



COMPLETE CLASSIFICATION OF THE DEHN FUNCTIONS OF BESTVINA–BRADY GROUPS

YU-CHAN CHANG, JERÓNIMO GARCÍA-MEJÍA, AND MATTEO MIGLIORINI

Abstract. We prove that the Dehn function of every finitely presented Bestvina–Brady group grows as a linear, quadratic, cubic, or quartic polynomial. In fact, we provide explicit criteria on the defining graph to determine the degree of this polynomial. As a consequence, we identify an obstruction that prevents certain Bestvina–Brady groups from admitting a CAT(0) structure.

1 Introduction

Bestvina and Brady introduced a class of subgroups of right-angled Artin groups, now known as *Bestvina–Brady groups*, that exhibit exotic finiteness properties [BB97]. For every finite simplicial graph Γ , they defined a group BB_Γ whose finiteness properties are determined by the homotopy type of the flag complex $\Delta(\Gamma)$ associated to Γ : they proved that BB_Γ is of (finiteness) type \mathcal{F}_n if and only if $\Delta(\Gamma)$ is $(n - 1)$ -connected. Recall that types \mathcal{F}_1 and \mathcal{F}_2 correspond to the more familiar notions of being finitely generated and finitely presented, respectively.

The family of Bestvina–Brady groups includes the Stallings–Bieri groups SB_n , which were the first known examples of groups of type \mathcal{F}_{n-1} but not of type \mathcal{F}_n . Stallings first constructed SB_3 [Sta63], and Bieri later generalised this to SB_n for $n > 3$ [Bie81]. Each SB_n can be realised as BB_Γ , where Γ is the join of n copies of the graph consisting of two non-adjacent vertices.

In this work, we study the Dehn functions of finitely presented Bestvina–Brady groups. The *Dehn function* of a finitely presented group is an important quasi-isometry invariant, which bounds the number of relations that must be applied to reduce a word representing the identity to the trivial word; see Sect. 2.5 for a precise definition. In this sense, the Dehn function can be interpreted as a quantitative version of finite presentability. Geometrically, the Dehn function of a finitely presented group G provides an upper bound on the area required to fill a loop in the universal cover of a finite 2-complex with fundamental group G . Notably, a group is hyperbolic if and only if its Dehn function is linear [Gro87], whereas any finitely presented non-hyperbolic group has at least quadratic Dehn function [Gro87, Bow95].

Keywords and phrases: Dehn functions, Filling invariants, Bestvina–Brady groups, Subgroups of right-angled Artin groups, Asymptotic cones

Mathematics Subject Classification: Primary 20F69, Secondary 20F05, 20F38, 20F65, 51F30

Dison proved that the Dehn function of a finitely presented Bestvina–Brady group is bounded above by a quartic polynomial [Dis08]; see [A+13] for a purely geometric proof. This naturally raises the question:

QUESTION 1. *Is the Dehn function of every finitely presented Bestvina–Brady group equivalent to a polynomial?*

Brady presented examples of quadratic, cubic, and quartic Dehn functions [BRS07]. Subsequently, Dison, Elder, Riley, and Young showed that Stallings’ group SB_3 has quadratic Dehn function [D+09]. Later, Carter and Forester, using geometric arguments, generalised this result by showing that SB_n has quadratic Dehn function for every $n \geq 3$ [CF17]. In fact, more generally, they showed that if Γ is a join of three nonempty graphs, then the Dehn function of BB_Γ is quadratic. More recently, Chang proved that the Dehn functions of certain Bestvina–Brady groups can be determined directly from the defining graphs [Cha21].

In the present work, we complete the picture giving a positive answer to Question 1, by showing that the Dehn function of every finitely presented Bestvina–Brady group is equivalent to a polynomial of degree $d \in \{1, 2, 3, 4\}$. Furthermore, we provide conditions on the defining graph of a finitely presented Bestvina–Brady group that determine the precise degree d ; see Theorem 1. In this sense, we obtain an effective classification of finitely presented Bestvina–Brady groups in terms of their Dehn functions.

Besides providing an answer to Question 1, our main result contributes to the understanding of the Dehn functions of subgroups of right-angled Artin groups; see [Bri13, BS19] for more examples other than Bestvina–Brady groups.

1.1 Main result. It is known that BB_Γ is hyperbolic if and only if Γ is a tree, as can be seen directly by looking at the presentation given in [DL99]. This characterises all the graphs Γ such that BB_Γ has linear Dehn function. However, providing a general criterion that distinguishes Bestvina–Brady groups with quadratic, cubic, and quartic Dehn functions is a much more challenging task.

Recall that a graph Γ is called *reducible* if it is a join $\Gamma_1 * \Gamma_2$, where Γ_1 and Γ_2 are nonempty subgraphs;¹ otherwise, it is called *irreducible*. Our main result gives a criterion on the defining graph Γ that determines the precise Dehn function of BB_Γ in terms of the *maximal reducible subgraphs* of Γ , that is, reducible subgraphs that are maximal with respect to inclusion among all reducible subgraphs of Γ . To state this full characterisation, we introduce the following definitions.

DEFINITION 1 (Essentially 2-reducible graph). A graph is called *essentially 2-reducible* if it is a join of two irreducible subgraphs, each with at least two vertices; equivalently, if it is neither a cone graph nor a join of three or more subgraphs.

¹ All subgraphs we consider are assumed to be *induced* subgraphs.

DEFINITION 2. Let Γ be a finite simplicial graph whose associated flag complex $\Delta(\Gamma)$ is simply connected. If this is the case, we say that Γ has property (\mathcal{D}_1) , and additionally, it has property

- (\mathcal{D}_2) if Γ is not a tree;
- (\mathcal{D}_3) if Γ contains a maximal reducible subgraph that is essentially 2-reducible;
- (\mathcal{D}_4) if Γ contains a maximal reducible subgraph whose associated flag complex is not simply connected.

We are now ready to state our main result.

Theorem 1. *Let Γ be a finite simplicial graph such that the associated flag complex $\Delta(\Gamma)$ is simply connected. Let $d(\Gamma)$ be the maximal $\alpha \in \{1, 2, 3, 4\}$ such that Γ has property (\mathcal{D}_α) . Then the Dehn function of the Bestvina–Brady group BB_Γ satisfies*

$$\delta_{BB_\Gamma}(n) \asymp n^{d(\Gamma)}.$$

Therefore, property (\mathcal{D}_α) is the obstruction to BB_Γ having Dehn function strictly less than a polynomial of degree α . This can be seen by combining Theorem 1 with the fact that property (\mathcal{D}_β) implies property (\mathcal{D}_α) for $\beta > \alpha$ on every graph Γ with $\Delta(\Gamma)$ simply connected. To see this, note that if Γ contains a reducible subgraph $\Lambda = \Lambda' * \Lambda''$ such that $\Delta(\Lambda)$ is not simply connected, then both Λ and Λ' are disconnected, and therefore irreducible with at least two vertices. This implies that Λ is essentially 2-reducible, so $(\mathcal{D}_4) \Rightarrow (\mathcal{D}_3)$. It is also clear that if Γ contains an essentially 2-reducible subgraph, then it contains a square and cannot be a tree, so $(\mathcal{D}_3) \Rightarrow (\mathcal{D}_2)$. Finally, (\mathcal{D}_1) is satisfied by definition.

If the defining graph Γ is reducible, then it coincides with its unique maximal reducible subgraph, meaning that Γ cannot have (\mathcal{D}_4) . Thus, Theorem 1 implies the following result.

COROLLARY 1. *If Γ is an essentially 2-reducible graph such that $\Delta(\Gamma)$ is simply connected, then BB_Γ has cubic Dehn function.*

If Γ is reducible but not essentially 2-reducible, then the Dehn function of BB_Γ is at most quadratic. While this is a simple application of Theorem 1, it also follows from previously known results. Indeed, a reducible graph Γ that is not essentially 2-reducible is either a cone graph, namely, $\Gamma = \{c\} * \Lambda$ for some subgraph Λ , or it splits as a join of three or more subgraphs. In the first case, the group BB_Γ is isomorphic to the right-angled Artin group A_Λ , which is a CAT(0) group² [Cha07, Theorem 2.6], and therefore has at most quadratic Dehn function [Bri02, Theorem 6.2.1]. In the second case, the Dehn function of BB_Γ being quadratic is a consequence of [CF17, Corollary 4.3]. Note that although we draw some inspiration from the work of Carter and Forester, our argument does not rely on their results.

² A CAT(0) group is a group admitting a proper and cocompact action on a CAT(0) space by isometries.

Theorem 1 is also consistent with [Cha21, Theorem 1.1]. Indeed, let $\Delta(\Gamma)$ be a 2-dimensional triangulated disc with square boundary. On the one hand, if Γ is reducible, then it is the suspension of a path. In this case, according to both Theorem 1 and [Cha21, Theorem 1.1], the Dehn function is cubic if the path has length greater than or equal to 3 and quadratic otherwise. On the other hand, if Γ is irreducible, then it has (\mathcal{D}_4) , as every reducible subgraph Λ containing the boundary of $\Delta(\Gamma)$ must be strictly contained in Γ . Therefore, the boundary of $\Delta(\Gamma)$ is not null-homotopic in $\Delta(\Lambda)$. Furthermore, the flag complex $\Delta(\Gamma)$ must contain an interior 2-simplex (otherwise, one gets a contradiction by [Cha21, Lemma 4.2]). Thus, both statements imply that BB_Γ has quartic Dehn function. However, note that the lower bound established in [Cha21, Proposition 1.5] is incorrect.³

1.2 Applications of the main result. Showing that the Dehn function of a group is strictly greater than quadratic yields geometric constraints on the group. In particular, since $\text{CAT}(0)$ groups are finitely presented and have at most quadratic Dehn function, condition (\mathcal{D}_3) provides an obstruction to a Bestvina–Brady group acting properly and cocompactly on a $\text{CAT}(0)$ space.

COROLLARY 2. *If a finite simplicial graph Γ has (\mathcal{D}_3) , then BB_Γ is not a $\text{CAT}(0)$ group.*

In general, not every finitely presented Bestvina–Brady group with quadratic Dehn function is a $\text{CAT}(0)$ group. Indeed, Stallings–Bieri groups have quadratic Dehn function, but they are not $\text{CAT}(0)$ since they are not of type \mathcal{F}_∞ [BB97]. This raises the following question:

QUESTION 2. *Does there exist a Bestvina–Brady group of type \mathcal{F}_∞ with quadratic Dehn function that is not a $\text{CAT}(0)$ group?*

To the best of our knowledge, no such examples are known in the literature. One possible approach to finding such examples would be looking at higher filling functions. Informally, the k -dimensional Dehn function of a group measures the difficulty of filling k -spheres with $(k+1)$ -balls in some suitable space on which the group acts; the 1-dimensional Dehn function is the usual Dehn function. We refer the reader to [A+13] for a precise definition of k -dimensional Dehn functions. For a $\text{CAT}(0)$ group, the k -dimensional Dehn function is at most $n^{(k+1)/k}$ [Gro83, Wen08]. In [A+13], the authors showed that for a Bestvina–Brady group of type \mathcal{F}_{k+1} , the k -dimensional Dehn function is bounded above by $n^{2(k+1)/k}$.

Finally, combining the result of [Pap96] with Theorem 1 yields the following result about the asymptotic cones of Bestvina–Brady groups.

³ For example, let $\Gamma = \{c\} * \Lambda$, where Λ is the suspension of the path of length three. Then the Dehn function of BB_Λ is cubic, whereas the Dehn function of BB_Γ is quadratic since Γ is a cone graph. The gap in the proof of [Cha21, Proposition 1.5] is due to [Cha21, Proposition 5.3], which is erroneous.




Figure 1: Brady’s examples with cubic (left) and quartic (centre) Dehn functions. On the right, a variation of Brady’s example whose Dehn function is quadratic.

COROLLARY 3. *Let Γ be a finite simplicial graph such that $\Delta(\Gamma)$ is simply connected. If Γ does not have (\mathcal{D}_3) , then the asymptotic cones of BB_Γ are simply connected.*

It would be interesting to know if there exist finitely presented Bestvina–Brady groups with non-quadratic Dehn functions that have simply connected asymptotic cones.

1.3 Examples. We present some examples to illustrate how to concretely apply Theorem 1 to determine the Dehn function of a Bestvina–Brady group. Consider the graphs Γ_1 , Γ_2 , and Γ_3 shown in Fig. 1. The graphs Γ_1 and Γ_2 were considered by Brady in [BRS07], where he constructed Bestvina–Brady groups with cubic and quartic Dehn functions, respectively. Theorem 1 recovers these cases. Although the graph Γ_3 is a slight variation of Γ_1 , the Dehn function of the corresponding Bestvina–Brady group was, to the best of our knowledge, not previously known. In particular, it was not addressed in [CF17] or [Cha21].

The graph Γ_1 . The graph Γ_1 is the suspension of the path of length three induced by the vertices C , D , E , and F . Therefore, it is essentially 2-reducible, so Corollary 1 implies that BB_{Γ_1} has cubic Dehn function.

The graph Γ_2 . In this case, the graph Γ_2 has (\mathcal{D}_4) . To see this, let Λ be the join of $\{C, F\}$ and the subgraph induced by the vertices G , A , B , and H . Then Λ is a maximal reducible subgraph of Γ_2 , and $\Delta(\Lambda)$ is not simply connected. Therefore, the graph Γ_2 has (\mathcal{D}_4) , and Theorem 1 implies that the Dehn function of BB_{Γ_2} is quartic.

The graph Γ_3 . Clearly, the graph Γ_3 has (\mathcal{D}_2) . Nevertheless, it does not have (\mathcal{D}_3) , since none of the maximal reducible subgraphs is essentially 2-reducible. In fact, all the maximal reducible subgraphs are cone graphs. The graph Γ_3 does contain a unique essentially 2-reducible subgraph, namely, the square induced by the vertices A , B , C , and E , but it is not maximal. Therefore, by Theorem 1, the Dehn function of BB_{Γ_3} is quadratic.

1.4 Strategy of the proof. In Theorem 1, the linear case corresponds to BB_Γ being hyperbolic, while the quadratic lower bound and the quartic upper bound

follow respectively from BB_Γ not being hyperbolic and from [Dis08, Theorem 1.1]. Therefore, to obtain Theorem 1, it suffices to show the following statement, which involves only (\mathcal{D}_3) and (\mathcal{D}_4) .

Theorem 2. *Let Γ be a finite simplicial graph such that $\Delta(\Gamma)$ is simply connected, and let $\alpha \in \{3, 4\}$. If Γ has property (\mathcal{D}_α) , then $\delta_{BB_\Gamma}(n) \succ n^\alpha$. Otherwise, $\delta_{BB_\Gamma}(n) \preccurlyeq n^{\alpha-1}$.*

The proof of Theorem 2 is divided into two parts: the quadratic and cubic upper bounds are given by Theorem 5.1, and the cubic and quartic lower bounds are given in Theorem 6.1.

From Bestvina and Brady’s construction [BB97], we can naturally view BB_Γ embedded in A_Γ as the 0-level set of the height function $\varphi_\Gamma: A_\Gamma \rightarrow \mathbb{Z}$, which sends every generator to 1. This geometric perspective allows us, via a standard argument, to compute the Dehn function by considering null-homotopic words and van Kampen diagrams in the right-angled Artin group that have heights close to 0. We formalise this by defining *alternating words* and *almost-flat van Kampen diagrams*. This is particularly useful for establishing the lower bound: we show that the height of a vertex in a van Kampen diagram can be computed by looking at *corridors* and *annuli*, and the restriction on the height yields a lower bound on the number of intersections between annuli, and hence, on the area of the van Kampen diagram.

The proof of the upper bound employs a refined version of Dison’s pushdown argument for the general quartic upper bound [Dis08]. Dison’s argument does not yield an optimal upper bound: it is possible for a null-homotopic word to have support contained in a cone subgraph, in which case it has at most quadratic area, although the “pushdown algorithm” may not recognise this and produce an inefficient filling.

We address this issue by introducing the notion of *k-coloured words*. A *k-coloured word* should be thought of as a word w in the right-angled Artin group A_Γ with the usual generating set given by the vertices of Γ , together with a decomposition $w = w_1 \cdots w_k$ and a sequence of generators a_1, \dots, a_k of A_Γ , such that every letter of w_i commutes with a_i . We think of a_i as colouring the subword w_i . We are particularly interested in the case where the coloured word has minimal k among all coloured words representing a fixed group element in BB_Γ .

The colouring information allows us to push coloured words down to the 0-level set and to exploit the fact that some subwords have support contained in cone subgraphs. We show that this technique produces the optimal upper bounds for the Dehn function of BB_Γ .

Throughout the proof of the upper bound, the case where Γ is irreducible is the most challenging. Indeed, if Γ is reducible, then any group element can be represented by a 2-coloured word. This is a useful property that can be used to produce concrete triangular diagrams, as demonstrated in [CF17]. However, for a general graph Γ , there is no uniform bound on the minimal number of colours required to represent an element. We overcome this by using corridors in a van Kampen diagram over A_Γ to produce a suitable subdivision of triangular diagrams into regions that have a uniformly bounded number of colours, while keeping a good control over the area.

1.5 Structure of the paper. Sect. 2 includes basic definitions and properties used throughout the paper. In Sect. 3, we define *alternating words* and *almost-flat van Kampen diagrams*, and show how they can be used to compute the Dehn function of Bestvina–Brady groups. In Sect. 4, we introduce *coloured words* and show how the colouring allows us to “push down” coloured words to the 0-level set in a canonical way to obtain alternating words. The entire Sect. 5 is devoted to the proof of Theorem 5.1, which establishes the upper bounds in Theorem 2. We establish the lower bounds in Theorem 2 by proving Theorem 6.1 in Sect. 6.

2 Preliminaries

In this section, we introduce basic definitions, present well-known results, and establish general terminology and notation used throughout.

2.1 Graphs. A *simplicial graph* is a graph without loops or multiple edges. Given a finite simplicial graph Γ , its vertex set and edge set are denoted by V_Γ and E_Γ , respectively. Two vertices u and v are *adjacent* if they are connected by an edge, in which case we write $\{u, v\} \in E_\Gamma$. An (*induced*) *subgraph* $\Lambda \subseteq \Gamma$ is a graph with $V_\Lambda \subseteq V_\Gamma$ and whose edge set E_Λ consists of all edges of Γ that connect vertices in V_Λ . Recall that all subgraphs considered in the present work are assumed to be induced subgraphs.

The flag complex associated to Γ , denoted by $\Delta(\Gamma)$, is the simplicial complex whose 1-skeleton is Γ , and in which every complete subgraph of Γ spans a simplex.

Let $u, v \in V_\Gamma$. A (*combinatorial*) *path* between u and v of length $k \in \mathbb{N}$ is a map $\gamma: \{0, \dots, k\} \rightarrow V_\Gamma$ such that $\gamma(i-1)$ is adjacent to $\gamma(i)$ for all $i \in \{1, \dots, k\}$. The *distance* between u and v in Γ is the length of a shortest path between them.

The *join* of two graphs Γ_1 and Γ_2 , denoted by $\Gamma_1 * \Gamma_2$, is the graph whose vertex set is $V_{\Gamma_1} \cup V_{\Gamma_2}$ and whose edge set is $E_{\Gamma_1} \cup E_{\Gamma_2} \cup \{\{u, v\} \mid u \in V_{\Gamma_1}, v \in V_{\Gamma_2}\}$. A graph is said to be *reducible* if it decomposes as a join of two nonempty subgraphs, and *irreducible* otherwise.

The *link* of a vertex $v \in V_\Gamma$, denoted by $\text{Lk}(v)$, is the subgraph induced by the vertices adjacent to v . The *star* of v , denoted by $\text{St}(v)$, is the subgraph $\{v\} * \text{Lk}(v)$.

The *complement* of Γ is the graph Γ^c with the same vertex set as Γ , such that two vertices are adjacent in Γ^c if and only if they are not adjacent in Γ . Note that Γ is irreducible if and only if its complement is connected. More generally, the graph Γ admits a unique decomposition (up to the order) $\Gamma = \Gamma_1 * \dots * \Gamma_k$, where $\Gamma_1, \dots, \Gamma_k$ are irreducible subgraphs of Γ ; these can be characterised as the connected components of the complement of Γ .

Let Λ be a reducible subgraph of Γ . We say that Λ is *maximal* if it is maximal with respect to inclusion among all reducible subgraphs of Γ . We say that Λ is *essentially 2-reducible* if it is a join $\Lambda_1 * \Lambda_2$ of exactly two irreducible subgraphs, each of which has at least two vertices. Equivalently, an essentially 2-reducible graph is a graph

whose complement has exactly two connected components, neither of which is a single vertex.

2.2 Words and free groups. For a set S , we denote by S^{-1} the set of formal inverses of the elements in S . By a *word w in S* we mean an element of the free monoid $(S \cup S^{-1})^*$ generated by $S \cup S^{-1}$; that is, w can be written as $s_1 \cdots s_k$ for some $s_i \in S \cup S^{-1}$. In this case, the elements s_1, \dots, s_k are called the *letters of w* , and k is called the *length* of w , denoted by $|w|$. A *subword of w* is a word of the form $s_i s_{i+1} \cdots s_j$ for some $1 \leq i \leq j \leq k$. A word in S is *freely reduced* if it does not contain subwords of the form ss^{-1} or $s^{-1}s$.

We denote the *free group on a set S* by $F(S)$. There is a natural epimorphism of monoids $(S \cup S^{-1})^* \rightarrow F(S)$ that sends a word w to its equivalence class under *free insertions and reductions*, that is, up to inserting or removing subwords of the form ss^{-1} or $s^{-1}s$ in w . Two words w and w' in S are said to be *freely equivalent*, written $w =_F w'$, if their images in $F(S)$ coincide. We write $w \equiv w'$ when words are actually identical “letter by letter”.

2.3 Groups and presentations. An (abstract) generating set for a group G is a set S equipped with a map $\iota: S \rightarrow G$ such that the induced map $F(S) \rightarrow G$ is an epimorphism. Note that we do not require the map ι to be injective. Let R be a set of words in S . We denote by $\langle S \mid R \rangle$ the group defined as the quotient of $F(S)$ by the normal closure of the image of R in $F(S)$. Given the canonical map $S \rightarrow \langle S \mid R \rangle$, the set S is naturally a generating set for $\langle S \mid R \rangle$. We say that $\langle S \mid R \rangle$ is *generated by S with relations R* .

If G is a group together with an isomorphism $G \cong \langle S \mid R \rangle$, we say that $\langle S \mid R \rangle$ is a *presentation for G* . When S and R are finite, we say that the *presentation is finite* or that G is *finitely presented*.

Given a group G with presentation $\langle S \mid R \rangle$, we say that two words w and w' in S *represent the same element in G* , written $w =_G w'$, if their images under the canonical surjective map $(S \cup S^{-1})^* \rightarrow G$ coincide.

2.4 Right-angled Artin groups and Bestvina–Brady groups. Let Γ be a finite simplicial graph, and let S_Γ be a set in one-to-one correspondence with the vertex set V_Γ ; we denote the latter as $\{v_s : s \in S_\Gamma\}$. The *right-angled Artin group A_Γ* associated to Γ is the group with finite presentation⁴

$$\mathcal{P}_\Gamma := \langle S_\Gamma \mid [s, t] = 1 \text{ for all } \{v_s, v_t\} \in E_\Gamma \rangle.$$

The set S_Γ and the presentation \mathcal{P}_Γ are called the *standard generating set* and the *standard presentation* of A_Γ , respectively.

⁴ We use the convention $[x, y] = xyx^{-1}y^{-1}$.

If Λ is a subgraph of Γ , then A_Λ is naturally a subgroup of A_Γ . We say that A_Γ is *irreducible* if there are no nontrivial subgraphs $\Lambda_1, \Lambda_2 \subseteq \Gamma$ such that A_Γ splits as $A_{\Lambda_1} \times A_{\Lambda_2}$; this is equivalent to saying that the graph Γ is irreducible.

Recall that associated to a right-angled Artin group A_Γ , there is a natural epimorphism $\varphi_\Gamma: A_\Gamma \rightarrow \mathbb{Z}$, called the *height function*, defined by sending every generator $s \in S_\Gamma$ to 1. The *Bestvina–Brady group* BB_Γ associated to Γ is the kernel of the height function.

Bestvina and Brady [BB97] showed that BB_Γ is finitely presented if and only if the flag complex $\Delta(\Gamma)$ is simply connected. In this case, Dicks and Leary [DL99] provided an algebraic proof of this result by exhibiting a finite presentation, which was further simplified by Papadima and Suciu [PS07]. We now describe their presentation.

Fix a total ordering \ll on the vertex set V_Γ , orient the edges of Γ increasingly, and choose a spanning tree T of Γ . The generating set for BB_Γ is the set of edges of T , denoted by S_T .

Let v_s and v_t be two vertices of Γ . The unique simple path in T from v_s to v_t defines a word $w_{s,t} := e_1^{\epsilon_1} \cdots e_k^{\epsilon_k}$ in S_T , where e_i is the i th edge in the path, and $\epsilon_i = 1$ if the traversed direction agrees with the orientation of e_i , and $\epsilon_i = -1$ otherwise. If $e = (v_s, v_t)$ is an oriented edge, we also write $w_e := w_{s,t}$.

A *directed triangle* is a triple of oriented edges (e, f, g) in Γ with $e = (v_{s_1}, v_{s_2})$, $f = (v_{s_2}, v_{s_3})$, and $g = (v_{s_1}, v_{s_3})$ for some vertices $v_{s_1} \ll v_{s_2} \ll v_{s_3}$. The presentation for BB_Γ provided by Papadima and Suciu can be expressed as

$$\mathcal{Q}_\Gamma = \langle S_T \mid [w_e, w_f] = 1 \text{ for all directed triangles } (e, f, g) \rangle.$$

2.5 Dehn functions. Let $\mathcal{P} := \langle S \mid R \rangle$ be a finite presentation for a group G . We say that a word w in S is *null-homotopic* if it represents the identity in G . The *area with respect to \mathcal{P}* of a null-homotopic word w in S is defined as

$$\text{Area}_{\mathcal{P}}(w) := \min \left\{ \ell \in \mathbb{N} : w =_F \prod_{i=1}^{\ell} x_i r_i^{\epsilon_i} x_i^{-1}, \text{ where } x_i \in (S \cup S^{-1})^*, r_i \in R, \text{ and } \epsilon_i = \pm 1 \right\}.$$

The *Dehn function of a finite presentation \mathcal{P}* is the function $\delta_{\mathcal{P}}: \mathbb{N} \rightarrow \mathbb{N}$ defined as

$$\delta_{\mathcal{P}}(n) := \max \{ \text{Area}_{\mathcal{P}}(w) : w \text{ null-homotopic word in } S \text{ such that } |w| \leq n \}.$$

Changing finite presentations of G does not alter the asymptotic behaviour of the Dehn function of the presentations. More precisely, the *Dehn function* δ_G of G is well-defined up to \asymp -equivalence of functions: two functions $f, g: \mathbb{N} \rightarrow \mathbb{N}$ are \asymp -equivalent, and we write $f \asymp g$, if $f \preceq g$ and $g \preceq f$, where $f \preceq g$ means that there exist constants $A, B > 0$ and $C, D, E \geq 0$ such that $f(n) \leq Ag(Bn + C) + Dn + E$ for all $n \geq 0$. If \mathcal{P}_1 and \mathcal{P}_2 are two finite presentations for G , then $\delta_{\mathcal{P}_1} \asymp \delta_{\mathcal{P}_2}$. Thus, the Dehn function of G is defined to be the \asymp -equivalence class $\delta_{\mathcal{P}}$ for some (hence any) finite presentation \mathcal{P} of G .

2.6 Van Kampen diagrams. The study of van Kampen diagrams over a finite group presentation provides a powerful geometric tool for computing the areas of null-homotopic words, and thus for establishing bounds on the Dehn function.

Let $\mathcal{P} = \langle S \mid R \rangle$ be a finite presentation; we may always assume that the relation set R is closed under taking cyclic conjugates and inverses.

Given a null-homotopic word w in \mathcal{P} , a *van Kampen diagram* for w is a connected, pointed, oriented, and labelled planar graph D such that each edge is labelled by an element of S , and the boundary word of each bounded region in $\mathbb{R}^2 \setminus D$ is freely equivalent to a word in R . The word is read either clockwise or counter-clockwise, from an arbitrary starting vertex in the boundary of the region. The label of each edge traversed this way has a $+1$ exponent (respectively, -1 exponent) if the reading direction coincides with (respectively, is opposite to) the orientation of the edge. Changing the starting point or direction alters the word by cyclic conjugation or inversion. The boundary word of D , when read from the base point, is labelled by w . The diagram D can also be given the structure of a 2-complex, where the graph is viewed as the 1-skeleton, and each bounded region corresponds to a 2-cell attached via the word labelling its boundary. The *area of a van Kampen diagram* D , denoted by $\text{Area}(D)$, is the number of bounded regions in $\mathbb{R}^2 \setminus D$, or equivalently, the number of 2-cells.

By van Kampen's Lemma [Bri02, Theorem 4.2.2], the area of a null-homotopic word w in a finite presentation \mathcal{P} satisfies

$$\text{Area}_{\mathcal{P}}(w) = \min\{\text{Area}(D) \mid D \text{ is a van Kampen diagram for } w \text{ over } \mathcal{P}\}.$$

A van Kampen diagram attaining this minimum is said to be a *minimal-area* van Kampen diagram.

2.7 Coarse diagrams. In the definition of van Kampen diagrams, we may weaken the condition on the boundary words of bounded regions by allowing them to be arbitrary null-homotopic words instead of relations. This yields a broader class of diagrams that are helpful for our work.

DEFINITION 2.1 (Coarse diagram). Let S a finite generating set for a group G . A *coarse diagram over S* for a null-homotopic word w in S is a van Kampen diagram for w , except that the bounded regions may be labelled by arbitrary null-homotopic words.

To be more precise, a coarse diagram can be defined as a van Kampen diagram over an *infinite* presentation, where the set of relations includes all null-homotopic words in a set of generators. Although coarse diagrams cannot be used directly to compute the Dehn functions of finitely presented groups, they can still be used to construct van Kampen diagrams over finite presentations in a “patchwork” manner, as follows.

Let $\mathcal{P} = \langle S \mid R \rangle$ be a finite presentation, and let D be a coarse diagram for a null-homotopic word w in S . By definition, each bounded region R_i of D is labelled

by a null-homotopic word w_i in S . In particular, for each w_i , there is a van Kampen diagram D_i over \mathcal{P} . Replacing each bounded region R_i in D with the corresponding van Kampen diagram D_i yields a van Kampen diagram D' for w over the finite presentation \mathcal{P} . In Sect. 3.3, we explain this “two-step filling” technique in greater detail for the case of Bestvina–Brady groups, which is the focus of this work.

2.8 Corridors. To estimate the area of van Kampen diagrams, it is often useful to study their geometric structure by means of *corridors* and *annuli*. We are interested in the corridors and annuli in the van Kampen diagrams over the standard presentation \mathcal{P}_Γ for A_Γ . For this reason, we give their definition only for this particular case.

Let D be a van Kampen diagram over \mathcal{P}_Γ for a null-homotopic word w in S_Γ , and let $a \in S_\Gamma$. An edge in D is called an *a-edge* if it is labelled by a . Let D^* be the graph dual to the 1-skeleton of D , together with a vertex v_∞ dual to the unbounded region of $D \subset \mathbb{R}^2$. Let λ be a loop in D^* whose edges are all dual to a -edges. The subdiagram C of D , consisting of all the closed 1-cells and 2-cells of D dual to $\lambda \setminus \{v_\infty\}$, is called an *a-corridor* if λ includes v_∞ , and an *a-annulus* if it does not. Moreover, if C is an *a-corridor*, then the path obtained from λ by removing its intersection with the interior of the unbounded region is called the *core of C*. We say that two corridors *cross* if they intersect in one or more 2-cells; equivalently, if their cores intersect.

The standard orientation of \mathbb{R}^2 induces an orientation on the core of a corridor C , so that a -edges intersect the core transversally from left to right. The paths consisting of 1-cells in C that are parallel to the core are called the *sides* of C . The *length* of a corridor is the number of its 2-cells, and the *label* of a corridor is the word read along either of its sides, following the orientation.

For $a, b \in S_\Gamma$, if an a -corridor and a b -corridor cross each other, then $[a, b] = 1$ must be a relation in \mathcal{P}_Γ . This tells us that the corresponding vertices $v_a, v_b \in \Gamma$ are adjacent. In particular, we have $a \neq b$, since, while a commutes with itself, the commutator $[a, a]$ is not a relation of \mathcal{P}_Γ . Thus, a corridor cannot cross itself. Moreover, we have the following standard result regarding minimal-area van Kampen diagrams in right-angled Artin groups.

LEMMA 2.2. *Let w be a null-homotopic word in the standard generating set S_Γ of A_Γ . If D is a minimal-area van Kampen diagram for w over \mathcal{P}_Γ , then it contains no annuli, and if two corridors intersect, then they do so at most once.*

Proof. If D contained an annulus, then removing it would yield a diagram of smaller area, a contradiction. Suppose instead that two corridors cross more than once, and call the region between them a bigon. By performing surgery on an innermost bigon, we can remove the double intersection and obtain a van Kampen diagram with two fewer 2-cells, again, a contradiction. \square

3 Alternating words and diagrams in Bestvina–Brady groups

Throughout the rest of the work, we denote by Γ a finite simplicial graph whose associated flag complex $\Delta(\Gamma)$ is simply connected.

As mentioned before, the group BB_Γ is naturally embedded in A_Γ as the 0-level set of the height function $\varphi_\Gamma: A_\Gamma \rightarrow \mathbb{Z}$. This perspective allows us to exploit our understanding of the ambient group A_Γ to study BB_Γ .

To this end, instead of working with the abstract presentation of BB_Γ , we use the presentation for A_Γ , but consider only words in S_Γ that are close to the 0-level set; we call these *alternating words* (see Definition 3.1). We also consider the van Kampen diagrams that are close to the 0-level set, leading to the definition of *almost-flat van Kampen diagrams* (see Definition 3.6). Almost-flat van Kampen diagrams are central in our estimates for the Dehn function of BB_Γ .

3.1 Alternating words and diagrams. We start by defining the set of alternating words in S_Γ , which represent elements in BB_Γ .

DEFINITION 3.1 (Alternating words). A word w in the standard generating set S_Γ for A_Γ is called an *alternating word* if it has even length and its letters alternate between elements of S_Γ and S_Γ^{-1} , starting with an element in S_Γ .

Geometrically, the height function $\varphi_\Gamma: A_\Gamma \rightarrow \mathbb{Z}$ extends naturally to a map $\tilde{\varphi}_\Gamma: \text{Cay}(A_\Gamma, S_\Gamma) \rightarrow \mathbb{R}$, where $\text{Cay}(A_\Gamma, S_\Gamma)$ denotes the Cayley graph of A_Γ with respect to the standard generating set S_Γ . Thus, a word is alternating if and only if it represents a path in $\tilde{\varphi}_\Gamma^{-1}([0, 1]) \subset \text{Cay}(A_\Gamma, S_\Gamma)$ that starts and ends at $\tilde{\varphi}_\Gamma^{-1}(0)$.

The following lemma shows that (up to free equivalence) there is a one-to-one correspondence between alternating words and words in the standard generating set S_T of BB_Γ .

LEMMA 3.2. *Let \ll be a total order on the vertices of Γ . Let T be a spanning tree for Γ , where the edges are oriented increasingly. Let S_T be the generating set for BB_Γ with respect to T , and let Alt_Γ be the set of alternating words in S_Γ . There exist canonical maps $\Psi: (S_T \cup S_T^{-1})^* \rightarrow \text{Alt}_\Gamma$ and $\Phi: \text{Alt}_\Gamma \rightarrow (S_T \cup S_T^{-1})^*$ that satisfy the following properties.*

- (1) *The composition $\Phi \circ \Psi$ is the identity map on $(S_T \cup S_T^{-1})^*$.*
- (2) *For every $u \in \text{Alt}_\Gamma$, we have $\Psi(\Phi(u)) =_F u$.*
- (3) *For every $u \in \text{Alt}_\Gamma$, the words $\Phi(u)$ and u represent the same element in BB_Γ . Similarly, for every $v \in (S_T \cup S_T^{-1})^*$, the words $\Psi(v)$ and v represent the same element in BB_Γ .*
- (4) *If r is a relation in the Papadima–Suciu presentation corresponding to a directed triangle with vertices $v_a \ll v_b \ll v_c$, then $\Psi(r)$ is freely equivalent to a conjugate of $ac^{-1}ba^{-1}cb^{-1}$.*

In particular, there exists a canonical isomorphism between $F(S_T)$ and the image of Alt_Γ in $F(S_\Gamma)$ via the natural quotient map $(S_T \cup S_T^{-1})^ \rightarrow F(S_\Gamma)$.*

Proof. The maps Ψ and Φ are defined as follows. If $w = e_1^{\epsilon_1} \cdots e_k^{\epsilon_k}$ is a word in S_T , where $e_i = (v_{s_i}, v_{t_i})$ is an oriented edge of T for $1 \leq i \leq k$, we define

$$\Psi(w) := (s_1 t_1^{-1})^{\epsilon_1} \cdots (s_k t_k^{-1})^{\epsilon_k}.$$

Conversely, let $u = s_1 t_1^{-1} \cdots s_n t_n^{-1}$ be an alternating word, where $s_i, t_i \in S_\Gamma$ for $1 \leq i \leq n$. Recall from the definition of the Papadima–Suciu presentation that w_{s_i, t_i} is the word in the generating set S_T obtained by taking the product of the edges along the path in T from v_{s_i} to v_{t_i} . We define

$$\Phi(u) := w_{s_1, t_1} \cdots w_{s_n, t_n}.$$

The statements (1), (2), and (3) are clear from the definition of Ψ and Φ , and the fact that $w_{s, t}$ represents the element st^{-1} . For (4), let (e, f, g) be a directed triangle with $e = (v_a, v_b)$ and $f = (v_b, v_c)$. Since $\Delta(\Gamma)$ is simply connected, the words $w_e e^{-1}$ and $w_f f^{-1}$ are null-homotopic. It then follows from the definition of Φ and Ψ that $\Phi(w_e) = \Phi(e) = ab^{-1}$; similarly $\Phi(w_f) = bc^{-1}$. Thus, for $r = [w_e, w_f]$, we have

$$\Phi(r) = ab^{-1}bc^{-1}ba^{-1}cb^{-1} =_F ac^{-1}ba^{-1}cb^{-1}. \quad \square$$

Using alternating words, we define the following notion of length for an element in BB_Γ .

DEFINITION 3.3. For $g \in BB_\Gamma$, we define the *flat norm* of g as

$$\|g\|_{\flat} := \min\{|w| : w \text{ is an alternating word representing } g\}.$$

For an element $g \in BB_\Gamma$, the flat norm $\|g\|_{\flat}$ and the usual word metric with respect to S_T , denoted by $\|g\|_{S_T}$, are Lipschitz equivalent, as shown by the following result.

LEMMA 3.4. *Let T be a spanning tree for Γ . There exists $C > 0$ such that for every $g \in BB_\Gamma$, we have*

$$\frac{1}{C} \|g\|_{S_T} \leq \|g\|_{\flat} \leq C \|g\|_{S_T}.$$

Proof. The result follows directly from the isomorphism between $F(S_T)$ and the image of Alt_Γ in $F(S_\Gamma)$; see Lemma 3.2. \square

Recall that coarse diagrams are van Kampen diagrams in which the bounded regions are labelled by arbitrary null-homotopic words. We now define a particular class of coarse diagrams, where the bounded regions are labelled by null-homotopic alternating words.

DEFINITION 3.5 (Alternating diagram). Let w be a null-homotopic alternating word in the standard generating set S_Γ of A_Γ , and let D be a coarse diagram over \mathcal{P}_Γ for w . We say that D is an *alternating diagram* for w if every loop based at a vertex of height 0 is labelled by an alternating word.

Definition 3.5 can also be interpreted geometrically as follows. The 1-skeleton of a coarse diagram D over \mathcal{P}_Γ admits a natural immersion into $\text{Cay}(A_\Gamma, S_\Gamma)$. To construct the immersion, for each vertex v of D , choose a combinatorial path connecting the base point of D to v , and send v to the group element g_v that is represented by the word labelling the chosen path. Since bounded regions are labelled by null-homotopic words, the element g_v is independent of the choice of path. In particular, every vertex v has a well-defined height $\varphi_\Gamma(g_v)$.

Alternating diagrams are the coarse diagrams whose images in the Cayley graph under the natural immersion are contained in $\tilde{\varphi}_\Gamma^{-1}([0, 1])$. Equivalently, they are coarse diagrams in which all vertices have height 0 or 1.

3.2 Almost-flat van Kampen diagrams. Note that a van Kampen diagram over the standard presentation \mathcal{P}_Γ of A_Γ cannot be an alternating diagram (unless the boundary word is freely trivial), since none of the defining relations are alternating words. To address this, we allow the van Kampen diagrams to be slightly farther away from the 0-level set.

DEFINITION 3.6 (Almost-flat van Kampen diagram). Let w be a null-homotopic alternating word in the standard generating set S_Γ of A_Γ . An *almost-flat van Kampen diagram* for w is a van Kampen diagram D for w over \mathcal{P}_Γ , such that every vertex $p \in D$ has height $\varphi_\Gamma(p) \in \{0, 1, 2\}$.

The reader should think of alternating diagrams as the coarse version of almost-flat van Kampen diagrams. All the regions of an alternating diagram D for a null-homotopic alternating word w are labelled by alternating words (assuming that we start reading each word from a vertex at height 0). So, by filling each bounded region of D with an appropriate almost-flat van Kampen diagram, we obtain an almost-flat van Kampen diagram for w .

We use almost-flat van Kampen diagrams to compute the asymptotic behaviour of the Dehn function of BB_Γ .

DEFINITION 3.7 (Almost-flat area). Let w be a null-homotopic alternating word in the standard generating set S_Γ of A_Γ . The *almost-flat area* of w is the quantity defined as

$$\text{Area}_b(w) = \min\{\text{Area}(D) : D \text{ is an almost-flat van Kampen diagram for } w\}.$$

PROPOSITION 3.8. Define $\delta_b : \mathbb{N} \rightarrow \mathbb{N}$ by

$$\delta_b(n) := \max\left\{\text{Area}_b(w) : w \text{ is a null-homotopic alternating word in } A_\Gamma \text{ with } |w| \leq n\right\}.$$

Then, we have $\delta_b \asymp \delta_{BB_\Gamma}$.

Proposition 3.8 follows from the following technical lemma.

LEMMA 3.9. *For every $H \in \mathbb{N}$, there exists $M_H \in \mathbb{N}$ such that the following holds. Let w be a null-homotopic alternating word in A_Γ . Suppose that w admits a van Kampen diagram D over \mathcal{P}_Γ such that the height of every vertex v satisfies $|\varphi_\Gamma(v)| \leq H$. Then w admits an alternating diagram with the same number of bounded regions as D , such that each bounded region is labelled by an alternating word of length at most M_H .*

Lemma 3.9 can be proved using a fairly standard “pushing down” argument; see [Dis08, A+13]. We postpone the proof of Lemma 3.9 until Sect. 4, where we introduce a variation of this technique to establish the upper bounds in Theorem 1. We now show how Lemma 3.9 implies Proposition 3.8.

Proof of Proposition 3.8. Fix a spanning tree T for Γ , and let the maps $\Phi: \text{Alt}_\Gamma \rightarrow (S_T \cup S_T^{-1})^*$ and $\Psi: (S_T \cup S_T^{-1})^* \rightarrow \text{Alt}_\Gamma$ be as in Lemma 3.2.

We start by proving $\delta_b \preceq \delta_{BB_\Gamma}$. Let u be a null-homotopic alternating word of length $2n$, and let D be a van Kampen diagram for the word $w = \Phi(u)$ over the Papadima–Suciu presentation \mathcal{Q}_Γ of BB_Γ ; its edges are labelled by elements of S_T , corresponding to oriented edges of T . Replacing each edge of D , labelled with an oriented edge of T oriented from v_s to v_t , with two consecutive edges labelled by st^{-1} yields an alternating diagram for $\Psi(w) =_F u$, where all bounded regions are labelled with words of the form $\Psi(r)$ for some relation r of \mathcal{Q}_Γ . By Lemma 3.2 (4), the relations of the presentation for BB_Γ are mapped to words that are, up to free equivalence and conjugation, of the form $ac^{-1}ba^{-1}cb^{-1}$. Such an alternating word admits an almost-flat van Kampen diagram of area 3, so u admits an almost-flat van Kampen diagram of area $3 \cdot \text{Area}(D)$. This shows that $\delta_b(2n) \leq 3\delta_{\mathcal{Q}_\Gamma}(n \cdot |\mathbb{E}_\Gamma|)$.

We now prove the inequality $\delta_b \succcurlyeq \delta_{BB_\Gamma}$. Let $H = 2$ and M_2 be the constants in Lemma 3.9. Let R be the set of null-homotopic words in S_T of length at most $M_2 \cdot |\mathbb{V}_\Gamma|$ and $\mathcal{P} := \langle S_T \mid R \rangle$. Consider a null-homotopic word w of length n in S_T , and let D be an almost-flat van Kampen diagram for $u := \Psi(w)$ with area $A \leq \delta_b(2n)$.

Since D is an almost-flat van Kampen diagram, the height of every vertex $v \in D$ satisfies $|\varphi_\Gamma(v)| \leq 2$. It follows from Lemma 3.9 that there exists a constant M_2 , independent of u , such that u admits an alternating diagram with A bounded regions, all labelled by alternating words of length at most M_2 . By van Kampen’s Lemma [Bri02, Theorem 4.2.2], there are alternating words α_i and u_i for $i \in \{1, \dots, A\}$, such that $|u_i| \leq M_2$ and

$$u =_F \prod_{i=1}^A \alpha_i u_i \alpha_i^{-1}.$$

By applying Φ on both sides of the identity we get

$$w =_F \prod_{i=1}^A \Phi(\alpha_i) \Phi(u_i) \Phi(\alpha_i)^{-1}.$$

A simple estimate on the length of $\Phi(u_i)$ shows that $\Phi(u_i) \in R$. This implies that \mathcal{P} is a presentation for BB_Γ , and $\delta_{\mathcal{P}}(n) \leq \delta_b(2n)$. \square

3.3 Estimating the almost-flat area. In this section, we develop some technical tools that are helpful to construct almost-flat van Kampen diagrams with good estimates on their areas.

As previously mentioned, one approach is to start with an alternating diagram D for a null-homotopic alternating word w and then fill each bounded region with an almost-flat van Kampen diagram. The following quantity, which is well-defined for any coarse diagram, allows us to control the area of the resulting almost-flat van Kampen diagram.

DEFINITION 3.10 (Density). Let G be a group with a finite generating set S , and let D be a coarse diagram over S for a null-homotopic word w in G . Denote the labels of the bounded regions of D by w_i , $i \in I$. The *density* of D is defined as

$$\rho(D) := \frac{\sum_{i \in I} |w_i|}{|w|}.$$

Equivalently, the density relates the total number of edges of D with the length of w .

LEMMA 3.11. *Let D be a coarse diagram for a null-homotopic word w , and denote by E_D the edge set of D . The density of D satisfies*

$$\rho(D) = \frac{2|E_D|}{|w|} - 1.$$

Proof. Let w_i , $i \in I$, denote the labels of the bounded regions of D . A double-counting argument on the lengths of w and the w_i yields

$$2|E_D| = |w| + \sum_{i \in I} |w_i|.$$

The result now follows. □

The next result illustrates how density comes into play when estimating the almost-flat area of the boundary alternating word in terms of the areas of the bounded regions.

LEMMA 3.12. *Let $\alpha \geq 1$ and $C > 0$ be real numbers. Suppose that w is a null-homotopic alternating word in S_Γ admitting an alternating diagram D , such that for every alternating word w_i labelling a bounded region of D , we have $\text{Area}_b(w_i) \leq C \cdot |w_i|^\alpha$. Then*

$$\text{Area}_b(w) \leq C \cdot \rho(D)^\alpha \cdot |w|^\alpha.$$

Proof. By assumption, we can fill each bounded region of D , labelled by w_i , with an almost-flat van Kampen diagram whose almost-flat area is at most $C \cdot |w_i|^\alpha$. After filling each bounded region of D , we obtain an almost-flat van Kampen diagram D'

for w with

$$\text{Area}(D') \leq \sum_i \text{Area}_b(w_i) \leq \sum_i C \cdot |w_i|^\alpha \leq C \cdot \left(\sum_i |w_i| \right)^\alpha = C \cdot \rho(D)^\alpha \cdot |w|^\alpha.$$

This implies $\text{Area}_b(w) \leq C \cdot \rho(D)^\alpha \cdot |w|^\alpha$. \square

Lemma 3.12 tells us that, whenever we can control the density of the alternating diagram, the almost-flat area of the boundary word has the same asymptotic behaviour as the almost-flat area of the labels of the bounded regions. We formalise this strategy as the following corollary.

COROLLARY 3.13. *Let W and W' be two families of null-homotopic alternating words. Let $\alpha \geq 1$ and $C > 0$ be constants such that each $w' \in W'$ has almost-flat area bounded above by $C \cdot |w'|^\alpha$. Suppose that for every $w \in W$, there exists an alternating diagram D_w such that*

- *all its bounded regions are labelled by words in W' , and*
- *its density is bounded above uniformly in w .*

Then there exists $C' > 0$ such that for all $w \in W$, we have $\text{Area}_b(w) \leq C' \cdot |w|^\alpha$.

Proof. Let $C' := C \cdot (\sup\{\rho(D_w) : w \in W\})^\alpha$, which is finite by hypothesis. The result is now a direct consequence of Lemma 3.12. \square

The two-step filling procedure described above can also be turned into a multistep procedure, by means of the following observation.

LEMMA 3.14. *Let D be an alternating diagram for a null-homotopic alternating word w in S_Γ , and denote by w_i , $i \in I$, the labels of the bounded regions of D . Let D' be the diagram where each bounded region is replaced with an alternating diagram D_i for w_i . Then*

$$\rho(D') \leq \rho(D) \cdot \max\{\rho(D_i) : i \in I\}.$$

Proof. Denote by $w_{i,j}$, $j \in J_i$, the labels on the bounded regions of D_i . Let $M := \max\{\rho(D_i) : i \in I\}$, we get

$$\sum_{i \in I} \sum_{j \in J_i} |w_{i,j}| = \sum_{i \in I} \rho(D_i) \cdot |w_i| \leq M \cdot \sum_{i \in I} |w_i| = M \cdot \rho(D) \cdot |w|. \quad \square$$

We conclude this section with a standard technique for obtaining upper bounds on the area of null-homotopic words. The idea behind this method can be traced back to Gromov [Gro93, 5.A₃]. Gersten and Short later formalised this idea and used it to provide upper bounds on the Dehn functions of certain subgroups of hyperbolic groups [GS02, Lemma 2.2]. In the next result, we present a generalised version due to Carter and Forester [CF17], applied to our setting of alternating words and almost-flat area. Intuitively, it tells us that to obtain an estimate of the Dehn function, it suffices to consider a certain class of triangular diagrams.

LEMMA 3.15 (Triangle Lemma). *Let $C > 0$ and $\alpha \geq 2$ be real numbers. For each $g \in BB_\Gamma$, choose an alternating word w_g representing g such that $|w_g| \leq C \cdot \|g\|_b$. Suppose that for every $g_1, g_2, g_3 \in BB_\Gamma$ satisfying $g_1 g_2 g_3 = 1$, we have*

$$\text{Area}_b(w_{g_1} w_{g_2} w_{g_3}) \leq C \cdot (\|g_1\|_b + \|g_2\|_b + \|g_3\|_b)^\alpha.$$

Then, $\delta_{BB_\Gamma}(n) \preceq n^\alpha$.

Proof. Since the proof is fairly standard and analogous to the proof of [CF17, Theorem 4.2], we give only the main idea and omit the details. Starting with an empty diagram for an alternating word w , we subdivide it into triangular regions whose sides are labelled by alternating words of the form w_g . By hypothesis, we can fill each triangular region with an almost-flat van Kampen diagram and estimate the almost-flat area of the whole diagram.

By exploiting the geometric series, one can see that the area is at most $C' \cdot |w|^\alpha$ for some $C' = C'(C, \alpha)$. Combining this with Proposition 3.8 yields the result. \square

4 Coloured words and coloured diagrams

Establishing an upper bound for the area of a null-homotopic alternating word w in BB_Γ is generally challenging. As mentioned in the introduction, there are cases where this is straightforward. For example, if we express w as a word in S_Γ and assume that all its letters commute with some fixed generator $a \in S_\Gamma$, then w represents an element in $BB_{\text{St}(v_a)}$, which is isomorphic to $A_{\text{Lk}(v_a)}$ via the homomorphism induced by $sa^{-1} \mapsto s$ for $v_s \in \text{Lk}(v_a)$. Thus, the area of w in BB_Γ is at most $|w|^2$. Concretely, to transform w into the empty word, we commute its letters different from a until they cancel out, using a as a “counterweight” to keep the word alternating at all times.

To handle the general case, the main idea is to split a null-homotopic alternating word w into subwords in such a way that each subword has its support contained in the star of some vertex of Γ . This provides a preferred way to manipulate each subword, as it belongs to the special case discussed above. We also want to keep track of the vertex used as a counterweight, as there might be multiple choices available.

This leads to the definition of *coloured words*, which can be intuitively understood as follows. In the example above, the counterweight $a \in S_\Gamma$ is now thought of as the colour of the word w . In general, the palette of colours consists of the standard generating set S_Γ that corresponds to the generators used as counterweights. Given a word w in S_Γ , we “paint” its letters so that each letter of w commutes with its colour.

This process splits w into subwords by grouping together the maximal substrings of consecutive letters in w that share the same colour. A subword with colour a defines an element in the right-angled Artin group $A_{\text{St}(v_a)}$. This gives rise to a canonical way to “push down” this coloured word to obtain an alternating word in $BB_{\text{St}(v_a)}$; see

Definition 4.10. Moreover, we exploit this idea to define a class of diagrams for null-homotopic coloured words, called *coloured diagrams*, and a way to “push down” these diagrams to obtain alternating diagrams.

4.1 Coloured words and pushdowns. We now give the formal definitions of coloured words and their pushdowns, and present some basic properties.

DEFINITION 4.1. The *coloured generating set* is the subset $S_\Gamma^{\text{col}} \subseteq S_\Gamma \times S_\Gamma$ consisting of all pairs (s, a) such that $[s, a] = 1$. It forms a generating set for A_Γ (with repeated generators) when equipped with the map $S_\Gamma^{\text{col}} \rightarrow A_\Gamma$ defined by $(s, a) \mapsto s$. A *coloured word* is a word in S_Γ^{col} , that is, an element in the free monoid generated by $S_\Gamma^{\text{col}} \cup (S_\Gamma^{\text{col}})^{-1}$.

The letter $(s, a) \in S_\Gamma^{\text{col}}$ should be thought of as the letter s coloured with the colour a . To emphasise this, we introduce the following notation.

NOTATION 4.2. Let $a \in S_\Gamma$. If $s \in S_\Gamma$ commutes with a , then we denote by $[s]_a$ the element $(s, a) \in S_\Gamma^{\text{col}}$ and refer to it as a *coloured letter*. The formal inverse of $[s]_a$ is denoted by $[s]_a^{-1}$. Given a word $w = s_1 \cdots s_k$ in S_Γ and $a \in S_\Gamma$ that commutes with s_i for all $i \in \{1, \dots, k\}$, we write $[w]_a$ to denote the coloured word $[s_1]_a \cdots [s_k]_a$; we refer to a as the *colour* of w . Moreover, whenever we write $[w]_a$, we implicitly assume that all the letters of w commute with a .

DEFINITION 4.3. Let $\mathbf{w} = [s_1]_{a_1} \cdots [s_k]_{a_k}$ be a coloured word. The *underlying word* of \mathbf{w} is the word $w = s_1 \cdots s_k$ in S_Γ obtained by forgetting all the colours.

By definition, if \mathbf{w} is a coloured word and w its underlying word, then $|\mathbf{w}| = |w|$.

REMARK 4.4. In the rest of the paper, we use bold letters for coloured words in S_Γ^{col} and the same letters in normal font for the corresponding underlying word in S_Γ .

A coloured word \mathbf{w} and its underlying word w represent the same element in A_Γ . In particular, a coloured word \mathbf{w} is null-homotopic if and only if its underlying word w is null-homotopic.

DEFINITION 4.5 (*k*-coloured word). A *k*-coloured word is a coloured word of the form

$$[w_1]_{a_1} \cdots [w_k]_{a_k},$$

where w_i is a nonempty word for all i , and a_i differs from a_{i+1} for all $i \in \{1, \dots, k-1\}$. We call 1-coloured words *monochromatic*.

REMARK 4.6. As an example, the word $[w_1]_a [w_2]_b [w_3]_a$ is 3-coloured, even though a colour is repeated.

Given a word $w = s_1^{\epsilon_1} \cdots s_k^{\epsilon_k}$ in S_Γ , where $\epsilon_i \in \pm 1$, there is a natural way to “colour it”, which we refer to as *self-colouring*, namely, the coloured word $[s_1]_{s_1}^{\epsilon_1} \cdots [s_k]_{s_k}^{\epsilon_k}$.

Coloured words are useful because the colours provide a preferred way to “push down” a word so that it becomes an alternating word in S_Γ . We formalise this as follows.

We fix, once and for all, a generator $s_0 \in S_\Gamma$.

DEFINITION 4.7 (Transition word). Let $a \in S_\Gamma$ and T a spanning tree of Γ . Let $\gamma_a: [0, \ell] \rightarrow T$ be the unique simple path in T from v_{s_0} to v_a . For $i \in \{0, \dots, \ell\}$, let $t_i \in S_\Gamma$ be such that $v_{t_i} = \gamma_a(i) \in V_\Gamma$; in particular, $t_0 = s_0$, $t_\ell = a$, and $[t_i, t_{i+1}] = 1$. For $h \in \mathbb{Z}$, we define the *transition word* $\tau_{a,h}$ as the alternating word

$$\tau_{a,h} := (t_0 t_1^{-1})^h \cdots (t_{\ell-1} t_\ell^{-1})^h.$$

Observe that $\tau_{a,h}$ represents the element $s_0^h a^{-h} \in A_\Gamma$. For every pair of vertices $v_x, v_y \in T$, Dison defined the word $p_h(v_x, v_y)$ in S_T ; see [Dis08, Sect. 4]. The image of $p_h(v_x, v_y)$ under the map $\Psi: (S_T \cup S_T^{-1})^* \rightarrow \text{Alt}_\Gamma$ from Lemma 3.2 is the word $\tau_{x,h}^{-1} \tau_{y,h}$.

We record the following straightforward observation about the length of the transition words.

LEMMA 4.8. *Let $a \in S_\Gamma$ and $h \in \mathbb{Z}$. If $a \neq s_0$, then $2|h| \leq |\tau_{a,h}| \leq 2|h| \cdot |V_\Gamma|$. Otherwise, $\tau_{a,h}$ is the trivial word.*

Proof. The result follows directly from the definition of $\tau_{a,h}$ and the observation that the length of the path γ_a defining $\tau_{a,h}$ is at most $|V_\Gamma|$. \square

NOTATION 4.9. For a word w in S_Γ and $a \in S_\Gamma$, we denote by $w(\Gamma a^{-1})$ the word obtained from w by replacing each letter $s \in S_\Gamma$ (respectively, $s^{-1} \in S_\Gamma^{-1}$) with sa^{-1} (respectively, as^{-1}).

DEFINITION 4.10 (Pushdown of coloured words). Let $h \in \mathbb{Z}$ and let $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ be a k -coloured word. Let $h_0 = h$ and $h_i = h + \sum_{j=1}^i \varphi_\Gamma(w_j)$ for $i \in \{1, \dots, k\}$. The h -pushdown of \mathbf{w} is the alternating word

$$\text{push}_h(\mathbf{w}) := \prod_{i=1}^k \tau_{a_i, h_{i-1}} w_i(\Gamma a_i^{-1}) \tau_{a_i, h_i}^{-1}.$$

We give an example for the definition of the pushdown of a coloured word.

EXAMPLE 4.11. Let $\Gamma = \{v_a\} * \Lambda$ be a cone graph and set $a \in S_\Gamma$ to be the generator corresponding to the vertex v_a . Fix the spanning tree of Γ consisting of all the edges that contain v_a as a vertex. Let $w = s_1^2 s_2^{-1} s_3$ be a word in S_Γ . Consider the monochromatic word $\mathbf{w} = [w]_a = [s_1]_a^2 [s_2]_a^{-1} [s_3]_a$. For $h \in \mathbb{Z}$, the h -pushdown of \mathbf{w} is the alternating word

$$\begin{aligned} \text{push}_h(\mathbf{w}) &= \tau_{a,h} w(\Gamma a^{-1}) \tau_{a, h + \varphi_\Gamma(w)}^{-1} \\ &= \tau_{a,h} (s_1 a^{-1})^2 (a s_2^{-1}) (s_3 a^{-1}) \tau_{a, h+2}^{-1} \end{aligned}$$

$$= (s_0 a^{-1})^h (s_1 a^{-1})^2 (a s_2^{-1}) (s_3 a^{-1}) (a s_0^{-1})^{h+2},$$

where s_0 is the chosen generator in S_Γ corresponding to a chosen base point in the spanning tree.

We now prove a series of lemmas about the pushdown of coloured words.

LEMMA 4.12. *Let $h \in \mathbb{Z}$. The h -pushdown has the following properties.*

(1) *Let $[s]_a^\epsilon$ be a coloured letter and $\epsilon = \pm 1$, then*

$$\text{push}_h([s]_a^\epsilon) = \tau_{a,h}(s a^{-1})^\epsilon \tau_{a,h+\epsilon}^{-1}.$$

(2) *Let \mathbf{w} and \mathbf{w}' be coloured words, then*

$$\text{push}_h(\mathbf{w}\mathbf{w}') =_F \text{push}_h(\mathbf{w}) \text{push}_{h+\varphi_\Gamma(\mathbf{w})}(\mathbf{w}').$$

(3) *If $\mathbf{w} =_F \mathbf{w}'$, then $\text{push}_h(\mathbf{w}) =_F \text{push}_h(\mathbf{w}')$.*

(4) *If the underlying word w of \mathbf{w} is alternating, then $\text{push}_0(\mathbf{w}) =_F w$.*

Proof. Statement (1) follows from Definition 4.10. Statement (2) also follows from Definition 4.10. In fact, if consecutive colours are different we get an identity between words, whereas if the last colour of \mathbf{w} coincides with the first colour of \mathbf{w}' , then a free insertion is required; we illustrate this in the case where \mathbf{w} and \mathbf{w}' are monochromatic. In this case, we have $\mathbf{w}\mathbf{w}' = [w]_a[w']_a = [ww']_a$, therefore

$$\begin{aligned} \text{push}_h([ww']_a) &= \tau_{a,h} \cdot (ww')(\Gamma a^{-1}) \cdot \tau_{a,h+\varphi_\Gamma(ww')}^{-1} \\ &=_F \tau_{a,h} w(\Gamma a^{-1}) \tau_{a,h+\varphi_\Gamma(w)}^{-1} \tau_{a,h+\varphi_\Gamma(w')} w'(\Gamma a^{-1}) \tau_{a,h+\varphi_\Gamma(ww')}^{-1} \\ &= \text{push}_h([w]_a) \text{push}_{h+\varphi_\Gamma(w)}([w']_a). \end{aligned}$$

Combining (1) and (2) shows that $\text{push}_h([s]_a[s]_a^{-1})$ is freely trivial, and (3) follows. Finally, to show (4), it suffices to check that the identity holds for a coloured word $\mathbf{w} = [a]_b[c]_d^{-1}$ of length 2. Here, we have

$$\text{push}_0(\mathbf{w}) = \tau_{b,0} a b^{-1} \tau_{b,1}^{-1} \tau_{d,1} d c^{-1} \tau_{d,0}^{-1} =_F a b^{-1} b s_0^{-1} s_0 d^{-1} d c^{-1} =_F a c^{-1},$$

where we use the fact that $\tau_{a',0}$ is the trivial word and $\tau_{a',1} =_F a' s_0^{-1}$ for $a' \in S_\Gamma$. \square

LEMMA 4.13. *Let $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ be a k -coloured word. Let $h_0 = h \in \mathbb{Z}$ and $h_i = h + \sum_{j=1}^i \varphi_\Gamma(w_j)$ for $i \in \{1, \dots, k\}$. Then*

$$|\text{push}_h(\mathbf{w})| \leq 2|\mathbf{w}| + 4 \left(\sum_{i=0}^k |h_i| \right) \cdot |V_\Gamma|.$$

Additionally, if \mathbf{w} is null-homotopic and $a_1 \neq a_k$, then $h_0 = h_k$ and

$$2|\mathbf{w}| + 2 \sum_{i=1}^k |h_i| \leq |\text{push}_h(\mathbf{w})|.$$

Proof. By Definition 4.10, we have

$$\begin{aligned} |\text{push}_h(\mathbf{w})| &= \sum_{i=1}^k |\tau_{a_i, h_{i-1}}| + \sum_{i=1}^k |w_i(\Gamma a_i^{-1})| + \sum_{i=1}^k |\tau_{a_i, h_i}^{-1}| \\ &= \sum_{i=1}^k |w_i(\Gamma a_i^{-1})| + \sum_{i=1}^{k-1} (|\tau_{a_{i+1}, h_i}| + |\tau_{a_i, h_i}|) + |\tau_{a_1, h_0}| + |\tau_{a_k, h_k}|. \end{aligned}$$

Since $|w_i(\Gamma a_i^{-1})| = 2|w_i|$, we have

$$\sum_{i=1}^k |w_i(\Gamma a_i^{-1})| = 2 \sum_{i=1}^k |w_i| = 2|\mathbf{w}|.$$

Using Lemma 4.8 to bound the lengths of the transition words gives

$$\sum_{i=1}^{k-1} (|\tau_{a_{i+1}, h_i}| + |\tau_{a_i, h_i}|) + |\tau_{a_1, h_0}| + |\tau_{a_k, h_k}| \leq \left(4 \sum_{i=1}^{k-1} |h_i| + 2|h_0| + 2|h_k| \right) \cdot |V_\Gamma|.$$

Putting all the identities and inequalities together yields the desired upper bound.

For the lower bound, note that \mathbf{w} being null-homotopic implies that $h_0 = h_k$. Moreover, for $i \in \{1, \dots, k-1\}$, since a_i and a_{i+1} are different, at least one of them is different from s_0 ; the same holds for a_1 and a_k . Therefore, Lemma 4.8 yields

$$\sum_{i=1}^{k-1} (|\tau_{a_{i+1}, h_i}| + |\tau_{a_i, h_i}|) + |\tau_{a_1, h_0}| + |\tau_{a_k, h_k}| \geq 2 \sum_{i=1}^{k-1} |h_i| + 2|h_k|. \quad \square$$

LEMMA 4.14. *Let $h \in \mathbb{Z}$. If $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ represents $g \in A_\Gamma$, then $\text{push}_h(\mathbf{w})$ represents the element $s_0^h g s_0^{-h - \varphi_\Gamma(g)}$ in BB_Γ .*

Proof. Let $h_0 = h \in \mathbb{Z}$ and $h_i = h + \sum_{j=1}^i \varphi_\Gamma(w_j)$ for $i \in \{1, \dots, k\}$. Since a_i commutes with w_i for $i \in \{1, \dots, k\}$, the word $w_i(\Gamma a_i^{-1})$ represents the element $w_i a_i^{-\varphi_\Gamma(w_i)} = a_i^{h_{i-1}} w_i a_i^{-h_i}$. Moreover, we have $\tau_{a_i, h} =_{A_\Gamma} s_0^h a_i^{-h}$. The result follows from Lemma 4.12. \square

4.2 Efficient coloured words. In the previous section, we showed that a coloured word can always be turned into an alternating word by means of the pushdown. In this section, we show that for each element of BB_Γ , there is a particularly nice coloured word whose pushdown represents it.

DEFINITION 4.15. Let $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ be a k -coloured word. A *proper block* of \mathbf{w} is a subword of the form $[w_i]_{a_i} [w_{i+1}]_{a_{i+1}} \cdots [w_j]_{a_j}$ for some $2 \leq i \leq j \leq k-1$.

DEFINITION 4.16. We call a k -coloured word \mathbf{w} representing $g \in A_\Gamma$:

- *geodesic*, if it is a minimal-length representative for g ;
- *chromatically minimal*, if there are no k' -coloured words representing g with $k' < k$;
- *efficient*, if it is geodesic and every proper block is chromatically minimal.

REMARK 4.17. Note that a coloured word \mathbf{w} is geodesic if and only if the underlying word w is a minimal-length representative for g with respect to the standard generating set S_Γ .

If a coloured word is chromatically minimal and geodesic, then it is also efficient because the proper blocks of a chromatically minimal word are chromatically minimal. The reason for using proper blocks in the definition of efficient coloured words is that being chromatically minimal does not pass to subwords. As an example, label the vertices of a path of length four, from left to right, by a , b , c , and d . Then $[ab]_a[cd]_c$ is chromatically minimal, whereas $[b]_a[cd]_c$ is not, since it represents the same element bcd as $[bcd]_c$. However, it is not difficult to see that a subword of an efficient word is also efficient, as the proper blocks of a subword are, in particular, proper blocks of the original coloured word.

LEMMA 4.18. *For each element $g \in BB_\Gamma$, there exists an efficient coloured word representing it.*

Proof. Let \mathbf{w} be a chromatically minimal representative for g . If it is not geodesic, then the underlying word w has a subword of the form $sw's^{-1}$, where w' is a word commuting with $s \in S_\Gamma \cup S_\Gamma^{-1}$. Removing the coloured letters corresponding to s and s^{-1} from \mathbf{w} yields another chromatically minimal word of shorter length. Repeating this procedure eventually yields a chromatically minimal and geodesic representative, and therefore, an efficient one. \square

In general, the length of the pushdown of a coloured word \mathbf{w} might be considerably greater than the length of \mathbf{w} because transition words contribute to its length. Since Bestvina–Brady groups can be quadratically distorted in right-angled Artin groups [Tra17, Theorem 1.1], the increase in the length of the pushdown cannot be avoided. Nevertheless, when \mathbf{w} is efficient, the length of its pushdown is always optimal (up to a multiplicative constant).

PROPOSITION 4.19. *Let \mathbf{w} be an efficient k -coloured word representing an element $g \in BB_\Gamma$. Then*

$$\|g\|_b \leq |\text{push}_0(\mathbf{w})| \leq (24|V_\Gamma| + 2)\|g\|_b.$$

To prove Proposition 4.19, we need the following two lemmas.

LEMMA 4.20. *Let w be a null-homotopic word in S_Γ and D a van Kampen diagram for w . Let p and q be two vertices on the boundary of D , and let g be the element represented by the word read from p to q along the boundary. Let C and C' be corridors containing p and q , respectively.*

- (1) *If C and C' cross, then g can be represented by a 2-coloured word.*
- (2) *If there is a corridor that crosses both C and C' , then g can be represented by a 3-coloured word.*

Proof. To prove (1), let C be an a -corridor and C' a b -corridor for some $a, b \in S_\Gamma$. If C and C' cross, then the side of C containing p intersects the side of C' containing

q in a vertex p' . Let w' be the word labelling the path along the side of C from p to p' , and let w'' be the word labelling the path along the side of C' from p' to q . It follows that $[w']_a[w'']_b$ is a 2-coloured word representing g .

For the proof of (2), let C'' be a c -corridor for some $c \in S_\Gamma$ that crosses both C and C' , and let p' and q' be the intersection vertices of one side of C'' with the sides of C and C' that contain p and q , respectively. Let w , w'' , and w' be the words labelling the paths from p to p' , p' to q' , and q' to q , respectively. Then, the element g is represented by the 3-coloured word $[w]_a[w'']_c[w']_b$. \square

LEMMA 4.21. *Let $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ be an efficient k -coloured word whose underlying word is $w = w_1 \cdots w_k$. Let w' be a word in A_Γ such that ww' is null-homotopic, and let D be a minimal-area van Kampen diagram for ww' . For $i \in \{1, \dots, k-1\}$, let p_i be the vertex on the boundary of D at the intersection of w_i and w_{i+1} , and let C_i be a corridor containing p_i . Then, for $i, j \in \{1, \dots, k-1\}$, the following statements hold.*

- (1) *If $|i - j| \geq 3$, then C_i and C_j do not cross.*
- (2) *If $|i - j| \geq 4$, no corridor crosses both C_i and C_j .*
- (3) *The total length of the corridors C_i satisfies*

$$\sum_i |C_i| \leq 2|ww'|.$$

Proof. Let $i, j \in \{1, \dots, k-1\}$ with $j \geq i+3$. If C_i and C_j cross, then by Lemma 4.20, the group element g represented by the proper block $[w_{i+1}]_{a_{i+1}} \cdots [w_j]_{a_j}$ can be represented by a 2-coloured word, contradicting the efficiency of \mathbf{w} . Similarly, let $j \geq i+4$, and suppose that C_i and C_j are crossed by the same corridor. Then the element g , defined as above, would be represented by a 3-coloured word, again a contradiction. This proves (1) and (2). In particular, the statement in (2) implies that a corridor C in D can cross at most four of the corridors C_1, \dots, C_{k-1} .

Since D is a minimal-area diagram, it contains no annuli, and every two corridors cross at most once. Therefore, the length of a corridor equals the number of corridors crossing it. There are $\frac{1}{2}|ww'|$ corridors in total, and each corridor can cross at most four of the corridors in the family $\{C_1, \dots, C_{k-1}\}$. That is, each of the $\frac{1}{2}|ww'|$ corridors contributes at most 4 to the overall sum $\sum_i |C_i|$. Statement (3) now follows from a double-counting argument. \square

We now proceed with the proof of Proposition 4.19.

Proof of Proposition 4.19. The first inequality $\|g\|_b \leq |\text{push}_0(\mathbf{w})|$ follows from the definition of $\|\cdot\|_b$; see Definition 3.3. For the second inequality, let $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ be an efficient k -coloured word with underlying word $w = w_1 \cdots w_k$, and let u be a shortest alternating word representing g , meaning that $|u| = \|g\|_b$. Since \mathbf{w} is geodesic, we have $|w| \leq |u|$.

Let D be a minimal-area van Kampen diagram for wu^{-1} , and let p_1, \dots, p_{k-1} be the points on the boundary path labelled by w as in Lemma 4.21. We also denote

by p_0 and p_k the points respectively at the start and end of w . Then the height h_i of p_i is $\varphi_\Gamma(w_1 \dots w_i)$.

Consider corridors C_i containing p_i for $i \in \{1, \dots, k-1\}$. Since \mathbf{w} is geodesic, the corridor C_i cannot have both ends on w , so it must begin at an edge of w and end at an edge of u . Since u is alternating, all the vertices of u have height 0 or 1, so p_i has distance at most $|C_i| + 1$ from a vertex of height 0. Therefore, we may estimate the height of p_i as $|h_i| \leq |C_i| + 1$ for every $i \in \{0, \dots, k-1\}$.

Finally, applying Lemmas 4.13 and 4.21 gives

$$\begin{aligned} |\text{push}_0(\mathbf{w})| &\leq 2|w| + 4 \left(\sum_{i=0}^k |h_i| \right) |V_\Gamma| \\ &\leq 2|w| + 4(2|wu^{-1}| + k - 1 + |h_k|) |V_\Gamma| \leq (24|V_\Gamma| + 2)|u|, \end{aligned}$$

where we use the fact that $\max\{k, |h_k|\} \leq |w| \leq |u|$. \square

4.3 Coloured diagrams. To compute the Dehn function of BB_Γ , we need to produce an algorithm that, given an arbitrary null-homotopic alternating word u in BB_Γ , constructs an almost-flat van Kampen diagram with small area. However, if u is the pushdown of a coloured word \mathbf{w} , it is more convenient to work directly with the coloured word. To this end, we define the notion of *coloured diagram*, which can be pushed down to obtain an alternating diagram for u .

DEFINITION 4.22 (Coloured diagram). A *coloured diagram* for a null-homotopic coloured word \mathbf{w} is a coarse diagram for \mathbf{w} over the generating set S_Γ^{col} .

Each edge of a coloured diagram is labelled with a coloured letter $[s]_a$, and we sometimes refer to a as the *colour* of the edge. In this sense, a coloured diagram is a coarse diagram over the standard generating set S_Γ , where every edge carries additional information given by the colour.

Just as the pushdown of a coloured word produces an alternating word, the pushdown of a coloured diagram produces an alternating diagram.

DEFINITION 4.23 (Pushdown of a coloured diagram). Let \mathbf{D} be a coloured diagram for a null-homotopic coloured word \mathbf{w} , and let $h \in \mathbb{Z}$. The *h -pushdown of \mathbf{D}* is the alternating diagram $\text{push}_h(\mathbf{D})$ obtained as follows.

- For each vertex p in \mathbf{D} , there is a vertex p' in $\text{push}_h(\mathbf{D})$.
- For each edge in \mathbf{D} , labelled by the coloured letter $[s]_a$ and oriented from p to q , we add a path from the corresponding vertices p' to q' in $\text{push}_h(\mathbf{D})$. This path is labelled with

$$\text{push}_{h+\varphi_\Gamma(p)}([s]_a) = \tau_{a,h+\varphi_\Gamma(p)} s a^{-1} \tau_{a,h+\varphi_\Gamma(q)}^{-1}.$$

- For every vertex p of \mathbf{D} , and for every pair of consecutive edges incident to p labelled with a letter of the same colour a , the corresponding paths in $\text{push}_h(\mathbf{D})$ have an initial segment labelled by the same word $\tau_{a,\varphi_\Gamma(p)}$. We *fold* these initial segments by identifying them together.

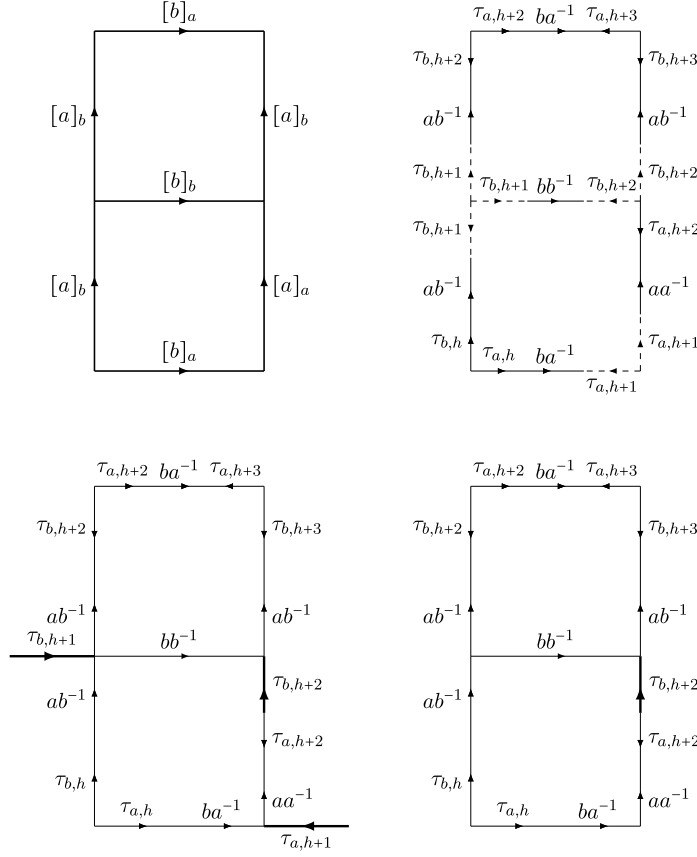


Figure 2: Example of the pushdown of a coloured diagram. The top left picture shows a coloured diagram \mathbf{D} . To obtain the h -pushdown, we first construct a diagram by replacing each edge of \mathbf{D} with a path labelled by the pushdown of the label on that edge, as shown in the top right picture. Next, we fold the paths originating from the same vertex that are labelled with the same transition word, as illustrated in the bottom left picture. Finally, wherever we have collapsed all the paths arising from a vertex, we can remove the resulting path; see the bottom right picture.

- If all the edges incident to p have the same colour a , the vertex p' corresponding to p in $\text{push}_h(\mathbf{D})$ has, after collapsing, a single path starting at it, which is labelled with $\tau_{a,\varphi_{\Gamma}(p)}$: we remove p' and the path altogether.

EXAMPLE 4.24. Fig. 2 illustrates a coloured diagram for $\mathbf{w} = [a^2]_b [b]_a [a^{-1}]_b [a^{-1}b^{-1}]_a$ together with its h -pushdown, which is an alternating diagram for $\text{push}_h(\mathbf{w})$ which is the alternating word

$$\tau_{b,h} ab^{-1} ab^{-1} \tau_{b,h+2}^{-1} \tau_{a,h+2} ba^{-1} \tau_{a,h+3}^{-1} \tau_{b,h+3} ba^{-1} \tau_{b,h+2}^{-1} \tau_{a,h+2} aa^{-1} ab^{-1} \tau_{a,h}^{-1}.$$

To ensure that the h -pushdown of a coloured diagram is well-defined, we must check that the labels of the h -pushdown of the bounded regions are null-homotopic words in S_{Γ} . This follows directly from the next result.

LEMMA 4.25. *Let $h \in \mathbb{Z}$. Let \mathbf{D} be a coloured diagram for a null-homotopic coloured word. Let $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ be the word read on a (possibly unbounded) region of \mathbf{D} , starting from a point p , and let u be the alternating word read on the corresponding region of $\text{push}_h(\mathbf{D})$. If $a_k \neq a_1$, then $u = \text{push}_{h+\varphi_\Gamma(p)}(\mathbf{w})$. Otherwise, if $a_k = a_1 = a$, then $\text{push}_{h+\varphi_\Gamma(p)}(\mathbf{w}) = \tau_{a,h+\varphi_\Gamma(p)} u \tau_{a,h+\varphi_\Gamma(p)}^{-1}$. In either case, the word u is null-homotopic.*

Proof. For $i \in \{0, \dots, k\}$, let $h_i = h + \varphi_\Gamma(p) + \varphi_\Gamma(w_1 \cdots w_i)$. Since \mathbf{w} is the label of a (possibly unbounded) region in \mathbf{D} , it is null-homotopic, so $h_0 = h_k = h + \varphi_\Gamma(p)$. From Definition 4.23, we get that, if $a_1 \neq a_k$, then

$$u = \prod_{i=1}^k \tau_{a_i, h_{i-1}} w_i (\Gamma a_i^{-1}) \tau_{a_i, h_i}^{-1} = \text{push}_{h+\varphi_\Gamma(p)}(\mathbf{w}).$$

The key point is that all the transition words appearing between letters of the same colour disappear after folding. The first and the last transition words are not folded because of the assumption $a_1 \neq a_k$. Otherwise, if they are folded, that is, if $a_k = a_1 = a$, we obtain

$$\text{push}_{h+\varphi_\Gamma(p)}(\mathbf{w}) = \tau_{a,h+\varphi_\Gamma(p)} u \tau_{a,h+\varphi_\Gamma(p)}^{-1}$$

In both cases, the fact that u is null-homotopic follows directly from Lemma 4.14, since \mathbf{w} is null-homotopic. \square

As highlighted in Lemma 4.25, it is often convenient to assume that the first and last colour of a null-homotopic coloured word \mathbf{w} are different. This is not a critical assumption, since monochromatic words are simple enough that we do not need to examine their coloured diagrams. If \mathbf{w} is not monochromatic, we can always achieve this assumption by considering a cyclic conjugate of \mathbf{w} . Geometrically, if we imagine \mathbf{w} drawn on a circle, this corresponds to reading the word starting from a vertex adjacent to edges of different colours. Therefore, we get the following consequence.

COROLLARY 4.26. *Let $h \in \mathbb{Z}$. If \mathbf{D} is a coloured diagram for a null-homotopic k -coloured word $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ with $a_k \neq a_1$, then $\text{push}_h(\mathbf{D})$ is an alternating diagram for $\text{push}_h(\mathbf{w})$. If the boundary word $\mathbf{w} = [w]_a$ is monochromatic, then $\text{push}_h(\mathbf{D})$ is an alternating diagram for $w(\Gamma a^{-1})$.*

We are now able to prove Lemma 3.9, which we recall here for the reader's convenience.

LEMMA 3.9. *For every $H \in \mathbb{N}$, there exists $M_H \in \mathbb{N}$ such that the following holds. Let w be a null-homotopic alternating word in A_Γ . Suppose that w admits a van Kampen diagram D over \mathcal{P}_Γ such that the height of every vertex v satisfies $|\varphi_\Gamma(v)| \leq H$. Then w admits an alternating diagram with the same number of bounded regions as D , such that each bounded region is labelled by an alternating word of length at most M_H .*

Proof. Let \mathbf{D} be the coloured diagram obtained from D by replacing each label s with $[s]_s$. The pushdown of \mathbf{D} is an alternating diagram D' for a word w' that is freely equivalent to w by Lemma 4.12 (4). After performing appropriate foldings in D' , we obtain an alternating diagram D_w for w . Since the bounded regions of \mathbf{D} are labelled with words of length 4, by Lemmas 4.13 and 4.25, the bounded regions of D_w are labelled with the words whose lengths are bounded by a constant M_H depending only on H . \square

Recall that our strategy for obtaining upper bounds on the Dehn function of BB_Γ is to produce an algorithm that takes an alternating word u as input and outputs an almost-flat van Kampen diagram whose almost-flat area satisfies the desired upper bound.

We introduced coloured diagrams to construct this algorithm in the case where $u = \text{push}_h(\mathbf{w})$ for some null-homotopic coloured word \mathbf{w} and some $h \in \mathbb{Z}$. In this case, if \mathbf{D} is a coloured diagram for \mathbf{w} , then $\text{push}_h(\mathbf{D})$ is an alternating diagram for u . By Lemma 4.25, the bounded regions of $\text{push}_h(\mathbf{D})$ are labelled by the pushdowns of the labels of the bounded regions of \mathbf{D} . Therefore, to estimate $\text{Area}_b(u)$, it suffices to estimate the almost-flat area of the label of each bounded region of $\text{push}_h(\mathbf{D})$.

By Lemma 3.12, the estimates are particularly nice if we have a uniform bound on the density of $\text{push}_h(\mathbf{D})$. Note that the coloured diagram \mathbf{D} , which is by definition a coarse diagram over the generating set S_Γ^{col} , also has a well-defined density. However, estimating the density of \mathbf{D} is not enough to obtain an estimate of the density of its pushdown, as one needs to take into account for the lengths of the transition words, which appear at every vertex incident to edges of different colours.

DEFINITION 4.27. Let \mathbf{D} be a coloured diagram. A vertex of \mathbf{D} is called *polychromatic* if it is incident to at least two edges of different colours. It is called *∂ -polychromatic* if it is on the boundary of \mathbf{D} and incident to two consecutive boundary edges of different colours.

PROPOSITION 4.28. *For every $M > 0$, there exists $C_M > 0$ such that the following holds. Let \mathbf{D} be a coloured diagram, and for every polychromatic vertex p , let γ_p be a path in \mathbf{D} from p to an arbitrary ∂ -polychromatic vertex (if p is ∂ -polychromatic, then we can choose γ_p to be the constant path). Suppose that*

- (1) *for every polychromatic vertex p in \mathbf{D} , we have $\deg p \leq M$,*
- (2) *for every edge e in \mathbf{D} , we have $|\{\gamma_p \mid e \in \gamma_p\}| \leq M$, and*
- (3) *$\rho(\mathbf{D}) \leq M$.*

Then $\rho(\text{push}_h(\mathbf{D})) \leq C_M$ for every $h \in \mathbb{Z}$.

Proof. Recall that by Lemma 3.11, it suffices to estimate the number of edges $|E_{\text{push}_h(\mathbf{D})}|$ for the pushdown diagram. We assume that there is at least one ∂ -polychromatic vertex; otherwise, there would not be any polychromatic vertices, so the labels of all bounded regions of \mathbf{D} would be monochromatic, and it would follow from Corollary 4.26 that $\rho(\mathbf{D}) = \rho(\text{push}_h(\mathbf{D}))$.

Let $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ be the boundary word of \mathbf{D} . Since the statement is independent of the choice of base point of \mathbf{D} , we may assume that the base point is ∂ -polychromatic, so $a_k \neq a_1$. In particular, the diagram \mathbf{D} has ∂ -polychromatic vertices p_1, \dots, p_k , where $k \geq 2$, and p_i is at the end of $[w_i]_{a_i}$.

Recall that an edge in \mathbf{D} from p to q labelled by $[s]_a$ is replaced by a path in $\text{push}_h(\mathbf{D})$ labelled by $\text{push}_{h+\varphi_\Gamma(p)}([s]_a) = \tau_{a, h+\varphi_\Gamma(p)} s a^{-1} \tau_{a, h+\varphi_\Gamma(q)}^{-1}$. Call the two edges in the middle labelled with $s a^{-1}$ *central edges*, and the ones in the paths labelled with $\tau_{a, \varphi_\Gamma(p)}$ and $\tau_{a, \varphi_\Gamma(q)}$ *transition edges*. Denote by Z and T the number of central and transition edges, respectively.

For every edge of \mathbf{D} , we get two central edges in $\text{push}_h(\mathbf{D})$, so

$$Z = 2|E_{\mathbf{D}}|.$$

If a vertex p is not polychromatic, then all the transition edges arising from it get removed. Otherwise, the vertex p gives rise to at most $2M \cdot |V_\Gamma| \cdot |h + \varphi_\Gamma(p)|$ transition edges by Lemma 4.8. Therefore, we obtain

$$T \leq 2M \cdot |V_\Gamma| \cdot \sum_{p \text{ polychr.}} |h + \varphi_\Gamma(p)|.$$

For $i \in \{1, \dots, k\}$, let P_i be the set of polychromatic vertices p connected to p_i by γ_p . It follows from the assumptions (1) and (2) that

$$|P_i| \leq M^2 + 1,$$

since there are at most M paths coming through each of the edges incident to p_i , plus the constant path from p_i to itself. For $p \in P_i$, we can estimate

$$|\varphi_\Gamma(p) - \varphi_\Gamma(p_i)| \leq |\gamma_p|.$$

Putting the estimates together yields

$$\begin{aligned} |E_{\text{push}_h(\mathbf{D})}| &= Z + T \\ &\leq 2|E_{\mathbf{D}}| + 2M \cdot |V_\Gamma| \cdot \sum_{p \text{ polychr.}} |h + \varphi_\Gamma(p)| \\ &= 2|E_{\mathbf{D}}| + 2M \cdot |V_\Gamma| \cdot \sum_{i=1}^k \sum_{p \in P_i} |h + \varphi_\Gamma(p)| \\ &\leq 2|E_{\mathbf{D}}| + 2M \cdot |V_\Gamma| \cdot \sum_{i=1}^k \sum_{p \in P_i} (|\gamma_p| + |h + \varphi_\Gamma(p_i)|) \\ &\leq 2|E_{\mathbf{D}}| + 2M \cdot |V_\Gamma| \cdot M|E_{\mathbf{D}}| + 2M \cdot |V_\Gamma|(M^2 + 1) \sum_{i=1}^k |h + \varphi_\Gamma(p_i)| \\ &= (1 + M^2|V_\Gamma|)(|\mathbf{w}| \rho(\mathbf{D}) + |\mathbf{w}|) + 2M \cdot |V_\Gamma|(M^2 + 1) \sum_{i=1}^k |h + \varphi_\Gamma(p_i)|, \end{aligned}$$

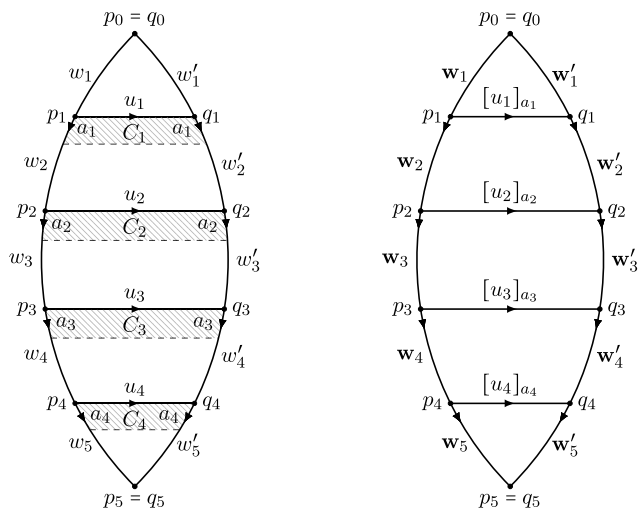


Figure 3: On the left is a minimal-area van Kampen diagram for ww'^{-1} , cut by a family of non-crossing corridors. On the right is the resulting coloured diagram for $\mathbf{w}\mathbf{w}'^{-1}$.

where the last equality follows from Lemma 3.11, which tells us $2|E_{\mathbf{D}}| = |\mathbf{w}| \cdot \rho(\mathbf{D}) + |\mathbf{w}'|$.

Finally, by Lemma 4.13, we have $|\text{push}_h(\mathbf{w})| \geq 2|\mathbf{w}| + 2\sum_{i=1}^k |h + \varphi_{\Gamma}(p_i)|$. Therefore, we can deduce

$$\begin{aligned} \rho(\text{push}_h(\mathbf{D})) &= \frac{2|E_{\text{push}_h(\mathbf{D})}|}{|\text{push}_h(\mathbf{w})|} - 1 \\ &\leq (1 + M^2|V_{\Gamma}|) \cdot (\rho(\mathbf{D}) + 1) + 2M(M^2 + 1)|V_{\Gamma}| - 1 \\ &\leq 3|V_{\Gamma}| \cdot M^3 + |V_{\Gamma}| \cdot M^2 + (1 + 2|V_{\Gamma}|) \cdot M, \end{aligned}$$

where the last estimate follows from the assumption (3). \square

4.4 Cutting along corridors. In this section, we develop a method to “cut” a coloured word along corridors to produce a coloured diagram.

Let \mathbf{w} and \mathbf{w}' be coloured words representing the same element in A_{Γ} , and let w and w' denote their underlying words, respectively. Let D be a minimal-area van Kampen diagram for ww'^{-1} , and let p_0 be the starting vertex of w , which we set as the base point of D . Let $k \in \mathbb{N}$, and for $i \in \{1, \dots, k-1\}$, let C_i be an a_i -corridor of D connecting an edge of w to an edge of w' . Assume that the corridors C_i are pairwise disjoint and are ordered by increasing distance from p_0 ; see the left diagram in Fig. 3.

Let γ_i be a side of C_i , and let p_i and q_i be its endpoints on w and w' , respectively. For consistency of notation, we also denote the base point p_0 as q_0 , and let $p_k = q_k$ be the other point where w and w' meet. Let w_i be the word read along the path from p_{i-1} to p_i , w'_i the word read along the path from q_{i-1} to q_i , and u_i the word read along γ_i from p_i to q_i .




Figure 4: Cutting a van Kampen diagram along corridors determined by an efficient coloured word on its boundary.

The decompositions $w = w_1 \cdots w_k$ and $w' = w'_1 \cdots w'_k$ naturally induce decompositions $\mathbf{w} = \mathbf{w}_1 \cdots \mathbf{w}_k$ and $\mathbf{w}' = \mathbf{w}'_1 \cdots \mathbf{w}'_k$ into coloured subwords, where w_i and w'_i are the underlying words of the coloured words \mathbf{w}_i and \mathbf{w}'_i , respectively. By the definition of corridor, all letters of u_i commute with a_i , so u_i is the underlying word of the monochromatic word $[u_i]_{a_i}$.

We can now construct a coloured diagram for $\mathbf{w}\mathbf{w}'^{-1}$ by cutting along the γ_i : this coloured diagram has k bounded regions, R_1, \dots, R_k , where the boundary word of R_i , when read starting from p_i , is $\mathbf{w}_i[u_i]_{a_i}\mathbf{w}'_i{}^{-1}[u_i^{-1}]_{a_i^{-1}}$, as shown in the right diagram in Fig. 3. We apply this construction to prove the following result.

LEMMA 4.29. *There exists a constant $C > 0$ such that the following holds. Let \mathbf{w} be an efficient coloured word, and let \mathbf{w}' be another coloured word representing the same element as \mathbf{w} . There exists a coloured diagram \mathbf{D} for $\mathbf{w}\mathbf{w}'^{-1}$ with $\rho(\text{push}_h(\mathbf{D})) \leq C$ for every $h \in \mathbb{Z}$, such that the label on every region of \mathbf{D} decomposes as a product $\mathbf{u}\mathbf{u}'^{-1}$, where \mathbf{u} is at most 5-coloured, and \mathbf{u}' is a subword of \mathbf{w}' .*

Proof. Let w and w' be the underlying words of \mathbf{w} and \mathbf{w}' , respectively. Assume that \mathbf{w} decomposes as $[w_1]_{a_1} \cdots [w_k]_{a_k}$, and let D be a van Kampen diagram for ww'^{-1} . Let $m = \lceil \frac{k}{3} \rceil$, and for each $i \in \{1, \dots, m-1\}$, let C_i be the corridor starting at the first edge of w_{3i+1} . Let γ_i be the side of C_i whose endpoint p_i on w separates w_{3i} and w_{3i+1} , and let q_i be the other endpoint of γ_i , and let u_i be the word read along γ_i from p_i to q_i . Since \mathbf{w} is efficient, the point q_i must belong to w' .

By Lemma 4.21, the corridors C_i do not cross each other, so we can construct a coloured diagram \mathbf{D} by cutting along the γ_i , where the corridor sides are labelled with $[u_i]_{a_{3i+1}}$; see Fig. 4. For each $i \in \{1, \dots, m\}$, the i th region of \mathbf{D} is labelled by

the word

$$[u_{i-1}]_{a_{3i-2}}^{-1} [w_{3i-2}]_{a_{3i-2}} [w_{3i-1}]_{a_{3i-1}} [w_{3i}]_{a_{3i}} [u_i]_{a_{3i+1}} \mathbf{w}'_i^{-1},$$

where we conventionally set $u_0 = u_m$ to be the empty word, and similarly for w_i when $k < i \leq 3m$. Then this word is of the form $\mathbf{u}\mathbf{u}'^{-1}$, where $\mathbf{u}' = \mathbf{w}'_i$, so it satisfies the desired properties.

It remains to establish the upper bound on the density of $\text{push}_h(\mathbf{D})$; we do so by checking the hypotheses for Proposition 4.28. By Lemma 4.21, part (3), we have $\sum_{i=1}^{m-1} |u_i| \leq 2|ww'|$, so the density of \mathbf{D} is at most 5. The only vertices of \mathbf{D} that may be polychromatic but not ∂ -polychromatic are the q_i ; this occurs when the two boundary edges incident to q_i have the same colour, which is different from the colour a_{3i+1} of the interior edge. In this case, we connect q_i to the ∂ -polychromatic vertex p_i via γ_i . Since the paths γ_i are all disjoint and all vertices of \mathbf{D} have degree at most 3, we can apply Proposition 4.28 and conclude. \square

5 Upper bounds

We now have all the necessary tools to establish upper bounds on the Dehn functions of Bestvina–Brady groups.

Theorem 5.1. *Let Γ be a finite simplicial graph such that the associated flag complex $\Delta(\Gamma)$ is simply connected, and let $\alpha \in \{3, 4\}$. If Γ does not have (\mathcal{D}_α) , then $\delta_{BB_\Gamma}(n) \preceq n^{\alpha-1}$.*

To prove Theorem 5.1, we need to show that, under the hypothesis above, every null-homotopic alternating word w in BB_Γ has an almost-flat area bounded above by $C \cdot |w|^{\alpha-1}$ for some constant $C > 0$. This is done by constructing an alternating diagram for w and applying our multi-step strategy to reduce the task to estimating the almost-flat area of simpler null-homotopic alternating words.

We begin by establishing upper bounds for these relatively simple words. Then, we explain how these pieces can be viewed as the building blocks of a general coloured diagram by means of cutting along corridors as discussed in Sect. 4.4.

5.1 Fundamental pieces. In the previous section, we showed that given a coloured diagram \mathbf{D} for a null-homotopic coloured word \mathbf{w} , its h -pushdown $\text{push}_h(\mathbf{D})$ is an alternating diagram for $\text{push}_h(\mathbf{w})$. This allows us to estimate the almost-flat area of $\text{push}_h(\mathbf{w})$ in terms of the almost-flat area of the bounded regions of $\text{push}_h(\mathbf{D})$.

We start with the case where \mathbf{w} belongs to one of the following three families of null-homotopic words, which we call *fundamental pieces*: these are *monochromatic words*, *coloured bigons* (Definition 5.4), and *coloured commutators* (Definition 5.5). These fundamental pieces are combinatorially easier to handle and can be used to estimate the almost-flat area of the h -pushdown of more complicated null-homotopic coloured words.

When dealing with the almost-flat area estimates of the fundamental pieces, we illustrate the power of working with diagrams: the filling of a coloured diagram, and thus of its pushdown, can be turned into a purely algebraic algorithm that tells us how to reduce the pushdown of the corresponding coloured word by applying relations at each step; see, for instance, Lemma 5.10, where estimating areas is equivalent to bounding the number of relations involved in proving an identity between words.

DEFINITION 5.2 (Support of a word). Let w be a word in S_Γ . The *support of w* , denoted by $\text{supp}(w)$, is the subgraph of Γ induced by the vertices v_s such that $s^{\pm 1}$ appears in w .

DEFINITION 5.3 (Palette of a word). The *palette of w* , denoted by $\text{Pal}(w)$, is the subgraph of Γ consisting of the intersection of the stars of all vertices in $\text{supp}(w)$. That is, the vertex set of $\text{Pal}(w)$ corresponds to the set of colours a such that $[w]_a$ is a well-defined coloured word.

DEFINITION 5.4 (Coloured bigon). A *coloured bigon* is a coloured word of the form $[w]_a[w]_b^{-1}$, where w is a word in S_Γ , and v_a and v_b are vertices of $\text{Pal}(w)$.

DEFINITION 5.5 (Coloured commutator). A *coloured commutator* is the commutator $[[w_1]_a, [w_2]_b]$ of two monochromatic words $[w_1]_a$ and $[w_2]_b$ such that $\{v_a\} \cup \text{supp}(w_2) \subset \text{Pal}(w_1)$ and $\{v_b\} \cup \text{supp}(w_1) \subset \text{Pal}(w_2)$.

By definition, the h -pushdown of coloured bigons and coloured commutators are null-homotopic alternating words. Observe that the h -pushdown of a coloured commutator is not in general the commutator of the h -pushdowns of the corresponding monochromatic coloured words.

To establish upper bounds for the areas of null-homotopic monochromatic words, coloured bigons, and coloured commutators, we first need some preliminary results. We begin by recalling a result from [Dis08], stated in our notation.

LEMMA 5.6 ([Dis08, Lemma 4.7]). *There exists a constant $K > 0$ such that the following holds. Let $h \in \mathbb{Z}$ and $\{v_a, v_b\} \in E_\Gamma$. The alternating word $\tau_{a,h}^{-1}\tau_{b,h}(ab^{-1})^{-h}$ is null-homotopic, and its almost-flat area satisfies*

$$\text{Area}_b(\tau_{a,h}^{-1}\tau_{b,h}(ab^{-1})^{-h}) \leq K|h|^2.$$

Throughout the rest of the section, we use the following notation.

NOTATION 5.7. Given two alternating words u and v in S_Γ that represent the same element in BB_Γ , we define the *almost-flat area of the identity $u =_{BB_\Gamma} v$* to be the almost-flat area of the null-homotopic word uv^{-1} . If A is the almost-flat area of the identity $u =_{BB_\Gamma} v$, we say that u can be rewritten as v with almost-flat area A .

PROPOSITION 5.8. *Let $h \in \mathbb{Z}$, $a \in S_\Gamma$, and $[w]_a$ be a null-homotopic monochromatic word. Then*

$$\text{Area}_b(\text{push}_h([w]_a)) \leq 3|w|^2.$$

Proof. Since the monochromatic word $[w]_a$ is null-homotopic, we have $\varphi_\Gamma(w) = 0$. It follows from the definition of the pushdown (Definition 4.10) that

$$\text{push}_h([w]_a) =_F \tau_{a,h} w(\Gamma a^{-1}) \tau_{a,h}^{-1}.$$

The word $w(\Gamma a^{-1})$ is null-homotopic in $BB_{\text{St}(v_a)}$. The latter is isomorphic to the right-angled Artin group $A_{\text{Lk}(v_a)}$ via the homomorphism defined by $sa^{-1} \mapsto s$ for all $v_s \in \text{Lk}(v_a)$, which sends $w(\Gamma a^{-1})$ to w . Since the area of w in $A_{\text{Lk}(v_a)}$ is at most $|w|^2$, it follows from Proposition 3.8 that the word $w(\Gamma a^{-1})$ admits an almost-flat van Kampen diagram of quadratic area.

To get the precise bound on the almost-flat area, note that a relation $[b, c] = 1$ in $A_{\text{Lk}(v_a)}$ is mapped to the null-homotopic alternating word $ba^{-1}cb^{-1}ac^{-1}$ in $BB_{\text{St}(v_a)}$, which admits an almost-flat diagram of area 3. Therefore, a van Kampen diagram for w can be transformed into an almost-flat van Kampen diagram for $w(\Gamma a^{-1})$ of area at most $3|w|^2$, as done in the proof of Proposition 3.8. \square

5.1.1 *Coloured bigons.* We now proceed to establish upper bounds for the almost-flat area of the h -pushdown of coloured bigons.

LEMMA 5.9. *There exists a constant $C > 0$ with the following property. Let $\mathbf{w} = [w]_a [w]_b^{-1}$ be a coloured bigon. Assume that v_a and v_b lie in the same connected component of $\text{Pal}(w)$. Then, for all $h \in \mathbb{Z}$, we have*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C \cdot (|w| + |h|)^2.$$

Proof. For any pair of adjacent vertices v_p and v_q in $\text{Pal}(w)$, we have that the alternating word

$$\tau_{p,h}^{-1} \text{push}_h([w]_p [w]_q^{-1}) \tau_{p,h} =_F w(\Gamma p^{-1}) \tau_{p,h+\varphi_\Gamma(w)}^{-1} \tau_{q,h+\varphi_\Gamma(w)} w(\Gamma q^{-1})^{-1} \tau_{q,h}^{-1} \tau_{p,h}$$

can be rewritten using Lemma 5.6 as

$$w(\Gamma p^{-1})(pq^{-1})^{h+\varphi_\Gamma(w)} w(\Gamma q^{-1})^{-1}(qp^{-1})^h.$$

This rewriting has almost-flat area at most $2K(|w| + |h|)^2$. This last word is a null-homotopic alternating word in the generating set of $BB_{\text{St}(v_p)}$, which is isomorphic to the right-angled Artin group $A_{\text{Lk}(v_p)}$. In particular, there is a constant $K' > 0$ such that it has almost-flat area bounded above by $K' \cdot (|w| + |h|)^2$. Putting everything together yields

$$\text{Area}_b(\text{push}_h([w]_p) \text{push}_h([w]_q^{-1})) \leq (2K + K') \cdot (|w| + |h|)^2.$$

Since a shortest path in $\text{Pal}(w)$ connecting v_a to v_b has length at most $|\mathbb{V}_\Gamma|$, by applying the above argument at most $|\mathbb{V}_\Gamma|$ times, we obtain

$$\begin{aligned} \text{Area}_b(\text{push}_h(\mathbf{w})) &= \text{Area}_b(\text{push}_h([w]_a) \text{push}_h([w]_b^{-1})) \\ &\leq |\mathbb{V}_\Gamma| \cdot (2K + K') \cdot (|w| + |h|)^2. \end{aligned} \quad \square$$




Figure 5: Filling of a coloured diagram for the coloured bigon when Λ decomposes as a join $\Lambda' * \Lambda''$.

In particular, if $\text{Pal}(w)$ is connected, then by Lemma 5.9, the almost-flat area of the pushdown of the coloured bigon is at most quadratic in its length.

LEMMA 5.10. *There exists a constant $C > 0$ with the following property. Let $\mathbf{w} = [w]_a[w]_b^{-1}$ be a coloured bigon. Assume that $\text{supp}(w)$ is reducible. Then, for all $h \in \mathbb{Z}$, we have*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C \cdot (|w| + |h|)^2.$$

Proof. Let $\Lambda := \text{supp}(w)$, and let $\Lambda = \Lambda' * \Lambda''$ be a join decomposition. Due to the decomposition $A_\Lambda = A_{\Lambda'} \times A_{\Lambda''}$, the element $g \in A_\Gamma$ represented by w can also be represented by a product of two words w' and w'' , whose supports lie respectively in Λ' and Λ'' , such that $|w'| + |w''| \leq |w|$.

This allows us to build a coloured diagram \mathbf{D} for $[w]_a[w]_b^{-1}$ with four bounded regions. Two of these regions are labelled by the null-homotopic monochromatic words $\mathbf{w}_1 = [w]_a[w'']_a^{-1}[w']_a^{-1}$ and $\mathbf{w}_2 = [w]_b[w'']_b^{-1}[w']_b^{-1}$; the other two are labelled by the coloured bigons $\mathbf{w}_3 = [w']_a[w']_b^{-1}$ and $\mathbf{w}_4 = [w'']_a[w'']_b^{-1}$; see Fig. 5.

Note that $\text{push}_h(\mathbf{D})$ is an alternating diagram for $\text{push}_h(\mathbf{w})$. We now proceed to estimate the almost-flat area of the four bounded regions in $\text{push}_h(\mathbf{D})$. These regions are labelled by the alternating words $u_i = \text{push}_{h_i}(\mathbf{w}_i)$ with $i \in \{1, 2, 3, 4\}$, where $h_1 = h_2 = h_3 = h$ and $h_4 = h + \varphi_\Gamma(w')$.

- (a) For the pushdown of the null-homotopic monochromatic words $\mathbf{w}_1 = [w]_a[w'']_a^{-1}[w']_a^{-1}$ and $\mathbf{w}_2 = [w]_b[w'']_b^{-1}[w']_b^{-1}$, we apply Proposition 5.8 to obtain $\text{Area}_b(u_i) \leq 12 \cdot |w|^2$ for $i \in \{1, 2\}$.
- (b) For the pushdown of the coloured bigon $\mathbf{w}_3 = [w']_a[w']_b^{-1}$, we can apply Lemma 5.9, since for every $v_c \in \Lambda' \subseteq \text{Pal}(w')$, we have $v_a, v_b \in \text{St}(v_c)$, and thus v_a and v_b lie in the same connected component of $\text{Pal}(w')$. The same argument holds for $\mathbf{w}_4 = [w'']_a[w'']_b^{-1}$. Therefore, we obtain $\text{Area}_b(u_i) \leq K \cdot (|w| + |h_i|)^2$ for $i \in \{3, 4\}$ and a sufficiently large constant $K > 0$.

Summing up, we find a constant $K' > 0$ such that $\text{Area}_b(u_i) \leq K' \cdot |u_i|^2$ for all $i \in \{1, \dots, 4\}$. For this, we are using the lower bound on $|u_i|$ given by Lemma 4.13.

There are now two different ways to conclude. The first is to apply Proposition 4.28 with $M = 4$, obtaining a uniform bound C_M on the density of $\text{push}_h(\mathbf{D})$, and then apply Lemma 3.12 to obtain

$$\text{Area}_b(\mathbf{w}) \leq K' \cdot C_M^2 \cdot |\text{push}_h(\mathbf{w})|^2 \leq C \cdot (|w| + |h|)^2$$

for large enough $C > 0$.

Alternatively, we can translate the coloured diagram \mathbf{D} into an algorithm to reduce the h -pushdown of $[w]_a[w]_b^{-1}$. Reducing this alternating word is equivalent to estimating the almost-flat area of the identity $\text{push}_h([w]_a) =_{BB_\Gamma} \text{push}_h([w]_b)$.

From (a), we get that the identity

$$\text{push}_{h_1}([w]_a) =_{BB_\Gamma} \text{push}_{h_1}([w']_a[w'']_a) \quad (5.1)$$

has almost-flat area at most $12 \cdot |w|^2$. Similarly, the almost-flat area of the identity

$$\text{push}_{h_2}([w]_b) =_{BB_\Gamma} \text{push}_{h_2}([w']_b[w'']_b) \quad (5.2)$$

is at most $12 \cdot |w|^2$. On the other hand, it follows from (b) that the identity

$$\text{push}_{h_3}([w']_a) =_{BB_\Gamma} \text{push}_{h_3}([w']_b) \quad (5.3)$$

holds with almost-flat area at most $K \cdot (|w| + |h_3|)^2$. Similarly, the almost-flat area of the identity

$$\text{push}_{h_4}([w'']_a) =_{BB_\Gamma} \text{push}_{h_4}([w'']_b) \quad (5.4)$$

is at most $K \cdot (|w| + |h_4|)^2$.

To assemble these identities together, we use the following property of the push-down

$$\text{push}_\ell(\mathbf{u}\mathbf{u}') =_F \text{push}_\ell(\mathbf{u}) \text{push}_{h+\varphi_\Gamma(u)}(\mathbf{u}'), \quad (5.5)$$

for any coloured words \mathbf{u} and \mathbf{u}' ; see Lemma 4.12, part (2). Recall $h_1 = h_2 = h_3 = h$ and $h_4 = h + \varphi_\Gamma(w')$. We obtain that the following identities

$$\begin{aligned} \text{push}_h([w]_a) &=_{BB_\Gamma} \text{push}_{h_1}([w']_a[w'']_a) && \text{by (5.1)} \\ &=_F \text{push}_{h_1}([w']_a) \text{push}_{h_4}([w'']_a) && \text{by (5.5)} \\ &=_{BB_\Gamma} \text{push}_{h_2}([w']_b) \text{push}_{h_4}([w'']_b) && \text{by (5.3) and (5.4)} \\ &=_F \text{push}_{h_2}([w']_b[w'']_b) && \text{by (5.5)} \end{aligned}$$




Figure 6: Filling of a coloured diagram for the coloured bigon $[w]_a[w]_b^{-1}$ with $w = s_1^{\epsilon_1} \cdots s_k^{\epsilon_k}$. The filling consists of k coloured bigons.

$$=_{BB_\Gamma} \text{push}_h([w]) \quad \text{by (5.2)}$$

hold with almost-flat area at most

$$\begin{aligned} & 24 \cdot |w|^2 + K \cdot \left((|w| + |h|)^2 + (|w| + |h + \varphi_\Gamma(w')|)^2 \right) \\ & \leq 24 \cdot |w|^2 + K \cdot \left((|w| + |h|)^2 + (2|w| + |h|)^2 \right) \\ & \leq 29 \cdot (K + 1) \cdot (|w| + |h|)^2. \end{aligned}$$

This second procedure yields in general a better constant, since we are estimating the areas of the regions of $\text{push}_0(\mathbf{D})$ more precisely without relying on the loose estimate of the density given by Proposition 4.28. \square

We saw that under certain conditions, the pushdown of a coloured bigon has quadratic area. We now show that in general we get a cubic upper bound.

LEMMA 5.11. *There exists a constant $C > 0$ with the following property. Let $\mathbf{w} = [w]_a[w]_b^{-1}$ be a coloured bigon. Then, for all $h \in \mathbb{Z}$, we have*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C \cdot (|w| + |h|)^3.$$

Proof. Let $k = |w|$, and write $w = s_1^{\epsilon_1} \cdots s_k^{\epsilon_k}$, where $s_i \in S_\Gamma$ and $\epsilon_i \in \{\pm 1\}$. By definition, we have

$$[w]_a = [s_1]_a^{\epsilon_1} \cdots [s_k]_a^{\epsilon_k}.$$

A coloured diagram for $\mathbf{w} = [w]_a[w]_b^{-1}$ can be constructed using k coloured bigons of the form $[s_i]_a^{\epsilon_i}[s_i]_b^{-\epsilon_i}$; see Fig. 6.

For every $i \in \{1, \dots, k\}$ and $\ell \in \mathbb{Z}$, the alternating word $\text{push}_\ell([s_i]_a^{\epsilon_i}[s_i]_b^{-\epsilon_i})$ is null-homotopic. Moreover, since $\text{Pal}(s_i^{\epsilon_i}) = \text{St}(s_i)$ is connected, Lemma 5.9 implies that

$$\text{Area}_b(\text{push}_\ell([s_i]_a^{\epsilon_i}[s_i]_b^{-\epsilon_i})) \leq K \cdot (1 + |\ell|)^2.$$

Therefore, for $h_0 = h$ and $h_i = h + \varphi_\Gamma(s_1 \cdots s_i)$, putting these k pieces together using the properties of the pushdown given in Lemma 4.12, we obtain the identities

$$\text{push}_h([w]_a) =_F \prod_{i=1}^k \text{push}_{h_{i-1}}([s_i]_a^{\epsilon_i}) =_{BB_\Gamma} \prod_{i=1}^k \text{push}_{h_{i-1}}([s_i]_b^{\epsilon_i}) =_F \text{push}_h([w]_b).$$

Thus, the almost-flat area of $\text{push}_h(\mathbf{w})$ satisfies

$$\begin{aligned}
\text{Area}_b(\text{push}_h(\mathbf{w})) &\leq K \cdot \sum_{i=1}^k (1 + |h_{i-1}|)^2 \\
&= K \cdot \sum_{i=1}^k 1 + 2|h_{i-1}| + |h_{i-1}|^2 \\
&\leq K \cdot |w| \left(1 + 2(|w| + |h|) + (|w| + |h|)^2\right) \\
&\leq 4K \cdot (|w| + |h|)^3. \quad \square
\end{aligned}$$

We now have all the tools to prove that coloured bigons satisfy Theorem 5.1.

PROPOSITION 5.12. *There exists a constant $C > 0$ with the following property. Assume that $\alpha \in \{3, 4\}$ and that Γ does not have (\mathcal{D}_α) . Let $\mathbf{w} = [w]_a [w]_b^{-1}$ be a coloured bigon. Then, for all $h \in \mathbb{Z}$, we have*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C \cdot |\text{push}_h(\mathbf{w})|^{\alpha-1}.$$

Proof. Observe that the case $\alpha = 4$ follows directly from Lemma 5.11, so we assume $\alpha = 3$. Since Γ does not have property (\mathcal{D}_3) , every maximal join is not essentially 2-reducible.

We may assume that v_a and v_b are not adjacent and that $\text{Pal}(w)$ is disconnected; otherwise, we could conclude by Lemma 5.9. In particular, the vertices v_a and v_b do not belong to $\text{supp}(w)$. If $\text{supp}(w)$ were reducible, we would also be done by Lemma 5.10, so we assume that $\text{supp}(w)$ is irreducible.

Since the subgraph $\text{supp}(w) * (\text{Pal}(w) \setminus \text{supp}(w))$ is reducible, it is contained in a maximal reducible subgraph $\Lambda \subseteq \Gamma$. Consider the decomposition $\Lambda = \Lambda_1 * \Lambda_2 * \cdots * \Lambda_k$, where $k \geq 2$ and Λ_i is irreducible for all $i \in \{1, \dots, k\}$.

Since $\text{supp}(w) \subseteq \Lambda$ is irreducible, it is contained in some Λ_i , say Λ_1 . Note that $\text{Pal}(w) = (\text{Pal}(w) \cap \Lambda_1) * \Lambda_2 * \cdots * \Lambda_k$. Since $\text{Pal}(w)$ is disconnected, it follows that $k = 2$ and $\text{Pal}(w) \cap \Lambda_1 = \emptyset$. Thus, the subgraph Λ_2 has at least two vertices v_a and v_b . Since Λ is not essentially 2-reducible, we have $\Lambda_1 = \{v_c\}$. In this case, we obtain $v_c \in \text{Pal}(w) \cap \Lambda_1$, which is a contradiction. \square

5.1.2 Coloured commutators. We now give an upper bound on the almost-flat area of the h -pushdown of a coloured commutator.

LEMMA 5.13. *There exists a constant $C > 0$ such that the following holds. Assume that $\alpha \in \{3, 4\}$ and that Γ does not have (\mathcal{D}_α) . Let $\mathbf{w} = [[w_1]_a, [w_2]_b]$ be a coloured commutator. If $\text{Pal}(w_1) \cap \text{Pal}(w_2) \neq \emptyset$, then for all $h \in \Gamma$, we have*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C \cdot (|w_1| + |w_2| + |h|)^{\alpha-1}.$$

Proof. Pick $v_c \in \text{Pal}(w_1) \cap \text{Pal}(w_2)$. A coloured diagram \mathbf{D} for \mathbf{w} consists of five bounded regions labeled by $\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3, \mathbf{w}_4$, and \mathbf{w}_5 . The first four coloured words




Figure 7: Filling of a coloured diagram for the coloured commutator $\mathbf{w} = [[w_1]_a, [w_2]_b]$ in the case where $v_c \in \text{Pal}(w_1) \cap \text{Pal}(w_2)$.

are coloured bigons of the form $\mathbf{w}_1 = [w_1]_a[w_1]_c^{-1}$, $\mathbf{w}_2 = [w_2]_b[w_2]_c^{-1}$, $\mathbf{w}_3 = [w_1]_a^{-1}[w_1]_c$, and $\mathbf{w}_4 = [w_2]_b^{-1}[w_2]_c$, respectively, while the last one is a null-homotopic monochromatic word $\mathbf{w}_5 = [[w_1]_c, [w_2]_c]$; see Fig. 7.

We now estimate the almost-flat area of the bounded regions in $\text{push}_h(\mathbf{D})$, which are labeled by the alternating word $u_i = \text{push}_{h_i}(\mathbf{w}_i)$ for $i \in \{1, \dots, 5\}$, where $h_1 = h_5 = h$, $h_2 = h + \varphi_\Gamma(w_1)$, $h_3 = h + \varphi_\Gamma(w_1 w_2)$, and $h_4 = h + \varphi_\Gamma(w_2)$.

On the one hand, for all the coloured bigons, it follows from Proposition 5.12 that there exists $C_1 > 0$ such that $\text{Area}_b(u_i) \leq C_1 \cdot |u_i|^{\alpha-1}$ for all $i \in \{1, \dots, 4\}$. On the other hand, it follows from Proposition 5.8 that for the null-homotopic monochromatic word \mathbf{w}_5 , we have $\text{Area}_b(u_5) \leq 12 \cdot (|w_1| + |w_2|)^2$.

All together, we find a constant $C_3 > 0$ such that for all $i \in \{1, \dots, 5\}$, we have

$$\text{Area}_b(u_i) \leq C_3 \cdot |u_i|^{\alpha-1}.$$

We conclude by first applying Proposition 4.28 with $M = 4$, obtaining a uniform bound C_M on the density of $\text{push}_h(\mathbf{D})$, and then by Lemma 3.12, we obtain

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C_3 \cdot C_M^2 \cdot |\text{push}_h(\mathbf{w})|^{\alpha-1} \leq C \cdot (|w_1| + |w_2| + |h|)^{\alpha-1}$$

for sufficiently large $C > 0$. \square

LEMMA 5.14. *There exists a constant $C > 0$ such that the following holds. Assume that $\alpha \in \{3, 4\}$ and that Γ does not have (\mathcal{D}_α) . Let $\mathbf{w} = [[w_1]_a, [w_2]_b]$ be a coloured commutator. If either $\text{supp}(w_1)$ or $\text{supp}(w_2)$ is reducible, then for all $h \in \mathbb{Z}$, we have*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C \cdot (|w_1| + |w_2| + |h|)^{\alpha-1}.$$

Proof. Suppose without loss of generality that $\text{supp}(w_1)$ decomposes as a join $\Lambda' * \Lambda''$. Set $\Lambda = \text{supp}(w_1)$. In this case, due to the decomposition $A_\Lambda = A_{\Lambda'} \times A_{\Lambda''}$, the element $g \in A_\Gamma$ represented by w can also be represented by a product of two words w' and w'' whose supports lie respectively in Λ' and Λ'' , and such that $|w'| + |w''| \leq |w|$. Thus, a coloured diagram \mathbf{D} for the alternating word $\text{push}_h(\mathbf{w})$ can be constructed with four bounded regions whose labels are coloured words \mathbf{w}_1 , \mathbf{w}_2 , \mathbf{w}_3 , and \mathbf{w}_4 . Two of these are null-homotopic monochromatic words $\mathbf{w}_1 = [w_1]_a[w'']_a^{-1}[w']_a^{-1}$ and




Figure 8: Filling of a coloured diagram for the coloured commutator $[[w_1]_a, [w_2]_b]$ in the case where $\text{supp}(w_1)$ is reducible.

$\mathbf{w}_3 = [w']_a [w''_a] [w_1]_a^{-1}$, and the others are coloured commutators $\mathbf{w}_2 = [[w''_a], [w_2]_b]$ and $\mathbf{w}_4 = [[w']_a, [w_2]_b]$; see Fig. 8.

For $i \in \{1, \dots, 4\}$, let $u_i = \text{push}_{h_i}(\mathbf{w}_i)$ be the alternating word that labels the corresponding bounded region of $\text{push}_h(\mathbf{D})$, where $h_1 = h_4 = h$, $h_2 = h + \varphi_\Gamma(w_1)$, and $h_3 = h + \varphi_\Gamma(w_2)$. By Proposition 5.8, for $i \in \{1, 3\}$, we have $\text{Area}_b(u_i) \leq 12 \cdot |w_1|^2$.

Since $\Lambda = \text{supp}(w_1)$ decomposes as a join $\Lambda' * \Lambda''$, the intersection $\text{Pal}(w'') \cap \text{Pal}(w_2)$ contains Λ' so it is non-empty. Therefore, it follows from Lemma 5.13 that

$$\text{Area}_b(u_2) \leq C_2 \cdot (|w''| + |w_2| + |h_2|)^{\alpha-1}.$$

for sufficiently large $C_2 > 0$. By symmetry, the same holds for u_4 , namely, $\text{Area}_b(u_4) \leq C_2 \cdot (|w'| + |w_2| + |h_4|)^{\alpha-1}$.

All together, we have $\text{Area}_b(u_i) \leq C_3 \cdot |u_i|^{\alpha-1}$ for sufficiently large $C_3 > 0$. A straightforward application of Proposition 4.28 with $M = 3$ gives us a uniform bound C_M on the density of the pushdown of the coloured diagram for \mathbf{w} . Thus, it follows from Lemma 3.12 that

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C_3 \cdot C_M^2 \cdot |\text{push}_h(\mathbf{w})|^{\alpha-1} \leq C \cdot (|w_1| + |w_2| + |h|)^{\alpha-1}$$

for sufficiently large $C > 0$. □

LEMMA 5.15. *There exists a constant $C > 0$ such that the following holds. Let $\mathbf{w} = [[w_1]_a, [w_2]_b]$ be a coloured commutator. Assume that Λ_1 and Λ_2 are subgraphs of Γ spanning a join with $\text{supp}(w_1) \cup \{v_b\} \subseteq \Lambda_1$ and $\text{supp}(w_2) \cup \{v_a\} \subseteq \Lambda_2$, such that either Λ_1 or Λ_2 is connected. Then for all $h \in \mathbb{Z}$, we have*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C \cdot (|w_1| + |w_2| + |h|)^3.$$

Proof. Assume without loss of generality that Λ_1 is connected. Let $w_1 = s_1^{\epsilon_1} \dots s_k^{\epsilon_k}$ with $\epsilon_i \in \{\pm 1\}$ for all $i \in \{1, \dots, k\}$. A coloured diagram for the coloured commutator $\mathbf{w} = [[w_1]_a, [w_2]_b]$ can be constructed using a total of $4k + 1$ bounded regions; see Fig. 9:




Figure 9: Filling of the coloured diagram for the coloured commutator $[[w_1]_a, [w_2]_b]$, where $w_1 = s_1^{\epsilon_1} s_2^{\epsilon_2} \dots s_k^{\epsilon_k}$.

- $k + 1$ bounded regions whose labels are coloured bigons of the form $[w_2]_b [w_2]_{s_1}^{-1}$ (corresponding to the top region in Fig. 9), of the form $[w_2]_{s_i} [w_2]_{s_{i+1}}^{-1}$ for $i \in \{1, \dots, k - 1\}$, and of the form $[w_2]_{s_k} [w_2]_b^{-1}$ (corresponding to the bottom region in Fig. 9);
- $2k$ bounded regions whose labels are coloured bigons of the form $[s_i]_a^{\epsilon_i} [s_i]_{s_i}^{-\epsilon_i}$ and $[s_i]_{s_i}^{\epsilon_i} [s_i]_a^{-\epsilon_i}$ (corresponding to the regions on the left and right sides in Fig. 9, respectively);
- k bounded regions whose labels are null-homotopic monochromatic words of the form $[[s_i]_{s_i}^{\epsilon_i}, [w_2]_{s_i}]$ for $i \in \{1, \dots, k\}$ (corresponding to the regions with four sides in Fig. 9).

Let $h_0 = h$ and $h_i = h + \varphi_\Gamma(s_1^{\epsilon_1} \dots s_i^{\epsilon_i})$ for $i \in \{1, \dots, k\}$. The pushdown of each of these $4k + 1$ pieces can be filled with at most quadratic almost-flat area:

- Since Λ_1 is connected, it follows from Lemma 5.9 that there exists $C_1 > 0$ such that the h -pushdown of the coloured bigon $[w_2]_b [w_2]_{s_1}^{-1}$ has almost-flat area at most $C_1 \cdot (|w_2| + |h|)^2$, the h_i -pushdown of the coloured bigon $[w_2]_{s_i} [w_2]_{s_{i+1}}^{-1}$ has almost-flat area at most $C_1 \cdot (|w_2| + |h_i|)^2$ for $i \in \{1, \dots, k - 1\}$, and the h_k -pushdown of the coloured bigon $[w_2]_{s_k} [w_2]_b^{-1}$ has almost-flat area at most $C_1 \cdot (|w_2| + |h_k|)^2$.
- For all $i \in \{1, \dots, k\}$, it follows from Lemma 5.9 that there exists $C_2 > 0$ such that the h_i -pushdown of the coloured bigon $[s_i]_a^{\epsilon_i} [s_i]_{s_i}^{-\epsilon_i}$ has almost-flat area at most $C_2 \cdot (1 + |h_{i-1}|)^2$. Similarly, we have that the almost-flat area of the $(h_i + \varphi_\Gamma(w_2))$ -pushdown of the coloured bigon $[s_i]_{s_i}^{\epsilon_i} [s_i]_a^{-\epsilon_i}$ is at most $C_2 \cdot (1 + |h_{i-1} + \varphi_\Gamma(w_2)|)^2$ for all $i \in \{1, \dots, k\}$.
- Finally, Proposition 5.8 implies that for all $i \in \{1, \dots, k\}$, the h_i -pushdown of the null-homotopic monochromatic word $[[s_i]_{s_i}^{\epsilon_i}, [w_2]_{s_i}]$ has almost-flat area at most $12 \cdot (1 + |w_2|)^2$.

By setting $C_3 = \max\{C_1, C_2, 16\}$, we obtain that $\text{Area}_b(\text{push}_h(\mathbf{w}))$ is bounded above by

$$C_3 \left(\sum_{i=0}^k (|w_2| + |h_i|)^2 + \sum_{i=1}^k (1 + |h_{i-1}|)^2 + (1 + |h_{i-1} + \varphi_\Gamma(w_2)|)^2 + (1 + |w_2| + |h_i|)^2 \right).$$

Since for all $i \in \{0, \dots, k\}$, we have $|h_i|, |h_i + \varphi_\Gamma(w_2)| \leq |w_1| + |w_2| + |h|$, and $k = |w_1|$, we obtain

$$\begin{aligned} \text{Area}_b(\text{push}_h(\mathbf{w})) &\leq C_3 \left(|w_1|(|w_2| + |w_1| + |h|)^2 + 3|w_1|(1 + |w_1| + |w_2| + |h|)^2 \right) \\ &\leq C (|w_1| + |w_2| + |h|)^3 \end{aligned}$$

for sufficiently large $C > 0$. □

We now conclude this section showing that coloured commutators satisfy Theorem 5.1.

PROPOSITION 5.16. *There exists a constant $C > 0$ with the following property. Assume that $\alpha \in \{3, 4\}$ and that Γ does not have (\mathcal{D}_α) . Let $h \in \mathbb{Z}$. The almost-flat area of the h -pushdown of the coloured commutator $\mathbf{w} = [[w_1]_a, [w_2]_b]$ satisfies*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C \cdot |\text{push}_h(\mathbf{w})|^{\alpha-1}.$$

Proof. Let $\mathbf{w} = [[w_1]_a, [w_2]_b]$ be a coloured commutator. Since $\text{supp}(w_1) \cup \{v_b\} \subseteq \text{Pal}(w_2)$ and $\text{supp}(w_2) \cup \{v_a\} \subseteq \text{Pal}(w_1)$, we may assume that

$$(\text{supp}(w_1) \cup \{v_b\}) \cap (\text{supp}(w_2) \cup \{v_a\}) = \emptyset,$$

otherwise, $\text{Pal}(w_1) \cap \text{Pal}(w_2) \neq \emptyset$, and the statement would follow from Lemma 5.13.

In particular, the subgraph induced by $\text{supp}(w_1) \cup \text{supp}(w_2) \cup \{v_a, v_b\}$ is reducible, and is contained in a maximal reducible subgraph $\Lambda \subseteq \Gamma$. Let $\Lambda = \Lambda_1 * \dots * \Lambda_k$ for $k \geq 2$, with each Λ_i irreducible.

Suppose that $\text{supp}(w_1) \cup \{v_b\}$ is reducible. Then, either $\text{supp}(w_1)$ is reducible or $v_b \in \text{Pal}(w_1)$, and we conclude by Lemma 5.14 or Lemma 5.13, respectively. So, we can assume that $\text{supp}(w_1) \cup \{v_b\}$ is irreducible and, symmetrically, that the same holds for $\text{supp}(w_2) \cup \{v_a\}$. This implies that $\text{supp}(w_1) \cup \{v_a\} \subseteq \Lambda_i$ and $\text{supp}(w_2) \cup \{v_a\} \subseteq \Lambda_j$ for some $i, j \in \{1, \dots, k\}$. We can also assume that $k = 2$ and that $\{i, j\} = \{1, 2\}$. If this were not the case, we would have $\text{Pal}(w_1) \cap \text{Pal}(w_2) \neq \emptyset$, and the statement would follow from Lemma 5.13. Therefore, we have $\Lambda = \Lambda_1 * \Lambda_2$ with $\text{supp}(w_1) \cup \{v_b\} \subseteq \Lambda_1$ and $\text{supp}(w_2) \cup \{v_a\} \subseteq \Lambda_2$.

If Γ does not have (\mathcal{D}_4) , then $\Delta(\Lambda)$ is simply connected and either Λ_1 or Λ_2 is connected. In this case, it follows from Lemma 5.15 that $\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C \cdot |\text{push}_h(\mathbf{w})|^3$ for some $C > 0$. If Γ does not have (\mathcal{D}_3) , then Λ is not essentially 2-reducible, so either Λ_1 or Λ_2 is a single vertex. In this case, the desired almost-flat area estimate now follows from Lemma 5.13. □

5.2 Putting the pieces together. Now that we have proven the upper bounds for the pushdown of the fundamental pieces, we combine them to obtain upper bounds for progressively more complicated words. We begin with the pushdown of null-homotopic k -coloured words whose underlying words admit van Kampen diagrams with well-behaved corridors. For now, we also allow constants to depend on the number of colours.

LEMMA 5.17. *For every $k \geq 2$, there exists a constant $C_k > 0$ such that the following holds. Let $\alpha \in \{3, 4\}$ and assume that Γ does not have (\mathcal{D}_α) . Let $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ be a null-homotopic k -coloured word with $a_k \neq a_1$, and $h \in \mathbb{Z}$. Assume that the underlying word $w = w_1 \cdots w_k$ admits a minimal-area van Kampen diagram D such that, for all $i \in \{1, \dots, k\}$, each corridor starting from an edge of w_i ends in an edge of either w_{i-2} , w_{i-1} , w_i , w_{i+1} , or w_{i+2} (indices intended mod k). Then*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C_k \cdot |\text{push}_h(\mathbf{w})|^{\alpha-1}.$$

Proof. We aim to construct a coloured diagram for \mathbf{w} whose bounded regions are fundamental pieces, and whose density can be bounded in terms of the number of colours.

For every $i \in \{1, \dots, k\}$, we choose a freely reduced coloured word $[w'_i]_{a_i}$ such that $w_i =_{A_\Gamma} w'_i$, and we choose a van Kampen diagram D' for $w'_1 \cdots w'_k$ such that corridors starting from an edge in w'_i end in an edge of w'_{i-2} , w'_{i-1} , w'_i , w'_{i+1} or w'_{i+2} . We make these choices so that the area of D' is the smallest possible. Note that it is always possible to make such choices, as one can always take $w'_i = w_i$ and $D' = D$.

Two corridors in D' that originate from the same w'_i cannot intersect. Otherwise, consider an innermost pair of intersecting corridors starting from w'_i ; these corridors must begin at adjacent letters. Exchanging those letters would yield a word w''_i such that $w'_1 \cdots w''_i \cdots w'_k$ admits a diagram with smaller area, a contradiction. Similarly, a corridor cannot connect two edges of the same w'_i . Since corridors starting from the same w'_i do not intersect, we can decompose each w'_i as

$$w'_i = u_i^{<} u_i^{<<} u_i^{>>} u_i^{>}$$

in such a way that the corridors of D' starting in $u_i^{>}$ end in $u_{i+1}^{<}$, and those starting in $u_i^{>>}$ end in $u_{i+2}^{<<}$. In particular, $u_i^{>} = (u_{i+1}^{<})^{-1}$ and $u_i^{>>} = (u_{i+2}^{<<})^{-1}$, and all letters of $u_i^{>>}$ commute with all letters of $u_{i+1}^{<<}$.

A coloured diagram \mathbf{D} for \mathbf{w} can be constructed, as shown in Fig. 10, using the following fundamental pieces:

- k monochromatic regions labelled by

$$[w_i]_{a_i} [w'_i]_{a_i}^{-1} = [w_i]_{a_i} [u_i^{>}]_{a_i}^{-1} [u_i^{>>}]_{a_i}^{-1} [u_i^{<<}]_{a_i}^{-1} [u_i^{<}]_{a_i}^{-1};$$

- k bigons labelled by $[u_i^{<}]_{a_i} [u_{i-1}^{>}]_{a_{i-1}}^{-1}$;
- k bigons labelled by $[u_{i+1}^{<<}]_{a_{i+1}} [u_{i-1}^{>>}]_{a_{i-1}}$;
- k coloured commutators $[[u_i^{>>}]_{a_i}, [u_{i+1}^{<<}]_{a_{i+1}}]$.




Figure 10: Filling of a coloured k -gon with coloured bigons and coloured commutators.

In this way, we obtain a coloured diagram \mathbf{D} which has $3k + 1$ polychromatic vertices, including the k ∂ -polychromatic vertices that are the vertices of the k -gon. By applying Proposition 4.28 with $M = 3k + 1$, the density of $\text{push}_h(\mathbf{D})$ is bounded above by some constant $C'_k > 0$.

Since every word w labelling the bounded regions of $\text{push}_h(\mathbf{D})$ has almost-flat area bounded above by $C \cdot |w|^{\alpha-1}$ by Propositions 5.12 and 5.16, for some constant $C > 0$, we conclude by applying Lemma 3.12. \square

By an induction argument, we can remove the hypothesis on the corridors in Lemma 5.17.

PROPOSITION 5.18. *Let $\alpha \in \{3, 4\}$ and assume that Γ does not have (\mathcal{D}_α) . For every $k \geq 1$, there exists a constant C_k such that for every null-homotopic k -coloured word \mathbf{w} and every $h \in \mathbb{Z}$, we have*

$$\text{Area}_b(\text{push}_h(\mathbf{w})) \leq C_k \cdot |\text{push}_h(\mathbf{w})|^{\alpha-1}.$$

Proof. We proceed by induction on k . The case $k = 1$ was handled in Proposition 5.8. Suppose that the statement holds for $\ell < k$.

By a cyclic conjugation, we may assume $\mathbf{w} = [w_1]_{a_1} \cdots [w_k]_{a_k}$ with $a_k \neq a_1$. Let D be a van Kampen diagram for the underlying word $w = w_1 \cdots w_k$. Assume that there is a corridor in D connecting a letter of w_i with a letter of w_j , with $3 \leq |i - j| \leq |n - 3|$; otherwise, we are under the hypotheses of Lemma 5.17 and can conclude directly.

By cutting along the boundary of the corridor, we obtain a coloured diagram \mathbf{D} with only two regions, labelled by a $(|i - j| + 2)$ -coloured word and a $(n - |i - j| + 2)$ -coloured word respectively; both have fewer than k colours. We have $\rho(\mathbf{D}) \leq 2$, and by Proposition 4.28, there is a constant $C > 0$ independent of k such that $\rho(\text{push}_h(\mathbf{D})) \leq C$.

By combining the inductive hypothesis with Lemma 3.12, we obtain that the statement holds with $C_k := \max\{C_\ell : \ell < k\} \cdot C^{\alpha-1}$. \square

Finally, we remove the dependence of the multiplicative constant C_k on the number of colours. The key to this is the following lemma.

LEMMA 5.19. *There exists a constant $C > 0$ such that the following holds. Let \mathbf{w}_1 , \mathbf{w}_2 , and \mathbf{w}_3 be efficient coloured words such that $\mathbf{w} = \mathbf{w}_1 \mathbf{w}_2 \mathbf{w}_3$ is null-homotopic. Then there exists a coloured diagram \mathbf{D} for \mathbf{w} whose bounded regions are at most 15-coloured, and such that $\rho(\text{push}_h(\mathbf{D})) \leq C$.*

Proof. We argue using Lemma 4.29 to construct a coloured diagram for the null-homotopic coloured word $\mathbf{w}_1 \mathbf{w}'^{-1}$, where $\mathbf{w}'^{-1} = \mathbf{w}_2 \mathbf{w}_3$. This yields a coloured diagram \mathbf{D}' such that $\rho(\text{push}_h(\mathbf{D}')) \leq C'$ for some constant $C' > 0$, and whose bounded regions are labelled by words of the form $\mathbf{u}'_1 \mathbf{w}'_2 \mathbf{w}'_3$, where \mathbf{u}'_1 is at most 5-coloured, and \mathbf{w}'_2 and \mathbf{w}'_3 are (possibly empty) subwords of \mathbf{w}_2 and \mathbf{w}_3 , respectively; they are, in particular, efficient.

By Lemma 3.14, it suffices to show that we can find a coloured diagram for the labels of the bounded regions of \mathbf{D}' , whose regions are at most 15-coloured and such that the density of the h -pushdown of \mathbf{D}' is at most C for some constant $C > 0$ independent of h .

To do so, we perform two more iterations of Lemma 4.29. First, we apply it to each bounded region of \mathbf{D}' to obtain coloured diagrams \mathbf{D}'' for the words $\mathbf{w}'_2 \cdot (\mathbf{w}'_3 \mathbf{u}'_1)$, such that $\rho(\text{push}_{h'}(\mathbf{D}'')) \leq C'$ for every $h' \in \mathbb{Z}$, and whose regions are labelled by words of the form $\mathbf{u}''_2 \mathbf{w}''_3 \mathbf{u}''_1$, where \mathbf{u}''_2 is at most 5-coloured, \mathbf{w}''_3 is a subword of \mathbf{w}'_3 , and \mathbf{u}''_1 is a subword of \mathbf{u}'_1 and therefore is at most 5-coloured. We replace each bounded region of \mathbf{D}' with the corresponding coloured diagram \mathbf{D}'' .

Second, we apply Lemma 4.29 once more to each bounded region of the coloured diagram constructed above, obtaining coloured diagrams \mathbf{D}''' for the words $\mathbf{w}''_3 \cdot (\mathbf{u}''_1 \mathbf{u}''_2)$ such that $\rho(\text{push}_{h''}(\mathbf{D}''')) \leq C''$ for every $h'' \in \mathbb{Z}$, and whose regions are labelled by words of the form $\mathbf{u}'''_3 \mathbf{u}'''_1 \mathbf{u}'''_2$, where \mathbf{u}'''_3 is at most 5-coloured, \mathbf{u}'''_1 is a subword of \mathbf{u}''_1 , and \mathbf{u}'''_2 is a subword of \mathbf{u}''_2 . We again replace each bounded region with the corresponding coloured diagram \mathbf{D}''' .

In this way, we obtain a coloured diagram for \mathbf{w} whose regions are labelled by words that are at most 15-coloured and whose pushdown has density at most C'^3 by Lemma 3.14. \square

We conclude the section by establishing the upper bound for the Dehn function of BB_Γ .

Proof of Theorem 5.1. For every $g \in BB_\Gamma$, choose a word u_g representing g that is the 0-pushdown of an efficient coloured word \mathbf{w}_g . By Proposition 4.19, the length of u_g is bounded linearly in terms of $\|g\|_b$, so we may apply Lemma 3.15. Now, it suffices to estimate the area of the word $u = u_{g_1} u_{g_2} u_{g_3}$, where $g_1, g_2, g_3 \in BB_\Gamma$ such that $g_1 g_2 g_3 = 1$.

To do so, consider the coloured diagram \mathbf{D} produced by Lemma 5.19 for the null-homotopic word $\mathbf{w} = \mathbf{w}_{g_1} \mathbf{w}_{g_2} \mathbf{w}_{g_3}$. Its 0-pushdown is an alternating diagram for the null-homotopic word

$$\text{push}_0(\mathbf{w}) = \text{push}_0(\mathbf{w}_{g_1}) \text{push}_0(\mathbf{w}_{g_2}) \text{push}_0(\mathbf{w}_{g_3}) = u;$$

here, we use the fact $\varphi_\Gamma(\mathbf{w}_{g_i}) = 0$ and apply Lemma 4.12. The density $\rho(\text{push}_0(\mathbf{D}))$ is bounded above by a constant C , and the regions of \mathbf{D} are labelled by alternating words u_i , which are the h_i -pushdown of coloured words that are at most 15-coloured. By Proposition 5.18, each u_i admits an almost-flat diagram of area at most $C_{15} \cdot |u_i|^{\alpha-1}$. We conclude by applying Lemma 3.12. \square

6 Lower bounds

We now turn our attention to the lower bounds. The aim of this section is to prove the following.

Theorem 6.1. *Let Γ be a finite simplicial graph such that the associated flag complex $\Delta(\Gamma)$ is simply connected, and let $\alpha \in \{3, 4\}$. If Γ has property (\mathcal{D}_α) , then $\delta_{BB_\Gamma}(n) \asymp n^\alpha$.*

To establish cubic and quartic lower bounds for the Dehn functions in Theorem 6.1, we need to produce a family of null-homotopic alternating words w_n such that we can find a cubic, respectively quartic, lower bound for the area of almost-flat van Kampen diagrams for w_n . To obtain these lower bounds, we analyse the behaviour of corridors and annuli inside such diagrams.

Recall that both corridors and annuli inherit a natural orientation: our convention is that the positive orientation of an a -corridor or a -annulus is the one such that the a -edges transverse to it are oriented from left to right.

Let D be an almost-flat van Kampen diagram over the standard presentation \mathcal{P}_Γ of A_Γ . Removing the interior of a corridor or annulus subdivides D into two connected components. We say that a corridor C *separates* two vertices p and q if they belong to different connected components after removing the interior of C , and that an annulus A *encloses* a vertex p if p is contained in the connected component that does not contain the boundary of D .

Let p_0 denote the base point of D . For every vertex $p \in D$, we define

$$\kappa_+(p) := |\{C : C \text{ is a corridor separating } p_0 \text{ and } p, \text{ with } p \text{ on the right of } C\}|$$

$$\kappa_-(p) := |\{C : C \text{ is a corridor separating } p_0 \text{ and } p, \text{ with } p \text{ on the left of } C\}|$$

$$\alpha_+(p) := |\{A : A \text{ is an annulus oriented clockwise enclosing } p\}|$$

$$\alpha_-(p) := |\{A : A \text{ is an annulus oriented counterclockwise enclosing } p\}|$$

and set

$$\alpha(p) := \alpha_+(p) - \alpha_-(p) \quad \text{and} \quad \kappa(p) = \kappa_+(p) - \kappa_-(p).$$

LEMMA 6.2. *For every vertex $p \in D$, we have $\varphi_\Gamma(p) = \alpha(p) + \kappa(p)$.*

Proof. Consider an arbitrary combinatorial path connecting p_0 and p . Every edge in the path crosses exactly one corridor or annulus. The height increases by one when crossing a corridor from left to right, when entering an annulus oriented clockwise, or when exiting an annulus oriented counterclockwise. Conversely, the height decreases by one when crossing a corridor from right to left, when exiting an annulus oriented clockwise, or when entering an annulus oriented counterclockwise. From this, we get the statement. \square

COROLLARY 6.3. *For every vertex $p \in D$, we have $|\alpha(p)| \geq |\kappa(p)| - 2$.*

Proof. The result follows from Lemma 6.2 and the fact that every vertex p in an almost-flat van Kampen diagram satisfies $\varphi_\Gamma(p) \in \{0, 1, 2\}$. \square

Thus, Corollary 6.3 tells us that in an almost-flat van Kampen diagram, a vertex p with large $|\kappa(p)|$ must be enclosed by many annuli. We show that, under certain assumptions, these annuli produce a lot of crossings, which in turn yield a large lower bound for the area of D .

DEFINITION 6.4. Let p be a vertex in D . We denote by $x(p)$ the set of pairs (A, A') , where A is an annulus enclosing p , A' is an annulus *not* enclosing p , and A crosses A' .

LEMMA 6.5. *Let p and q be vertices of D , and let γ be a (combinatorial) path connecting them. Assume that no annulus encloses both p and q , and that every vertex in γ is enclosed by at least k annuli of D . Then*

$$|x(p)| \geq \frac{k(k+1)}{2}.$$

Proof. Assume that $\gamma: [0, \ell] \rightarrow D$ is parametrised by length, with $\gamma(0) = p$ and $\gamma(\ell) = q$. For a vertex p' of D , denote by $\Theta(p')$ the set of annuli enclosing p' , and by $\Omega(p')$ the set of annuli not enclosing p' .

For $A \in \Theta(p)$, we call the *exit time* of A the number

$$\epsilon(A) := \min\{t \in \mathbb{N} \cap [0, \ell] : A \in \Omega(\gamma(s)) \text{ for } s \in \mathbb{N} \cap [t, \ell]\}.$$

That is, the number $\epsilon(A)$ is the time after which γ last crosses A . Note that this number is well-defined since $A \in \Omega(q)$. Moreover, an annulus $A \in \Theta(p)$ crosses, by definition, the edge between $\gamma(\epsilon(A) - 1)$ and $\gamma(\epsilon(A))$, so different annuli have distinct exit times.

Choose the ordering $\Theta(p) = \{A_1, \dots, A_n\}$ according to the exit times so that $\epsilon(A_i) < \epsilon(A_j)$ whenever $i < j$. For example, if the annuli are pairwise disjoint, then A_1 is the innermost annulus, and A_n the outermost. For $i \in \{1, \dots, n\}$, let $p_i := \gamma(\epsilon(A_i))$ denote the exit point of A_i .

By the chosen ordering, we have $A_j \in \Omega(p_i)$ for $i \geq j$. In particular, $|\Omega(p_i) \cap \Theta(p)| \geq i - 1$, and therefore $|\Theta(p_i) \cap \Theta(p)| \leq n - (i - 1)$. By the assumption $|\Theta(p_i)| \geq k$, it follows that $|\Theta(p_i) \cap \Omega(p)| \geq k - n + i - 1$. Every annulus $A' \in \Theta(p_i) \cap \Omega(p)$ must cross A_i , since p is enclosed by A_i but not by A' , p_i is enclosed by A' but not by A_i , and the vertex after p_i (moving towards p) is enclosed by both A' and A_i . Therefore, $(A_i, A') \in x(p)$.

Since $n \geq k$, we conclude that $x(p)$ contains at least $k + (k - 1) + \dots + 1 = \frac{k(k+1)}{2}$ pairs of crossing annuli. \square

We have the following consequence.

COROLLARY 6.6. *Assume that there exists a family $P = \{p_1, \dots, p_m\}$ of vertices in D such that no annulus encloses two or more vertices of P . Furthermore, suppose that for every $i \in \{1, \dots, m\}$, there is some integer $k_i \geq 0$ and a vertex q_i in D such that:*

- *no annulus encloses both p_i and q_i ;*
- *there exists a path connecting p_i and q_i such that every vertex of the path is enclosed by at least k_i annuli.*

Then

$$\text{Area}(D) \geq \frac{1}{2} \sum_{i=1}^m k_i^2.$$

Proof. The area of D coincides with the total number of crossings between annuli and/or corridors. In particular, it is bounded below by the cardinality of

$$X := \{(A, A') : A \neq A' \text{ are annuli and } A \text{ crosses } A'\}.$$

Note that we do not divide by two to take the symmetry into account, since if two annuli cross, they must do so at least twice.

Denote by L_1, \dots, L_m the families of annuli enclosing p_1, \dots, p_m , respectively. By assumption, the sets L_i are pairwise disjoint, so $x(p_1), \dots, x(p_m) \subseteq X$ are also pairwise disjoint. By Lemma 6.5, we have $|x(p_i)| \geq \frac{k_i^2}{2}$ for every $i \in \{1, \dots, m\}$. Therefore, we obtain $|X| \geq \frac{1}{2} \sum_{i=1}^m k_i^2$, as desired. \square

6.1 Cubic lower bound. Throughout this subsection, we assume that Γ has property (\mathcal{D}_3) ; that is, it admits a maximal reducible subgraph Λ that is essentially

2-reducible. In particular, the subgraph Λ decomposes as a join $\Lambda' * \Lambda''$, where Λ' and Λ'' are irreducible and each has at least two vertices. We aim to show that, in this case, the Dehn function of BB_Γ is bounded below by a cubic polynomial.

REMARK 6.7. The results in this subsection also apply when Γ has (\mathcal{D}_4) , since (\mathcal{D}_4) implies (\mathcal{D}_3) . In fact, the proof of the quartic lower bound in Sect. 6.2 builds on the work done for the cubic lower bound in Sect. 6.1, using the additional hypothesis given by property (\mathcal{D}_4) to promote the cubic lower bound into a quartic one.

Denote by \mathcal{A} and \mathcal{B} the sets of generators corresponding to the vertices of Λ' and Λ'' , respectively, and by $\mathcal{C} := S_\Gamma \setminus (\mathcal{A} \cup \mathcal{B})$ the generators corresponding to the vertices outside the join.

We begin by noting the following fact.

LEMMA 6.8. *Let $s \in S_\Gamma$ be a generator. If s commutes with all $a \in \mathcal{A}$, then $s \in \mathcal{B}$. Similarly, if s commutes with all $b \in \mathcal{B}$, then $s \in \mathcal{A}$.*

Proof. We prove the first assertion; the other follows by symmetry. Suppose that $s \in S_\Gamma$ commutes with all $a \in \mathcal{A}$, so the vertex v_s belongs to $\bigcap_{a \in \mathcal{A}} \text{St}(v_a)$. If $v_s = v_a$ for some $a \in \mathcal{A}$, then Λ' would be a cone graph, contradicting the fact that Λ' is irreducible. If $v_s \neq v_a$ for all $a \in \mathcal{A}$, then v_s must belong to Λ'' ; otherwise, the graph $\Lambda' * (v_s \cup \Lambda'')$ would be a reducible graph containing Λ , contradicting the maximality of Λ . Thus, we have $v_s \in \Lambda''$, and therefore $s \in \mathcal{B}$. \square

Lemma 6.8 has an immediate corollary.

COROLLARY 6.9. *There is no generator that commutes with all $s \in \mathcal{A} \cup \mathcal{B}$.*

Proof. Such a generator should belong to both \mathcal{A} and \mathcal{B} , but $\mathcal{A} \cap \mathcal{B} = \emptyset$. \square

LEMMA 6.10. *There exist $k \in \mathbb{N}$ and a sequence $a_0, a_1, \dots, a_k = a_0$ of non-necessarily distinct elements of \mathcal{A} such that $\{a_1, \dots, a_k\} = \mathcal{A}$, and $[a_{i-1}, a_i] \neq 1$ for all $i \in \{1, \dots, k\}$. Similarly, there exist $\ell \in \mathbb{N}$ and a sequence $b_0, b_1, \dots, b_\ell = b_0$ of non-necessarily distinct elements of \mathcal{B} such that $\{b_1, \dots, b_\ell\} = \mathcal{B}$, and $[b_{j-1}, b_j] \neq 1$ for all $j \in \{1, \dots, \ell\}$.*

Proof. We prove the first statement; the other follows from a similar argument. Since Λ' is irreducible, its complement $(\Lambda')^c$ is connected. Thus, the graph $(\Lambda')^c$ admits a combinatorial closed path $\gamma: [0, k] \rightarrow (\Lambda')^c$ (that is, $\gamma(k) = \gamma(0)$) that visits all the vertices at least once. Notice that we do not require the path γ to be injective, for instance, it may backtrack. Let $a_0, a_1, \dots, a_k \in \mathcal{A}$ be the sequence of generators corresponding to the vertices of γ , so that $\gamma(i) = v_{a_i}$. Since consecutive vertices on γ are not adjacent in Λ' , the corresponding generators do not commute. \square

Let $a_1, \dots, a_k = a_0 \in \mathcal{A}$ and $b_1, \dots, b_\ell = b_0 \in \mathcal{B}$ be any choice of elements satisfying the assumptions of Lemma 6.10. For $n \in \mathbb{N}$, consider the alternating

REMARK 6.11. Following the notation in Sect. 4, we have

$$\begin{aligned} w'_n &= \text{push}_0([(a_1 \cdots a_k)^n (a_1^{-1} \cdots a_k^{-1})]_{b_1}) \\ w''_n &= \text{push}_0([(b_1 \cdots b_\ell)^n (b_1^{-1} \cdots b_\ell^{-1})]_{a_1}). \end{aligned}$$

So, the word $w_n = [w'_n, w''_n]$ is the pushdown of the coloured commutator $[[w'_n]_{b_1}, [w''_n]_{a_1}]$ (recall Definition 5.5). While this fact is not used in the following arguments, it is still worthy of note: we have seen in Sect. 5 that upper bounds on the almost-flat area of the pushdowns of fundamental pieces (monochromatic words, coloured bigons, and coloured commutators) yield an upper bound on the Dehn function of BB_Γ . This suggests that the Dehn functions of BB_Γ should be witnessed by a family of such words.

DEFINITION 6.12. We call a corridor of D_n an \mathcal{A} -corridor, \mathcal{B} -corridor or \mathcal{C} -corridor, if it is an s -corridor for some $s \in \mathcal{A}, \mathcal{B}, \mathcal{C}$ respectively. Define \mathcal{A} -annulus, \mathcal{B} -annulus and \mathcal{C} -annulus similarly.

REMARK 6.13. Since w_n is a word in $\mathcal{A} \cup \mathcal{B}$, there are no \mathcal{C} -corridors in D_n ; however, \mathcal{C} -annuli may be present.

The properties described in Lemma 6.10 place strong restrictions on how w_n can be transformed into the trivial word by applying relations. This allows us to describe precisely the behaviour of corridors inside D_n .

DEFINITION 6.14. A corridor in D_n is called a *bridge* if it connects two i th letters on the opposite sides of the square (counting from left to right and from bottom to top). It is called a *half-annulus* if it connects the i th letter with the $(4n - i)$ th letter on the same side of the square.

LEMMA 6.15. *Every \mathcal{A} -corridor in D_n is either a half-annulus that starts and ends on the left or right side of the square, or a bridge connecting the top and bottom sides. Analogously, every \mathcal{B} -corridor in D_n is either a half-annulus that starts and ends on the top or bottom side of the square, or a bridge connecting the left and right sides; see Fig. 11.*

Proof. We prove the first statement; the second statement follows from a similar argument. Consider the j th leftmost a_i -edge e on the bottom side of D_n , with $2 \leq i \leq n$. The \mathcal{A} -corridor starting at e must end at another a_i -edge with opposite orientation. This corridor divides the other edges of D_n into two sets, one on each side of the corridor. Since two a_i -corridors cannot cross, the total number of boundary a_i -edges on each side of the corridor, counted with sign, must be zero. To satisfy this condition, there are only two possibilities: the corridor must end either at the j th leftmost a_i -edge e' on the top side, or at the j th rightmost a_i -edge e'' on the bottom side.

Suppose, by contradiction, that there is a corridor C connecting e to e'' . Then there are $(n - j)$ many a_{i-1} -edges below C oriented to the right and $(n - j + 1)$ many a_{i-1} -edges below C oriented to the left. Since a_{i-1} and a_i do not commute, every

corridor that starts at a letter a_{i-1} below C cannot intersect C . Thus, there is no way to pair all the a_{i-1} edges below C . Therefore, a corridor that starts at e must end at e' , that is, the corridor C is a bridge.

For a_1 -edges, the reasoning is similar: let e be an a_1 -edge on the bottom side that is not the leftmost one. There is only one other edge with the same label that can be reached by a corridor C without intersecting any a_2 -corridors or a_k -corridors, namely the a_1 -edge directly above on the top side. Thus, the corridor C is a bridge.

The leftmost a_1 -edges on the top and bottom sides, together with all the a_1 -edges on the left and right sides, are oriented in such a way that the only possible a_1 -corridors that do not pairwise intersect and do not intersect any a_2 -corridors or a_k -corridors are bridges and half-annuli, respectively, as shown in Fig. 11. \square

Lemma 6.15 shows that bridges in D_n look almost like a grid pattern, in the sense that \mathcal{A} -bridges are parallel and do not intersect, and the same holds for \mathcal{B} -bridges. Every \mathcal{A} -bridge crosses every \mathcal{B} -bridge at least once, and possibly multiple times, as shown in Fig. 11. Moreover, the half-annuli must stay “close” to the boundary, as they cannot cross a bridge: for example, the bottom half-annuli labelled by b_1 would need to cross a b_1 -bridge, and the top ones would have to cross a b_k -bridge, and both are not allowed. This gives a complete description of the corridors inside D_n .

We now proceed to investigate the behaviour of the annuli. To this end, we introduce the following definition.

DEFINITION 6.16. For $-(n-1) \leq i \leq n-1$, define the \mathcal{A} -gap, denoted by $\Sigma(\mathcal{A}, i)$, to be the region between the $(n+i)$ th a_k -bridge and the $(n+i+1)$ th a_1 -bridge, counted from left to right. Similarly, define the \mathcal{B} -gap, denoted by $\Sigma(\mathcal{B}, j)$, to be the region between the $(j+n)$ th b_ℓ -bridge and the $(j+n+1)$ th b_1 -bridge, counted from bottom to top.

LEMMA 6.17. *Let A be an annulus that intersects two distinct \mathcal{A} -gaps. Then A is a \mathcal{B} -annulus. Similarly, if it intersects two distinct \mathcal{B} -gaps, then it is an \mathcal{A} -annulus.*

Proof. Let $s \in S_\Gamma$. If A is an s -annulus intersecting two distinct \mathcal{A} -gaps, it must also cross every \mathcal{A} -bridge in between. Then s commutes with all the generators of \mathcal{A} , so it belongs to \mathcal{B} by Lemma 6.8. The argument for the second statement is similar. \square

For $i, j \in \mathbb{Z}$ with $|i|, |j| \leq n-1$, denote by $R_{i,j}$ the intersection $\Sigma(\mathcal{A}, i) \cap \Sigma(\mathcal{B}, j)$, which is non-empty as it contains at least intersections of the sides of the bridges. However, it may be disconnected. For example, in Fig. 11, the region $R_{0,0}$ has three connected components.

LEMMA 6.18. *For every vertex $p \in R_{i,j}$, we have $\kappa(p) = k(n - |i|) + \ell(n - |j|)$*

Proof. The bottom and top half-annuli, which are the b_1 -corridors, do not intersect \mathcal{B} -gaps, as to do so they would need to intersect a b_i -bridge for every $i \in \{1, \dots, \ell\}$. Similarly, the left and right half-annuli do not intersect \mathcal{A} -gaps. In particular, they cannot separate p from the base point, so only the bridges contribute to $\kappa(p)$. The statement now follows from the definition of κ . \square

For $|i| \leq n-1$, choose a point p_i inside $R_{i,i}$. We verify in the following two lemmas that the family $P = \{p_i : |i| \leq n-1\}$ satisfies the two conditions in the hypotheses of Corollary 6.6.

LEMMA 6.19. *No annulus can enclose both p_i and p_j for $i \neq j$.*

Proof. If such an annulus exists, then it would intersect two distinct \mathcal{A} -gaps and two distinct \mathcal{B} -gaps. By Lemma 6.17, it would be both an \mathcal{A} -annulus and a \mathcal{B} -annulus, which is a contradiction. \square

LEMMA 6.20. *Let $p \in R_{i,j}$ be a vertex, with $|i|, |j| \leq n-2$. There exists $q \in R_{i',j'}$, for some i', j' satisfying $|i-i'| = |j-j'| = 1$, and a (combinatorial) path connecting p and q that intersects only the four gaps $\Sigma(\mathcal{A}, i)$, $\Sigma(\mathcal{A}, i')$, $\Sigma(\mathcal{B}, j)$, and $\Sigma(\mathcal{B}, j')$. In particular, for every point x on the path between p and q , we have $|\kappa(x)| \geq k(n-|i|-1) + \ell(n-|j|-1)$.*

Proof. Consider a path, starting from p and ending on the boundary of D_n , that is entirely contained in $\Sigma(\mathcal{A}, i)$. This path must intersect a \mathcal{B} -gap different from $\Sigma(\mathcal{B}, j)$. Let p' be the first vertex inside such a gap $\Sigma(\mathcal{B}, j')$. Clearly, $|j-j'| = 1$.

Now, repeat the same argument by taking a path inside $\Sigma(\mathcal{B}, j')$ starting at p' and ending on the boundary of D_n ; let q be the first vertex on this path that lies in a different gap $\Sigma(\mathcal{A}, i')$. Then, $|i-i'| = 1$. The union of the two paths, from p to p' and from p' to q , intersects only the four gaps in the statement. Since this path lies entirely within these four gaps, the estimate for $\kappa(x)$ follows analogously to the proof of Lemma 6.18. \square

Now, we can apply Corollary 6.6 to obtain the cubic lower bound.

PROPOSITION 6.21. *If Γ has property (\mathcal{D}_3) , then $\delta_{BB_\Gamma}(n) \gtrsim n^3$.*

Proof. We prove that there exists a constant $C > 0$, independent of the choice of the almost-flat diagram D_n for w_n , such that the area satisfies

$$\text{Area}(D_n) \geq \frac{1}{C}n^3 - C.$$

Consider the family $P = \{p_i \in R_{i,i} : |i| \leq n-1\}$. Lemma 6.19 says that there are no annuli enclose p_i and p_j for $i \neq j$. By Lemma 6.20, for each i , we get a point $q_i \in R_{i\pm 1, i\pm 1}$ and a path from p_i to q_i , and every point x on this path is enclosed by at least

$$\alpha_+(x) + \alpha_-(x) \geq |\alpha(x)| \geq |\kappa(x)| - 2 \geq (k + \ell)(n - |i| - 1) - 2$$

many annuli. Thus, the family P satisfies the hypotheses of Corollary 6.6, and applying the conclusion yields

$$\begin{aligned} \text{Area}(D_n) &\geq \frac{1}{2} \sum_{i=-(n-1)}^{n-1} ((k+\ell)(n-|i|-1)-2)^2 \\ &\geq \frac{1}{C} n^3 - C \end{aligned}$$

for some constant $C > 0$ sufficiently large and independent of n . \square

6.2 Quartic lower bound. In this section, we assume that Γ has property (\mathcal{D}_4) ; that is, it has a maximal reducible subgraph $\Lambda \subseteq \Gamma$ such that $\Delta(\Lambda)$ is not simply connected. Thus, the subgraph Λ decomposes as a join $\Lambda' * \Lambda''$, where Λ' and Λ'' are both disconnected. In particular, the subgraph Λ is essentially 2-reducible.

In this case, we may choose the generators a_1, \dots, a_k and b_1, \dots, b_ℓ that satisfy the hypotheses of Lemma 6.10, and with the additional constraint that the vertices v_{a_1} and v_{a_k} in Λ' , corresponding to a_1 and a_k , lie in different connected components of Λ' , and the same holds for v_{b_1} and v_{b_ℓ} in Λ'' .

With this additional assumption, we show that the area of an almost-flat van Kampen diagram D_n for the word w_n , defined in Sect. 6.1, has a quartic lower bound. We begin with the following standard topological result.

LEMMA 6.22. *Let C be a 2-dimensional contractible cell complex, and let $A \sqcup B$ be a partition of its vertices. Assume that for every 2-dimensional cell σ of C , the full subcomplexes of σ whose vertices belong to A and to B are connected or empty. Let $\gamma: S^1 \rightarrow C$ be a combinatorial loop (not necessarily simple) in C , obtained by concatenating four combinatorial paths $\gamma_1, \gamma_2, \gamma_3$, and γ_4 . Then, either there exists a combinatorial path connecting the images of γ_1 and γ_3 , with all vertices in A , or there exists a combinatorial path connecting the images of γ_2 and γ_4 , with all vertices in B .*

REMARK 6.23. Note that Lemma 6.22 does not assume that the vertices of γ_1, γ_3 belong to A , nor that the vertices of γ_2, γ_4 belong to B . However, a counterexample cannot be constructed by requiring γ_1, γ_3 to be proper subsets of B and γ_2, γ_4 to be proper subsets of A , since $(\gamma_1 \cup \gamma_3) \cap (\gamma_2 \cup \gamma_4) \neq \emptyset$ (it contains the four points where the paths are concatenated).

Proof. Suppose that this is not the case. We claim that there exists $f: C \rightarrow S^1$ such that $f \circ \gamma$ is homotopic to a homeomorphism. This leads to a contradiction, as C being contractible implies that the composition $f \circ \gamma$ must be null-homotopic. For the following, we view S^1 as the unit circle in \mathbb{C} .

To define f , we first note that a vertex $v \in A$ cannot be connected inside A to both γ_1 and γ_3 . For $v \in A$, define $f(v) = 1$ if it is connected to γ_1 , and $f(v) = -1$ if it is connected to γ_3 or to neither. Similarly, for $v \in B$, define $f(v) = i$ if v is connected to γ_2 , and $f(v) = -i$ otherwise.

If two vertices v and v' share an edge, then $f(v) \neq -f(v')$. So there is a canonical way to extend f to the 1-skeleton by choosing the shortest path in S^1 connecting the images. Then, if σ is a 2-cell, by hypothesis the vertices $\partial\sigma \cap A$ are all in the same connected component, so $f|_{\partial\sigma}$ is not surjective. Therefore, the map f can be extended to σ continuously.

It remains to show that $f \circ \gamma$ is homotopic to a homeomorphism. This follows since f never assumes the value -1 on γ_1 , $-i$ on γ_2 , 1 on γ_3 , or i on γ_4 , so f is a degree-one map. \square

We now use Lemma 6.22 to obtain information about the structure of the \mathcal{A} -gaps and \mathcal{B} -gaps.

LEMMA 6.24. *In every \mathcal{A} -gap (respectively, \mathcal{B} -gap) of D_n , there is a combinatorial path connecting the two opposite boundary components that does not cross any \mathcal{A} -annuli (respectively, \mathcal{B} -annuli).*

Proof. Let p be a vertex in D_n that is enclosed in an \mathcal{A} -annulus. Consider the set of outermost \mathcal{A} -annuli that enclose p , that is, those not enclosed by any larger \mathcal{A} -annulus. Since they pairwise cross, all the corresponding generators must commute, so their associated vertices all belong to the same connected component of Λ' . Denote this connected component by $\text{Comp}(p)$. If p and p' are adjacent vertices in D_n and each is enclosed by at least one \mathcal{A} -annulus, then there is an \mathcal{A} -annulus that encloses both, and therefore $\text{Comp}(p) = \text{Comp}(p')$.

We now apply Lemma 6.22 to the \mathcal{A} -gap, where the partition of the vertices is given by those enclosed by some \mathcal{A} -annulus and those that are not. Suppose by contradiction that there is no path that connects the boundary components of an \mathcal{A} -gap and lies outside every \mathcal{A} -annulus. Then, there is a path L connecting the a_k -corridor to the a_1 -corridor such that every vertex of L lies in at least one \mathcal{A} -annulus. Therefore, repeating the argument in the previous paragraph on the vertices of L gives $\text{Comp}(p) = \text{Comp}(q)$, where p and q are the endpoints of L .

However, all \mathcal{A} -annuli enclosing p must cross the a_k -corridor, so $\text{Comp}(p)$ is the connected component containing v_{a_k} . Similarly, the connected component $\text{Comp}(q)$ contains v_{a_1} . Since we assumed that v_{a_1} and v_{a_k} belong to different connected components of Λ' , we obtain a contradiction. \square

We apply Lemma 6.24 to the gaps $\Sigma(\mathcal{A}, i)$ and $\Sigma(\mathcal{B}, j)$ and deduce the existence of two paths $\gamma \subseteq \Sigma(\mathcal{A}, i)$ and $\gamma' \subseteq \Sigma(\mathcal{B}, j)$, whose vertices are not enclosed by any \mathcal{A} -annuli and any \mathcal{B} -annuli, respectively. By intersecting γ and γ' , we obtain a vertex $p_{i,j} \in R_{i,j} = \Sigma(\mathcal{A}, i) \cap \Sigma(\mathcal{B}, j)$ that is not enclosed by any \mathcal{A} -annulus or any \mathcal{B} -annulus.

We are now ready to apply Corollary 6.6 to establish the quartic lower bound.

PROPOSITION 6.25. *If Γ has property (\mathcal{D}_4) , then $\delta_{BB\Gamma}(n) \gtrsim n^4$.*

Proof. We claim that the area of the almost-flat van Kampen diagram D_n for w_n is bounded below by a quartic polynomial. First, note that there is no annulus enclosing two distinct points $p_{i,j}$ and $p_{i',j'}$. Suppose, for instance, that $i \neq i'$ (the case $j \neq j'$

is similar). Then an annulus enclosing both $p_{i,j}$ and $p_{i',j'}$ would have to intersect different \mathcal{A} -gaps. By Lemma 6.17, it must be a \mathcal{B} -annulus, but this is ruled out by the definition of $p_{i,j}$.

By Lemma 6.20, we obtain a path connecting $p_{i,j}$ to some other point $q_{i,j} \in R_{i\pm 1, j\pm 1}$, with every point along the path contained in at least $k(n - |i| - 1) + \ell(n - |j| - 1) - 2$ many annuli. By the argument above, no annulus can enclose both $p_{i,j}$ and $q_{i,j}$. Therefore, we can apply Corollary 6.6 to the family $p_{i,j}$ and obtain

$$\begin{aligned} \text{Area}(D_n) &\geq \frac{1}{2} \sum_{i=-(n-1)}^{n-1} \sum_{j=-(n-1)}^{n-1} (k(n - |i| - 1) + \ell(n - |j| - 1) - 2)^2 \\ &\geq \frac{1}{C} n^4 - C \end{aligned}$$

for some constant $C > 0$ sufficiently large and independent of n . \square

Acknowledgements

We are thankful for the environment provided by the 2023 thematic program *Geometric Group Theory* at the Centre de Recherches Mathématiques, Université de Montréal. We are grateful to Claudio Llosa Isenrich and Matt Zaremsky for helpful comments and suggestions on an earlier version of this work. We thank Claudio Llosa Isenrich for also raising interesting questions related to Corollary 2 and Corollary 3. We thank the anonymous referee for their careful reading and for their valuable comments and suggestions, including the remark about higher Dehn functions in the discussion about Question 2.

Funding information

The second author is supported by the Royal Society of Great Britain. The second and third authors acknowledge funding from the DFG 281869850 (RTG 2229).

Open Access. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- [A+13] ABRAMS, A., BRADY, N., DANI, P., DUCHIN, M., YOUNG, R.: Pushing fillings in right-angled Artin groups. *J. Lond. Math. Soc. (2)* **87**(3), 663–688 (2013).
- [BB97] BESTVINA, M., BRADY, N.: Morse theory and finiteness properties of groups. *Invent. Math.* **129**(3), 445–470 (1997).
- [Bie81] BIERI, R.: Homological Dimension of Discrete Groups, 2nd edn. Queen Mary College Mathematics Notes. Queen Mary College, Department of Pure Mathematics, London (1981).
- [Bow95] BOWDITCH, B.H.: A short proof that a subquadratic isoperimetric inequality implies a linear one. *Mich. Math. J.* **42**(1), 103–107 (1995).
- [Bri02] BRIDSON, M.R.: The geometry of the word problem. In: *Invitations to Geometry and Topology*, vol. 7, pp. 29–91 (2002).
- [Bri13] BRIDSON, M.R.: On the subgroups of right-angled Artin groups and mapping class groups. *Math. Res. Lett.* **20**(2), 203–212 (2013).
- [BRS07] BRADY, N., RILEY, T., SHORT, H.: The Geometry of the Word Problem for Finitely Generated Groups. *Advanced Courses in Mathematics*. CRM Barcelona. Birkhäuser, Basel (2007). Papers from the Advanced Course held in Barcelona, July 5–15, 2005.
- [BS19] BRADY, N., SOROKO, I.: Dehn functions of subgroups of right-angled Artin groups. *Geom. Dedic.* **200**(1), 197–239 (2019).
- [CF17] CARTER, W., FORESTER, M.: The Dehn functions of Stallings–Bieri groups. *Math. Ann.* **368**(1), 671–683 (2017).
- [Cha07] CHARNEY, R.: An introduction to right-angled Artin groups. *Geom. Dedic.* **125**, 141–158 (2007).
- [Cha21] CHANG, Y.-C.: Identifying Dehn functions of Bestvina-Brady groups from their defining graphs. *Geom. Dedic.* **214**, 211–239 (2021).
- [D+09] DISON, W., ELDER, M., RILEY, T.R., YOUNG, R.: The Dehn function of Stallings’ group. *Geom. Funct. Anal.* **19**(2), 406–422 (2009).
- [Dis08] DISON, W.: An isoperimetric function for Bestvina-Brady groups. *Bull. Lond. Math. Soc.* **40**(3), 384–394 (2008).
- [DL99] DICKS, W., LEARY, I.J.: Presentations for subgroups of Artin groups. *Proc. Am. Math. Soc.* **127**(2), 343–348 (1999).
- [Gro83] GROMOV, M.: Filling Riemannian manifolds. *J. Differ. Geom.* **18**(1), 1–147 (1983).
- [Gro87] GROMOV, M.: Hyperbolic groups. In: *Essays in Group Theory*. Math. Sci. Res. Inst. Publ., vol. 8, pp. 75–263. Springer, New York (1987).
- [Gro93] GROMOV, M.L.: Asymptotic invariants of infinite groups. In: *Geometric Group Theory*, Vol. 2 (Sussex, 1991), London Math. Soc. Lecture Note Ser., vol. 182, pp. 1–295. Cambridge University Press, Cambridge (1993).
- [GS02] GERSTEN, S., SHORT, H.: Some isoperimetric inequalities for kernels of free extensions. *Geom. Dedic.* **92**(1), 63–72 (2002).
- [Pap96] PAPASOGLU, P.: On the asymptotic cone of groups satisfying a quadratic isoperimetric inequality. *J. Differ. Geom.* **44**(4), 789–806 (1996).
- [PS07] PAPADIMA, S., SUCIU, A.: Algebraic invariants for Bestvina-Brady groups. *J. Lond. Math. Soc. (2)* **76**(2), 273–292 (2007).
- [Sta63] STALLINGS, J.: A finitely presented group whose 3-dimensional integral homology is not finitely generated. *Am. J. Math.* **85**, 541–543 (1963).
- [Tra17] TRAN, H.C.: Geometric embedding properties of Bestvina-Brady subgroups. *Algebr. Geom. Topol.* **17**(4), 2499–2510 (2017).

[Wen08] WENGER, S.: A short proof of Gromov's filling inequality. *Proc. Am. Math. Soc.* **136**(8), 2937–2941 (2008).

Publisher's note. Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Yu-Chan Chang

Department of Mathematics, The Ohio State University, 231 W. 18th Ave., Columbus, OH 43210, USA.

chang.2628@osu.edu

Jerónimo García-Mejía

Mathematical Institute, Andrew Wiles Building, Observatory Quarter, University of Oxford, Oxford OX2 6GG, UK.

jeronimo.garcia-mejia@maths.ox.ac.uk

Matteo Migliorini

Faculty of Mathematics, Karlsruhe Institute of Technology, Englerstraße 2, 76131 Karlsruhe, Germany.

matteo.migliorini@kit.edu

Received: 23 September 2025

Revised: 17 December 2025

Accepted: 23 December 2025