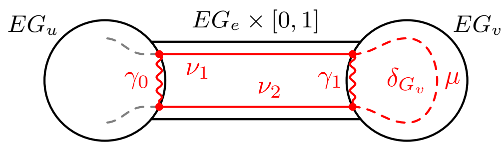

FIGURE 3. Lift of the collapse of an edge $\{u, v\}$ to a homotopy in $\widehat{X}$. The loop $\mu \cdot \gamma_1$ is filled using the Dehn function δ_{G_v} of the vertex stabilizer G_v.

3. A BROWN THEOREM FOR ONE-DIMENSIONAL DEHN FUNCTIONS

Throughout this section, let X be a G-simplicial complex that is 1-connected, cocompact, and locally finite in degree 2. We further assume that all vertex stabilizers are finitely presented and all edge stabilizers are finitely generated.

Because X is cocompact, $G\backslash X$ has finite 1-skeleton. Let $\widehat{X}$ and $\pi\colon \widehat{X} \to X$ be as described in Remark 2.22. We write edist_X for the maximum of the distortion functions $\operatorname{dist}_{G_e}^G$ over the edges e in X. Recall that, as in Construction 2.14, we fix a representative of the conjugacy class of every cell stabilizer, which we denote by G_σ. Since there are only finitely many orbits of vertices and edges, edist_X is well-defined, and similar considerations apply to the other maxima below.

3.1. **Lifting fillings of tree van Kampen diagrams.** We start with the special case in which the projection of a combinatorial loop in $\widehat{X}$ to X admits a van Kampen diagram in X with no 2-cells or, equivalently, a van Kampen diagram whose underlying complex is a tree. Note that if X is 1-dimensional, i.e., a Bass–Serre tree, then this holds for all van Kampen diagrams. Thus, the results in this section produce an upper bound on the Dehn function of a graph of groups.

Lemma 3.1 (Edge collapse)**.** *Let $\gamma\colon [0,1] \to \widehat{X}$ be a combinatorial path of length n. Let $e = \{u, v\}$ be an edge of X. Assume γ decomposes as $\gamma = \nu_1 \cdot \mu \cdot \nu_2$, where $\pi(\gamma(0)) = \pi(\gamma(1)) = u$, $\pi \circ \mu = \mathrm{const}_v$, and $\ell(\nu_1) = \ell(\nu_2) = 1$. Then γ is homotopic to a path $\gamma'\colon [0,1] \to \widehat{X}$ via a homotopy h fixing endpoints such that*

1. *$\ell(\gamma') = d_{G_e}(\gamma(0), \gamma(1))$;*
2. *h has at most $\delta_{G_v}\big(\ell(\gamma) + \ell(\gamma')\big) + \ell(\gamma')$ nonsingular 2-cells.*

Proof. By assumption, ν_1 and ν_2 define horizontal edges in $EG_e \times [0,1]$. Let $\gamma_0'\colon [0,1] \to EG_e \times \{0\}$ be a combinatorial path of minimal length from $\nu_1(0) = \gamma(0)$ to $\nu_2(1) = \gamma(1)$. Then $\ell(\gamma_0') = d_{G_e}(\gamma(0), \gamma(1))$. Denote by γ_1' the parallel copy of γ_0' inside $EG_e \times \{1\}$. Then $\gamma_1'(0) = \mu(0)$ and $\gamma_1'(1) = \mu(1)$. Thus, there is a homotopy from μ to γ_1' with at most $\delta_{G_v}\big(\ell(\gamma) + \ell(\gamma_0')\big)$ 2-cells. Moreover, parallel translation of γ_1' to γ_0' inside $EG_e \times [0,1]$ defines a homotopy from $\nu_1 \cdot \gamma_1' \cdot \nu_2$ to γ_0' with at most $\ell(\gamma_0')$ 2-cells (see Fig. 3). Composing these two homotopies and choosing $\gamma' \coloneqq \gamma_0'$ completes the proof. $\square$

Remark 3.2. Let g_0 and g_1 be the elements in G_e represented by $\gamma'(0) = \gamma(0)$ and $\gamma'(1) = \gamma(1)$ respectively. Then

$$d_{\widehat{X}_e}(\gamma'(0), \gamma'(1)) = d_{G_e}(g_0, g_1) \leq \operatorname{dist}_{G_e}^G\big(d_G(g_0, g_1)\big).$$

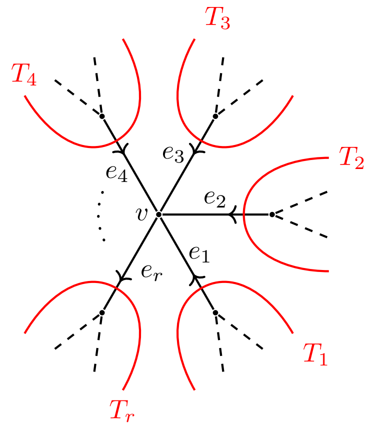

FIGURE 4. The subtrees T_i around a vertex v. Arrows indicate the direction of the collapse of the adjacent edges.

Moreover, G and $\widehat{X}^{(1)}$ are quasi-isometric, so there exists a constant $D = D(X) \geq 1$ such that $d_G(g_0, g_1) \leq D \cdot d_{\widehat{X}}(\gamma(0), \gamma(1))$. Consequently,

$$\ell(\gamma') \leq \operatorname{dist}^G_{G_e}\big(D \cdot d_{\widehat{X}}(\gamma(0), \gamma(1))\big) \leq \operatorname{dist}^G_{G_e}\big(D \cdot \ell(\gamma)\big). \tag{3.1}$$

Theorem 3.3. *Let $\gamma\colon [0,1] \to \widehat{X}$ be a combinatorial path of length n. Assume that $\pi \circ \gamma$ admits a van Kampen diagram (T, f, ι) where T is a tree. Then there exists a filling of γ consisting of at most $n \cdot \big(\operatorname{edist}_X(Dn) + \max_{v \in X^{(0)}}\big(\delta_{G_v}(\overline{\operatorname{edist}_X}(Dn))\big)\big)$ 2-cells.*

Proof. For ease of notation, we identify the vertices and edges of T with their images under f in X.

Fix a sequence of edges along which T is collapsed to a single vertex v_0. Let $v \in T^{(0)}$ be a vertex, let $e_1, \ldots, e_r$ be the edges adjacent to v in the order in which they are collapsed, and denote by T_i the subtree extending outwards from v along the edge e_i. This situation is depicted in Fig. 4.

For $1 \leq i \leq r-1$ let γ_{e_i} be the subpath of γ corresponding to T_i and γ'_{e_i} be the path from $\gamma_{e_i}(0)$ to $\gamma_{e_i}(1)$ obtained when applying Lemma 3.1 to collapse the edge e_i connecting T_i to v in our iterative procedure. Then

$$\ell(\gamma'_{e_i}) \leq \operatorname{dist}^G_{G_{e_i}}\big(D \cdot \ell(\gamma_{e_i})\big),$$

by Eq. (3.1). The same holds for the path γ'_{e_r} when collapsing e_r. Therefore, the homotopy we obtain from Lemma 3.1 when we collapse the edge e_r from v has at most

$$\begin{aligned}
\delta_{G_v}\Big(\ell(\gamma) + \sum_{i=1}^{r} \operatorname{dist}^G_{G_{e_i}}\big(D \cdot \ell(\gamma_{e_i})\big)\Big) &\leq \delta_{G_v}\Big(\ell(\gamma) + \sum_{i=1}^{r} \operatorname{edist}_X\big(D \cdot \ell(\gamma_{e_i})\big)\Big) \\
&\leq \delta_{G_v}\Big(\ell(\gamma) + \overline{\operatorname{edist}_X}\Big(\sum_{i=1}^{r} D \cdot \ell(\gamma_{e_i})\Big)\Big) \\
&\leq \delta_{G_v}\Big(n + \overline{\operatorname{edist}_X}(Dn)\Big)
\end{aligned}$$

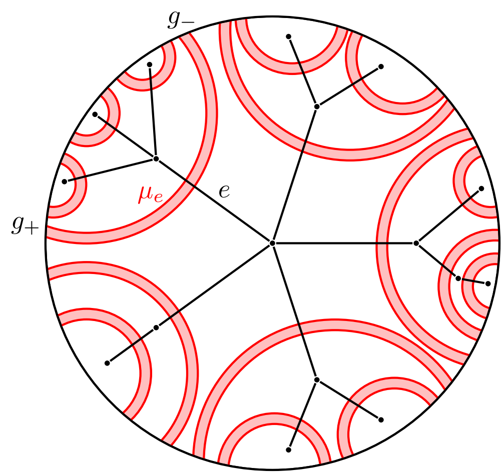

FIGURE 5. A tree van Kampen diagram and the corresponding disk partition obtained by the connecting paths of the pairings $\{g_+, g_-\}$.

2-cells coming from the filling inside EG_v and at most

$$\ell(\gamma'_{e_r}) \leq \mathrm{dist}^G_{G_{e_r}}\big(D \cdot \ell(\gamma_{e_r})\big) \leq \mathrm{edist}_X(D \cdot \ell(\gamma_{e_r})) \leq \mathrm{edist}_X(D \cdot \ell(\gamma))$$

2-cells coming from the collar inside EG_{e_r}. After iteratively collapsing T to a single vertex v_0, we have obtained a homotopy h of γ to a path γ' such that $\pi \circ \gamma' = \mathrm{const}_{v_0}$. The same analysis as above shows that $\ell(\gamma') \leq n + \overline{\mathrm{edist}_X}(n)$. Filling this loop inside EG_{v_0} completes the homotopy h to a filling of γ. Because T has at most $\frac{n}{2}$ edges, the resulting filling has area at most

$$\frac{n}{2} \cdot \Big(\max_{v \in X^{(0)}} \delta_{G_v}\big(n + \overline{\mathrm{edist}_X}(Dn)\big) + \overline{\mathrm{edist}_X}(Dn)\Big) + \delta_{G_{v_0}}\big(n + \overline{\mathrm{edist}_X}(Dn)\big)$$
$$\leq n \cdot \Big(\max_{v \in X^{(0)}} \delta_{G_v}\big(n + \overline{\mathrm{edist}_X}(Dn)\big) + \overline{\mathrm{edist}_X}(Dn)\Big). \qquad \square$$

3.2. **Intuition for the case of trees.** Before proceeding, we give an intuitive explanation of the proof when the projection $\pi \circ \gamma$ to X of a combinatorial loop $\gamma \colon [0,1] \to \widehat{X}$ admits a van Kampen diagram with no 2-cells. Due to our choice of $\widehat{X}$, we can partition the edges defining γ into horizontal edges, which project onto edges in X under $\pi \circ \gamma$, and vertical edges, which project onto vertices in X under $\pi \circ \gamma$.

If we choose a van Kampen diagram for $\pi \circ \gamma$ with no 2-cells, then its underlying complex is a tree and $\pi \circ \gamma$ corresponds to a parametrization of its boundary. This parametrization traverses every edge of the van Kampen diagram precisely twice in opposite directions, inducing a pairing between the edges that define $\pi \circ \gamma$ and thus between the horizontal edges of γ.

For every pair of edges e_+, e_- of γ, there is an edge e of X and elements $g_+, g_- \in G_e$ such that $\gamma|_{e_+}$ and $\gamma|_{e_-}$ parametrize edges of the form $\{g_+\} \times [0,1]$ and $\{g_-\} \times [0,1]$ in $EG_e \times [0,1]$. There is a path $\mu_e \colon [0,1] \to EG_e$ between g_+ and g_- of length $\ell(\mu_e) \leq \mathrm{dist}^G_{G_e}(g_+, g_-)$.

To build a van Kampen diagram for γ, we label the boundary of a 2-disc by the edges defining γ and subdivide it along corridors between paired edges e_+, e_- labeled by the paths μ_e (see Fig. 5). The boundary of every bounded region in this disc describes a word of length $\leq \operatorname{edist}(\ell(\gamma))$ in some vertex group G_v. This implies that there is a van Kampen diagram for γ satisfying the upper area bound in Theorem 3.3.

3.3. **Reduction to trees.** If X is not 1-dimensional, we have to consider loops in X that do not admit a van Kampen diagram with no 2-cells. We will deal with this case by reducing it to the case of a van Kampen diagram that is a tree. For this, we use Lemma 2.12, which states that any finite contractible planar simplicial 2-complex can be collapsed onto a tree by a sequence of elementary collapses of 2-simplices along codimension-one faces.

Lemma 3.4 (Face collapse)**.** *Let $\gamma\colon [0,1] \to \widehat{X}$ be a combinatorial path. Let $e = \{v_1, v_2\}$ be an edge of X contained in a 2-simplex $\sigma = \{v_1, v_2, v_3\}$. Assume that $\pi(\gamma(0)) = v_1$, $\pi(\gamma(1)) = v_2$, and $\ell(\gamma) = 1$.*

Then γ is homotopic to a path $\gamma' = \mu_1 \cdot \lambda_1 \cdot \nu_1 \cdot \lambda_3 \cdot \nu_2 \cdot \lambda_2 \cdot \mu_2 \colon [0,1] \to \widehat{X}$ via a homotopy h fixing endpoints such that

1. *$\pi \circ \mu_i = \mathrm{const}_{v_i}$ for $i \in \{1,2\}$, and $\pi \circ \lambda_i = \mathrm{const}_{v_i}$ for $i \in \{1,2,3\}$;*
2. *$\pi \circ \nu_1$ parametrizes $\{v_1, v_3\}$ and $\pi \circ \nu_2$ parametrizes $\{v_3, v_2\}$;*
3. *ν_1 and ν_2 have length 1;*
4. *$\ell(\lambda_i) \leq 1$ for $i \in \{1,2,3\}$;*
5. *$\ell(\mu_1),\ \ell(\mu_2) \leq [G_e : G_\sigma] - 1$;*
6. *h has at most $[G_e : G_\sigma]$ many 2-cells.*

Proof. Since $[G_e : G_\sigma] < \infty$, there is a path $\mu_1 \colon [0,1] \to EG_e \times \{0\}$ of length $\ell(\mu_1) \leq [G_e : G_\sigma] - 1$ such that $\mu_1(0) = \gamma(0)$ and $\mu_1(1)$ lies in the image of $EG_\sigma^{(0)}$ in $EG_e^{(0)}$. We choose μ_2 to be the path parallel to μ_1 in $EG_e \times \{1\}$ parametrized in the opposite direction. By Remark 2.22, there is, for some g, a 2-cell gP_σ over σ whose boundary contains the edge from $\mu_1(1)$ to $\mu_2(0)$, edges ν_1 and ν_2 projecting onto $\{v_1, v_3\}$ and $\{v_3, v_2\}$, respectively, and possibly edges λ_i in the vertex fibers, where λ_i is constant if this edge is absent. By construction, parallel translation of μ_1 inside $EG_e \times [0,1]$, together with gP_σ, defines the required homotopy from γ to γ' with $\ell(\mu_1) + 1 \leq [G_e : G_\sigma]$ 2-cells. $\square$

Remark 3.5. If X is a cocompact G-simplicial complex that is locally finite in degree 2, then

$$C(X) \coloneqq \max\{[G_e : G_\sigma] \mid \ \sigma \text{ is a 2-cell and } e \subset \sigma \text{ is an edge}\} - 1$$

is finite.

Proposition 3.6. *Let $\gamma\colon [0,1] \to \widehat{X}$ be a combinatorial loop of length n. Let (Δ, f, ι) be a van Kampen diagram for $\pi \circ \gamma$. Let $C \coloneqq C(X)$. Then γ is homotopic to a loop $\gamma' \colon [0,1] \to \widehat{X}$ via a homotopy h that fixes the basepoint such that*

1. *$\pi \circ \gamma'$ has a van Kampen diagram that is a tree;*
2. *$\ell(\gamma') \leq \ell(\gamma) + (2C + 4) \cdot \mathrm{Area}(f)$;*
3. *h has at most $\mathrm{Area}(f) \cdot (C + 1)$ 2-cells.*

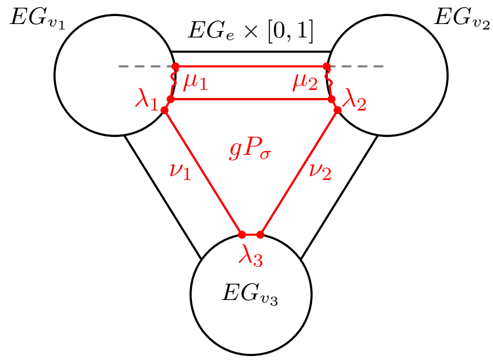

FIGURE 6. Lift of the collapse of a 2-simplex $\{v_1, v_2, v_3\}$ to a homotopy in $\widehat{X}$.

Proof. Since Δ is a contractible planar simplicial 2-complex, it is collapsible. Thus, by Lemma 2.10, there is a sequence of $\mathrm{Area}(f)$ elementary collapses

$$\Delta = \Delta_0 \searrow \Delta_1 \searrow \dots \searrow \Delta_{\mathrm{Area}(f)}.$$

Set $\Delta' \coloneqq \Delta_{\mathrm{Area}(f)}$. This sequence transforms Δ into a tree Δ' and leaves the vertex set of the diagram invariant. By Lemma 3.4, it induces a sequence of basepoint-preserving homotopies $h_1, \dots, h_{\mathrm{Area}(f)}$ between loops $\gamma = \gamma_0, \gamma_1, \dots, \gamma_{\mathrm{Area}(f)}$. Set $\gamma' \coloneqq \gamma_{\mathrm{Area}(f)}$. Then

- $\pi \circ \gamma'$ parametrizes the boundary of Δ' and thus admits a van Kampen diagram that is a tree;
- $\ell(\gamma_{i+1}) \leq \ell(\gamma_i) + 2C + 4$;
- h_i has at most $C+1$ many 2-cells.

Choosing h to be the concatenation of the homotopies $h_1, \dots, h_{\mathrm{Area}(f)}$ thus completes the proof. $\square$

We can now prove Theorems B and C.

Proof of Theorem C. Up to $\approx$-equivalence we may assume that all functions δ_X, δ_{G_v}, and edist_X are at least linear. It follows from Proposition 3.6 and Theorem 3.3 that $\widehat{X}$ is simply connected and by construction it admits a free cocompact cellular G-action. Hence, $\delta_{\widehat{X}} \approx \delta_G$.

Together, Proposition 3.6 and Theorem 3.3 yield that

$$\delta_G(n) \leq m_n \cdot \Big(\max_{v \in X^{(0)}} \delta_{G_v}\big(m_n + \overline{\mathrm{edist}_X}(Dm_n)\big) + \mathrm{edist}_X(Dm_n) \Big) + \delta_X(n)(C+1),$$

where $m_n = n + (2C+4) \cdot \delta_X(n)$. As $m_n \preccurlyeq \delta_X(n)$, it follows that

$$m_n + \overline{\mathrm{edist}_X}(Dm_n) \preccurlyeq \overline{\mathrm{edist}_X}(Dm_n) \preccurlyeq \overline{\mathrm{edist}_X}(\delta_X(n))$$

and

$$\delta_{G_v}\big(m_n + \overline{\mathrm{edist}_X}(Dm_n)\big) + \mathrm{edist}_X(Dm_n) \preccurlyeq \delta_{G_v}\big(\overline{\mathrm{edist}_X}(\delta_X(n))\big),$$

so the claimed upper bound follows. $\square$

Proof of Theorem B. This is the special case of Theorem C when X is a tree. $\square$

Remark 3.7. Theorem C generalizes to actions on polyhedral 2-complexes. Indeed, the construction in Remark 2.22 works for arbitrary n-gons by attaching a $2n$-cell instead of a hexagon and similarly Lemmas 3.4 and 2.12 generalize to arbitrary n-gons up to enlarging the constants. We restricted ourselves to the case of simplicial complexes to simplify the exposition.

4. Higher Dehn functions

In this section, we explain how our proof can be adapted to obtain upper bounds on higher Dehn functions of graphs of groups. As a consequence, we prove Theorem D and Corollary E.

4.1. **Lifting fillings of spheres and the second Dehn function.** We now provide a proof of the upper bound on the second Dehn function of a group acting cocompactly on a tree. The main result of this subsection is originally due to Barnard, Brady, and Dani (see [BBD12, Proposition 6.1]). We include a proof for the reader's convenience.

Throughout this section, let X be a cocompact G-simplicial tree whose vertex stabilizers are of type F_3 and whose edge stabilizers are of type F_2. Let $\widehat{X}$ and $\pi\colon \widehat{X} \to X$ be the complex and map obtained from Proposition 2.21 for $n = 3$, that is, $\widehat{X}^{(3)}$ is a cocompact G-polyhedral complex and $\pi|_{\widehat{X}^{(3)}}$ is a polyhedral map.

Definition 4.1. Let $f\colon S \cong S^2 \to X$ be a polyhedral map from a finite polyhedral 2-sphere. A *tree van Kampen diagram* for f is a triple (T, F, ι) consisting of a finite tree T and polyhedral maps $\iota\colon S \to T$, $F\colon T \to X$ such that

(i) $F \circ \iota = f$;
(ii) $\iota^{-1}(v)$ is connected for every vertex v of T;
(iii) for every edge e of T, there is a polyhedral homeomorphism $\iota^{-1}(e^\circ) \cong S^1 \times e^\circ$ under which ι corresponds to projection onto e°.

Remark 4.2. The definition of a singular combinatorial map $f\colon S^2 \to \widehat{X}$ implies that singular 1-cells are mapped onto 0-cells and that singular 2-cells are mapped onto the image of their attaching maps in the 1-skeleton, which is a polyhedral subcomplex of $\widehat{X}^{(1)}$. A suitable simplicial subdivision of the singular 2-cells of S^2 thus yields a polyhedral map $f'\colon S^2 \to \widehat{X}$. By construction, f and f' have the same image. Consequently, any filling $F'\colon D^3 \to \widehat{X}$ of f' induces a filling $F\colon D^3 \to \widehat{X}$ of f with $\mathrm{Vol}(F) = \mathrm{Vol}(F')$. Thus, we can restrict to bounding the filling volume of polyhedral maps $S^2 \to \widehat{X}$ to derive an upper bound on δ^2_G.

The same line of argument allows us to replace singular combinatorial fillings of combinatorial loops by polyhedral fillings.

Proposition 4.3. *Let X be a polyhedral complex and let $f\colon S \cong S^2 \to X$ be a polyhedral map from a finite polyhedral 2-sphere. If $f(S)$ is a tree, then f admits a tree van Kampen diagram.*

Proof. Define an equivalence relation on S by declaring $x \sim y$ if $f(x) = f(y)$ and x and y lie in the same connected component of $f^{-1}(f(x))$. Let $T \coloneqq S/\sim$ and denote the quotient map by $\iota\colon S \to T$. Since f is constant on equivalence classes, there is an induced map $F\colon T \to X$ such that $F \circ \iota = f$.

Since S is compact and f is polyhedral, T is a finite graph. Hence, with the induced polyhedral structure on T, the maps ι and F are polyhedral.

Let e be an edge of T. The restriction $\iota^{-1}(e^\circ) \to e^\circ$ is a PL bundle whose fibers are compact 1-manifolds. Its fibers have empty boundary because S has no boundary, and they are connected by construction. Hence, every fiber is a circle. Since e° is an interval, the bundle is trivial, so $\iota^{-1}(e^\circ) \cong S^1 \times e^\circ$, with ι corresponding to projection onto e°.

It remains to show that T is a tree. It is connected because S is connected. Suppose that T contains a cycle and choose a point t in the interior of an edge on this cycle. Then $\iota^{-1}(t)$ is an embedded circle in $S \cong S^2$, so it separates S by the Jordan–Schönflies theorem. On the other hand, $T \setminus \{t\}$ is connected. Since ι is monotone, the preimage of this connected subspace is connected. Thus, $S \setminus \iota^{-1}(t)$ is connected, a contradiction. It follows that T has no cycles and hence is a tree. □

Lemma 4.4. *Let (T, F, ι) be a tree van Kampen diagram for a polyhedral map $f\colon S \cong S^2 \to X$. Let $e = [u, v]$ be an edge of T such that v is a leaf of T. Then $\overline{\iota^{-1}((u, v])}$ is a polyhedral 2-disk.*

Proof. Choose $t \in (u, v)$ and set $C \coloneqq \iota^{-1}(t)$. By Definition 4.1, C is an embedded circle in S. Thus, by the Jordan–Schönflies theorem, $S \setminus C$ has two components, both open 2-discs.

Since v is a leaf, $(t, v]$ is a component of $T \setminus \{t\}$. The map ι is monotone, so $\iota^{-1}((t, v])$ is connected. It follows that $\iota^{-1}((t, v])$ is one of the two components of $S \setminus C$, and hence is an open 2-disc. Its closure is therefore a closed 2-disc with boundary C.

The product structure over the open edge identifies $\iota^{-1}((u, t])$ with an annular collar of C. Hence, $\overline{\iota^{-1}((u, v])}$ is obtained from the closed disc $\overline{\iota^{-1}((t, v])}$ by extending this collar. Our injectivity assumption on the attaching maps implies that this is again a closed 2-disc. Since ι is polyhedral, this disc is polyhedral. □

Lemma 4.5. *Let $f\colon D \cong D^2 \to \widehat{X}$ be a polyhedral map. Let $e = \{u, v\}$ be an edge in X and assume that $(\pi \circ f)^{-1}(u) = \partial D$ and $\pi(f(D)) = e$. Then f is homotopic to a polyhedral map $f'\colon D' \cong D^2 \to \widehat{X}_u$ via a homotopy h fixing ∂D such that*

1. *$\mathrm{Vol}(f') \le \delta_{G_e}(\mathrm{Vol}(f|_{\partial D}))$;*
2. *h has at most $\delta^2_{G_v}\big(\mathrm{Vol}(f) + \mathrm{Vol}(f')\big) + \mathrm{Vol}(f')$ nonsingular 3-cells.*

Proof. It follows from the assumptions that the boundary loop $f|_{\partial D}$ is contained in $\widehat{X}_e \times \{u\}$. Let $f'_0\colon D' \to \widehat{X}_e \times \{0\}$ be a polyhedral filling of this loop with $\mathrm{Vol}(f'_0) \le \delta_{G_e}(\mathrm{Vol}(f|_{\partial D}))$. Denote by $f'_1\colon D' \to \widehat{X}_e \times \{1\}$ the parallel copy of f'_0 at the other end of the cylinder $\widehat{X}_e \times [0, 1]$.

Collapse $\widehat{X}_e \times [0, 1]$ to the end attached to the vertex fiber over v. Composing f with this collapse gives a polyhedral map $f_v\colon D \to \widehat{X}_v$ whose boundary agrees with the boundary of f'_1. Hence, $f_v \cup f'_1$ is a polyhedral map from a 2-sphere to $\widehat{X}_v$ of volume at most $\mathrm{Vol}(f) + \mathrm{Vol}(f'_1)$. We fill this sphere inside $\widehat{X}_v$ using at most

$$\delta^2_{G_v}\big(\mathrm{Vol}(f) + \mathrm{Vol}(f'_1)\big)$$

nonsingular 3-cells.

Finally, parallel translation of f'_1 back to f'_0 inside $\widehat{X}_e \times [0, 1]$ gives a collar with $\mathrm{Vol}(f'_1)$ many nonsingular 3-cells. Gluing this collar to the filling inside $\widehat{X}_v$ gives a homotopy from f to $f' \coloneqq f'_0$ fixing ∂D, with the claimed bound. □

Theorem 4.6. *Let $f\colon S \cong S^2 \to \widehat{X}$ be a polyhedral map with $\mathrm{Vol}(f) = n$. Assume that $\pi \circ f$ admits a tree van Kampen diagram (T, F, ι). Then there exists a filling of f with*

$$\mathrm{Vol}(f) \leq \overline{\max_e \delta_{G_e}}(n) + \overline{\max_v \delta^2_{G_v}}\big(n + 2 \cdot \overline{\max_e \delta_{G_e}}(n)\big).$$

Proof. The proof works analogously to the proof of Theorem 3.3. We iteratively apply Lemma 4.5 to the map $f|_D\colon D \to \widehat{X}$, where $D = \overline{\iota^{-1}((u,v])}$ is the disk from Lemma 4.4. Doing so gives a homotopy to a polyhedral map $f'\colon S' \to \widehat{X}$ such that $\pi \circ f'$ admits a tree van Kampen diagram with one fewer edge than T.

Using the notation from Lemma 4.5, let $e_1, \dots, e_{r_v}$ be the edges incident to v in the order in which they are collapsed (see again Fig. 4). Denote by $f_{v,i}\colon D_{v,i} \to \widehat{X}$ the polyhedral map to which we apply Lemma 4.5 when collapsing the edge e_i, and by $f'_{v,i}\colon D'_{v,i} \to \widehat{X}$ the resulting polyhedral map. When collapsing the last edge e_{r_v}, we apply Lemma 4.5 to a map $f_{v,r_v}\colon D_{v,r_v} \to \widehat{X}$ defined on the 2-disc $D_{v,r_v} = S_v \cup D'_{v,1} \cup \ldots \cup D'_{v,r_v-1}$. The resulting homotopy therefore has at most

$$\begin{aligned}\delta^2_{G_v}\Big(\mathrm{Vol}(f|_{S_v}) + \sum_{i=1}^{r_v} \mathrm{Vol}(f'_{v,i})\Big) &\leq \delta^2_{G_v}\Big(\mathrm{Vol}(f|_{S_v}) + \sum_{i=1}^{r_v} \max_e \delta_{G_e}\big(\mathrm{Vol}(f|_{\partial D_{v,i}})\big)\Big)\\ &\leq \delta^2_{G_v}\Big(\mathrm{Vol}(f|_{S_v}) + \overline{\max_e \delta_{G_e}}\Big(\sum_{i=1}^{r_v} \mathrm{Vol}(f|_{\partial D_{v,i}})\Big)\Big)\end{aligned}$$

nonsingular 3-cells. The resulting filling therefore has at most

$$\overline{\max_v \delta^2_{G_v}}\Big(\sum_v \mathrm{Vol}(f|_{S_v}) + 2 \cdot \overline{\max_e \delta_{G_e}}\Big(\sum_v \sum_{i=1}^{r_v} \mathrm{Vol}(f|_{\partial D_{v,i}})\Big)\Big) \tag{4.1}$$

nonsingular 3-cells coming from fillings inside the EG_v. Observe that $\mathrm{Vol}(f|_{\partial D_{v,i}})$ coincides with the number of nonsingular 2-cells in the collar above e_i that extends $\partial D_{v,i}$ into the original sphere S. In particular, the 2-cells of the S_v and the 2-cells of the collars defined by the $\partial D_{v,i}$ are pairwise disjoint subcomplexes of S. Thus, $\sum_v \mathrm{Vol}(f|_{S_v}) \leq n$ and $\sum_v \sum_{i=1}^{r_v} \mathrm{Vol}(f|_{\partial D_{v,i}}) \leq n$, and hence Eq. (4.1) is bounded above by $\overline{\max_v \delta^2_{G_v}}\big(n + 2 \cdot \overline{\max_e \delta_{G_e}}(n)\big)$. Similarly, one sees that the fillings inside the EG_e contribute at most $\overline{\max_e \delta_{G_e}}(n)$ nonsingular 3-cells. □

We can now deduce [BBD12, Proposition 6.1] from Theorem 4.6 and Remark 4.2:

Corollary 4.7 ([BBD12, Proposition 6.1])**.** *If X is a tree, then*

$$\delta^2_G(n) \preccurlyeq \overline{\max_v \delta^2_{G_v}}\big(\overline{\max_e \delta_{G_e}}(n)\big).$$

4.2. **Homological filling functions.**

We now turn to upper bounds for the higher homological Dehn functions and the proof of Theorem D.

As before, let X be a cocompact G-simplicial tree and $\widehat{X}$ and $\pi\colon \widehat{X} \to X$ be the complex and map obtained from Proposition 2.21 for $n = k+1$. The proof follows a similar strategy to those in Sections 3 and 4.1.

Definition 4.8. If $\alpha = \sum_\sigma \alpha_\sigma \sigma$ is a cellular chain and A is a subspace, write $\alpha|_A \coloneqq \sum_{\sigma \subseteq A} \alpha_\sigma \sigma$ for the restriction of α to A.

Let $k \geq 1$, let $T \subseteq X$ be a finite subtree, and let $z \in Z_k(\widehat{X})$ be supported in $\pi^{-1}(T)$. For each cell σ of T, set $c(\sigma) \coloneqq z|_{\pi^{-1}(\sigma)} - z|_{\pi^{-1}(\partial\sigma)}$. Thus, $c(\sigma)$ consists precisely of the terms supported on cells that project onto σ. We call the family

$(c(\sigma))_{\sigma\subseteq T}$ the *T-decomposition* of z. The T-decomposition of z is unique and satisfies $z = \sum_{\sigma\subseteq T} c(\sigma)$.

For an edge e of T incident to a vertex v, set $b(e,v) \coloneqq \partial c(e)|_{\widehat{X}_e\times\{v\}}$.

Lemma 4.9. *Let $(c(\sigma))_{\sigma\subseteq T}$ be a T-decomposition of a cycle $z \in Z_k(\widehat{X})$ supported on a tree $T \subseteq X$. Then*

1. *$\partial c(e) = b(e,u) + b(e,v)$ for every edge $e = \{u,v\}$;*
2. *$\partial b(e,v) = 0$ for every edge e with vertex v;*
3. *$\partial c(v) + \sum_{e\ni v} b(e,v) = 0$ for every vertex v;*
4. *$\sum_{\sigma\subseteq T} \|c(\sigma)\| = \|z\|$;*
5. *$\|b(e,v)\| \le \|c(e)\|$ for every edge e with vertex v.*

Proof. By Proposition 2.21 (3), every term of $\partial c(e)$ not supported over a vertex projects onto e°. No boundary of another term in the T-decomposition projects onto e°, so these terms form the part of ∂z projecting onto e° and therefore vanish, proving (1). By (1), $\partial b(e,u) + \partial b(e,v) = 0$. Because $b(e,u)$ and $b(e,v)$ have disjoint supports, no cancellation can occur, implying (2). To see (3), observe that $(\partial z)|_{\pi^{-1}(v)} = \partial c(v) + \sum_{e\ni v} b(e,v) = 0$.

The chains $c(\sigma)$ have pairwise disjoint cellular supports. Their masses therefore add without cancellation, implying (4).

Finally, by Proposition 2.21 (3), every cell occurring in $c(e)$ is a product with e and has one face over each endpoint. The simplicial embeddings from Proposition 2.21 (2) make the induced endpoint maps norm non-increasing, proving (5). □

Remark 4.10. Let z be a cellular k-cycle with a T-decomposition, and let $e = \{u,v\}$ be an edge of T such that v is a leaf. Set $E \coloneqq \widehat{X}_e \times e$ and $E_u \coloneqq \widehat{X}_e \times \{u\}$. Then $a \coloneqq c(e) + c(v)$ is supported in $E \cup \widehat{X}_v$ and satisfies $a|_{E_u} = 0$. Moreover, by Lemma 4.9 we have $\partial a = b(e,u) + b(e,v) + \partial c(v) = b(e,u) \in C_{k-1}(E_u)$.

Lemma 4.11. *Let $k \ge 2$ and let $e = \{u,v\}$ be an edge of X. Set $E \coloneqq \widehat{X}_e \times e$, $E_u \coloneqq \widehat{X}_e \times \{u\}$, and $E_v \coloneqq \widehat{X}_e \times \{v\}$. Let $a \in C_k(\widehat{X})$ be a cellular k-chain whose support is contained in $E \cup \widehat{X}_v$. Assume that $a|_{E_u} = 0$ and that $b \coloneqq \partial a \in C_{k-1}(E_u)$. Then there are chains $a' \in C_k(E_u)$ and $H \in C_{k+1}(\widehat{X})$ such that*

1. *$\partial a' = b$;*
2. *$\partial H = a - a'$;*
3. *$\|a'\| \le \mathrm{FV}^k_{G_e}(\|b\|)$;*
4. *$\|H\| \le \mathrm{FV}^{k+1}_{G_v}\big(\|a\| + \|a'\|\big) + \|a'\|$.*

Proof. Since $b = \partial a$, the chain b is a $(k-1)$-cycle in E_u. Choose a filling $a' \in C_k(E_u)$ with $\partial a' = b$ and $\|a'\| \le \mathrm{FV}^k_{G_e}(\|b\|)$.

Put $Y \coloneqq E \cup \widehat{X}_v$. The attaching map $E_v \to \widehat{X}_v$ is cellular and injective, so its induced map on cellular chains is norm non-increasing.

Let $r\colon Y \longrightarrow \widehat{X}_v$ be the cellular retraction obtained by collapsing the interval factor of $E = \widehat{X}_e \times e$ to the endpoint v and then using the attaching map $E_v \to \widehat{X}_v$. Let $a_v \coloneqq r_*a'$. Then $\zeta \coloneqq r_*a - a_v$ is a k-cycle in $\widehat{X}_v$, since $\partial\zeta = r_*b - r_*b = 0$. Moreover, $\|\zeta\| \le \|a\| + \|a'\|$. Hence, there is a chain $B \in C_{k+1}(\widehat{X}_v)$ such that $\partial B = \zeta$ and $\|B\| \le \mathrm{FV}^{k+1}_{G_v}\big(\|a\| + \|a'\|\big)$.

Let $j\colon \widehat{X}_v \hookrightarrow Y$ be the inclusion. Then $j \circ r$ is homotopic to id_Y relative to $\widehat{X}_v$. Let P_* be the chain homotopy induced by this homotopy. It vanishes on $j_*C_*(\widehat{X}_v)$.

For every q-cell σ of $\widehat{X}_e$, it satisfies $P_q(\sigma \times \{u\}) = \pm(\sigma \times e)$ and $P_{q+1}(\sigma \times e) = 0$. Thus, $P_k(a) = 0$ because $a|_{E_u} = 0$. Moreover, P_k is norm non-increasing on $C_k(E_u)$. Consequently,

$$\| P_k(a - a') \| = \| P_k(a') \| \leq \| a' \|.$$

Since $a - a'$ is a cycle, we have $\partial P_k(a - a') = (j \circ r)_*(a - a') - (a - a')$. Moreover, $(j \circ r)_*(a - a') = j_*(r_* a - r_* a') = j_* \zeta$. Therefore, $H \coloneqq j_* B - P_k(a - a')$ satisfies $\partial H = a - a'$ and $\|H\| \leq \|B\| + \|a'\|$, which is the claimed bound. □

Remark 4.12. In the setting of Remark 4.10, suppose that $a' \in C_k(E_u)$ satisfies $\partial a' = b(e, u)$, and let T' be obtained from T by deleting e and its leaf v. Then $z' \coloneqq z - a + a'$ is a cycle supported over T'. Its T'-decomposition is obtained by replacing $c(u)$ with $c(u) + a'$. Moreover, $\|a\| + \|z'\| \leq \|z\| + \|a'\|$. Indeed, a and $z - a$ have disjoint supports, while $z' = (z - a) + a'$.

Theorem 4.13. *Let $k \geq 2$ and let $z \in C_k(\widehat{X})$ be a cycle with $\|z\| \leq n$. Assume that z is supported in $\pi^{-1}(T)$ for a finite subtree $T \subseteq X$. Then z bounds a cellular $(k+1)$-chain H such that*

$$\| H \| \leq \overline{\max_e \mathrm{FV}^k_{G_e}}(n) + \overline{\max_v \mathrm{FV}^{k+1}_{G_v}}\big(n + 2 \cdot \overline{\max_e \mathrm{FV}^k_{G_e}}(n)\big),$$

where the maxima are taken over all edges and vertices of X, respectively.

Proof. Let r denote the number of edges in T and fix a vertex v_0 of T. We use the iterative collapse procedure from the proofs of Theorems 3.3 and 4.6, which treat the homotopical cases $k = 1, 2$. Let $e_1, \ldots, e_r$ be the sequence of edges we use to collapse T onto v_0 iteratively and let $z_0 \coloneqq z$.

We write $e_i = \{u_i, v_i\}$ for the edge we remove in the i-th step by collapsing its leaf v_i, z_{i-1} for the cycle we obtained from the $(i-1)$-th step, and c_{i-1} and b_{i-1} for the corresponding chains of the T-decomposition of z_{i-1}. Let $a_i \coloneqq c_{i-1}(e_i) + c_{i-1}(v_i)$ and $\beta_i \coloneqq b_{i-1}(e_i, u_i)$. By Remark 4.10 and Lemma 4.11, there are chains a'_i and H_i such that $\partial H_i = a_i - a'_i$. Setting $z_i \coloneqq z_{i-1} - a_i + a'_i = z_{i-1} - \partial H_i$ yields the decomposition after the collapse of e_i by Remark 4.12. In particular, $b_i(e_j, w) = b_0(e_j, w)$ for every edge e_j of T with $j > i$ and every vertex w of such an edge e_j. After all r edges have been collapsed, z_r is supported in $\widehat{X}_{v_0}$ and $\partial(\sum_i H_i) = z_0 - z_r$.

Each β_i is therefore equal to $b_0(e, w)$ for some edge e of T and some vertex w of e. By Lemma 4.9,

$$\sum_i \| \beta_i \| \leq \sum_{e \subseteq T} \| c_0(e) \| \leq \| z_0 \| \leq n.$$

Set $A \coloneqq \sum_i \|a'_i\|$. By Lemma 4.11 and superadditivity, we obtain

$$A \leq \overline{\max_e \mathrm{FV}^k_{G_e}}\Big(\sum_i \| \beta_i \|\Big) \leq \overline{\max_e \mathrm{FV}^k_{G_e}}(n).$$

For the i-th collapse, the inequality in Remark 4.12 gives

$$\| a_i \| + \| z_i \| \leq \| z_{i-1} \| + \| a'_i \|.$$

Summing these inequalities for all i and cancelling the terms $\|z_1\|, \ldots, \|z_{r-1}\|$ gives

$$\| z_r \| + \sum_i \| a_i \| \leq \| z_0 \| + \sum_i \| a'_i \| \leq n + A,$$

hence $\|z_r\| + \sum_i \big(\|a_i\| + \|a'_i\|\big) \leq n + 2A$.

Choose a filling H_0 of z_r in $\widehat{X}_{v_0}$ such that $\|H_0\| \leq \mathrm{FV}^{k+1}_{G_{v_0}}(\|z_r\|)$. By Lemma 4.11, each H_i with $1 \leq i \leq r$ satisfies $\|H_i\| \leq \mathrm{FV}^{k+1}_{G_{v_i}}\big(\|a_i\| + \|a_i'\|\big) + \|a_i'\|$. Set $H \coloneqq H_0 + \sum_i H_i$. Then

$$\begin{aligned}\|H\| &\leq \|H_0\| + \sum\nolimits_i \|H_i\| \\ &\leq \overline{\max_v \mathrm{FV}^{k+1}_{G_v}}\Big(\|z_r\| + \sum\nolimits_i \big(\|a_i\| + \|a_i'\|\big)\Big) + A \\ &\leq \overline{\max_v \mathrm{FV}^{k+1}_{G_v}}(n + 2A) + A \\ &\leq \overline{\max_v \mathrm{FV}^{k+1}_{G_v}}(n + 2 \cdot \overline{\max_e \mathrm{FV}^{k}_{G_e}}(n)) + \overline{\max_e \mathrm{FV}^{k}_{G_e}}(n).\end{aligned}$$

Finally, since $\partial H_0 = z_r$ and $\partial(\sum_i H_i) = z - z_r$ we have $\partial H = z$. This completes the proof. $\square$

Proof of Theorem D. This is now a direct consequence of Theorem 4.13. $\square$

5. Special cases and applications

In this section, we discuss some special cases and applications of our methods to group extensions and amalgamated products of nilpotent groups.

5.1. **Group extensions.** Let $1 \to H \to G \to Q \to 1$ be a group extension, where H and Q are finitely presented.

Throughout, fix a finite presentation of Q and let X be the corresponding Cayley complex. It becomes a G-polyhedral complex via the quotient map $G \twoheadrightarrow Q$, and all stabilizers are H. Moreover, fix symmetric generating sets S_H and S_G of H and G, respectively, such that $S_H \subseteq S_G$.

Since all vertex and edge stabilizers of this action are equal to H and $\delta_X \approx \delta_Q$, Theorem C yields the bound

$$\delta_G(n) \preccurlyeq \delta_Q(n) \cdot \delta_H\big(\overline{\mathrm{dist}^G_H}(\delta_Q(n))\big).$$

This bound is not optimal in general; see Remark 5.9. The goal of this subsection is to give a better upper bound for group extensions by establishing a tighter control on the lengths of the paths in Lemma 3.1.

Remark 5.1. Fix a choice of the data in Construction 2.14. If $e = (v, w)$ is an oriented edge of $G\backslash X$, define

$$m_e \coloneqq h_{w,e} h_{v,e}^{-1}.$$

With this convention, transporting a point $a \in \widehat{X}_v$ along $\widehat{X}_e \times [0, 1]$ results in the element $m_e^{-1} a m_e \in \widehat{X}_w$. We have $m_{\bar{e}} = m_e^{-1}$, where $\bar{e} = (w, v)$ is the same edge with the opposite orientation.

Definition 5.2 (Monodromy distortion)**.** An $\widehat{X}$*-path* is a sequence

$$\widehat{\gamma} = (h_0, e_1, h_1, \ldots, e_k, h_k),$$

where $h_i \in H$ and the $e_i = (v_{i-1}, v_i)$ are oriented edges in X which form an oriented path. It is called an $\widehat{X}$*-loop* if $v_k = v_0$. Its length is defined as

$$\ell_{\widehat{X}}(\widehat{\gamma}) \coloneqq k + \sum_{i=0}^{k} \ell_H(h_i).$$

For an oriented edge e_i in X, let m_{e_i} denote the element associated to the oriented edge $p(e_i)$ of $G\backslash X$ as in Remark 5.1. In this way, we can associate to each $\widehat{X}$-loop $\widehat{\gamma}$ an element

$$W(\widehat{\gamma}) \coloneqq h_0 m_{e_1} h_1 \cdots m_{e_k} h_k \in H.$$

The element $W(\widehat{\gamma})$ is obtained from the vertical pieces h_i under the monodromy map defined by the path $e_1 \cdots e_i$, that is, by pushing each h_i to the fiber over v_0 along this path.

We define the *monodromy distortion* function $\mathrm{Mon}_{\widehat{X}} \colon \mathbb{N} \to \mathbb{R}$ by

$$\mathrm{Mon}_{\widehat{X}}(n) \coloneqq \max\{\ell_H(W(\widehat{\gamma})) \mid \ell_{\widehat{X}}(\widehat{\gamma}) \le n\}.$$

Lemma 5.3. *Let $\widehat{X}$ and $\widehat{X}'$ be constructed using two different choices of elements $\{h_{v,e}\}$ and $\{h'_{v,e}\}$. Then $\mathrm{Mon}_{\widehat{X}} \approx \mathrm{Mon}_{\widehat{X}'}$.*

Proof. For every incidence $v \subset e$ in $G\backslash X$, set $a_{v,e} \coloneqq h_{v,e}(h'_{v,e})^{-1} \in H$. If $e = (v,w)$, then the corresponding elements from Remark 5.1 satisfy

$$m_e = a_{w,e} m'_e a_{v,e}^{-1}.$$

Let $C \coloneqq \max\{\ell_H(a_{v,e})\}$. Write $W_{\widehat{X}}$ and $W_{\widehat{X}'}$ for the words formed using the elements m_e and m'_e, respectively. Given a $\widehat{X}'$-loop $\widehat{\gamma} = (h_0, e_1, h_1, \ldots, e_k, h_k)$, where $e_i = (v_{i-1}, v_i)$, define an $\widehat{X}$-loop $\widehat{\eta} = (g_0, e_1, g_1, \ldots, e_k, g_k)$ by

$$g_i \coloneqq \begin{cases} h_0 a_{v_1,e_1}^{-1} & i = 0; \\ a_{v_{i-1},e_i} h_i a_{v_{i+1},e_{i+1}}^{-1} & 1 \le i \le k-1; \\ a_{v_{k-1},e_k} h_k & i = k. \end{cases}$$

Then $W_{\widehat{X}}(\widehat{\eta}) = W_{\widehat{X}'}(\widehat{\gamma})$ and

$$\ell_{\widehat{X}}(\widehat{\eta}) \le \ell_{\widehat{X}'}(\widehat{\gamma}) + 2Ck \le (2C+1)\ell_{\widehat{X}'}(\widehat{\gamma}).$$

Hence, $\mathrm{Mon}_{\widehat{X}'}(n) \le \mathrm{Mon}_{\widehat{X}}((2C+1)n)$. Because the argument is symmetric, we obtain $\mathrm{Mon}_{\widehat{X}} \approx \mathrm{Mon}_{\widehat{X}'}$. □

Definition 5.4. The *monodromy distortion* of the extension $1 \to H \to G \to Q \to 1$ is defined as $\mathrm{Mon}_H^G \coloneqq \mathrm{Mon}_{\widehat{X}}$ for some choice of elements $\{h_{v,e}\}$.

Lemma 5.5. *Let $\widehat{\gamma}$ be an $\widehat{X}$-loop based at v_0. Then*

$$d_{\widehat{X}_{v_0}}\big(\widehat{\gamma}(0), \widehat{\gamma}(1)\big) \le \mathrm{Mon}_{\widehat{X}}\big(\ell_{\widehat{X}}(\widehat{\gamma})\big).$$

Proof. Let $\widehat{\gamma} = (h_0, e_1, h_1, \ldots, e_k, h_k)$, $n = \ell_{\widehat{X}}(\widehat{\gamma})$, and set $u_i \coloneqq m_{e_1} \cdots m_{e_i}$. Note that $u_k \in H$, as $\widehat{\gamma}$ is a loop. The associated word $W(\widehat{\gamma})$ can be written as

$$\begin{aligned} W(\widehat{\gamma}) &= h_0 m_{e_1} h_1 \cdots m_{e_k} h_k \\ &= h_0 (u_1 h_1 u_1^{-1})(u_2 h_2 u_2^{-1}) \cdots (u_k h_k u_k^{-1}) u_k. \end{aligned} \tag{5.1}$$

The formula (5.1) describes the path in $\widehat{X}_{v_0}$ obtained by pulling the vertical pieces h_i back along $e_1, \ldots, e_i$ to the base vertex fiber. Hence, the word $W(\widehat{\gamma})$ labels a path in $\widehat{X}_{v_0}$ from $\widehat{\gamma}(0)$ to $\widehat{\gamma}(1)$. Therefore,

$$d_{\widehat{X}_{v_0}}\big(\widehat{\gamma}(0), \widehat{\gamma}(1)\big) \le \ell_H(W(\widehat{\gamma})) \le \mathrm{Mon}_{\widehat{X}}(n). \qquad \square$$

Theorem 5.6. *Let $1 \to H \to G \to Q \to 1$ be a group extension, where Q and H are finitely presented. Then $\delta_G(n) \preccurlyeq \delta_Q(n) \cdot \delta_H\big(\overline{\mathrm{Mon}_H^G}(\delta_Q(n))\big)$.*

Proof. Let S_H be the chosen finite generating set of H. Recall that, for an incidence $v \subset e$ in $G\backslash X$, $c_{v,e}\colon H \to H$ denotes conjugation by $h_{v,e}$, so that the attaching map $\psi_{v,e}\colon \widehat{X}_e \to \widehat{X}_v$ is induced by $c_{v,e}$ on vertices. Set

$$C \coloneqq \max\bigl\{1, \ell_H(c_{v,e}^{-1}(s)) \mid v \subset e \subseteq G\backslash X \text{ and } s \in S_H\bigr\}.$$

Then $\ell_H(c_{v,e}^{-1}(h)) \leq C\ell_H(h)$ for every $h \in H$. Let γ_{e_i} be the subpaths from the proof of Theorem 3.3 and denote their endpoints by $a_i, b_i \in H$ identified as vertices in $\widehat{X}_v$. The endpoints of the paths $\gamma'_{e_i} \subset \widehat{X}_e$ are therefore $a'_i = \psi_{v,e}^{-1}(a_i)$ and $b'_i = \psi_{v,e}^{-1}(b_i)$, respectively. Then

$$\ell(\gamma'_{e_i}) = d_{\widehat{X}_e}(a'_i, b'_i) \leq C \cdot d_{\widehat{X}_v}(a_i, b_i) \leq C \cdot \mathrm{Mon}_{\widehat{X}}(\ell(\gamma_{e_i})), \tag{5.2}$$

where the last inequality is due to Lemma 5.5. Replacing the inequality $\ell(\gamma'_{e_i}) \leq \overline{\mathrm{edist}}_X(D \cdot \ell(\gamma_{e_i}))$ by Eq. (5.2) one obtains the same upper bound on fillings as in Theorem 3.3 with $\overline{\mathrm{edist}}_X$ replaced by $C \cdot \overline{\mathrm{Mon}}_{\widehat{X}}$. One now proceeds as in the proof of Theorem C, keeping in mind that $\delta_X \approx \delta_Q$. □

Lemma 5.7. $\mathrm{Mon}_H^G \preccurlyeq \mathrm{dist}_H^G$.

Proof. Let $C \coloneqq \max\{\ell_G(m_e) \mid e \in G\backslash X\}$ and $\widehat{\gamma} = (h_0, e_1, h_1, \ldots, e_k, h_k)$ be an $\widehat{X}$-loop with $\ell_{\widehat{X}}(\widehat{\gamma}) \leq n$. Then $\ell_G(h_i) \leq \ell_H(h_i)$, so

$$\ell_G(W(\widehat{\gamma})) \leq Ck + \sum_{i=0}^{k} \ell_G(h_i) \leq C\Bigl(k + \sum_{i=0}^{k} \ell_H(h_i)\Bigr) \leq Cn.$$

As $W(\widehat{\gamma}) \in H$, we have $\ell_H(W(\widehat{\gamma})) \leq \mathrm{dist}_H^G(Cn)$. □

The following special case recovers a result on the Dehn function of central extensions by Conner.

Corollary 5.8 ([Con95])**.** *Let $1 \to H \to G \to Q \to 1$ be a central group extension where Q and H are finitely presented. Then $\delta_G(n) \preccurlyeq \delta_Q(n)^3$.*

Proof. Because H is central, it is abelian, so δ_H is at most quadratic. Moreover, using the description of the words $W(\widehat{\gamma})$ from Eq. (5.1), one sees that Mon_H^G is linear, as the conjugation action by any element of G on H is trivial. □

Remark 5.9. In general, for an extension $1 \to H \to G \to Q \to 1$, we have $\mathrm{Mon}_H^G \not\approx \mathrm{dist}_H^G$. Indeed, the proof of Corollary 5.8 shows that, in the case of a central extension, $\mathrm{Mon}_H^G(n) \approx n$. However, there are central extensions of nilpotent groups with polynomial distortion of any integer degree $d \geq 1$.

Example 5.10. An example of an extension for which the bound in Theorem 5.6 cannot be improved is the solvable group

$$\mathrm{Sol}_3(\mathbb{Z}) = \mathbb{Z}^2 \rtimes_A \mathbb{Z} = \bigl\langle x, y, t \mid [x, y], txt^{-1} = x^2y, tyt^{-1} = xy\bigr\rangle \text{ for } A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$$

Indeed, it is well known that the Dehn function of $\mathrm{Sol}_3(\mathbb{Z})$ is exponential and that the same is true for the distortion of $\mathbb{Z}^2$ in $\mathrm{Sol}_3(\mathbb{Z})$ [ECH+92]. Since Lemma 5.7 implies that $\mathrm{Mon}_{\mathbb{Z}^2}^{\mathrm{Sol}_3(\mathbb{Z})}$ is bounded above by an exponential function, the optimality follows.

Note that the discussion in the previous paragraph also implies that the growth rate of $\mathrm{Mon}_{\mathbb{Z}^2}^{\mathrm{Sol}_3(\mathbb{Z})}$ is precisely exponential. One can also give a direct proof that

$\mathrm{Mon}_{\mathbb{Z}^2}^{\mathrm{Sol}_3(\mathbb{Z})}$ satisfies an exponential lower bound. For this, observe that the Bass–Serre tree X of $\mathrm{Sol}_3(\mathbb{Z})$ is the real line with edges corresponding to unit intervals with integer endpoints. There is a single orbit of edges. Since t acts by $x \mapsto x+1$, we can choose $m_e = t^{-1}$ for all positively oriented edges e of X. Let e_i be the edge from $i-1$ to i and $\overline{e_i}$ the inverse edge from i to $i-1$.

With these choices, consider the $\widehat{X}$-loop

$$\widehat{\gamma} = (h_0, e_1, h_1, \dots, e_k, h_k, \overline{e_k}, h_{k+1}, \overline{e_{k-1}}, \dots, \overline{e_1}, h_{2k}),$$

with $h_k = x$ and $h_i = 1$ for $i \neq k$. Then $W(\widehat{\gamma}) = t^{-k} x t^k = x^{f_{2k-1}} y^{-f_{2k}}$, where f_i denotes the i-th Fibonacci number. In particular, $W(\widehat{\gamma})$ grows exponentially in $\ell_{\widehat{X}}(\widehat{\gamma}) = 2k+1$. We deduce that $\mathrm{Mon}_{\mathbb{Z}^2}^{\mathrm{Sol}_3(\mathbb{Z})}(n) \succcurlyeq e^n$.

5.2. **Amalgams of nilpotent groups along cyclic subgroups.** Hidber proved in [Hid99] that if $G = G_1 *_H G_2$ is an amalgam of two nilpotent groups of class $\leq c$ along an infinite cyclic subgroup H, then $\mathrm{dist}_H^G(n) \preccurlyeq n^c$ and $\delta_G(n) \preccurlyeq n^{4c^2}$. Combining Hidber's distortion bound, the fact that Dehn functions of nilpotent groups of class c are bounded by n^{c+1} [GHR03], and Theorem B, we obtain the following improved upper bound.

Corollary 5.11. *Let $G = G_1 *_H G_2$ be an amalgamated product of finitely generated nilpotent groups G_1 and G_2 along a cyclic subgroup H. Then $\delta_G(n) \preccurlyeq n^{c(c+1)+1}$.*

References

[ABDY13] Aaron Abrams, Noel Brady, Pallavi Dani, and Robert Young, *Homological and homotopical Dehn functions are different*, Proc. Natl. Acad. Sci. USA **110** (2013), no. 48, 19206–19212. MR3153947 ↑3

[Alo94] Juan M. Alonso, *Finiteness conditions on groups and quasi-isometries*, J. Pure Appl. Algebra **95** (1994), no. 2, 121–129. MR1293049 ↑1

[AWP99] J. M. Alonso, X. Wang, and S. J. Pride, *Higher-dimensional isoperimetric (or Dehn) functions of groups*, J. Group Theory **2** (1999), no. 1, 81–112. MR1670329 ↑1, 3, 4

[BBD12] Josh Barnard, Noel Brady, and Pallavi Dani, *Super-exponential 2-dimensional Dehn functions*, Groups Geom. Dyn. **6** (2012), no. 1, 1–51. MR2888944 ↑3, 16, 18

[BBFS09] Noel Brady, Martin R. Bridson, Max Forester, and Krishnan Shankar, *Snowflake groups, Perron-Frobenius eigenvalues and isoperimetric spectra*, Geom. Topol. **13** (2009), no. 1, 141–187. MR2469516 ↑2, 3, 4

[BC00] Stephen G. Brick and Jon M. Corson, *On Dehn functions of amalgamations and strongly undistorted subgroups*, Internat. J. Algebra Comput. **10** (2000), no. 5, 665–681. MR1781853 ↑2

[BC98] ———, *Dehn functions and complexes of groups*, Glasgow Math. J. **40** (1998), no. 1, 33–46. MR1612141 ↑2

[BH99] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Springer Berlin Heidelberg, 1999. ↑2, 7

[Bri02a] Martin R. Bridson, *The geometry of the word problem*, Invitations to geometry and topology, 2002, pp. 29–91. MR1967746 ↑3, 5

[Bri02b] ———, *Polynomial Dehn functions and the length of asynchronously automatic structures*, Proc. London Math. Soc. (3) **85** (2002), no. 2, 441–466. MR1912057 ↑3, 4, 5

[Bri93] Stephen G. Brick, *On Dehn functions and products of groups*, Trans. Amer. Math. Soc. **335** (1993), no. 1, 369–384. MR1102884 ↑2

[Bro87] Kenneth S. Brown, *Finiteness properties of groups*, Proceedings of the Northwestern conference on cohomology of groups (Evanston, Ill., 1985), 1987, pp. 45–75. MR885095 ↑1

[BRS07] Noel Brady, Tim Riley, and Hamish Short, *The geometry of the word problem for finitely generated groups*, Advanced Courses in Mathematics. CRM Barcelona, Birkhäuser Verlag, Basel, 2007. Papers from the Advanced Course held in Barcelona, July 5–15, 2005. MR2281936 ↑3, 5

[Con95] Gregory R. Conner, *Isoperimetric functions for central extensions* (1995), 73–78. ↑23

[DK18] C. Druţu and M. Kapovich, *Geometric group theory*, Colloquium Publications, American Mathematical Society, 2018. ↑7, 8, 9

[ECH+92] 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. MR1161694 ↑23

[Ger92] S. M. Gersten, *Dehn functions and l_1-norms of finite presentations*, Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989), 1992, pp. 195–224. MR1230635 ↑2

[GHR03] S. M. Gersten, D. F. Holt, and T. R. Riley, *Isoperimetric inequalities for nilpotent groups*, Geom. Funct. Anal. **13** (2003), no. 4, 795–814. MR2006557 ↑24

[Hae92] André Haefliger, *Extension of complexes of groups*, Annales de l'Institut Fourier **42** (1992), no. 1-2, 275–311 (en). MR93j:20080 ↑7

[Hid99] Christian Hidber, *Isoperimetric functions of finitely generated nilpotent groups*, J. Pure Appl. Algebra **144** (1999), no. 3, 229–242. MR1731333 ↑24

[LM22] Xingzhe Li and Fedor Manin, *Homological filling functions with coefficients*, Groups Geom. Dyn. **16** (2022), no. 3, 889–907. MR4506540 ↑2

[Ril03] T. R. Riley, *Higher connectedness of asymptotic cones*, Topology **42** (2003), no. 6, 1289–1352. MR1981358 ↑4, 5

[SW26] Roman Sauer and Jannis Weis, *Polynomial homological Dehn functions from non-proper actions*, arXiv preprint arXiv:2606.25897 (2026). ↑3

[Whi39] J. H. C. Whitehead, *Simplicial spaces, nuclei and m-groups*, Proceedings of the London Mathematical Society **s2-45** (1939), no. 1, 243–327, available at `https://academic.oup.com/plms/article-pdf/s2-45/1/243/4244250/s2-45-1-243.pdf`. ↑6

[You11] Robert Young, *Homological and homotopical higher-order filling functions*, Groups Geom. Dyn. **5** (2011), no. 3, 683–690. MR2813532 ↑3, 4, 5, 6

[Zar26] M. C. B. Zaremsky, *Some open problems*, 2026. Dated July 12, 2026. Available on `https://zaremsky.github.io/open_problems.pdf`. ↑2

Department of Mathematics, University of Luxembourg, 6 Avenue de la Fonte, 4364 Esch-sur-Alzette, Luxembourg
Email address: `claudio.llosaisenrich@uni.lu`

Fakultät für Mathematik, Karlsruher Institut für Technologie, 76131 Karlsruhe, Germany
Email address: `jannis.weis@kit.edu`