Homological Dehn Functions

by

Jakub Tucker

MA4K9 Dissertation

Submitted to The University of Warwick

Mathematics Institute

April, 2025

Contents

1 Introduction

Since the time of Gromov's fundamental work [Gro87], hyperbolic groups have been a key focus of work in geometric group theory. A useful characterisation is that a group G is hyperbolic if and only if it is finitely presentable and satisfies a linear isoperimetric inequality. In general, there is much less known about subgroups of hyperbolic groups: for example, it has long been known that hyperbolicity is not a property that extends to all subgroups. In [Rip82], predating Gromov's introduction of the notion, Rips constructs finitely generated subgroups of small cancellation groups – which are hyperbolic – that are not finitely presentable. In [Bra99], Brady constructs the first example of a finitely presented subgroup of a hyperbolic group that is not of type FP_3, and therefore not hyperbolic.

Brady's subgroup appears as the kernel of a surjective map from a hyperbolic group into $\mathbb{Z}$. Since this group is not hyperbolic, it does not satisfy a linear isoperimetric inequality. Knowledge of the *isoperimetric spectrum*, the subset of reals x which appear as Dehn functions $n \mapsto n^x$ for some finitely presented group: that there is only one gap [BB00], which is the open interval $(1, 2)$ – a result of Gromov, proven by Bowditch in [Bow95], tells us that the Dehn function of Brady's subgroup is at least quadratic. The question of an upper bound prompted the work of Gersten and Short in [GS02], which resulted in:

Theorem. *Given a split extension*

$$1 \to K \to H \to F_n \to 1$$

Where F_n is free and K is finitely presented, then if H is hyperbolic, K satisfies a polynomial isoperimetric inequality.

The methods described in [GS02] rely on the existence of certain constants depending on an explicit presentation for H, which in practice are difficult to compute. An explicit polynomial upper bound for Brady's subgroup of n^{96} was computed in [KIS25] by Kropholler, Isenrich and Soroko using a more geometric approach. On a similar note, recent work of Isenrich in [Ise24] shows that we can exchange F_n for $\mathbb{Z}^n$ in the above statement, prompted by the existence of non-hyperbolic finitely presented subgroups appearing as such kernels.

In [Ger96a], Gersten proves that if a hyperbolic group is of cohomological dimension 2, then all its subgroups of type FP_2 are hyperbolic and therefore finitely presented. However, this property is unique to dimension 2, as shown by Kropholler and Vigolo in [KV21], where a group which is of type FP_2 but not finitely presented (and so not hyperbolic) appears as the kernel of a homomorphism $\phi : H \to \mathbb{Z}$ from a hyperbolic group H of cohomological dimension 3. The existence of such a subgroup prompts the author's work.

There is a notion of a generalisation of the isoperimetric inequality for groups of type FP_2 that are not finitely presented. Many foundational results are given in [BKS21] by Brady, Kropholler and Soroko. Notice that the circumstances of this subgroup's appearance as a kernel are not dissimilar to that of Brady's, it is a natural question to ask whether the subgroups constructed by Kropholler and Vigolo also satisfy a polynomial upper bound to their *homological Dehn function*.

Indeed, such an upper bound is achieved as an immediate corollary of a generalisation of Gersten and Short's methods in [GS02] to homological filling functions:

Theorem 4.5. *Let H be a split extension of a group K of type FP_2 by a finitely generated free group F_n, so that we have the short exact sequence*

$$1 \to K \to H \to F_n \to 1.$$

If (f, g) is a homological area-radius pair for H with $f(n) \geq n$ for all n, then there is a constant $C > 1$ such that $FA_K \preccurlyeq C^g f$.

To obtain this result, we generalise the definition of an area-radius pair as used in [GS02] by introducing a *homological radius function* analogously to the area functions defined in [BKS21]. Fundamental properties of the radius function also hold for the homological counterpart, most importantly the independence of choice of (homological) finite presentation up to $\cong$-equivalence. We conclude with a discussion of homological isoperimetric inequalities for hyperbolic groups to obtain a result which applies to the subgroup described in [KV21].

1.1 Acknowledgements

The author would like to thank their supervisor, Dr Robert Kropholler, for bringing these problems to their attention, and for many helpful discussions without which the following would not have been written.

2 The Homological Dehn Function

The following consists of definitions and basic results from [Bri02] for homotopical Dehn functions, and [BKS21] for their homological counterpart. Before we introduce the notion of a homological Dehn function, it is natural to first understand the homotopical Dehn function. We want to meaningfully compare these two functions, observing which properties are common and which are not.

2.1 Homotopical Dehn Functions

Given a group G with finite presentation $\langle A|R\rangle$, a classical question to ask is whether G has a *solvable word problem*, that is, given some word w over $A \cup A^{-1}$, is there some finite algorithm that will determine whether or not $w = 1$ in G? The (homotopical) Dehn function is a well-studied complexity measure for the word problem, and is typically thought of both algebraically and geometrically. In order to define the Dehn function, we first need to have a reasonable notion of the *area* of a word representing the identity.

For the algebraic perspective, given some word w that represents the identity in G, this word can be reduced to the empty word using repeated applications of 3 different processes, namely, *free reductions* – the removal of a subword aa^{-1} for $a \in A$ wherever it may appear in w, *free expansions* – the addition of a subword aa^{-1} for $a \in A$ at any point in w, and, most importantly, *applications of defining relators* – the addition of any word r in $R \cup R^{-1}$, or one of its rotations, at any point in w. A sequence of such moves which reduces a word w to the empty word is called a *null sequence* for w. We define the $\mathrm{Area}_{\langle A|R\rangle}(w) \in \mathbb{N}$ to be the minimum number N of applications of defining relators among all null sequences for w. From this, if we denote by $|w|$ the length of the word w, we can define the Dehn function with respect to the presentation $\langle A|R\rangle$ to be

Definition 2.1. The Dehn function $\delta_{\langle A|R\rangle}$ with respect to the presentation $\langle A|R\rangle$ is

$$\delta_{\langle A|R\rangle}(n) = \sup\left\{\mathrm{Area}_{\langle A|R\rangle}(w) \mid w \in F(A), w =_G 1, \text{ and } |w| \leq n\right\}$$

It turns out that if we look at the Dehn function of any finite presentation of the group G, then they are *equivalent* in the following sense:

Given two functions $f, g : [0,\infty) \to [0,\infty)$, we say $f \preccurlyeq g$ if there exists a constant $C \geq 1$ with $f(x) \leq Cg(Cx+C)+Cx+C$ for all $x \in [0,\infty)$, and we say $f \simeq g$ if $f \preccurlyeq g$ and $g \preccurlyeq f$. Note that we restrict this definition to functions $\mathbb{N} \to \mathbb{N}$ when applying it to Dehn functions. It is common practice to consider Dehn functions only up to $\simeq$-equivalence, in which case the notation δ_G to denote the Dehn function of a finitely presented group G is well-defined. That is to say, up to $\simeq$-equivalence, the Dehn function is independent of choice of (finite) presentation.

To view the Dehn function of a group geometrically, we construct a space $X = X(A,R)$ called the Cayley 2-complex of the presentation $\langle A|R\rangle$. As the name suggests, this space is a 2-dimensional CW complex, whose 1-skeleton is the Cayley graph $\Gamma(G,A)$ – a labelled directed graph with one vertex for every element g of the group G, and two edges at each vertex for each generator $a \in A$: one incoming edge from $a^{-1} \cdot g$ to g labelled a at each vertex, and one outgoing edge from g to $a \cdot g$ labelled a. The 2-cells of X correspond to relations: at each vertex g, for each $r \in R$, a 2-cell is attached along the edge loop labelled r in $X^{(1)}$. This space is the universal cover of the presentation complex for $\langle A|R\rangle$, and as such has trivial fundamental group. Since our presentation is finite, the Cayley 2-complex X is locally finite.

The Cayley graph $\Gamma = \Gamma(G,A)$ is often equipped with the graph edge metric, where every edge in the graph has length 1 and for any two vertices $g, h \in G$, $d_\Gamma(g,h)$ is the length of the shortest path from g to h in Γ. With this metric, Γ is a geodesic metric space - meaning that there is always a path realising the distance between any two points.

Fixing any vertex x_0 in X as a basepoint, loops in the 1-skeleton based at x_0, called *edge loops*, correspond bijectively to words w with $w =_G 1$, by reading off the labels of the edges the loops traverses. Given any such edge loop γ, since $\pi_1(X)$ is trivial, this loop can be filled by a disc in X. A filling of an edge loop labelled by a word $w =_G 1$ by 2-cells in X is called a *Van Kampen diagram* for w. We define $Area_X(\gamma)$ to be the minimum number of 2-cells contained in any filling of γ in X. An *isoperimetric inequality* for X is a function $f : \mathbb{N} \to \mathbb{N}$ such that for any edge loop γ traversing at most n edges, $\mathrm{Area}_X(\gamma) \leq f(n)$. The Dehn function $\delta_{\langle A|R\rangle}$ is a minimal isoperimetric inequality for $X(A, R)$, that is, it maps $n \in \mathbb{N}$ to the maximum number of 2-cells among least-area Van Kampen diagrams for all words $w =_G 1$ with $|w| \leq n$.

A fundamental result concerning Dehn functions is that they are a *quasi-isometry invariant*, that is, if two groups G and H are quasi-isometric with respect to some word metrics, then the Dehn functions δ_G and δ_H are $\simeq$-equivalent. A priori, the Dehn function is only defined on *finitely presented* groups, if we try to extend the definition above to groups with infinite sets of relations, then we no longer have independence of presentation. For an example, one might consider a presentation of a finitely presented group which has a non-linear Dehn function in which every word $w =_G 0$ is contained in the set of relators, which has an isoperimetric inequality given by the constant function 1, which is $\simeq n \mapsto n$. In order to generalise the notion of a Dehn function to groups which are not finitely presented, it is necessary to consider different spaces and different area functions. For example, the Cayley 2-complex of a presentation with an infinite set of relations is not locally finite, since each vertex is on the boundary of at least as many 2-cells as there are relations.

2.2 Homological Dehn Functions

The homological Dehn function is defined for groups of type FP_2, this is a finiteness condition defined below. Notably, every finitely presented group is of type FP_2, but there exist groups of type FP_2 which are not finitely presentable, as shown by Bestvina and Brady in [BB97]. Given a group H with presentation $\langle A|R\rangle$, we will not assume that R is finite. We will see that the homological Dehn function is well-defined for finitely presented groups and is bounded above by the superadditive closure of the homotopical Dehn function.

Definition 2.2. A group H is of type $FP_2(\mathbb{Z})$ – which we will denote by FP_2 for brevity – if there exists a resolution of $\mathbb{Z}$ by finitely generated projective $\mathbb{Z}H$-Modules

$$P_2 \to P_1 \to P_0 \to \mathbb{Z} \to 0$$

We will be mostly concerned with an equivalent condition: the existence of a *homological finite presentation* for H, that is,

Definition 2.3 ([BKS21])**.** A homological finite presentation for $H = \langle A|R\rangle$, where A is finite, is a pair $\langle A \| R_0 \rangle$, where R_0 is a finite subset of R such that $H_1(X(A, R_0)) = 0$ – the homological Cayley complex has trivial first homology group. The homological Cayley complex is a CW complex constructed similarly to the Cayley 2-complex, one begins with the Cayley graph $\Gamma(H, A)$, and for each vertex $x \in H$ and each relation $r \in R_0$, one attaches a 2-cell D_r^x with boundary the loop labelled by r beginning at x.

Remark. The homological Cayley complex retains a key property of the full Cayley 2-complex: that every 1-cycle is the boundary of a 2-chain – though it no longer holds that every loop bounds a disc. The homological Cayley complex has beneficial properties that the Cayley 2-complex does not, for example, it is locally finite where the Cayley 2-complex may not be.

Remark. The group H acts freely, cellularly, cocompactly and vertex transitively on the homological Cayley complex $X(A, R_0)$ on the left as follows: on vertices, an element $h \in H$ maps the vertex $x \in H$ to the vertex hx; on edges, an edge from x to xa with label a is mapped by h to the edge from hx to hxa with the same label; on 2-cells, the 2-cell D_r^x is mapped to the 2-cell D_r^{hx}. Note that we require R_0 to be finite to ensure that this action is cocompact.

The equivalence of these characterisations of groups of type FP_2 is Proposition 2.7 of [BKS21]. Note that if H is finitely presented, $\langle A \| R \rangle$ is a homological finite presentation for H, since $\pi_1(X(A, R)) = 1$ and so its abelianisation $H_1(X(A, R))$ is also trivial.

There are two competing definitions for a generalisation of the Dehn function to groups H of type FP_2, those being the *homological Dehn function* FA_H and the *abelianised Dehn function* δ_H^{ab} as follows. These functions differ by the construction of their corresponding filling functions, which are defined more generally on certain 2-dimensional CW complexes. As such, some new definitions are warranted.

Definition 2.4 ([BKS21])**.** A map between CW complexes is called *combinatorial* if each open cell of the domain is mapped homeomorphically onto an open cell of the target. A CW complex is called *combinatorial*, or simply a *combinatorial complex*, if all of its attaching maps are combinatorial.

Importantly, the Cayley 2-complex of a homological finite presentation $\langle A \| R_0 \rangle$ is combinatorial, and Proposition 2.6 of [BKS21] asserts a partial converse, that is,

Proposition 2.5 ([BKS21])**.** *Let X be a connected 2-dimensional combinatorial complex with trivial first homology on which H acts freely, cellularly, cocompactly and vertex-transitively, then there exists a homological finite presentation $\langle A \| R_0 \rangle$ for H whose homological Cayley complex is X.*

Definition 2.6 ([BKS21])**.** Given a 2-dimensional combinatorial complex X with $H_1(X) = 0$, then for any 1-cycle γ in the 1-skeleton $X^{(1)}$, there exists a 2-chain $c = \sum_i a_i \sigma_i$ for 2-cells σ_i and $a_i \in \mathbb{Z}$ with $\partial c = \gamma$. We define the *homological area* of γ to be

$$\mathrm{HArea}_X(\gamma) \coloneqq \min\left\{\textstyle\sum_i |a_i| \,\middle|\, \gamma = \partial c \text{ for } c = \sum_i a_i \sigma_i\right\},$$

the *homological filling function*

$$\mathrm{FA}_X(n) \coloneqq \sup\left\{\mathrm{HArea}_X(\gamma) \,\middle|\, \gamma \text{ is a 1-cycle in } X^{(1)} \text{ with } |\gamma| \le n\right\},$$

and the *abelianised filling function*

$$\delta_X^{ab}(n) \coloneqq \sup\left\{\mathrm{HArea}_X(\gamma) \,\middle|\, \gamma \text{ is a loop in } X^{(1)} \text{ with } |\gamma| \le n\right\},$$

where since γ is a 1-cycle, it can be realised as a combinatorial map $\bigsqcup_i S^1 \to X$ (when γ is a loop this map is $S^1 \to X$) and $|\gamma|$ is defined to be the minimal number of 1-cells in $\bigsqcup_i S^1$ among all representations of $\gamma : \bigsqcup_i S^1 \to X$.

Remark. Observe that the supremum taken in the definition of FA_X also considers every loop considered in the supremum in the definition of δ_X^{ab}, so $\delta_X^{ab} \preccurlyeq \mathrm{FA}_X$.

Definition 2.7 ([BKS21])**.** Consider a group H of type FP_2 and a homological finite presentation $\langle A \| R_0 \rangle$, the *homological Dehn function* of the triple $\langle H, A, R_0 \rangle$ is defined by

$$\mathrm{FA}_{\langle H,A,R_0\rangle}(n) \coloneqq \mathrm{FA}_{X(A,R_0)}(n),$$

and the *abelianised Dehn function* is simlarly

$$\delta^{ab}_{\langle H,A,R_0\rangle}(n) \coloneqq \delta^{ab}_{X(A,R_0)}(n).$$

Like the homotopical Dehn function, Proposition 2.24 of [BKS21] tells us that up to $\simeq$-equivalence, the homological and abelianised Dehn functions are independent of choice of homological finite presentation, so we can write FA_H and δ_H^{ab} without specifying a choice of homological finite presentation. There are many such parallels with homotopical Dehn functions, for example, Theorem 2.36 tells us that the homological Dehn function is a quasi-isometry invariant; or if a group H splits as a free product $H_1 * H_2$, then its homological Dehn function $\mathrm{FA}_H \simeq \max(\mathrm{FA}_{H_1}, \mathrm{FA}_{H_2})$, and the corresponding result holds for homotopical Dehn functions of finitely presented groups. However, the homological Dehn function differs from its homotopical counterpart in many regards. For one, it is proven in Proposition 2.26 of [BKS21] that every homological Dehn function is $\simeq$-equivalent to a superadditive function, while this is still an open conjecture for homotopical Dehn functions, as well as for the abelianised Dehn function. A more significant difference is the following: the link between the homological Dehn function and the word problem is weaker than its homotopical counterpart, as discussed in Corollary 6.1 of [BKS21], there exist groups of type FP_2 with unsolvable word problems whose homological Dehn functions are $\simeq$-equivalent to $n \mapsto n^4$, whereas it is well-known that a finitely presented group has solvable word problem if and only if its Dehn function is bounded above by a recursive function.

The homological and abelianised Dehn functions can be compared, Proposition 2.21 of [BKS21] shows that for a locally finite combinatorial complex X with $H_1(X) = 0$, $\mathrm{FA}_X \simeq$

$\overline{\delta_X^{ab}}$, where $\overline{\delta_X^{ab}}$ is the *superadditive closure* of δ_X^{ab}:

$$\overline{f}(n) \coloneqq \max\{f(n_1) + \cdots + f(n_r) \,|\, r \geq 1, n_i \in \mathbb{N}, n_1 + \cdots + n_r = n\}.$$

Since the Cayley complex for a homological finite presentation satisfies the above conditions, we can say that $\mathrm{FA}_H \simeq \overline{\delta_H^{ab}}$.

This comparison is very useful, for example, Proposition 2.22 of [BKS21] tells us that, when the complex X is locally finite, as is the case when X is a homological Cayley complex, the suprema in the definitions of δ_X^{ab} and FA_X are attained. This is done for δ_X^{ab} by arguing that there are only finitely many edge paths of a given length, then using that FA_X is equivalent to the superadditive closure of δ_X^{ab}.

We can also compare, when a group H of type FP_2 is also finitely presented, the homological and homotopical Dehn functions.

Lemma 2.8. *If $H \cong \langle A|R\rangle$ is a finitely presented group, then $FA_H(n) \simeq \overline{\delta_H^{ab}} \preccurlyeq \overline{\delta_H}(n)$.*

Since its Cayley 2-complex X is also a homological Cayley complex, one can consider some loop γ of length n in the 1-skeleton. This loop can be filled with a disc, as in the definition of $\mathrm{Area}_H(\gamma)$; or a 2-chain, as in the definition of $\mathrm{HArea}_X(\gamma)$. Note that since a disc filling is already a 2-chain filling, we immediately have $\mathrm{HArea}_H(\gamma) \leq \mathrm{Area}_H(\gamma)$, and thus $\delta_H^{ab} \preccurlyeq \delta_H$, the result follows by taking the superadditive closure.

3 Finitely Presented Kernels of Free Extensions

Given a group H with finite presentation $\mathcal{P} = \langle A|R\rangle$: $F(A)$ denotes the free group on A, $\overline{A}$ denotes the set $A \cup A^{-1}$ and $\overline{A}^*$ denotes the free monoid on $\overline{A}$ - like the free group but we consider elements which differ by a sequence of free reductions and expansions to be different, *i.e.* in $F(A)$, $abb^{-1} = a$, but in $\overline{A}^*$, abb^{-1} and a are distinct.

We consider words $w \in H$ with $w =_H 1$ as elements of $F(A)$ or $\overline{A}^*$ and their Van Kampen diagrams over $\mathcal{P}$ – or $\mathcal{P}$-diagrams – as connected, oriented, labelled, planar graphs of which the bounded regions of the complement have boundaries labelled by elements of R and the unbounded region has boundary labelled by w. An *area function* $f : \mathbb{N} \to \mathbb{R}$ for $\mathcal{P}$ is a function such that every word $w =_H 1$ in $F(A)$ with length at most n, $Area_{\mathcal{P}}(w) \leq f(n)$ – such a function is also called an *isoperimetric inequality* and the Dehn function $\delta_{\mathcal{P}}$ is a minimal such function. A *radius function* $g : \mathbb{N} \to \mathbb{R}$ for $\mathcal{P}$ is a function such that for each word $w =_H 1$ in $F(A)$, there is a Van Kampen diagram $D_{\mathcal{P}}(w)$ such that for every vertex v in $D_{\mathcal{P}}(w)$, there is a path in the 1-skeleton from v to the boundary $\partial D_{\mathcal{P}}(w)$ of length at most $g(n)$. A pair of such functions is called an *area-radius* (AR) pair.

Gersten and Short's general result, which implies the specific case Theorem 5.2, is

Theorem 3.1. *Given a split extension*

$$1 \to K \to H \to F_n \to 1$$

Where F_n is free and K is finitely presented, if (f,g) is an area-radius pair for H with $f(n) \geq n$ for all $n \in \mathbb{N}$, then there exists a constant $C > 1$ such that $C^g f$ is an isoperimetric inequality for K.

Like we used above for Dehn functions, the AR pairs are often considered up to equivalence. Gersten and Short define for $f, \overline{f} : \mathbb{N} \to \mathbb{R}$, $f \simeq \overline{f}$ if there exist four integer constants A, B, C, D such that both $f(n) \leq A\overline{f}(Bn) + Cn + D$ and $\overline{f}(n) \leq Af(Bn) + Cn + D$, and for $g, \overline{g} : \mathbb{N} \to \mathbb{R}$, $g \cong \overline{g}$ if there exist three integer constants A, B, C such that both $g(n) \leq A\overline{g}(Bn) + C$ and $\overline{g}(n) \leq Ag(Bn) + C$, each for all $n \in \mathbb{N}$. Note that $\cong$ is a strictly stronger relation than $\simeq$, and this definition of $\simeq$ is equivalent to the one we defined earlier: observe that if $f(n) \leq A\overline{f}(Bn) + Cn + D$, then $f(n) \leq A\overline{f}(Bn + 0) + Cn + D$, and if $f(n) \leq A\overline{f}(Bn + B') + Cn + D$, then $f(n) \leq A\overline{f}\left((B + B')n\right) + Cn + D$.

When considering equivalence of AR pairs (f,g) and $(\overline{f}, \overline{g})$, we say that these are equivalent when $f \simeq \overline{f}$ and $g \cong \overline{g}$. Proposition 2.1 of [GS02] tells us that the equivalence class of an AR pair is a group invariant, if $\mathcal{P}$ and $\mathcal{Q}$ are finite presentations for H, then there are equivalent AR pairs (f,g) for $\mathcal{P}$ and $(\overline{f}, \overline{g})$ for $\mathcal{Q}$.

To prove Gersten and Short's result, it is necessary to understand how the area of a Van Kampen diagram can be affected by the application of an automorphism, which leads to Lemmas 3.1 and 3.2 of [GS02]: Let $\mathcal{P} = \langle A | R \rangle$ be a finite presentation for a group H and let $\phi : H \to H$ be a group automorphism. For each $a \in A$, choosing a word representing $\phi(a)$ in $F(A)$ induces a monoid homomorphism $\Phi : \overline{A}^* \to \overline{A}^*$ with $\Phi(a) =_H \phi(a), \Phi(a^{-1}) =_H \phi(a^{-1}) = \phi(a)^{-1}$. Since ϕ is an automorphism, its inverse also induces a monoid homomorphism $\Psi : \overline{A}^* \to \overline{A}^*$ where $\Phi(\Psi(a^{\pm 1})) =_H a^{\pm 1} =_H \Psi(\Phi(a^{\pm 1}))$.

Lemma 3.2. *There is a constant $S > 0$ such that if $D_{\mathcal{P}}(w)$ is a Van Kampen diagram for a word $w =_H 1$, then there is a Van Kampen diagram $D'_{\mathcal{P}}(\Phi(w))$ over $\mathcal{P}$ for $\Phi(w)$ with*

$$\operatorname{Area}\left(D'_{\mathcal{P}}(\Phi(w))\right) \leq S \cdot \operatorname{Area}\left(D_{\mathcal{P}}(w)\right),$$

and conversely,

Lemma 3.3. *There are constants $S', S'' > 0$ such that if $D'_{\mathcal{P}}(\Phi(w))$ is a Van Kampen diagram for $\Phi(w)$, then there is a Van Kampen diagram $D''_{\mathcal{P}}(w)$ for w such that*

$$\operatorname{Area}\left(D''_{\mathcal{P}}(\Phi(w))\right) \leq S' \cdot \operatorname{Area}\left(D'_{\mathcal{P}}(\Phi(w))\right) + S'' \cdot \ell(w).$$

Gersten and Short's proof of our Theorem 3.2 begins by taking a finite presentation $\mathcal{P}_K$ for the kernel of the split extension. We will have to adapt this to the case where K is of type FP_2. Similarly, the notion of an equivalence class of AR pairs is proved to be an invariant among finite presentations of groups, we will have to work around this to attempt to apply the main ideas in the proof to groups of type FP_2. Their proof of the main theorem is as follows:

Proof of Theorem 3.2. Let $\mathcal{P}_K = \langle x_1, \ldots, x_m | r_1, \ldots, r_k \rangle$ be a finite presentation for K and let the generators of F_n be $\{t_1, \ldots, t_n\}$. In the split extension $1 \to K \to H \to F_n \to 1$, for each t_i generating F_n, there is an automorphism $\phi_i : K \to K$ given by conjugation of K by t_i in H. Define the monoid endomorphism $\Phi_i : \overline{X}^* \to \overline{X}^*$ as before, then H has finite presentation

$$\mathcal{P}_H = \left\langle x_1, \ldots, x_m, t_1, \ldots, t_n \,\middle|\, r_1, \ldots, r_k, \{t_i^{-1} x_j t_i \Phi_i(x_j)^{-1}\} \right\rangle.$$

Given a word in K $w =_K 1$, we consider a Van Kampen diagram D over $\mathcal{P}_H$. Since there are no t-edges on the boundary, [Ger96b] tells us each edge labelled t is in a unique annulus formed of 2-cells corresponding to relations $t_i^{-1} x_j t_i \Phi(x_j)$ which lies in the interior of D, called a 't-ring' – as shown in Fig. 1. Each t-ring is associated to one of the t_i, which we call *stable letters*.

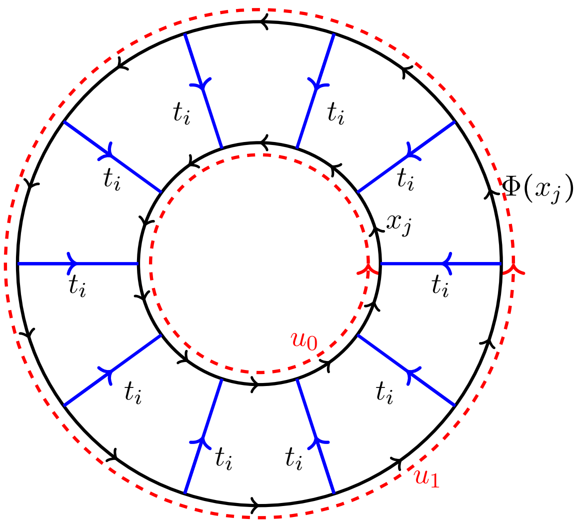

Figure 1: A t-ring associated to t_i

Consider the two components of a $\mathcal{P}$-diagram minus a t-ring A, let D_0 be the inner component, with the inner boundary of A, and D_1 the outer component with the outer boundary of A. Let the words read around the inner and outer boundaries of A be called u_0 and u_1 respectively, observe u_0, u_1 are words in $\overline{X}^*$ and each $u_i =_K 1$, with D_0 being a $\mathcal{P}_G$-diagram for u_0. Since t-rings can be nested only as deep as the radius of the diagram, if A is an innermost t-ring, then D_0 is a $\mathcal{P}_K$-diagram for u_0. Depending on the orientation, we might have $u_1 = \Phi_i(u_0)$ or $u_0 = \Phi_i(u_1)$. In the former case, applying Φ_i to the $\mathcal{P}_K$-diagram D_0 gives us a $\mathcal{P}_K$-diagram $D(u_1)$ for u_1, which Lemma 3.2 tells us has an area bounded above by $S \cdot \mathrm{Area}(D_0)$ for some positive constant S. Replacing $A \cup D_0$ with $D(u_0)$ here, and doing analogously to all innermost rings, at most multiplies the area of D by S. In the latter case, we can instead apply Lemma 3.3 to get a $\mathcal{P}_K$-diagram $D(u_1)$ for u_1 with area bounded by $S' \cdot \mathrm{Area}(D_0) + S'' \ell(u_1)$. Since each edge of D can lie in at most 2 t-rings, we can bound above $\ell(u_1)$ by 2x the number of edges in D, so

$\ell(u_1) \leq \ell(w) + 2\rho\text{Area}(D)$, where ρ is the length of the longest relation in the presentation $\mathcal{P}_G$. Replacing $D_0 \cup A$ by $D(u_1)$ on all innermost t-rings at most multplies the area by S' and adds $S''(\ell(w) + 2\rho(D) \leq Mf(\ell(w))$ for some $M > \max(S, S', 1)$.

We can apply the above procedures no more than $g(\ell(w))$ times in order to obtain a $\mathcal{P}_K$-diagram for w, whose area is bounded by

$$M(\ldots(M(Mf(\ell(w)) + Mf(\ell(w))) + Mf(\ell(w))) + \cdots + Mf\ell(w)) \leq M^{g(\ell(w))+1}f(\ell(w)).$$

Replacing M by a larger constant $C > M^2$ gives Theorem 3.1. □

4 Kernels of type FP_2

Gersten and Short's proof regards finitely presented groups, but there few instances where the existence of finite presentations is assumed. For the remainder of this section, we will be working under the hypothesis that we have a split extension

$$1 \to K \to H \to F_n \to 1$$

Where the groups K and H are of type FP_2, and F_n is free of rank n.

In this section we prove homological analogues of key results from [GS02] used in the proof of Theorem 3.1, beginning with expanding the notion of an area-radius pair, adapting the results of Lemma 3.2 and Lemma 3.3, before fitting the results together as Gersten and Short did.

4.1 Homological Radius Functions

Gersten and Short's upper bound is in terms of an area-radius pair. A priori, these are defined only for finitely presented groups. We already have a notion of isoperimetric inequality for groups H of type FP_2, that being the homological Dehn function FA_H. This satisfies key properties of the area function used by Gersten and Short, most importantly being independent of homological finite presentation up to $\simeq$-equivalence. Consider the definition of the radius function in [GS02], and observe that it still holds when using a homological filling[1] in place of a Van Kampen Diagram.

Definition 4.1. Let X be a 2-dimensional combinatorial complex with $H_1(X) = 0$. For any 1-cycle γ in the 1-skeleton $X^{(1)}$, there exists a 2-chain $c = \sum_i a_i\sigma_i$ for 2-cells σ_i and $a_i \in \mathbb{Z}$ with $\partial c = \gamma$. We define the *homological radius* of γ to be

$$\text{HRad}_X(\gamma) := \min\left\{\text{Rad}_X(c) \,\middle|\, \gamma = \partial c \text{ for } c = \textstyle\sum_i a_i\sigma_i\right\},$$

[1] This subsection is retreaded in Appendix A using surface diagrams.

where

$$\mathrm{Rad}_X(c) = \max_{v \text{ a vertex in } c}\left\{\min\left\{|\alpha_v| \,\middle|\, \alpha_v \text{ is a path from } v \text{ to } \partial c \text{ in } c^{(1)}\right\}\right\}.$$

Analogously to the homological filling function, we define the *homological radius function* for the space X to be

$$\mathrm{FR}_X(n) = \sup\left\{\mathrm{HRad}_X(\gamma) \,|\, \gamma \text{ is a 1-cycle in } X^{(1)} \text{ with } |\gamma| \leq n\right\}.$$

For a group H of type FP_2 with homological finite presentation $\langle A \| R_0\rangle$ and homological Cayley complex $X(A, R_0)$, we denote

$$\mathrm{FR}_{\langle H,A,R_0\rangle}(n) \coloneqq \mathrm{FR}_{X(A,R_0)}(n).$$

Remark. If H is a finitely presented group, then every Van Kampen diagram is a homological filling, so the usual radius function is an upper bound for the homological radius function for any fixed finite presentation.

To use this new function in place of the radius function in our analogue of Gersten and Short's proof, we will require that it is invariant of presentation. This homological radius function shares some of the same key properties as the homological filling function, we will argue similarly to Propositions 2.22 and 2.24 of [BKS21] to show this property holds.

Lemma 4.2. *Let X be a connected combinatorial complex which is locally finite with $H_1(X) = 0$ which admits a cellular cocompact group action. Then the supremum in the definition of homological radius function is attained.*

Proof. Since X is locally finite, and there are only finitely many orbits of vertices under the group action, there exists a natural number k such that every vertex of X has valence at most k.

For any fixed vertex v of X, there are at most k^n paths of length n based at v, it follows that there are only finitely many orbits of loops of a given length n, so the supremum

$$\sup\left\{\mathrm{HRad}_X(\gamma) \,|\, \gamma \text{ is a loop in } X^{(1)} \text{ with } |\gamma| \leq n\right\}$$

is attained.

If a 1-cycle γ is not a loop, then it can be expressed as a formal combination of loops $\gamma = \gamma_1 + \cdots + \gamma_\ell$. Each of the γ_i can be optimally filled by a 2-chain c_i, the formal combination $c_1 + \cdots + c_\ell$ is therefore a filling of γ by a 2-chain, the radius of which is the maximum of the radii of the c_i, that is, $\mathrm{Rad}_X(c) = \max_i\{\mathrm{HRad}_X(c_i)\}$. It follows that $\mathrm{HRad}_X(\gamma) \leq \mathrm{Rad}_X(c) = \mathrm{HRad}_X(c_j)$ for one of the the loops c_j – the homological radii of the individual loops of any 1-cycle bound above the homological radius of the 1-cycle, so

$$\mathrm{FR}_X(n) = \sup\left\{\mathrm{HRad}_X(\gamma) \,|\, \gamma \text{ is a loop in } X^{(1)} \text{ with } |\gamma| \leq n\right\}$$

is attained by a loop. $\square$

With this property, we are able to proceed as the authors of [BKS21] in their Proposition 2.24.

4.1.1 Independence of Presentation

Proposition 4.3. *Up to $\cong$-equivalence, the homological radius function is independent of a homological finite presentation for H.*

Proof. As in [BKS21], we will first prove that two homological finite presentations which share a generating set A have $\cong$-equivalent radius functions.

Let $\mathcal{P}_0 = \langle A \| R_0 \rangle, \mathcal{P}_1 = \langle A \| R_1 \rangle$ be homological finite presentations for H. Consider the homological Cayley complex X_0 for $\mathcal{P}_0$. Note that this space shares its 1-skeleton with the homological Cayley complex X_1 for $\mathcal{P}_1$ – both are the Cayley graph for H with respect to generating set A.

Let γ be a loop in $X_0^{(1)}$ with homological filling $c = \sum_i a_i \sigma_i$ with a radius $\mathrm{HRad}_{X_0}(c)$. Each 2-cell σ_i is $g_i \cdot r_i$ for some $g_i \in H$ and some relation $r_i \in R_0$, so each boundary $\partial \sigma_i$ is a loop labelled r_i, so can be filled by a 2-chain in $\mathcal{P}_1$ with radius n_i. Filling all the σ_i in this way gives a filling of γ in X_1, the radius of which is bounded above by

$$\mathrm{HRad}_{X_0}(\gamma) + \max_{r_i \in R_0}\{n_i\},$$

since for any vertex v, the distance in the 1-skeleton from v to the boundary $\partial \sigma_i$ is at most $\max_i\{n_i\}$, and the distance from that vertex to the boundary ∂c is at most $\mathrm{HRad}_{X_0}(c)$. Since $\max_i\{n_i\}$ is a constant, we have that

$$\mathrm{FR}_{\langle H,A,R_0\rangle} \cong \mathrm{FR}_{\langle H,A,R_1\rangle}$$

by applying the same argument, starting with X_1 in place of X_0.

This brings us to the general case. Suppose now that $\langle A \| R_0 \rangle$ and $\langle A_1 \| S_0 \rangle$ are two homological finite presentations for H. For each generator b in A_1, there is a word $v(A)$ in $F(A)$ with $b =_H v(A)$. Fix a $v(A)$ for each $b \in A_1$ and let R_1 be the set of relations $\{b = v(A) \,|\, b \in A_1\}$. As in [BKS21], we construct a new complex X_1 from the homological Cayley complex $X(A, R_0)$ as follows: for each generator $b \in A_1$ and for each vertex $x \in X(A, R_0)$, we attach an edge e_b^x labelled by b with initial vertex x and terminal vertex $x \cdot v(A)$; then for each such edge, we attach a 2-cell D_b^x to the closed path reading $b^{-1}v(A)$ from x. So the complex X_1 is

$$X_1 \coloneqq X(A, R_0) \cup \bigcup_{\substack{x \in H \\ b \in A_1}} e_b^x \cup \bigcup_{\substack{x \in H \\ b \in A_1}} D_b^x$$

The H-action on $X(A, R_0)$ is extended to a free, vertex transitive and cocompact action on this new complex by defining $g \cdot e_b^x = e_b^{g \cdot x}$ and $g \cdot D_b^x = D_b^{g \cdot x}$. [BKS21] also gives us a deformation retraction $r : X_1 \to X(A, R_0)$ which sends edges labelled b to paths labelled $v(A)$ with the same initial and terminal vertices, and collapses the 2-cells D_b^x. The existence of this deformation retraction tells us $H_1(X_1) \cong H_1(X(A, R_0))$ is trivial, so $\langle A \sqcup A_1 \| R_0 \sqcup R_1 \rangle$ is a homological finite presentations for H.

Let w be a 1-cycle in X_1. Every vertex in w lies in $X(A, R_0) \subset X_1$ by construction. For every edge of w labelled by a generator of A_1, we can apply a relator from R_1 to replace this edge with edges labelled by A. Applying no more than $|w|$ such relators gives us a 1-cycle entirely contained within $X(A, R_0)$, we call this 1-cycle w'. Observe that $|w'| \leq k|w|$, where $k = \max\left\{|v(A)| \,\middle|\, b^{-1}v(A) \text{ is a relator in } R_1\right\}$. The 1-cycle w' admits a homological filling c using only relators from R_0. Given any vertex in this filling, its distance to the boundary ∂c is at most $\mathrm{Rad}_{X(A,R_0)}(c)$, and this boundary vertex has distance at most $k/2$ from a vertex in the 1-cycle w. Combining this w'-filling c with the R_1-relators applied earlier gives us a filling of w whose radius is bounded above by $\mathrm{Rad}_{X(A,R_0)}(c) + k/2$. It follows that

$$\mathrm{HRad}_{X_1}(w) \leq \mathrm{HRad}_{X(A,R_0)}(w') + k/2 \leq \mathrm{FR}_{X(A,R_0)}(k|w|) + k/2.$$

Since the choice of 1-cycle w was arbitrary, we may conclude that

$$\mathrm{FR}_{X_1}(n) \leq \mathrm{FR}_{X(A,R_0)}(kn) + k/2.$$

For the converse direction, let w be a 1-cycle in $X(A, R_0)$ whose edges are labelled only by letters from A, $|w| \leq n$ and $\mathrm{HRad}_{X(A,R_0)}(w) = \mathrm{FR}_{X(A,R_0)}(n)$. We know that such a 1-cycle realising the supremum in the definition of $\mathrm{FR}_{X(A,R_0)}$ exists by Lemma 4.2. Let c be a filling of w in X_1. Consider the image of c under the retraction r, this is a filling for w in $X(A, R_0)$. Since w is a 1-cycle whose optimal $X(A, R_0)$-filling has radius $\mathrm{FR}_{X(A,R_0)}(n)$, we must have that the radius $\mathrm{Rad}_{X(A,R_0)}(r(c)) \geq \mathrm{FR}_{X(A,R_0)}(n)$. The retraction r sends each A_1-edge to a word $v(A)$ of length at most k, it follows that the radius $\mathrm{Rad}_{X(A,R_0)}(r(c)) \leq k \cdot \mathrm{Rad}_{X_1}(c) \leq k \cdot \mathrm{FR}_{X_1}(n)$, therefore

$$\mathrm{FR}_{X(A,R_0)}(n) \leq k \cdot \mathrm{FR}_{X_1}(n)$$

Bringing the two inequalities together, we obtain that

$$\mathrm{FR}_{X_1} \cong \mathrm{FR}_{X(A,R_0)}.$$

Working analogously to the above, for each $a \in A$, we can find words $u(A_1)$ in (A_1) with $a =_H u(A_1)$, calling the set of such relations S_1. We similarly construct a homological finite presentation $\langle A \sqcup A_1 \| S_0 \sqcup S_1 \rangle$ with homological Cayley complex Y_1, and apply the same argument to see $\mathrm{FR}_{Y_1} \cong \mathrm{FR}_{X(A_1,S_0)}$. Note that the homological finite presentations for X_1 and Y_1 share a generating set, so we can apply the first part of the proof to see

$$\mathrm{FR}_{X(A,R_0)} \cong \mathrm{FR}_{X(A_1,S_0)}.$$

i.e. Up to $\cong$-equivalence, the homological radius function is independent of homological finite presentation. $\square$

Notation. *Considering homological radius functions up to $\cong$-equivalence, the notion of a homological radius function for a group, $FR_H \coloneqq FR_{X(A,R_0)}$ for any homological finite presentation $\langle A \| R_0 \rangle$ of H, is well-defined.*

With these results in place, we can now extend the notion of an area-radius pair to groups H of type FP_2, using the functions FA_H and FR_H. Since these share the same key properties of the area and radius functions of Gersten and Short, we can use FA_H and FR_H in their place.

Definition 4.4. Given a group H of type FP_2 with homological finite presentation $\mathcal{P} = \langle A \| R_0 \rangle$, we call the pair (f, g) of functions $f, g : \mathbb{N} \to \mathbb{R}$ a *homological area-radius pair* if for any 1-cycle γ of length at most n in $X(A, R_0)$, there exists a 2-chain $c = \sum_i a_i \sigma_i$ with $\partial c = \gamma$ and $\sum_i |a_i| \leq f(n)$ and $\mathrm{Rad}_{X(A,R_0)}(c) \leq g(n)$.

4.2 A Homological Finite Presentation for H

Now we are able to state a generalised version of Gersten and Short's result for groups of type FP_2:

Theorem 4.5. *Let H be a split extension of a group K of type FP_2 by a finitely generated free group F_n, so that we have the short exact sequence*

$$1 \to K \to H \to F_n \to 1.$$

If (f, g) is a homological area-radius pair for H with $f(n) \geq n$ for all n, then there is a constant $C > 1$ such that $FA_K \preccurlyeq C^g f$.

A priori, we don't know whether the group H is necessarily of type FP_2. Gersten and Short's proof opens by constructing a finite presentation for H from one for K, in the following result, we claim that the analogous construction yields a homological finite presentation for H. First we will consider the case $n = 1$, before adapting to the general case.

4.2.1 The Case $n = 1$

Lemma 4.6. *Let H be a split extension of a group $K = \langle A | R \rangle$ of type FP_2 by $\mathbb{Z} = \langle t | \rangle$, so that we have the short exact sequence*

$$1 \to K \to H \to \mathbb{Z} \to 1.$$

There is an automorphism $\phi : K \to K$ given by conjugation by t in H, let Φ be a lift of ϕ to a semigroup endomorphism as in Section 4.4. If $\langle A \| R_0 \rangle$ is a homological finite presentation for K, then $\left\langle A \cup \{t\} \,\middle\|\, R_0 \cup \left\{t^{-1}a_j t\Phi(a_j)^{-1} \,|\, a_j \in A\right\}\right\rangle$ is a homological finite presentation for H.

Remark. Note that that existence of a homological finite presentation for H proves that H is of type FP_2.

Proof. Note first that [GS02] tells us that $\left\langle A \cup \{t\} \,\middle|\, R \cup \left\{t^{-1}a_j t\Phi(a_j)^{-1} \,|\, a_j \in A\right\}\right\rangle$ is a presentation for H. The existence of the split extension sequence tells us that H can be realised as a semidirect product $K \rtimes_\phi \mathbb{Z}$, importantly, each element of H can be written in the form kt^n for some $k \in K$.

Consider the Cayley graph $\Gamma(H, A \cup \{t\})$ for H. Each coset Kt^n for $n \in \mathbb{Z}$ corresponds to a unique copy of the Cayley graph $\Gamma(K, A)$ for K as a subgraph of $\Gamma(H, A \cup \{t\})$. All edges labelled by generators $a_j \in A$ have initial and terminal vertices in the same coset, and all t-edges have initial vertex in a coset Kt^n and terminal vertex in an 'adjacent' coset Kt^{n+1}. We can picture the Cayley graph $\Gamma(H, A \cup \{t\})$ as a tower of copies of $\Gamma(K, A)$, each connected to the copies above and below only by t-edges (See Fig. 2).

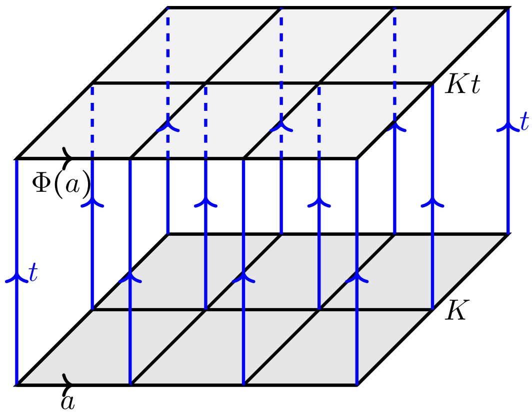

Figure 2: Sketch of part of $\Gamma(H, A \cup \{t\})$

We construct a 2-complex X from $\Gamma(H, A \cup \{t\})$ by attaching for each vertex $x \in H$ and for each of the finitely many relations $r \in R_0 \cup \left\{t^{-1}a_j t\Phi(a_j)^{-1} \,|\, a_j \in A\right\}$ a 2-cell with boundary the loop labelled by r beginning at x. The construction of this space mirrors that of a homological Cayley complex for H, we claim that $H_1(X) = 1$ and therefore that

$$\left\langle A \cup \{t\} \,\middle\|\, R_0 \cup \left\{t^{-1}a_j t\Phi(a_j)^{-1} \,|\, a_j \in A\right\}\right\rangle$$

is a homological finite presentation for H.

By construction, each copy of K in the Cayley graph for H corresponds to a copy of the homological Cayley complex $X(A, R_0)$ for K in the larger space X. To prove that $H_1(X)$ is trivial, we show that every 1-cycle admits a filling by a 2-chain. It suffices to prove the result for a loop γ in $X^{(1)}$.

If γ lies entirely in one of the cosets Kt^n, that is, γ contains no t-edges, then we can view γ as a loop in $X(A, R_0)$. $H_1(X(A, R_0))$ is trivial, so γ admits a filling by a 2-chain c in $X(A, R_0)$, which we can embed into X by identifying $X(A, R_0)$ with the subcomplex of X corresponding to the coset Kt^n.

For a general loop γ in $X^{(1)}$, γ can only intersect finitely many of the cosets Kt^n. Let $m = \min\{n \in \mathbb{Z} \mid \gamma \cap Kt^n \neq \varnothing\}$, $M = \max\{n \in \mathbb{Z} \mid \gamma \cap Kt^n \neq \varnothing\}$. Consider the part of γ which lies in Kt^m. If γ traverses no edges in Kt^m, then γ must go down and up the same t-edge which can be ignored in the filling. Suppose not, then for each edge labelled a_j that γ traverses in Kt^m, we can attach a 2-cell with boundary label $t^{-1}a_j t\Phi(a_j)^{-1}$ to *push up* the part of the loop γ in Kt^m into Kt^{m+1}. We can proceed as such recursively until the loop γ is pushed up to a loop γ' which lies in Kt^M. As above, there exists a filling c' of γ', which we can extend to a filling of γ by adding all the 2-cells $t^{-1}a_j t\Phi(a_j)^{-1}$ used to push up γ. □

4.2.2 The General Case

Many ideas used in the above will still hold in the general case, however, we need to refine what we mean to 'push down' or 'push up' fillings when working in the Cayley graph for $K \rtimes F_n$. In the case $F_n \cong \mathbb{Z}$, there was only one stable letter to apply. In the general case, we must be careful which we use to construct the partial 2-chain filling which allows us to reduce the task of finding a 2-chain filling in H to a filling in K.

Lemma 4.7. *Let H be a split extension of a group $K = \langle A|R\rangle$ of type FP_2 by a finitely generated free group F_n, so that we have the short exact sequence*

$$1 \to K \to H \to F_n \to 1.$$

There is an automorphism $\phi_i : K \to K$ given by conjugation by t_i in H, let Φ_i be a lift of ϕ_i to a semigroup endomorphism as in Section 4.4. If $\langle A \| R_0\rangle$ is a homological finite presentation for K, then $\left\langle A \cup \{t_1, \dots, t_n\} \,\middle\|\, R_0 \cup \left\{t_i^{-1}a_j t_i \Phi_i(a_j)^{-1} \,\middle|\, a_j \in A, i \in \{1, \dots, n\}\right\}\right\rangle$ is a homological finite presentation for H.

Proof. We construct X analogously to the above. We can picture the Cayley graph $\Gamma(H, A \cup \{t_i, \dots, t_n\})$ as a copy of the Cayley graph $\Gamma(K, A)$ for K for each vertex of the Cayley graph of the free group $F(t_1, \dots, t_n)$ (the $2n$-valent tree), each connected to the $2n$ adjacent copies by edges labelled t_i. X is obtained from $\Gamma(H, A \cup \{t_i, \dots, t_n\})$ by gluing for each vertex $x \in H$ and for each of the finitely many relations $r \in R_0 \cup$

$\{t_i^{-1} a_j t_i \Phi_i(a_j)^{-1} \mid a_j \in A, i \in \{1, \dots, n\}\}$ a 2-cell with boundary the loop labelled by r beginning at x. Under this construction, the copies of $\Gamma(K, A)$ corresponding to cosets Kw in H have become copies of the homological Cayley complex $X(K, A, R_0)$. Importantly, each copy of the homological Cayley complex for K has trivial first homology. It suffices to prove that $H_1(X)$ is trivial, for which we show that every 1-cycle admits a filling by a 2-chain. It suffices to check the result for a loop γ in $X^{(1)}$.

A loop γ in $X^{(1)}$ can only intersect finitely many of the cosets Kw where w is a reduced word in $F(t_1, \dots, t_n)$. A continuous loop γ can only traverse from a coset Kw to an adjacent one Kwt_i or Kwt_i^{-1} by an edge labelled t_i. Consider the set S of the words w such that γ intersects Kw non-trivially. S must contain a unique word of minimal length. Suppose that $u \neq v$ are two words of minimal length, then there must exist a path along γ from Ku to Kv. If this path traverses an edge labelled t_i, then it moves into the coset Kut_i. We must have $|ut_i| = |u| + 1$ by minimality of $|u|$. Any subsequent t-edges traversed can only move to Ku' for $|u'| > |u|$ or back into Ku, where the above argument applies again. There is no way of changing any of the first $|u|$ letters of u by traversing t-edges without moving into a coset Ku' with $|u'| < |u|$, contradicting minimality. Set $\overline{w}$ to be the unique word of minimal length in S.

Pick a word $w \in S$ of maximal length. If the last letter of w is t_i, that is, $w = w' * t_i = w't_i$ for some strictly shorter word w' of length $|w| - 1$. Let Ψ_i be as in Lemma 4.12. For each edge labelled a_j that γ traverses in Kw, there are 2-chains in Kw with boundary $a_j(\Phi_i \circ \Psi_i)(a_j)^{-1}$, for each edge a_k in $\Psi_i(a_j)$, there is a 2-cell $t^{-1} a_k t \Phi_i(a_k)^{-1}$. Taking all such 2-chains and 2-cells allows us to 'push down' the part of the loop γ to a homologous loop γ' which is the same as γ except the section in Kw is replaced by a homologous section in Kw'. If instead the last letter of w is t_i^{-1}, that is , $w = w' * t_i^{-1} = w't_i^{-1}$. Take, for all edges a_j traversed by γ in Kw, 2-cells of the form $t^{-1} a_j t \Phi_i(a_j)^{-1}$. The collection of such 2-cells allows us to 'push up' the part of the loop γ in Kw to a homologous part in Kw' to obtain a homologous loop γ' which does not intersect Kw.

Applying the process above repeatedly to all words in S from longest to shortest eventually yields a loop $\overline{\gamma}$ contained in $K\overline{w}$. $\overline{\gamma}$ is the boundary of a 2-chain in $K\overline{w}$. By taking the union of this 2-chain with all the 2-cells required to push γ into $K\overline{w}$, we obtain a 2-chain whose boundary is γ. Thus X is a homological Cayley complex for H. □

4.3 Surface Diagrams

In [GS02], Gersten and Short do not work directly in the Cayley 2-complex for H, but rather on *disc diagrams*, 2-discs with an associated combinatorial structure that map into the Cayley 2-complex as a filling. We want to use a similar method to work with a space with simpler geometry than an arbitrary 2-chain realised in the homological Cayley complex, while still obtaining the same information about the group.

There is a more general notion of this: a *surface diagram*. To define these surface diagrams, we will take cues from Arenas' definition in [AW22], as well as the definitions of (combinatorial) disc and annular diagrams and their filling maps in [BKS21]. Like Arenas, we will define initially our surface diagrams as having a specific genus, but in practice, we will not know which genus of surface diagram is required and will simply refer to an arbitrary surface diagram.

Definition 4.8 (Surface Diagram)**.** A (genus g) *surface diagram* S is a compact combinatorial 2-complex defined as

$$\overline{S} \smallsetminus (\textstyle\bigsqcup_{i=1}^{n} e_i)$$

For some closed, possibly disconnected, genus g surface $\overline{S}$, where the e_i are finitely many disjoint open 2-cells in a combinatorial cell structure on $\overline{S}$. The surface diagram has n *boundary paths*, corresponding to the attaching maps $\phi_i : \partial D^2_{e_i} \hookrightarrow \overline{S}$. The boundary of S is the union of such paths.

We define the *area* $\mathrm{Area}(S)$ of a surface diagram S to be the number of 2-cells in S.

Definition 4.9 (Surface Diagram filling of a 1-cycle)**.** Let X be a combinatorial complex with a 1-cycle γ in the 1-skeleton $X^{(1)}$. A *filling* for γ is a pair (S, π) consisting of a surface diagram S with a combinatorial map $\pi : S \to X$ where $\gamma = \pi \circ \bigsqcup_{i=1}^{n} \phi_i$, and each $\pi \circ \phi_i \subseteq \gamma$ is a closed loop in $X^{(1)}$. *i.e.* each boundary loop maps to a loop in the 1-cycle γ.

Lemma 4.10. *Let $X = X(A, R_0)$ be a a homological Cayley complex for a group H of type FP_2 with homological finite presentation $\langle A \| R_0 \rangle$, let γ be a 1-cycle in the 1-skeleton $X^{(1)}$. Then there exists a surface diagram S that is a filling for γ.*

Proof. X is a homological Cayley complex so $H_1(X) = 0$ and there exists a (finite) 2-chain $c = \sum_{i=1}^{k} a_i \sigma_i$ in X with boundary $\partial c = \gamma$. Each 2-cell σ_i has boundary a loop in $X^{(1)}$ labelled by a relator $r_i \in R_0$ beginning at a vertex $h_i \in H$.

We construct a combinatorial structure of a several-times punctured surface S which maps combinatorially to c as follows. For a 2-cell σ_i, consider its boundary word r_i starting at g_i in c. Observe that every edge in the loop r_i starting at g_i either lies on a loop of the boundary γ – a *boundary edge* – or is also contained in the boundary of another 2-cell σ_j – an *internal edge.*

We begin with a collection of all the 2-cells in c, with multiplicity given by the a_i in the sum. Pick some 2-cell σ_i to add to our surface diagram, if σ_i has no internal edges, then its boundary must form a loop in γ so we move on to another 2-cell in c. If σ_i has an internal edge with label g and initial vertex h, pick some other 2-cell σ_j that shares this edge in c, and glue along it – note that another 2-cell with nothing glued along this edge always exists since this edge does not appear in the boundary, so must appear in c an even number of times so that its value is zero when the boundary is taken. Return to the vertex h in S, there are 2 non-glued edges with h at their endpoint, if they are both boundary edges, then we move on to another vertex in the boundary of S. If one of the edges containing h in S is internal, then pick another 2-cell σ_k in c that shares this edge, and glue along it – note that we might choose a σ_k that is already in S, thus identifying in S two (possibly previously separate) vertices corresponding to the same group element. There might be other 2-cells that share this edge in c, but in the surface diagram S we do not necessarily identify their endpoints labelled h in c. Compare with Fig. 3.

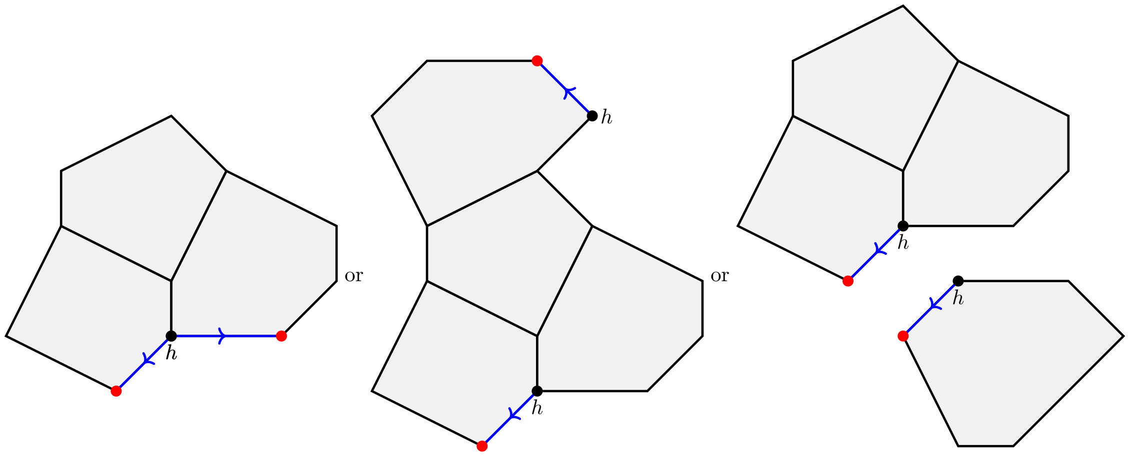

Figure 3: Gluing edges around h

This is not a cause for concern, we can return to this vertex h in S and glue 2-cells until its only non-glued edges correspond to those in γ, or there are no non-glued edges. As in Fig. 4: in the former case, locally around h we have a neighbourhood homeomorphic to the upper half plane in $\mathbb{R}^2$; in the latter, we have a neighbourhood homeomorphic to $\mathbb{R}^2$. Notice that any point in an internal edge we have glued is locally homeomorphic to $\mathbb{R}^2$, and any point on a boundary edge is locally homeomorphic to the upper half plane, and any point internal to a 2-cell is locally homeomorphic to $\mathbb{R}^2$. Once this vertex h is locally a surface – possibly with boundary – we turn our attention to a different vertex in S, and repeat the process outlined above.

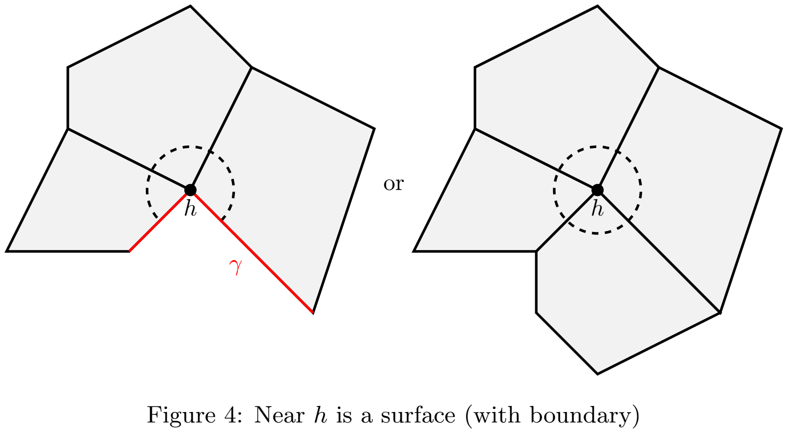

Figure 4: Near h is a surface (with boundary)

If this process appears to terminate early, that is, we have a closed surface with boundary S, but there are still 2-cells in c unaccounted for in S, then it must be the case that these will correspond to a separate connected component in S. We can continue to build our surface diagram by picking any unaccounted-for 2-cell in c and repeating the construction.

Since the 2-chain c is finite, we have only finitely many edges, vertices and 2-cells, so this process must terminate. Note that it is impossible for the process to terminate without yielding a (possibly disconnected) surface with boundary, since any point internal to a 2-cell has a neighbourhood homeomorphic to $\mathbb{R}^2$, is does any point on a glued internal edge, and each vertex has a neighbourhood homeomorphic to either $\mathbb{R}^2$ or the upper half plane, as do the boundary edges. If any point in S does not satisfy this property, then it must lie on an unglued internal edge, or a vertex which appears as the endpoint of an unglued internal edge, in which case there must be a 2-cell in c sharing this edge which has not yet been glued, since the boundary of c does not contain this edge. The above condition implies that, after this construction is complete, every point of S has a neighbourhood homeomorphic to $\mathbb{R}^2$ or the upper half plane, that is, S is a *surface with boundary.*

We have constructed the surface diagram S in such a way that the filling map is easy to define. We define $\pi : S \to c \subseteq X$ to be the map that sends each point in S that is a vertex to the corresponding vertex in c, every point on an edge in S to the corresponding point in the corresponding edge in c, and every point in a 2-cell to the corresponding point in the corresponding 2-cell in c. By construction, n-cells are mapped homeomorphically onto n-cells (π restricts to the identity map on any given n-cell), and thus the map is combinatorial. The boundary components of S correspond exactly to γ by construction, so $\gamma = \pi(\partial S)$ and π is a surface diagram filling for γ. □

Remark. By construction, the area of a surface diagram S in a filling (S, π) is precisely the area of the 2-chain that is its image – they have the same number of 2-cells.

Remark. If (S, π) is a surface diagram filling for a closed loop γ in the Cayley graph of a group H of type FP_2, then S has one boundary component, and it is a loop combinatorially homeomorphic to γ. This follows since a boundary component in S is a loop, and must map to a loop under the combinatorial map π, so each boundary component on S surjects onto γ. If S has multiple boundary components, then under π, the loop γ will be several times counted, but $\pi(\partial S) = \partial c = \gamma$, where c is some 2-chain in X the homological Cayley complex, and γ is only once counted by assumption.

4.4 New Filling from an Automorphism

Section 3 of [GS02] gives upper bounds for the area of a Van Kampen diagram after the application of an automorphism, we will need a similar result for surface diagram fillings of 1-cycles. Fortunately, the methods of Gersten and Short's proofs for their Lemmas 3.1

and 3.2 apply with minimal adjustments to the case of 2-chain and surface diagram fillings of 1-cycles.

Let $\mathcal{P}_0 = \langle A \| R_0 \rangle$ be a homological finite presentation for a group H of type FP_2, and let $\phi : H \to H$ be a group automorphism. For each $a \in A$, fixing a choice of word representing $\phi(a)$ in $F(A)$ induces a monoid homomorphism $\Phi : \overline{A}^* \to \overline{A}^*$ with $\Phi(a) =_H \phi(a), \Phi(a^{-1}) =_H \phi(a^{-1}) = \phi(a)^{-1}$. Since ϕ is an automorphism, its inverse also induces a monoid homomorphism $\Psi : \overline{A}^* \to \overline{A}^*$ where $\Phi(\Psi(a^{\pm 1})) =_H a^{\pm 1} =_H \Psi(\Phi(a^{\pm 1}))$. For a 1-cycle γ in $X(A, R_0)$, we denote by $\Phi(\gamma)$ the 1-cycle in $X(A, R_0)$ obtained by mapping each vertex $x \in \gamma$ to $\Phi(x)$ and each edge labelled a with initial vertex x and terminal vertex y to the path labelled by the word $\Phi(a)$ from $\Phi(x)$ to $\Phi(y)$. If $\gamma = \gamma_1 + \cdots + \gamma_\ell$ is the decomposition of a 1-cycle gamma into distinct loops γ_i, then $\Phi(\gamma) = \Phi(\gamma_1) + \cdots + \Phi(\gamma_\ell)$, and if a loop γ_i has boundary word w_i, $\Phi(\gamma_i)$ has boundary word $\Phi(w_i)$.

Lemma 4.11. *Let $H = \langle A | R \rangle$ be a group of type FP_2 with homological finite presentation $\langle A \| R_0 \rangle$. Let γ be a 1-cycle in $X(A, R_0)^{(1)}$ and let $c = \sum_i a_i \sigma_i$ be a 2-chain with $\partial c = \gamma$. There exists a constant C and a 2-chain c' with $\partial c' = \Phi(\gamma)$ and $\mathrm{Area}(c') \leq C \cdot \mathrm{Area}(c)$.*

Proof. Let $c^{(1)}$ denote the 1-skeleton of the subcomplex of $X(A, R_0)$ consisting of the edges bounding each 2-cell in c. Subdivide and relabel each edge in $c^{(1)}$, so that the edge previously labelled by a generator a_i is now labelled by a word $\Phi(a_i)$. The outside boundary of this complex is the image $\Phi(\gamma)$, the compact regions have boundaries labelled by words $\Phi(r)$ for $r \in R_0$. Each $\Phi(r)$ is a relation in H, so can be filled by a 2-chain. Fix a 2-chain $c_{\Phi(r)}$ for each $r \in R_0$, set

$$C = \max \left\{ \mathrm{Area}(c_{\Phi(r)}) \,\middle|\, r \in R_0 \right\},$$

then by filling each of the compact regions with boundary labelled $\Phi(r)$ by the 2-chain $c_{\Phi(r)}$, we obtain a 2-chain filling c' for $\Phi(w)$ with

$$\mathrm{Area}(c') \leq C \cdot \mathrm{Area}(c).$$

□

Remark. It follows that since the above holds for any filling c, we must have

$$\mathrm{HArea}(\Phi(\gamma)) \leq C \cdot \mathrm{HArea}(\gamma).$$

Lemma 4.12. *Let $H = \langle A | R \rangle$ be a group of type FP_2 with homological finite presentation $\langle A \| R_0 \rangle$. Let γ be a 1-cycle in $X(A, R_0)^{(1)}$ and let $c' = \sum_i a_i \sigma_i$ be a 2-chain with $\partial c' = \Phi(\gamma)$. There exist constants C', C'' and a 2-chain c'' with $\partial c'' = \gamma$ and $\mathrm{Area}(c'') \leq C' \cdot \mathrm{Area}(c') + C'' \cdot |\gamma|$.*

Proof. Analogous to the above, we obtain a filling $c_{\Psi\circ\Phi(\gamma)}$ for $\Psi\circ\Phi(\gamma)$ with area at most $C'\cdot\mathrm{Area}(c')$, where we fix fillings of the relations $\Psi\circ\Phi(r)$ and set

$$C' = \max\left\{\mathrm{Area}(c_{\Psi\circ\Phi(r)})\,|\,r \in R_0\right\}.$$

For each generator $a \in A$, fix a 2-chain filling c_a for the relation $a = \Psi\circ\Phi(a)$. Let

$$C'' = \max\left\{\mathrm{Area}(c_a)\,|\,a \in A\right\}.$$

Attaching such 2-chains to the boundary of $c_{\Psi\circ\Phi(\gamma)}$ - there will be one for each edge in the boundary of γ - yields a 2-chain filling c'' for γ with

$$\mathrm{Area}(c'') \leq C'\cdot\mathrm{Area}(c') + C''\cdot|\gamma|.$$

□

Remark. As above, we must then have

$$\mathrm{HArea}(\gamma) \leq C'\cdot\mathrm{HArea}(\Phi(\gamma)) + C''\cdot|\gamma|.$$

Corollary 4.13. *The above lemmas can be applied to surface diagram fillings as follows. For a* 1*-cycle* γ *in* $X(A,R_0)^{(1)}$*, there exists a surface diagram filling* (S,π) *for* γ*, then* $\pi(S)$ *is exactly some* 2*-chain* c *with* $\partial c = \gamma$*, which can be exchanged using Lemma 4.11 for a* 2*-chain* c' *with* $\partial c' = \Phi(\gamma)$*, which has a corresponding surface diagram filling* (S',π') *by Lemma 4.10, noting that the area of the surface diagram is exactly the area of the* 2*-chain by construction. The process for Lemma 4.12 is similar.*

4.5 t-rings

4.5.1 In the Homological Cayley Complex

We justify the use of surface diagrams as opposed to working within c itself since the placement of t-rings on a surface diagram is 'nicer' than those in a 2-chain. One possible concern regarding the structure of an arbitrary 2-chain filling regards the uniqueness of annuli containing a given t-edge. Consider the diagram Fig. 5 where the 2-chain is not a surface, but has the cross-section of a Θ-graph. In such a case, Gersten's argument in section 5 of [Ger96b] fails - there exists a t-edge which is contained in two distinct annuli. However, in a surface diagram S for this 2-chain filling c, these become either two disjoint loops or one larger loop – thus removing this potential cause for concern.

Another problem with working directly with c is as follows. In [GS02], the proof of Theorem B asserts that each edge in a disc diagram for a loop w in H can be contained in at most 2 t-rings, consider the case of Fig. 6. In c, this is no longer the case, since the 2-chain c is not necessarily a surface.

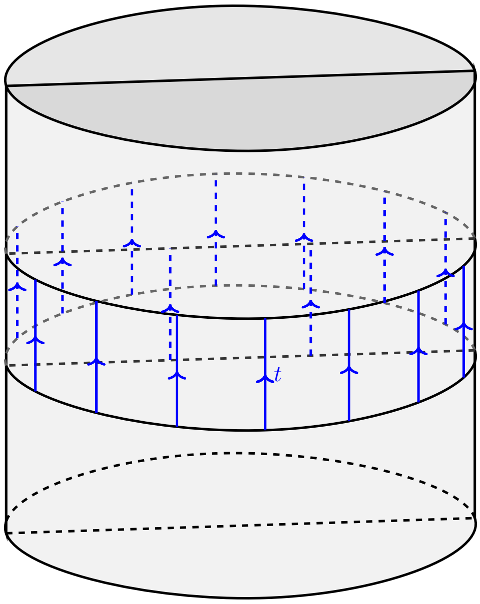

Figure 5: A Θ-shaped t-cycle – the 2-cells in the centre are doubly counted

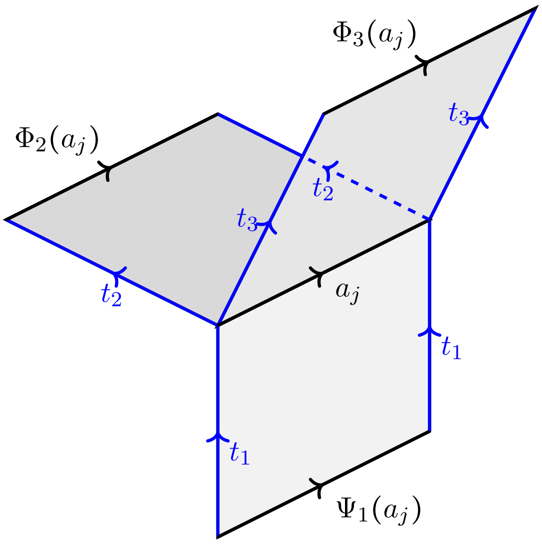

Figure 6: An edge contained in 3 t-rings

It turns out that one can still bound above the number of t-rings containing a given edge by $2n$, where n is the rank of the free group in the split extension. However, using surface diagrams allows us to use Gersten's initial stronger bound of 2.

4.5.2 On Surface Diagrams

We consider what a t-ring we might encounter in an arbitrary surface diagram filling S of a 1-cycle γ could look like. When K is finitely presented, a loop always admits a filling by a disc diagram, and one can always 'push down' a t-ring by exchanging a disc diagram filling of a word w for a disc diagram filling of $\Phi(w)$, or vice versa. However, when the filling is by a surface diagram, there is a concern that a t-ring may not be nullhomotopic in S – consider the case pictured in Fig. 7a, where a t-ring wraps around a handle. A priori, in this case we cannot apply Lemma 4.11 or Lemma 4.12 to push down our loop, since there is no surface diagram filling for either w or $\Phi(w)$.

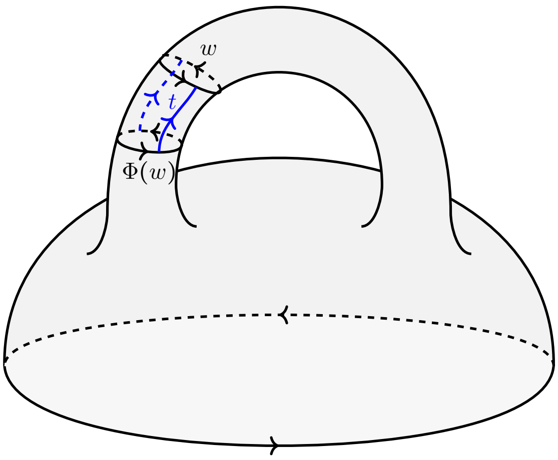

(a) A concerning t-ring placement

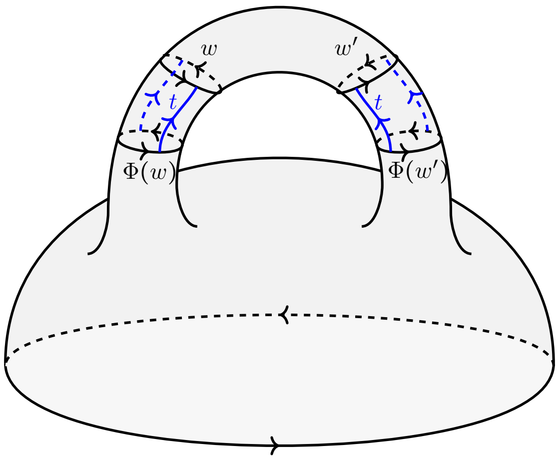

(b) The actual case

Figure 7

Fortunately, such a situation cannot occur - for the same reason as the proofs in Lemma 4.6 and Lemma 4.7 work. Recall that the Cayley graph of H consists of many copies of K, each corresponding to an element of F_n, attached to one another only by t-edges when the corresponding vertices are adjacent in the Cayley graph of F_n generated by the stable letters t_i. As such, the coarse structure of the Cayley graph of H looks like the Cayley graph of F_n but with copies of the Cayley graph for K at each vertex. The situation in Fig. 7a cannot occur as such since it would require a path traversing t-edges to have the same endpoints as a path which does not traverse t-edges. This is not possible, since the latter path must be contained in a single coset and the former must be contained in a different one. Since no path can exist in the Cayley graph, no path can exist in a surface diagram, and it follows that we must have the situation as in Fig. 7b, where we can apply Lemma 4.11 to remove these two t-rings simultaneously.

The two t-rings pictured in Fig. 7b are both components of the boundary of a sub-surface of S, and it is impossible to separate the two when ‘pushing down’, that is, no sub-surface of S exists whose boundary is only one of these two ‘linked’ t-rings. Hence it is natural to think of these not as two separate t-rings but as one t-*cycle*. Note that the generality of Lemma 4.11 and Lemma 4.12 allow us to ‘push up’ or ‘push down’ general 1-cycles, not just loops.

In the Cayley graph, a t-cycle that we work with has all its initial vertices in one coset Kw, and all its terminal vertices in the coset Kwt_i for some $i \in \{1, \dots, n\}$. We call the collection of loop in S corresponding to the boundary in Kw the *lower boundary* and that for Kwt_i the *upper boundary.* More precisely, we obtain our t-cycles as follows.

Construction 4.14 (Finding t-cycles)**.** *Let c be a finite 2-chain in*

$$X = X(A \cup \{t_1, \dots, t_n\}, R_0 \cup \{t_i^{-1} a_j t_i \Phi_i(a_j)^{-1}\})$$

whose boundary is a 1-cycle γ with no t-edges. Then every edge labelled t_i in c has associated to it a t_i-cycle whose two boundaries are 1-cycles contained in subcomplexes without t-edges.

Consider a surface diagram S for c. Consider an arbitrary t-edge with label t_i in S. The initial vertex of this edge lies in the sub-surface corresponding to the coset Kw for some reduced word $w \in F(t_1, \dots, t_n)$ and the terminal vertex lies in the subcomplex corresponding to Kwt_i. These two sub-surfaces are connected only by edges labelled t_i, since the corresponding subcomplexes in c have this property.

Arguing as Gersten in [Ger96b], since there are no t-edges on the boundary γ, each t-edge must be contained in an annulus in the interior of the surface diagram S, whose 2-cells correspond to stringing together relators of the form $t_i^{-1} a_j t_i \Phi_i(a_j)^{-1}$ in a circuit. This argument still holds in the more general case of a surface diagram, and we can also deduce that each t-edge can only be contained in one annulus, since S is a surface and cannot have features such as those pictured in Fig. 5 or Fig. 6. The t-cycle we can use with Lemma 4.11 and Lemma 4.12 consists of all t_i-annuli with one boundary in Kw and the other in Kwt_i.

4.6 Proving the Generalised Theorem

Recall the hypotheses of Theorem 4.5.

Theorem 4.5. *Let H be a split extension of a group K of type FP_2 by a finitely generated free group F_n, so that we have the short exact sequence*

$$1 \to K \to H \to F_n \to 1.$$

If (f, g) is a homological area-radius pair for H with $f(n) \geq n$ for all n, then there is a constant $C > 1$ such that $FA_K \preccurlyeq C^g f$.

Proof. Let $K \cong \langle A \,|\, R \rangle$ have a homological finite presentation $\langle A \| R_0 \rangle$. Then by Lemma 4.6, $\left\langle A \cup \{t_1, \dots, t_n\} \,\middle\|\, R_0 \cup \left\{t_i^{-1} a_j t_i \Phi_i(a_j)^{-1} \,|\, a_j \in A, i \in \{1, \dots, n\}\right\}\right\rangle$ is a homological finite presentation for H, with corresponding homological Cayley complex X. Let γ be a 1-cycle in $\Gamma(K, A)$. We embed γ into X by identifying $\Gamma(K, A)$ with the subcomplex corresponding to K in the 1-skeleton $X^{(1)} = \Gamma(H, A \cup \{t_1, \dots, t_n\})$.

Let $c = \sum_i a_i \sigma_i$ be a 2-chain with $\partial c = \gamma$ in X with area no more than $f(|\gamma|)$. Let (S, π) be a surface diagram filling for c. In S, the 2-cells corresponding to relations $t^{-1} a_j t \Phi(a_j)^{-1}$

form t-cycles as in Construction 4.14. Each t-cycle is associated to one of the letters t_i and connects regions of S corresponding to cosets Kw and Kwt_i. Let W_c be the set of all reduced words $w \in F(t_1, \ldots, t_n)$ such that the filling c has non-empty intersection with the subcomplex of X corresponding to the coset Kw. Note that the set W_c contains the empty word, since the loop γ, the boundary of c, lies in the subcomplex corresponding to K.

Let w be a word of maximal length k in W_c, we consider 2 possible cases.

If the last letter of w is a stable letter, that is, $w = w' * t_i$ for some $0 \leq i \leq n$, then the region of S corresponding to Kw is a surface diagram filling S_u for the upper boundary u of the t_i-cycle connecting Kw' to Kw which has no interior t-edges, that is, a K-filling. The labels on the upper boundary u can be obtained by applying Φ_i to the lower boundary ℓ of the t_i-cycle. Applying Lemma 4.12 as described in Corollary 4.13, we can exchange the K-filling S_u of u and the t_i-cycle itself for a K-filling S_ℓ of ℓ whose area $\mathrm{Area}(S_\ell) \leq C' \cdot \mathrm{Area}(S_u) + C'' \cdot |u|$, where C' and C'' are constants as in Lemma 4.12. Any edge in S can occur in no more than 2 t-cycles, so twice the number of edges in S is an upper bound on the length $|u|$. Let ρ be the length of the longest relation in $R_0 \cup \left\{t_i^{-1} a_j t_i \Phi_i(a_j)^{-1}\right\}$, then $|u| \leq |\gamma| + 2\rho \mathrm{Area}(S)$. Applying this process simultaneously to all maximal words at worst multiplies the area of S by a constant C' and adds at most $C''(|\gamma| + 2\rho \mathrm{Area}_X(c)) \leq Mf(|\gamma|)$, for some positive constant M since $f(n) \geq n$. To avoid degeneracies, we take $M \geq \max\{C, C', C''(2\rho+1), 1\}$.

If instead the last letter of w is the inverse of a stable letter, *i.e.* $w = w' * t_i^{-1}$ for some $0 \leq i \leq n$, then the part of S corresponding to Kw is a K-filling S_ℓ for the lower boundary ℓ of the t_i-cycle connecting Kw' and Kw. As above, we have that $u = \Phi(\ell)$ and applying Lemma 4.11 as described in Corollary 4.13 tells us that we can exchange the K-filling S_ℓ for ℓ and the t_i-cycle itself with a K-filling S_u for u whose area $\mathrm{Area}(S_u) \leq C \cdot \mathrm{Area}(S_\ell)$, where C is a constant as in Lemma 4.11. Applying this process simultaneously to all maximal words at worst this multiplies the area of S by a constant $C \leq M$. In both cases, we have removed all words of length k in W_c from the filling by eliminating their corresponding t_i-cycles, that is, we have obtained a new surface diagram filling for γ that intersects only the cosets corresponding to words in W_c of length strictly less than k.

The length k of a maximal word $w \in W_c$ is bounded above by the radius $g(|\gamma|)$, hence one can apply the above procedure $k \leq g(|\gamma|)$ times to obtain a surface diagram filling (S', π') for γ whose image is contained in K and whose area is no more than

$$\mu^{g(\gamma)}\left(f\left(|\gamma|\right)\right) \leq M^{g(|\gamma|)+1} f(|\gamma|) \leq (M^2)^{g(|\gamma|)} f(|\gamma|)$$

where $\mu(x) = Mx + Mf(|\gamma|)$ is the upper bound on the area increase given by each iteration of Lemma 4.12. Thus it suffices to take the constant $C \geq M^2$ in the statement of the theorem. $\square$

5 Hyperbolic Groups

A key motivation for Gersten and Short's work was to better understand isoperimetric inequalities for subgroups of hyperbolic groups. Hyperbolic groups are a well-studied class of groups since Gromov's seminal work [Gro87].

Definition 5.1. A group H is called *word hyperbolic* if its Cayley graph Γ with respect to a fixed finite generating set, endowed with the graph metric – each edge is given length 1, and the distance between two vertices is given by the length of the shortest path between them – satisfies the property that every geodesic triangle is δ-slim. Meaning that there exists some $\delta > 0$ where if $[x, y]$ denotes a geodesic path between vertices x and y in Γ, then for any $a, b, c \in H$, any two of the δ-neighbourhoods $B_\delta([a,b]), B_\delta([a,c]), B_\delta([b,c])$ covers the geodesic triangle $\Delta(a, b, c)$ (see Fig. 8).

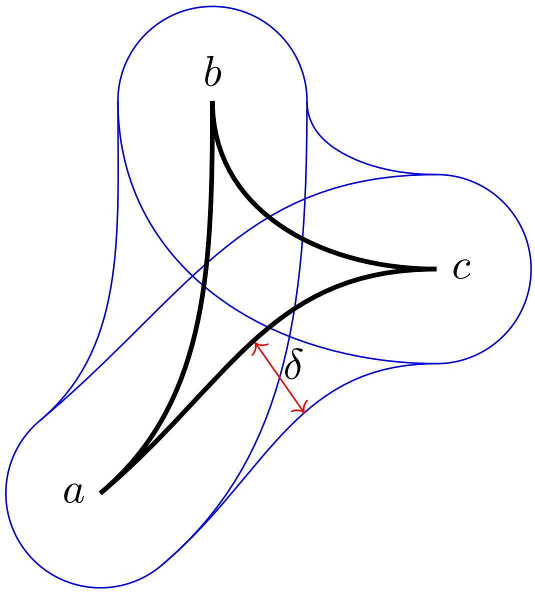

Figure 8: A δ-slim triangle

Hyperbolic groups are characterised by satisfying a linear isoperimetric inequality, that is, a finitely presented group H is hyperbolic if and only if $\delta_H(n) \simeq n$, but this property does not necessarily hold for subgroups of H.

Gersten and Short's general result is applied to these subgroups in Theorem A of [GS02]:

Theorem 5.2. *Given a split extension*

$$1 \to K \to H \to F_n \to 1$$

Where F_n is free and K is finitely presented, then if H is hyperbolic, K satisfies a polynomial isoperimetric inequality.

For hyperbolic groups, Gersten and Short give, in Lemma 2.2 of [GS02], an area-radius pair (the proof is contained in Appendix B).

Lemma 5.3. ([GS02]) *Let $\mathcal{P} = \langle A|R\rangle$ be a finite presentation of a hyperbolic group H, then there are constants $B, C > 0$ such that for any nullhomotopic word $w \in \overline{A}^*$ with $\ell(w) \geq 1$, there is a Van Kampen diagram over $\mathcal{P}$ with area at most $B\ell(w)(\log_2(\ell(w)) + 1)$ and radius at most $C(\log_2(\ell(w)) + 1)$.*

In fact, Corollary 4.2 of [GR02] tells us that all minimal area Van Kampen diagrams of finite presentations of hyperbolic groups satisfy a uniform logarithmic upper bound on their radii. Applying this AR pair to Theorem 3.1 gives us Theorem 5.2.

Applying this Area-Radius pair to Theorem 4.5, after noting that an AR pair in the sense of [GS02] is also a homological AR pair, we can obtain the following corollary:

Theorem 5.4. *Let H be a split extension of a group K of type FP_2 by a finitely generated free group F_n, so that we have the short exact sequence*

$$1 \to K \to H \to F_n \to 1.$$

If H is hyperbolic, then FA_K is bounded above by a polynomial.

5.1 Linear Homological Isoperimetric Inequality

A key property of hyperbolicity in groups is that all hyperbolic groups (which are necessarily finitely presented) satisfy a linear isoperimetric inequality, and any finitely presented group satisfying one is necessarily hyperbolic. This link extends to bounds on the homological area as well, as illustrated by the following results:

Lemma 5.5. *Let H be a hyperbolic group, then $FA_H(n) \preccurlyeq n$.*

Proof. First one notes that all hyperbolic groups are finitely presented, argued in Chapter III.Γ of [BH99]. We discussed in Lemma 2.8 the relationship between homological and homotopical Dehn functions, notably, $\mathrm{FA}_H(n) \preccurlyeq \overline{\delta_H(n)}$ for any finitely presented group H. Since hyperbolic groups are known to satisfy a linear isoperimetric inequality, observing that the map $n \mapsto n$ is (super)additive, we can say that $\mathrm{FA}_H(n) \preccurlyeq n$. □

For the converse direction, we take cue from Theorem 2.9 of chapter III.H of [BH99] for our proof, adapting this proof to hold for groups of type FP_2.

Lemma 5.6. *Let H be a group of type FP_2 which satisfies a linear homological isoperimetric inequality, then H is hyperbolic.*

Proof. First note that H satisfies a linear homological isoperimetric inequality, so we can say $\mathrm{HArea}_H(n) \leq Cn$ for some constant C. We will show that the Cayley graph for H is δ-hyperbolic for some δ by studying a homological Cayley complex for H. Let $S \subseteq H$ be a generating set, $H \cong \langle A|R\rangle$ a presentation and $\langle A\|R_0\rangle$ a homological finite presentation for H.

We will show that there exists a $\delta > 0$ such that $X(A, R_0)$ is δ-hyperbolic by bounding above integers $n > 0$ so that there is a geodesic triangle not $(n+1)$-slim.

Fix n, and consider a geodesic triangle $\Delta = \Delta(a,b,c)$ in $\Gamma(H,A)$ for $a,b,c \in H$. Suppose that there is a point $v \in [a,b]$ the geodesic in Δ so that

$$d_\Gamma(v, [a,c] \cup [b,c]) \coloneqq \inf\{ d_\Gamma(v,x) | x \in [a,c] \cup [b,c]\} \geq n+1$$

i.e. Δ is not $(n+1)$-slim. If v lies on an edge, we exchange it for its nearest vertex in $[a,b]$, in which case we now have $d_\Gamma(v, [a,c] \cup [b,c]) \geq n$.

Define

- $L = \max\{|r| \mid r \in R_0\}$ – the maximum length of a relator in R_0;
- $N = \max\{d_\Gamma(x,y) \mid x, y \in r, r \in R_0\}$ – the maximum diameter of a relator in R_0;
- $k = LCN$, $m = LC$ – two constants, note $k = mN$. To avoid degenerate cases, we will assume that $n > 6k$.

After possibly renaming a and b, We consider two possible cases, the *hexagonal* and *quadrilateral* cases. In the hexagonal case, $[a,v]$ in disjoint from the $4k$-neighbourhood of $[b,c]$ and $[v,b]$ is disjoint from the $4k$-neighbourhood of $[a,c]$. In the quadrilateral case, there is some $w \in [v,b]$, $w' \in [a,c]$ so that $d_\Gamma(w,w') = 4k$.

In the hexagonal case, there is a minimal subarc $[u,w] \subset [a,b]$ containing v with endpoints at distance exactly k from the other sides of Δ. There are points realising this distance k, denote by u' the point closest to a at distance k from u, w' the point closest to b at distance k from w. We also define u'', w'' to be so that $[u',u''] \subset [a,c]$, $[w',w''] \subset [b,c]$ are maximal subarcs so that the k-neighbourhoods of their interiors are disjoint, so $d_\Gamma(u'', w'') = 2k$. In the quadrilateral case, we define u' similarly. Compare with Fig. 9 below.

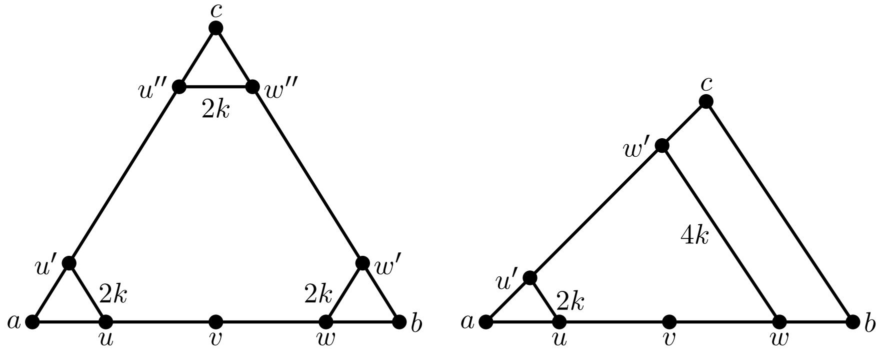

(a) The Hexagonal Case (b) The Quadrilateral Case

Figure 9

We will deal first with the hexagonal case. Let $\mathscr{H}$ denote a geodesic hexagon with vertices u, u', u'', w, w', w'' as illustrated, where the sides $[u,w], [u',u''], [w',w'']$ are subarcs of the edges in $\Delta(a,b,c)$ and the other three sides have length $2k$. $\mathscr{H}$ is a loop in the Cayley graph, so admits a surface diagram filling (S,π) as defined in Section 4.3. Note that the boundary of each 2-cell in the combinatorial structure on S has at most L edges. We identify points in $\mathscr{H}$ with their preimage under π.

In order to derive a contradiction, we aim to estimate the number of 2-cells in the k-neighbourhood of $\mathscr{H}$ in S. Since the k-neighbourhoods of the arcs $[u,w],[u',u''],[w',w'']$ are disjoint in S, we will obtain a lower bound for the number of faces in each k-neighbourhood.

Let $\alpha = d_\Gamma(u,w)$, $\beta = d_\Gamma(u',u'')$, and $\gamma = d_\Gamma(w',w'')$. Consider first the 2-cells in S that intersect the subarc $[u,w]$ of $\mathscr{H}$. We define the subsurface D_m inductively: let D_1 denote the minimal combinatorial subsurface of S containing the union of these two cells. Let D_{n+1} denote the minimal combinatorial subsurface of S containing the union of the 2-cells in S which intersect D_n. Observe that D_m lies in the k-neighbourhood of $[u,w]$ in S, since $k = mN$ where N is the maximum diameter of any 2-cell.

We estimate the number of 2-cells in D_m as follows. Each 2-cell has a perimeter of length no more than L, it follows that there are at least $\frac{\alpha}{L}$ 2-cells in D_1. Fo each $i \in \{1,2,\dots,m-1\}$, there is a unique injective edge path connecting u to w in ∂D_i with no edges in $[u,w]$, this has combinatorial length at least α. Since $m = \frac{k}{N}$, if any edges of the path lie in $\mathscr{H}$, they can only lie on the edges $[u,u']$, $[w,w']$ – each of length $2k$, and the intersection with the k-neighbourhood of $[u,w]$ (containing D_m) is at most k. It follows that there are at least $\frac{\alpha-2k}{L}$ 2-cells in $S \smallsetminus D_i$ which abut this path, so at least as many faces in $D_{i+1} \smallsetminus D_i$. This yields a lower bound of

$$\frac{m}{L}(\alpha - 2k) = C(\alpha - 2k)$$

for the number of 2-cells in the k-neighbourhood of $[u,w]$ in S.

Applying the same argument to the edges $[u',u''],[w',w'']$, we get a lower bound

$$\mathrm{Area}(S) \geq C(\alpha+\beta+\gamma) - 6kC.$$

This lower bound can be improved upon, consider the k-neighbourhood of v, our partial filling D_m must intersect this, and we know that $d_\Gamma(v, \mathscr{H} \smallsetminus [u,w]) \geq n-2k$, so there is some combinatorial arc $\mathbf{A}_+$ on the boundary $\partial D_m \smallsetminus \mathscr{H}$ of length at least $(n-3k)$ – corresponding to the bold blue region in Fig. 10 – contained in the $(n-2k)$-neighbourhood of v. Noting that we took n to be at least $6k$, this is a positive length. There are at least $\frac{n-3k}{L}$ 2-cells in S which abut $\mathbf{A}_+$ which are not contained in either D_m or the k-neighbourhoods of $[u',u''],[w',w'']$ considered previously. It follows that we obtain the lower bound

$$\mathrm{Area}(S) \geq C(\alpha+\beta+\gamma) - 6kC + \frac{n-3k}{L}$$

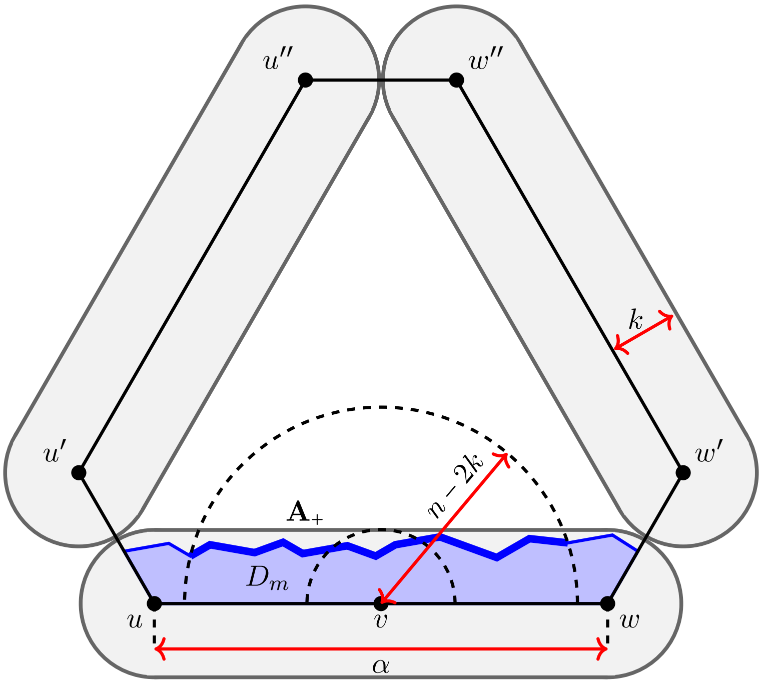

Figure 10: The Hexagonal Case $\mathscr{H}$

for the number of 2-cells in S. However, since $\mathrm{HArea}_H(n) \leq Cn$, we can note that since the length of $\mathscr{H}$ is $\alpha + \beta + \gamma + 6k$, that

$$\mathrm{Area}(S) \leq C(\alpha + \beta + \gamma) + 6kC.$$

So we must have the case that

$$\tfrac{n-3k}{L} \leq 12kC,$$

but the right hand side of this expression is a constant – it is clear that n is bounded.

We now turn our attention to the quadrilateral case. Let $\mathscr{Q}$ denote a geodesic quadrilateral with vertices u, w, w', u' as pictured in Fig. 9, where the sides $[u, w], [u', w']$ are subarcs of the sides $[a, b], [a, c]$ respectively, $[u, u']$ has length $2k$ and $[w, w']$ has length $4k$.

Let (S, π) be a surface diagram filling for $\mathscr{Q}$, and identify points in $\mathscr{Q}$ with their preimage under π. The k-neighbourhoods of the arcs $[u, w], [u', w']$ are disjoint in S, we argue similarly to the hexagonal case. Let α denote the length of $[u, w]$ and define D_m analogously to the above. With the same reasoning, we can see that there are at least $K(\alpha - 2k)$ 2-cells in the k-neighbourhood of $[u, w]$ is S. If β denotes the length of the arc $[u', w']$, we can deduce

$$\mathrm{Area}(S) \geq C(\alpha + \beta) - 6kC.$$

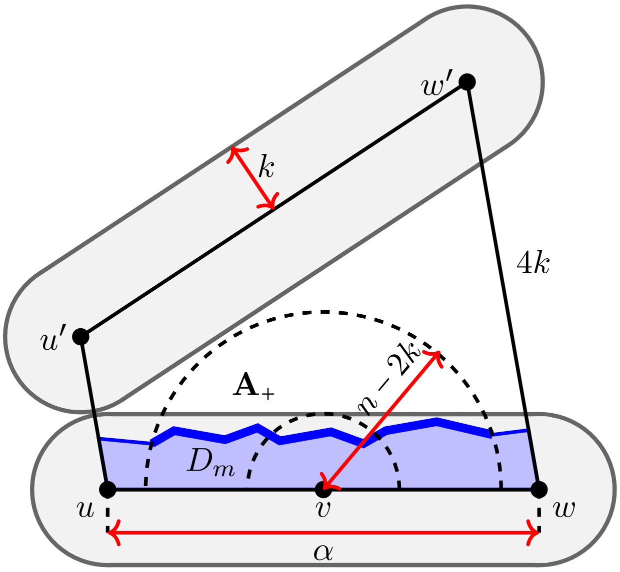

Figure 11: The Quadrilateral Case $\mathscr{Q}$

As above, the $(n-2k)$-neighbourhood of v does not intersect the k-neighbourhood of $[u', w']$, thus we can deduce that there is a corresponding combinatorial arc $\mathbf{A}_+$ of length $n-3k$ in this neighbourhood – as shown in Fig. 11, which must abut at least $\frac{n-3k}{L}$ 2-cells in S. So we obtain the sharper lower bound of

$$\text{Area}(S) \geq C(\alpha+\beta) - 6kC + \tfrac{n-3k}{L}.$$

Similarly noting that the linear isoperimetric inequality yields an upper bound of

$$\text{Area}(S) \leq C(\alpha+\beta) + 6kC,$$

we can deduce that n is bounded as before.

This bound on n says exactly that there is an upper bound on the integers n so that there exists a geodesic triangle that is not $(n+1)$-slim (recalling our hypothesis). One can deduce that H is δ-hyperbolic, and see that $\delta \leq 13kCL + 3k$ is a sufficient bound for both cases. $\square$

Corollary 5.7. *A group H of type FP_2 is hyperbolic if and only if it satisfies a linear homological isoperimetric inequality.*

5.2 Subquadratic Homological Isoperimetric Inequality

The above results can be strengthened further. In [Gro87], Gromov uses analytic techniques to prove that any a group satisfying a subquadratic isoperimetric inequality is in fact hyperbolic. This result proves the existence of the so-called *Gromov gap* in the isoperimetric spectrum – any subquadratic function that appears as a Dehn function must in fact be linear. In fact, Gromov's result applies more broadly to isoperimetric inequalities in other areas, such as Riemannian manifolds. We aim to extend this result to apply to groups H of type FP_2, by applying Bowditch's result from [Bow95].

5.2.1 Groups of type FP_2 Satisfy Bowditch's Axioms

Bowditch's result is very general, imposing only two requirements on a geodesic metric space (X, d), where we consider only loops and paths of finite length:

(A1) The Triangle inequality for theta curves. If loops $\gamma_1, \gamma_2, \gamma_3$ form a theta curve, then $\mathrm{Area}(\gamma_1) \leq \mathrm{Area}(\gamma_2) + \mathrm{Area}(\gamma_3)$,

(A2) The Rectangle inequality. If a loop γ is split into four subpaths $\gamma = \alpha_1 \cup \alpha_2 \cup \alpha_3 \cup \alpha_4$, then $\mathrm{Area}(\gamma) \geq d_1 d_2$, where $d_1 = \min\{\, d(x,y) \,|\, x \in \alpha_1, y \in \alpha_3\}$, and $d_2 = \min\{\, d(x,y) \,|\, x \in \alpha_2, y \in \alpha_4\}$,

where a *theta curve* is a collection of 3 loops formed by taking three paths $\alpha_1, \alpha_2, \alpha_3$ from x to y, and defining the loop $\gamma_i \coloneqq \alpha_{i+1} \cup (-\alpha_{i+2})$, where $-\alpha$ is the path α traversed backwards, and the indices are taken mod 3. Note that theta curve can be obtained from a loop $\gamma : S^1 \to X$ for any two points $t \neq u \in S^1$ by cutting γ into two paths α_1, α_2, one being the image of the segment $[t, u]$ and the other being the image of $[u, t]$, and setting α_3 to be a geodesic $[\gamma(t), \gamma(u)]$, as pictured in Fig. 12.

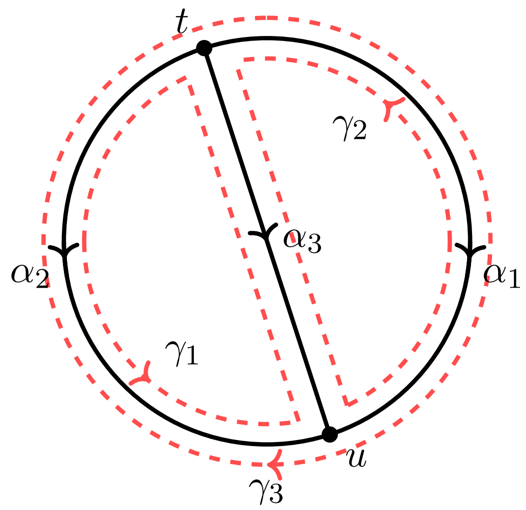

Figure 12: A Theta Curve

We first check that groups H of type FP_2, with distance from the edge metric on the Cayley graph and area being the minimal number of 2-cells in a surface diagram filling for a loop, satisfy these two axioms.

We suppose that we have a *triangular presentation* for H, that is, one in which the length of every relation is 3. Given any presentation $\langle A|R\rangle$ for H, one can turn it into a triangular presentation as follows. If there is a relator of length 1, then this means that the corresponding generator is trivial, so remove it. If there is a relator of length 2, then this means that either two generators are equal, one generator is inverse to another, or there is a generator a which is self-inverse. In the former cases, we can exchange one generator for the other (or its inverse) whenever it appears in a relator, and remove it from the generating set; in the latter case, we can pass to the index 2 subgroup $\langle A \smallsetminus \{a\}\rangle$ (since hyperbolicity is a quasi-isometry invariant, it suffices to prove that H has a finite-index subgroup which is hyperbolic). If there is a relator $r_1 r_2 \cdots r_n$ of length greater than 3, we can introduce a new generator a, and exchange the relator $r_1 r_2 \cdots r_n$ for two relators: $r_1 r_2 a$ – of length 3, and $a^{-1} r_3 r_4 \cdots r_n$ – of length $n-1$. One repeats this process until all relators are of length 3.

Geometrically, a triangular presentation for H means that every 2-cell in the surface diagram has 3 edges on its boundary, this property is useful for proving (A2). Without assuming triangularity, Bowditch's axiom holds up to a constant multiple, and the proof of hyperbolicity works with minimal adaptation.

Lemma 5.8. *Groups of type FP_2 satisfy Bowditch's axioms (A1), (A2).*

Note in the below proof, since we restrict ourselves to loops, this holds for both the abelianised Dehn function δ_H^{ab} and the homological Dehn function FA_H.

Proof. For (A1), observe that for any three loops forming a theta curve, given surface diagram fillings for two of the loops γ_2, γ_3, gluing these along the shared boundary yields a surface diagram whose area is $\mathrm{Area}(\gamma_2) + \mathrm{Area}(\gamma_3)$ and whose boundary is the third loop γ_1, thus it follows that $\mathrm{Area}(\gamma_1) \le \mathrm{Area}(\gamma_2) + \mathrm{Area}(\gamma_3)$ since this surface diagram may not be minimal.

For (A2), consider a loop γ and a minimal surface diagram filling (S, π). Split γ into four subpaths $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ as in the statement of the axiom. Define

$$d_1' = \min\{|\delta| \mid \delta \text{ is a path from } x \text{ to } y \text{ in } S\ x \in \alpha_1, y \in \alpha_3\},$$
$$d_2' = \min\{|\delta| \mid \delta \text{ is a path from } x \text{ to } y \text{ in } S\ x \in \alpha_2, y \in \alpha_4\}.$$

Note that $d_i' \ge d_i$, since there may exist a shorter path in the Cayley graph than exists in the surface diagram, and d_i' maps under π to a path of the same length in the Cayley graph.

Each edge along the path α_2 is contained in the boundary of a triangular 2-cell in S. Suppose that fewer than d_1' triangles abut α_2, then since the diameter of each triangle is 1, this gives a path of length strictly less than d_1' from α_1 to α_3, contradicting the definition of d_1', so there is a layer of at least d_1' triangles abutting α_2. We consider the edges of these triangles which do not lie in α_2, there is a path from α_1 to α_3 consisting

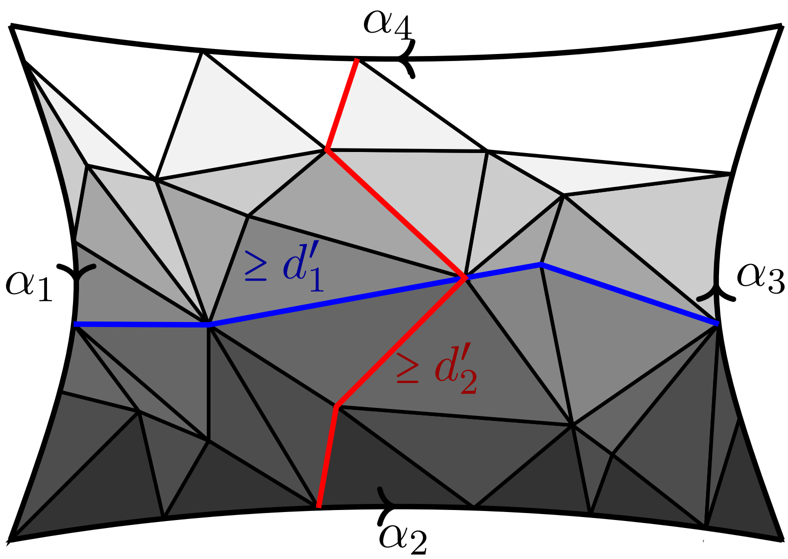

Figure 13: At least d'_2 layers, each of at least d'_1 triangular 2-cells contained in S

of such edges of length at least d'_1 (beginning with taking all of the non-boundary edges and observing that there must be a connected subpath from α_1 to α_3, which has length at least d'_1 by assumption). Consider the subsurface of S obtained by removing all of the triangles in the first layer. By the same argument, the new boundary must abut at least d'_1 triangles. We repeat this argument, until a layer containing a triangle with a vertex in α_4. Suppose that this is the n-th layer. Since the triangles each have diameter 1, this gives a path traversing n edges from α_2 to α_4, so $n \geq d'_2$. It follows that the surface diagram contains at least d'_2 layers of d'_1 triangles, so $\text{Area}(\gamma) = \text{Area}(S) \geq d'_1 d'_2 \geq d_1 d_2$ as desired. Compare with Fig. 13, where each layer is illustrated by a different shade. □

5.2.2 Bowditch's Proof

Now we can proceed with applying Bowditch's proof in [Bow95] to the specific case of groups of type FP_2. Bowditch splits the result into two lemmas, the first of which we modify to suit the case the H is a group of type FP_2. Recall that $\delta_H^{ab} \preccurlyeq \text{FA}_H$ (the set that the supremum in the definition of FA_H is taken over contains the corresponding set in the definition of δ_H^{ab}), so if FA_H is subquadratic, so is δ_H^{ab}. Suppose that δ_H^{ab} is subquadratic, and suppose that we prove δ_H^{ab} is in fact linear. Then it follows that, since $\text{FA}_H \simeq \overline{\delta_H^{ab}}$, FA_H is linear (linear functions $n \mapsto Cn + C$ are $\simeq$-equivalent to superadditive functions $n \mapsto Cn$) – so it suffices to prove Bowditch's result for δ_H^{ab}.

Lemma 5.9. *Fix a homological finite presentation $\langle A \| R_0 \rangle$ for H, denote $\delta^{ab}_{\langle H,A,R_0\rangle}$ by δ_H^{ab}. For all $n \in \mathbb{N}$, there exist $p, q \in \mathbb{N}$ so that*

1. $\delta_H^{ab}(n) \leq \delta_H^{ab}(p) + \delta_H^{ab}(q)$,
2. $p, q \leq \frac{3}{4}n + 3\sqrt{\delta^{ab}(n)}$,

3. $p + q \leq n + 6\sqrt{\delta_H^{ab}(n)}$.

Proof. Let $n \in \mathbb{N}$. By Proposition 2.22 of [BKS21], there exists a loop γ in the Cayley graph $\Gamma(H, A)$ which attains the supremum in the definition of δ_H^{ab}, that is, $\mathrm{HArea}_{X(A,R_0)}(\gamma) = \delta_H^{ab}(n)$ and $|\gamma| \leq n$.

Let $\gamma : S^1 \to \Gamma(H, A)$ be a parameterisation. Let $\Delta \subseteq S^1 \times S^1$ be the set of pairs (t, u) so that $\gamma(t), \gamma(u) \in H$ cut γ into two paths, each of length at most $\frac{3}{4}|\gamma|$. Set

$$\ell = \min\{d_\Gamma(\gamma(t), \gamma(u) \,|\, t, u \in \Delta\}.$$

Pick a pair $(t_0, u_0) \in \Delta$ so that $d_\Gamma(\gamma(t_0), \gamma(u_0)) = \ell$. Set $a = \gamma(t_0), b = \gamma(u_0)$ and call the two subpaths β_0 and β_1 of γ. Without loss of generality, we can say that $|\beta_0| \leq |\beta_1| \leq \frac{3}{4}|\gamma|$. Let $[a, b]$ denote a geodesic from a to b in $\Gamma(H, A)$, and pick $a', b' \in [a, b]$ (possibly not vertices) so that the segments $\delta = [a', b'], \alpha_1 = [a, a'], \alpha_2 = [b', b]$ each have length $\ell/3$.

By construction, $d_\Gamma(\alpha_1, \alpha_2) = \min\{d_\Gamma(x, y) \,|\, x \in \alpha_1, y \in \alpha_2\} = d_\Gamma(a', b') = \ell/3$. We claim that $d_\Gamma(\delta, \beta_0) \geq \ell/3$ also.

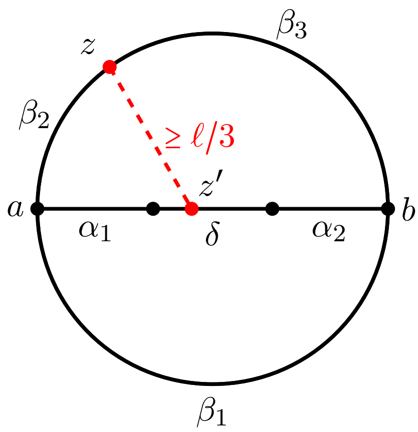

Figure 14

Let $z = \gamma(v) \in \beta_0$, and $z' \in \delta$. z cuts β_0 into two subpaths β_2, β_3, where $a \in \beta_2$ and $b \in \beta_3$ as shown in Fig. 14. Without loss of generality, we may assume that $|\beta_2| \leq |\beta_3|$. Note that $|\beta_1| + |\beta_2| + |\beta_3| = |\gamma|$, $|\beta_1| \leq |\beta_2| + |\beta_3|$ and $|\beta_2| \leq |\beta_3|$. It follows from these inequalities that $|\beta_1| \leq \frac{1}{2}|\gamma|$, $|\beta_2| \leq \frac{1}{2}(|\gamma| - |\beta_1|)$, so $|\beta_1| + |\beta_2| \leq \frac{3}{4}|\gamma|$. We also have that $|\beta_3| \leq \frac{3}{4}|\gamma'|$, which means that $(u_0, v) \in \Delta$. By minimality of $d_\Gamma(a, b)$ and the triangle inequality, we can say that

$$d_\Gamma(z, z') + d_\Gamma(z', b) \geq d_\Gamma(z, b) \geq \ell = d_\Gamma(a, b) = d_\Gamma(a, z') + d_\Gamma(z', b).$$

Then observe that

$$d_\Gamma(z, z') \geq d_\Gamma(a, z') \geq d_\Gamma(a, a') = \ell/3,$$

as claimed.

The geodesic $[a, b]$ forms 2 new curves γ_1, γ_2, where $\gamma_1 = \beta_0 \cup \alpha_1 \cup \delta \cup \alpha_2$ and $\gamma_2 = \beta_1 \cup \bar{\alpha_2} \cup \bar{\delta} \cup \bar{\alpha_1}$, where the bar denotes the reverse path. Together with γ, these form a theta curve. By the axiom (A1) we have that

$$\delta^{ab}(n) = \mathrm{HArea}(\gamma) \leq (\mathrm{HArea}(\gamma_1) + \mathrm{HArea}(\gamma_2)) \leq \delta^{ab}(p) + \delta^{ab}(q),$$

– condition 1. By the axiom (A2), we can see that $\mathrm{HArea}(\gamma_1) \leq (\ell/3)^2$, it follows that $\ell \leq 3\sqrt{\mathrm{HArea}(\gamma_1)} \leq 3\sqrt{\delta_H^{ab}(\gamma_1)} \leq 3\sqrt{\delta_H^{ab}(n)}$. Then $p, q \leq \frac{3}{4}|\gamma| + \ell \leq \frac{3}{4}n + 3\sqrt{\delta_H^{ab}(n)}$ – condition 2. We can similarly deduce that $p + q \leq |\gamma| + 2\ell \leq n + 6\sqrt{\delta_H^{ab}(n)}$ – condition 3. □

With the above proven for groups of type FP_2, Bowditch's Lemma 2 in [Bow95] gives precisely the following.

Lemma 5.10. *Suppose $f : [0, \infty) \to [0, \infty)$ is increasing, suppose there are constants $K > 0$ and $0 < \lambda < 1$ such that for all $x \in [0, \infty)$, there exist $p, q \in [0, \infty)$ such that*

1. *$f(x) \leq f(p) + f(q)$,*
2. *$p, q \leq \lambda x + K\sqrt{f(x)}$,*
3. *$p + q \leq x + K\sqrt{f(x)}$,*

then if $f(x) = o(x^2)$, then $f(x) = O(x)$.

From the above, and the earlier remarks, we can deduce the result.

Corollary 5.11. *If a group H of type FP_2 satisfies a subquadratic homological isoperimetric inequality, that it in fact satisfies a linear one, and is therefore hyperbolic.*

Remark. Moreover, one may also recall that all hyperbolic groups are in fact finitely presented, so this result also implies the strictly weaker result that if a group H is of type FP_2 and satisfies a subquadratic homological isoperimetric inequality, then it is finitely presented.

6 Conclusion

To summarise, we have argued that, with some slight modification, the methods of Gersten and Short in [GS02] still apply in the more general scenario of homological isoperimetric inequalities. This improved result is motivated by interest in subgroups of hyperbolic groups which are of type FP_2 but not finitely presented, known examples of which appear as kernels of surjective homomorphisms into $\mathbb{Z}$ (or free abelian groups $\mathbb{Z}^n$). We also looked at hyperbolic groups themselves, and showed that homological and homotopical Dehn functions share some key properties concerning such groups, namely that satisfying a linear (or even subquadratic) homological isoperimetric inequality is equivalent to hyperbolicity.

The proof given for Theorem 4.5 is not entirely constructive, actually obtaining such a presentation, and computing the constants used is a difficult task. While it is theoretically possible to compute the polynomial bound in the hyperbolic case explicitly, one might assume that an alternative approach would prove more practical, as was the case when Kropholler, Isenrich, and Soroko computed an explicit bound for a finitely presented analogue of this case in [KIS25].

The motivation to study this result of Gersten and Short also raises some other questions. As mentioned in the introduction, there is a generalisation of [GS02] in [Ise24], one might ask whether this result also holds when the kernel is not necessarily finitely presented but is of type FP_2? Isenrich's work uses more elaborate techniques of morse theory and BNSR invariants, rather than the hands-on approach of Gersten and Short. As such, adapting this result is beyond the scope of this report. There is a worry in attempting to follow the proof in [Ise24] with the weakened assumption, since in the comparison of Theorems 2.6 and 2.7, Isenrich uses that the kernel being of type F_2 and $FP_n(\mathbb{Z})$ is equivalent to being of type F_n. If we assume only that our kernel is of type $FP_2(\mathbb{Z})$ without the added property of being finitely presented, or equivalently of type F_2, then we cannot echo this argument.

The closing remark of section 5 suggests another question: which functions appear as homological Dehn functions of groups of type FP_2 that are not finitely presented? This question mirrors that of the isoperimetric spectrum. We have proven that '1' only appears as an exponent of a polynomial homological Dehn function when the group is finitely presented, so a natural question to ask is whether this (strictly) homological isoperimetric spectrum is also dense in the interval $[2, \infty)$, or if it may contain other gaps.

Appendices

A Homological Radius Functions with Surface Diagrams

As this project progressed, initially the work in Section 4 was done all in terms of the homological Cayley complex, with the surface diagrams content added retroactively. As such, Section 4.1 on the homological radius function was not originally written with surface diagrams in mind, however, the content seems to suit this approach quite nicely.

This appendix will follow along very closely with 4.1 using surface diagrams.

Definition A.1. Let X be a 2-dimensional combinatorial complex with $H_1(X) = 0$. For any 1-cycle γ in the 1-skeleton $X^{(1)}$, there exists a surface diagram filling (S, π) for γ. We define the *homological radius* of γ to be

$$\mathrm{HRad}_X(\gamma) \coloneqq \min\Big\{\mathrm{Rad}(S)\,\Big|\,(S,\pi) \text{ a surface diagram filling for } \gamma\Big\},$$

where

$$\mathrm{Rad}(S) = \max_{v \text{ a vertex in } S}\Big\{\min\Big\{|\alpha_v|\,\Big|\,\alpha_v \text{ is a path from } v \text{ to } \partial S \text{ in } S^{(1)}\Big\}\Big\}.$$

Analogously to the homological filling function, we define the *homological radius function* for the space X to be

$$\mathrm{FR}_X(n) = \sup\big\{\mathrm{HRad}_X(\gamma)\,|\,\gamma \text{ is a 1-cycle in } X^{(1)} \text{ with } |\gamma| \le n\big\}.$$

For a group H of type FP_2 with homological finite presentation $\langle A \| R_0\rangle$ and homological Cayley complex $X(A, R_0)$, we denote

$$\mathrm{FR}_{\langle H,A,R_0\rangle}(n) \coloneqq \mathrm{FR}_{X(A,R_0)}(n).$$

Remark. If H is a finitely presented group, then every Van Kampen diagram is a homological filling, so the usual radius function is an upper bound for the homological radius function for any fixed finite presentation.

Remark. With the original homological radius in mind, one may note that since S surjects onto a 2-chain c viewed geometrically in the Cayley 2-complex, the new homological radius is bounded below by the old one.

Lemma 4.2 holds still, by the same proof as before. Also, by a similar argument, we can see that the supremum is attained by a loop as in Section 4.1.

A.1 Independence of Presentation

Proposition A.2. *Up to $\cong$-equivalence, the homological radius function is independent of a homological finite presentation for H.*

Proof. As in [BKS21] (and as we did in Section 4), we will first prove that two homological finite presentations which share a generating set A have $\cong$-equivalent radius functions.

Let $\mathcal{P}_0 = \langle A \| R_0 \rangle, \mathcal{P}_1 = \langle A \| R_1 \rangle$ be homological finite presentations for H. Consider the homological Cayley complex X_0 for $\mathcal{P}_0$. Note that this space shares its 1-skeleton with the homological Cayley complex X_1 for $\mathcal{P}_1$ – both are the Cayley graph for H with respect to generating set A.

Let γ be a loop in $X_0^{(1)}$ with surface diagram filling (S, π) with a radius $\mathrm{HRad}_{X_0}(S)$. The boundary of each 2-cell $\partial\sigma_i$ is a loop labelled by $r_i \in R_0$, so can be filled by a surface diagram in $\mathcal{P}_1$ with radius n_i. Filling all the σ_i in this way, and gluing all these surfaces together corresponding to the gluings in S gives a surface diagram filling for γ with edge labels from $\mathcal{P}_1$, the radius of which is bounded above by

$$\mathrm{HRad}_{X_0}(\gamma) + \max_{r_i \in R_0}\{n_i\},$$

since for any vertex v in this diagram, the distance in the 1-skeleton from v to the boundary $\partial\sigma_i$ is at most $\max_i\{n_i\}$, and the distance from that vertex to the boundary ∂S is at most $\mathrm{HRad}_{X_0}(S)$. Since $\max_i\{n_i\}$ is a constant, we have that

$$\mathrm{FR}_{\langle H,A,R_0\rangle} \cong \mathrm{FR}_{\langle H,A,R_1\rangle}$$

by applying the same argument, starting with X_1 in place of X_0.

This brings us to the general case. Suppose now that $\langle A \| R_0 \rangle$ and $\langle A_1 \| S_0 \rangle$ are two homological finite presentations for H. For each generator b in A_1, there is a word $v(A)$ in $F(A)$ with $b =_H v(A)$. Fix a $v(A)$ for each $b \in A_1$ and let R_1 be the set of relations $\{b = v(A) \mid b \in A_1\}$. As in [BKS21], we construct a new complex X_1 from the homological Cayley complex $X(A, R_0)$ as follows: for each generator $b \in A_1$ and for each vertex $x \in X(A, R_0)$, we attach an edge e_b^x labelled by b with initial vertex x and terminal vertex $x \cdot v(A)$; then for each such edge, we attach a 2-cell D_b^x to the closed path reading $b^{-1}v(A)$ from x. So the complex X_1 is

$$X_1 \coloneqq X(A, R_0) \cup \bigcup_{\substack{x \in H \\ b \in A_1}} e_b^x \cup \bigcup_{\substack{x \in H \\ b \in A_1}} D_b^x$$

The H-action on $X(A, R_0)$ is extended to a free, vertex transitive and cocompact action on this new complex by defining $g \cdot e_b^x = e_b^{g \cdot x}$ and $g \cdot D_b^x = D_b^{g \cdot x}$. [BKS21] also gives us a deformation retraction $r : X_1 \to X(A, R_0)$ which sends edges labelled b to paths labelled $v(A)$ with the same initial and terminal vertices, and collapses the 2-cells D_b^x. The existence of this deformation retraction tells us $H_1(X_1) \cong H_1(X(A, R_0))$ is trivial, so $\langle A \sqcup A_1 \| R_0 \sqcup R_1 \rangle$ is a homological finite presentations for H.

Let w be a 1-cycle in X_1. Every vertex in w lies in $X(A, R_0) \subset X_1$ by construction. For every edge of w labelled by a generator of A_1, we can apply a relator from R_1 to replace this edge with edges labelled by A. Applying no more than $|w|$ such relators gives us a 1-cycle entirely contained within $X(A, R_0)$, we call this 1-cycle w'. Observe that $|w'| \leq k|w|$, where $k = \max\left\{|v(A)| \,\middle|\, b^{-1}v(A) \text{ is a relator in } R_1\right\}$. The 1-cycle w' admits a surface diagram filling S which uses only relators from R_0. Given any vertex in this filling, its distance to the boundary ∂S is at most $\mathrm{Rad}(S)$, and this boundary vertex has distance at most $k/2$ from a vertex in the 1-cycle w. Combining the image of this w'-diagram-filling $\pi(S)$ with the R_1-relators applied earlier gives us a 2-chain with boundary w whose corresponding surface diagram's radius is bounded above by $\mathrm{Rad}(S) + k/2$. It follows that

$$\mathrm{HRad}_{X_1}(w) \leq \mathrm{HRad}_{X(A,R_0)}(w') + k/2 \leq \mathrm{FR}_{X(A,R_0)}(k|w|) + k/2.$$

Since the choice of 1-cycle w was arbitrary, we may conclude that

$$\mathrm{FR}_{X_1}(n) \leq \mathrm{FR}_{X(A,R_0)}(kn) + k/2.$$

For the converse direction, let w be a 1-cycle in $X(A, R_0)$ whose edges are labelled only by letters from A, $|w| \leq n$ and $\mathrm{HRad}_{X(A,R_0)}(w) = \mathrm{FR}_{X(A,R_0)}(n)$. We know that such a 1-cycle realising the supremum in the definition of $\mathrm{FR}_{X(A,R_0)}$ exists by Lemma 4.2. Let (S, π) be a filling of w in X_1. Consider the image of S under the retraction r (by applying r to its image under π and reconstructing a new surface diagram), this is a surface diagram filling for w with labels from (A, R_0). Since w is a 1-cycle whose optimal $X(A, R_0)$-filling has radius $\mathrm{FR}_{X(A,R_0)}(n)$, we must have that the radius $\mathrm{Rad}(r(S)) \geq \mathrm{FR}_{X(A,R_0)}(n)$. The retraction r sends each A_1-edge to a word $v(A)$ of length at most k, it follows that the radius $\mathrm{Rad}(r(S)) \leq k \cdot \mathrm{Rad}(S) \leq k \cdot \mathrm{FR}_{X_1}(n)$, therefore

$$\mathrm{FR}_{X(A,R_0)}(n) \leq k \cdot \mathrm{FR}_{X_1}(n)$$

Bringing the two inequalities together, we obtain that

$$\mathrm{FR}_{X_1} \cong \mathrm{FR}_{X(A,R_0)}.$$

Working analogously to the above, for each $a \in A$, we can find words $u(A_1)$ in (A_1) with $a =_H u(A_1)$, calling the set of such relations S_1. We similarly construct a homological finite presentation $\langle A \sqcup A_1 \| S_0 \sqcup S_1 \rangle$ with homological Cayley complex Y_1, and apply the same argument to see $\mathrm{FR}_{Y_1} \cong \mathrm{FR}_{X(A_1,S_0)}$. Note that the homological finite presentations for X_1 and Y_1 share a generating set, so we can apply the first part of the proof to see

$$\mathrm{FR}_{X(A,R_0)} \cong \mathrm{FR}_{X(A_1,S_0)}.$$

i.e. Up to $\cong$-equivalence, the homological radius function is independent of homological finite presentation. □

Notation. *Considering homological radius functions up to $\cong$-equivalence, the notion of a homological radius function for a group, $FR_H := FR_{X(A,R_0)}$ for any homological finite presentation $\langle A \| R_0 \rangle$ of H, is well-defined.*

With these results in place, we can now extend the notion of an area-radius pair to groups H of type FP_2, using the functions FA_H and FR_H. Since these share the same key properties of the area and radius functions of Gersten and Short, we can use FA_H and FR_H in their place.

Definition A.3. Given a group H of type FP_2 with homological finite presentation $\mathcal{P} = \langle A \| R_0 \rangle$, we call the pair (f, g) of functions $f, g : \mathbb{N} \to \mathbb{R}$ a *homological area-radius pair* if for any 1-cycle γ of length at most n in $X(A, R_0)$, there exists a surface diagram filling (S, π) for γ and $\mathrm{Area}(S) \leq f(n)$ and $\mathrm{Rad}(S) \leq g(n)$.

B A (Homological) Area-Radius Pair for Hyperbolic Groups

In the main body of this report, we cite [GS02] for our Lemma 5.3. In this appendix, we will go through Gersten and Short's proof of this result (their Lemma 2.2), for completeness.

Let us first recall the result:

Lemma 5.3 *Let $\mathcal{P} = \langle A | R \rangle$ be a finite presentation of a hyperbolic group H, then there are constants $B, C > 0$ such that for any nullhomotopic word $w \in \overline{A}^*$ with $\ell(w) \geq 1$, there is a Van Kampen diagram over $\mathcal{P}$ with area at most $B\ell(w)(\log_2(\ell(w)) + 1)$ and radius at most $C(\log_2(\ell(w)) + 1)$.*

Proof. [GS02] Gersten and Short use a different but equivalent definition of hyperbolicity: that for a geodesic triangle Δ, fibres of the *tripod map* have diameter at most δ. Where the tripod map is defined by first mapping $\Delta \to \Delta'$ a triangle in Euclidean space with the same side lengths, then to the tripod graph T whose 3-valent vertex is at the centre of the inscribed circle in Δ'. In this proof, this 'δ-fineness' condition is used very similarly to our δ-slimness definition.

Given a word $w =_H 1$, a disc diagram is constructed as follows. Place $n = |w|$ vertices around S^1 the circle, labelled by the integers 0 to $n-1$, map this into $\Gamma = \Gamma(H, A)$ so that the boundary reads w. The 0th and $\lfloor n/2 \rfloor$th vertices are joined by a straight line, mapped to a geodesic joining the corresponding vertices in Γ.

Join now, for each $j = 2, 3, \ldots, \lfloor \log_2(n) + 1 \rfloor$, $i = 1, 2, \ldots, 2^j$, the vertices labelled $\lfloor (i-1)n/2^j \rfloor$, $\lfloor in/2^j \rfloor$ by geodesics γ_j^i, we call these 'level j edges'. These form many triangles, some in the last layer will have degenerate edges, and we always take the geodesic γ_j^i to be the edge in the loop when it joins two adjacent vertices. We call each triangle with two level j edges and one level $(j-1)$ edge a 'level j triangle'. Compare with Fig. 15. A key property is that, since the level j geodesics form a shorter loop from the basepoint to itself, the total length of level j geodesics γ_j^i is at most n for each j.

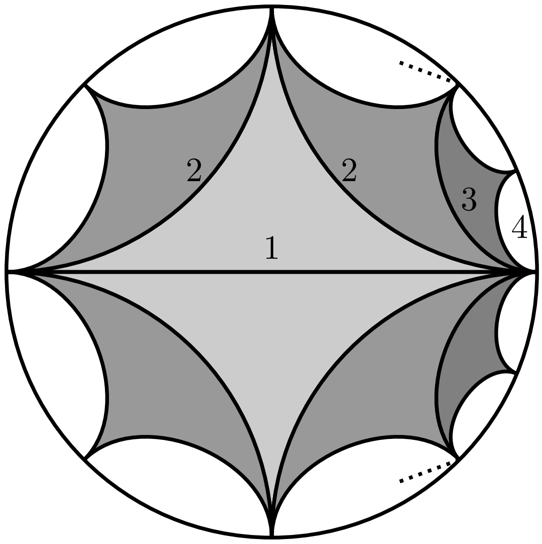

Figure 15: Splitting the loop into triangles

By hyperbolicity, in any level j triangle T, the level $j-1$ edge is always at most δ distance away from one of the level j edges. This means that we can divide up a level j triangle into three triangles, several rectangles and one central polygon as pictured in Fig. 16. In this triangle, every rectangle has boundary length at most $2+2\delta$, each of the corner triangles, which are really degenerate rectangles, have boundary length at most $2+\delta$ and the central polygon has boundary length at most $3+3\delta$. If the central polygon is a triangle, we remove one of its boundary edges that has a vertex in the level $j-1$ edge.

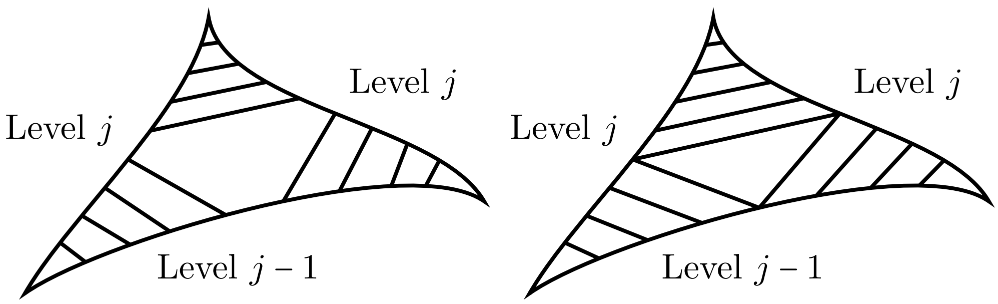

Figure 16: Inside a level j triangle (recreation of [GS02] Fig. 2)

We then finally obtain our disc diagram by gluing disc diagrams for each of these rectangles and central polygons. Let C_1, C_2 be the maximum radius, area among all disc diagram fillings of loops of radius at most $3+3\delta$ respectively. It remains to study this filling and deduce the AR pair.

For the radius, observe that, for any vertex in the disc diagram, there is a path of length at most C_1 to reach a triangle as above, and then one can follow a path along the δ-edges to the boundary. Since each vertex has a δ-edge whose other endpoint is on a higher-level geodesic, this path has length at most $C_1 + \delta\lfloor \log_2(n) + 1 \rfloor \leq C(\log_2(\ell(w)) + 1)$ for a slightly larger constant C. For the area, since the total length of level j edges is always at most n, we know that there are at most n of our δ-polygons at each level. It follows that the area of our disc diagram is at most $nC_2(\lfloor \log_2(n) + 1 \rfloor) \leq B\ell(w)(\log_2(\ell(w)) + 1)$ for a slightly larger constant B. □

Remark. Observe that, in the above, we could replace each instance of 'disc diagram' with 'surface diagram' and the proof would work exactly the same. Thus we were justified in using this result as a homological area-radius pair (using our new surface diagram definition of AR pair).

Bibliography

[AW22] Macarena Arenas and Daniel T. Wise. Linear isoperimetric functions for surfaces in hyperbolic groups, 2022.

[BB97] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Inventiones mathematicae*, 129(3):445–470, Aug 1997.

[BB00] N. Brady and M. R. Bridson. There is only one gap in the isoperimetric spectrum. *Geometric and Functional Analysis*, 10(5):1053–1070, Dec 2000.

[BH99] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer Berlin, Heidelberg, 1999.

[BKS21] Noel Brady, Robert Kropholler, and Ignat Soroko. Homological Dehn functions of groups of type FP_2, 2021.

[Bow95] B. H. Bowditch. A short proof that a subquadratic isoperimetric inequality implies a linear one. *Michigan Mathematical Journal*, 42(1):103 – 107, 1995.

[Bra99] Noel Brady. Branched coverings of cubical complexes and subgroups of hyperbolic groups. *Journal of the London Mathematical Society*, 60(2):461–480, 1999.

[Bri02] Martin R. Bridson. *The geometry of the word problem*, volume 7, chapter 2, pages 29–91. Oxford Graduate Texts in Mathematics, 2002.

[Ger96a] S. M. Gersten. Subgroups of word hyperbolic groups in dimension 2. *Journal of the London Mathematical Society*, 54(2):261–283, 1996.

[Ger96b] S.M. Gersten. Preservation and distortion of area in finitely presented groups. *Geometric and functional analysis*, 6(2):301–345, 1996.

[GR02] Steve M. Gersten and Tim R. Riley. Filling radii of finitely presented groups. *The Quarterly Journal of Mathematics*, 53(1):31–45, 03 2002.

[Gro87] M. Gromov. *Hyperbolic Groups*, pages 75–263. Springer New York, New York, NY, 1987.

[GS02] Steve Gersten and Hamish Short. Some isoperimetric inequalities for kernels of free extensions. *Geometriae Dedicata*, 92(1):63–72, Jul 2002.

[Ise24] Claudio Llosa Isenrich. From the second BNSR invariant to Dehn functions of coabelian subgroups, 2024.

[KIS25] Robert Kropholler, Claudio Llosa Isenrich, and Ignat Soroko. Explicit polynomial bounds on dehn functions of subgroups of hyperbolic groups, 2025.

[KV21] Robert Kropholler and Federico Vigolo. Hyperbolic groups with almost finitely presented subgroups, 2021.

[Rip82] E. Rips. Subgroups of small cancellation groups. *Bulletin of the London Mathematical Society*, 14(1):45–47, 01 1982.