

PUSHING FILLINGS IN RIGHT-ANGLED ARTIN GROUPS

ABRAMS BRADY DANI DUCHIN YOUNG

ABSTRACT. We define a family of quasi-isometry invariants of groups called higher divergence functions, which measure isoperimetric properties “at infinity.” We give sharp upper and lower bounds on the divergence functions for right-angled Artin groups, using different *pushing maps* on the associated cube complexes. In the process, we define a class of RAAGs we call *orthoplex groups*, which have the property that their Bestvina-Brady subgroups have hard-to-fill spheres. Our results give sharp bounds on the higher Dehn functions of Bestvina-Brady groups, a complete characterization of the divergence of geodesics in RAAGs, and an upper bound for filling loops at infinity in the mapping class group.

1. INTRODUCTION

This paper focuses on a new construction for right-angled Artin groups that we call a *pushing map*. This map can be used to study the difference between efficient filling and “obstructed” filling of cycles, by pushing chains into constrained parts of the complex associated to the group. This will say something about the flexibility of the geometry. We note at the outset that the main results and techniques in this paper, though their properties are stated and proved in the homological category, work just as well with homotopical definitions.

Right-angled Artin groups (or RAAGs) are given by presentations in which each relator is a commutator of two generators. These groups have been studied extensively, and much is known about them both algebraically and geometrically. For instance, they have automatic structures and useful normal forms, and they act geometrically on CAT(0) cube complexes. Many tools are available for their study, in part because these complexes contain many flats arising from mutually commuting elements (see [Cha07]). Frequently, topological invariants of RAAGs can be read off of the defining graph, and along these lines we will relate properties of the graph to the filling functions and divergence functions.

The divergence function grew from the desire to study the geometry of a group “at infinity” through the use of filling functions. Recall that the most basic filling function in groups is the *Dehn function*, which measures the area necessary to fill a closed loop by a disk; these functions have been a key part of geometric group theory since Gromov used them to characterize hyperbolic groups [Gro87] (or arguably longer, since Dehn used related ideas to find fast solutions to the word problem). This can naturally be generalized to higher-order Dehn functions, which describe the difficulty of filling k -spheres by $(k + 1)$ -balls or k -chains by $(k + 1)$ -cycles.

Topology at infinity is the study of the asymptotic structure of groups by attaching topological invariants to the complements of large balls; the theorems of

Date: April 27, 2010.

This work was supported by a SQuaRE grant from the American Institute of Mathematics.

Hopf and Stallings about ends of groups were early examples of major results of this kind.

One can bring the two together by defining *higher divergence functions*, which measure rates of filling in complements of balls in groups and other metric spaces. With respect to a fixed basepoint x_0 in a space, we will describe a map whose image is disjoint from the ball of radius r about x_0 as being *r-avoidant*. Roughly, the k -dimensional divergence function is a filling invariant for avoidant spheres and balls; it measures the volume needed to fill an avoidant sphere by an avoidant ball. (We will make this precise in §2.2.)

As with Dehn functions, the divergence functions become meaningful for finitely generated groups by adding an appropriate equivalence relation to make the definition invariant under quasi-isometry.

For $k \geq 0$, our functions Div^k are closely related to the higher divergence functions defined by Brady and Farb in [BF98] for the special case of Hadamard manifolds. Using this definition, combined results of Leuzinger and Hindawi prove that the higher divergence functions detect the real-rank of a symmetric space, as Brady-Farb had conjectured [Leu00, Hin05]. Thus the geometry and the algebra are connected. Wenger generalized this, showing that higher divergence functions detect the Euclidean rank of any CAT(0) space [Wen06].

These divergence functions are interesting in part because they unify a number of concepts in coarse geometry and geometric group theory. For instance, in the $k = 0$ case, this is the classical divergence of geodesics, which relates to the curvature and in particular detects the hyperbolicity of the space. Gromov showed that a space is δ -hyperbolic if and only if it has exponential divergence of geodesics in a certain precise sense. More recently, many papers in geometric group theory have studied polynomial divergence of geodesics, including but not limited to [Ger94b, Ger94a, KL98, DR09, Mac08, DMS09]. Much of this work arose to explore an expectation offered by Gromov in [Gro93] that nonpositively curved spaces should, like symmetric spaces, exhibit a gap in the possible rates of divergence of geodesics (between linear and exponential). On the contrary, it is now clear that quadratic divergence of geodesics often occurs in groups where many “chains of flats” are present, and Macura has produced examples of groups with polynomial divergence of every degree.

Avoidant fillings can often be constructed by perturbing an efficient filling to avoid a ball. When this is “hard” (i.e., when the perturbation blows up volume), then there can be a large difference between the Dehn functions and the divergence functions of a space. For instance, in a nonpositively curved space of rank k , efficient fillings of $(k - 1)$ -spheres are often essentially unique, so modifying them to avoid a ball is difficult; Leuzinger’s result cited above showed that in a symmetric space of rank k , one can find avoidant $(k - 1)$ -spheres that require exponential volume to fill by an avoidant k -ball.

At the other extreme, fillings in Euclidean space are extremely flexible; if $k \leq d - 2$, there are many ways to fill a k -sphere in \mathbb{R}^d with essentially the same volume, and filling a sphere avoidantly takes not much more volume than the most efficient filling. In the next section, we will study this example in more detail. Similarly, in [Youa], Young constructs efficient horizontal fillings of spheres in a class of nilpotent groups. These fillings satisfy certain tangency conditions, and the set of horizontal maps in these groups satisfies an h-principle. This flexibility means that there are

many ways to fill a sphere by a ball and makes it likely that Young's constructions can be modified to produce a bound on the divergence functions as well.

In this paper we develop several applications of pushing maps (defined in §4), which are “singular retractions” defined from the complex associated to a RAAG onto various subsets of this complex, such as the exterior of a ball or a Bestvina-Brady subgroup. We will use these maps to obtain upper bounds on higher Dehn functions of Bestvina-Brady groups by pushing fillings into these subgroups (§5), and we will use them in a different way to find special examples where the upper bounds are achieved. In §6 we study higher divergence functions by pushing fillings out of balls: if a RAAG A_Γ is k -connected at infinity, we can guarantee that avoidant fillings satisfy a polynomial bound, namely $\text{Div}^k(A_\Gamma) \preceq r^{2k+2}$. Next, these upper bounds are shown to be sharp by using the earlier estimates for Bestvina-Brady subgroups.

Although the upper and lower bounds are sharp in every dimension, we are not able to specify which rates of divergence occur between the two extremes for $k \geq 1$. However, for $k = 0$, we show in §7 that every RAAG must have either linear or quadratic divergence, and this depends only on whether the group is a direct product (a property that can be read off of the defining graph).

Finally, several of the techniques developed to deal with RAAGs have applications in other groups. We will show that a related pushing construction can be applied in the mapping class group $\text{Mod}(S)$, which also has many flats arising from Dehn twists about disjoint curves. It is already known that $\text{Div}^0(\text{Mod}(S)) \asymp r^2$, the same as the highest growth rate realized by right-angled Artin groups. We show in §8 that $\text{Div}^1(\text{Mod}(S)) \preceq r^4$, obtaining again the same upper bound as for RAAGs.

2. DEHN FUNCTIONS AND DIVERGENCE FUNCTIONS

In this section, we will define the higher-order Dehn functions and the higher divergence functions and illustrate the basic methods of this paper by using a pushing map to bound the divergence functions of \mathbb{R}^d .

2.1. Higher Dehn functions. We will primarily use homological Dehn functions, following [ECH⁺92, Gro93, Wen08]. Homological Dehn functions describe the difficulty of filling cycles by chains. In contrast, some other papers ([AWP99, BBFS09]) use homotopical Dehn functions, which measure fillings of spheres by balls.

In low dimensions, these functions may differ, but they are essentially the same for high-dimensional fillings in highly-connected spaces. When $k \geq 3$, the topologies of the boundary and of the filling are irrelevant, and the homological and homotopical k -th order Dehn function are the same. When $k = 2$, the topology of the boundary is relevant, but the topology of the filling is not: a homological filling of a sphere guarantees a homotopical one of the same volume, so that the homotopical Dehn function is bounded above by the homological one [Groa, Grob]. (See also [Gro83, App.2.(A')], [BBFS09, Rem.2.6(4)].) However this bound may not be sharp, that is, there may exist cycles that are “harder to fill” than spheres [Youb].

We will define the higher Dehn function in two ways, one better-suited to dealing with complexes, and one better for dealing with manifolds.

We define a polyhedral complex to be a CW-complex in which each cell is isometric to a convex polyhedron in Euclidean space and the gluing maps are isometries.

If X is a polyhedral complex, we can define filling functions of X based on cellular homology. Assume that X is k -connected and let $C_k^{\text{cell}}(X)$ be the set of cellular k -chains of X with integer coefficients. If $a \in C_k^{\text{cell}}(X)$, then $a = \sum_i a_i \sigma_i$ for some integers a_i and distinct k -cells σ_i . Set $\|a\| = \sum |a_i|$. If $Z_k^{\text{cell}}(X)$ is the set of cellular k -cycles and $a \in Z_k^{\text{cell}}(X)$, then the fact that X is k -connected implies that $a = \partial b$ for some $b \in C_{k+1}^{\text{cell}}(X)$. Define the *filling volume* of a to be

$$\delta_X^{k;\text{cell}}(a) = \min_{\substack{b \in C_{k+1}^{\text{cell}}(X) \\ \partial b = a}} \|b\|,$$

and define the k -th order Dehn function of X to be

$$\begin{aligned} \delta_X^{k;\text{cell}} : \mathbb{R}^+ &\longrightarrow \mathbb{R}^+ \\ l &\longmapsto \max_{\substack{a \in Z_{k+1}^{\text{cell}}(X) \\ \|a\| \leq l}} \delta_X^{k;\text{cell}}(a). \end{aligned}$$

Alonso, Wang, and Pride [AWP99] showed that if G and G' are quasi-isometric groups acting geometrically (i.e., properly discontinuously, cocompactly, and by isometries) on certain associated k -connected complexes X and Y respectively, then $\delta_X^{k;\text{cell}}$ and $\delta_Y^{k;\text{cell}}$ grow at the same rate; in particular, this shows that the growth rate of $\delta_X^{k;\text{cell}}$ depends only on G , so we can define $\delta_G^{k;\text{cell}}$. This is made rigorous by defining an equivalence relation \asymp on functions, as follows. There is a partial order on the set of functions $\mathbb{R}^+ \rightarrow \mathbb{R}^+$ given by the following symbol: $f \preceq g$ means there exists $A > 0$ such that

$$f(t) \leq Ag(At + A) + At + A$$

for all $t \geq 0$, and the same property may be written $g \succeq f$. Then $f \asymp g$ if and only if $f \preceq g$ and $f \succeq g$. This is the standard notion of equivalence for coarse geometry, because it amounts to allowing a linear rescaling of domain and range, as in a quasi-isometry. Note that the equivalence relation \asymp identifies all linear and sublinear functions into one class, but distinguishes polynomials of different degrees.

Another way to define a homological higher-order Dehn function, somewhat better suited to Riemannian manifolds and CAT(0)-spaces, is to use singular Lipschitz chains, as in [ECH⁺92], [Gro83], and [Wen08]. A full introduction to this approach can be found in Chapter 10.3 of [ECH⁺92]. Let X be a k -connected Riemannian manifold or locally finite polyhedral complex (more generally, a local Lipschitz neighborhood retract). Singular Lipschitz k -chains (sometimes simply called Lipschitz k -chains) are formal sums of Lipschitz maps from the standard simplex Δ^k to X . As in the cellular case, we will consider chains with integral coefficients. The boundary operator is defined as for singular homology. Since a Lipschitz map is differentiable almost everywhere, we can define the k -volume of a Lipschitz map $\Delta^k \rightarrow X$, and we define the *mass* of a Lipschitz chain to be the total volume of its summands, weighted by the coefficients. For k -connected X , if $C_k^{\text{Lip}}(X)$ is the set of Lipschitz k -chains in X and $Z_k^{\text{Lip}}(X)$ is the set of Lipschitz k -cycles then we can define filling functions by

$$\delta_X^{k;\text{Lip}}(a) := \inf_{\substack{b \in C_{k+1}^{\text{Lip}}(X) \\ \partial b = a}} \text{mass}(b),$$

for all $a \in Z_k^{\text{Lip}}(X)$, and

$$\begin{aligned} \delta_X^{k;\text{Lip}} : \mathbb{R}^+ &\longrightarrow \mathbb{R}^+ \\ l &\longmapsto \sup_{\substack{a \in Z_k^{\text{Lip}}(X) \\ \text{mass } a \leq l}} \delta_X^{k;\text{Lip}}(a). \end{aligned}$$

These two definitions of Dehn functions are very similar, and if X is a polyhedral complex with bounded geometry (or if X is a space which can be approximated by such a polyhedral complex), one can show that they grow at the same rate by using the Federer-Fleming Deformation Theorem. We briefly explain the notation before stating the theorem: we will be approximating a Lipschitz chain a by a cellular chain $P(a)$. This may necessitate changing the boundary, and $R(a)$ interpolates between the old and new boundaries. Finally, $Q(a)$ interpolates between a and $P(a) + R(a)$. Note that if X is a polyhedral complex, then there is an inclusion $C_k^{\text{cell}}(X) \hookrightarrow C_k^{\text{Lip}}(X)$.

Theorem 2.1 (Federer-Fleming [FF60]). *Let X be a polyhedral complex with finitely many isometry types of cells. There is a constant c depending on X such that if $a \in C_k^{\text{Lip}}(X)$, then there are $P(a) \in C_k^{\text{cell}}(X)$, $Q(a) \in C_{k+1}^{\text{Lip}}(X)$, and $R(a) \in C_k^{\text{Lip}}(X)$ such that*

- (1) $\|P(a)\| \leq c \cdot \text{mass}(a)$;
- (2) $\|Q(a)\| \leq c \cdot \text{mass}(a)$;
- (3) $\|R(a)\| \leq c \cdot \text{mass}(\partial a)$;
- (4) $\partial Q(a) = a - P(a) - R(a)$; and
- (5) $\partial R(a) = \partial a - \partial P(a)$.

If $\partial a \in C_k^{\text{cell}}(x)$, we can take $R(a) = 0$. Furthermore, $P(a)$ and $Q(a)$ are supported in the smallest subcomplex of X which contains the support of a , and $R(a)$ is supported in the smallest subcomplex of X which contains the support of ∂a .

This version of the Federer-Fleming theorem is close to the one in [ECH⁺92], which addresses the case that a is a cycle.

As an application, it is straightforward to prove that if X is as above, then $\delta_X^{k;\text{Lip}}(l) \asymp \delta_X^{k;\text{cell}}(l)$. We will thus generally refer to δ_X^k or δ_G^k , using cellular or Lipschitz methods as appropriate.

Another (very important) application of the Federer-Fleming theorem is the isoperimetric inequality in Euclidean space [FF60]: if $1 \leq k \leq d - 1$, then

$$\delta_{\mathbb{R}^d}^k(l) \asymp l^{\frac{k+1}{k}}.$$

This is extended to general CAT(0) spaces, obtaining the same upper bound, in [Gro83, Wen08]. We state the version we will need for our filling results; it describes the key properties of Wenger's construction.

Proposition 2.2 (CAT(0) isoperimetric inequality [Wen08]). *If X is a CAT(0) polyhedral complex and $k \geq 1$, then the k -th order Dehn function of X satisfies*

$$\delta_X^k(l) \preceq l^{\frac{k+1}{k}}.$$

In fact, a slightly stronger condition is satisfied: there is a constant m such that if $a \in Z_k^{\text{Lip}}(X)$, then there is a chain $b \in C_{k+1}^{\text{Lip}}(X)$ such that $\partial b = a$,

$$\text{mass } b \leq m(\text{mass } a)^{\frac{k+1}{k}},$$

and $\text{supp } b$ is contained in a $m(\text{mass } a)^{\frac{1}{k}}$ -neighborhood of $\text{supp } a$.

2.2. Higher divergence functions. We will define divergence invariants $\text{Div}^k(X)$ for spaces X with sufficient connectivity at infinity. Our goal is to study the divergence functions of groups. We will solve filling problems in model spaces, such as $K(G, 1)$ spaces and other cell complexes with a geometric G -action. To make this meaningful, we therefore want Div^k to be invariant under quasi-isometries. The somewhat complicated equivalence relation defined in this section is designed to achieve this.

$\text{Div}^k(X)$ will basically generalize the definitions of divergence found in Gersten's work for $k = 0$ and Brady-Farb for $k \geq 1$ [Ger94b, BF98]. However, ours is not quite the same notion of equivalence. In particular, ours distinguishes polynomials of different degrees, whereas Brady-Farb identifies all polynomials into a single class by the equivalence relation used to define Div^k . The equivalence classes here are strictly finer than theirs. Moreover, there is a subtle error in the definition of \preceq found in Gersten's original paper (making the equivalence classes far larger than intended) that propagated into the rest of the literature.

Let X be a metric space with basepoint x_0 . Recall from above that a map to X is r -avoidant if its image is disjoint from the ball of radius r about x_0 . We say that a chain (Lipschitz singular or cellular) in X is r -avoidant if its support is disjoint from the ball of radius r about x_0 . For $\rho \leq 1$, we say that X is (ρ, k) -acyclic at infinity if for every r -avoidant n -cycle a in X , where $0 \leq n \leq k$, there exists a ρr -avoidant $(k + 1)$ -chain b with $\partial b = a$. For fixed k , we sometimes write $\bar{\rho}$ for the supremum of the values for which (ρ, k) -acyclicity at infinity holds. Note that if X is (ρ, k) -acyclic at infinity for any ρ then it is k -acyclic at infinity (cf. [BM01] for the definition of acyclicity at infinity). The converse is false in general, and we will discuss the special case of right-angled Artin groups in more detail in the next section.

For a metric space X , define the *divergence dimension* $\text{divdim}(X)$ to be the largest whole number $k \leq \dim X - 2$ such that X is (ρ, k) -acyclic at infinity for some $0 < \rho \leq 1$. For instance, $\text{divdim}(\mathbb{R}^d) = \text{divdim}(\mathbb{H}^d) = d - 2$. The reason for the dimension bound is that if X has dimension less than or equal to k , then there is no $(k + 1)$ -volume to measure, and if X has dimension $k + 1$ then the functions $\text{div}_{\rho, \alpha}^k$ defined below may exist for some values of α and ρ but not others. We will define Div^k for $k \leq \text{divdim}(X)$.

The definition of Div^k will be a bit special when $k = 0$, so we deal with that case later. Suppose first that $1 \leq k \leq \text{divdim}(X)$. Given a k -cycle a , we define

$$\text{div}_{\rho}^k(a, r) := \inf \text{mass } b,$$

where the inf is over ρr -avoidant Lipschitz $(k + 1)$ -chains b such that $\partial b = a$. We then define

$$\text{div}_{\rho}^k(l, r) := \sup \text{div}_{\rho}^k(a, r), \quad (k > 0)$$

where the sup is over r -avoidant Lipschitz k -cycles a of mass at most l .

In order to see the effect of removing a ball from the space, consider what happens as r and l go to infinity simultaneously. In the nonpositively curved setting, the difficulty of filling spheres that arise as the intersection of a large ball around the basepoint with a flat of rank $k + 1$ tends to be a distinguishing feature of the asymptotic geometry (as in the results for symmetric spaces referenced above). These spheres have $l = O(r^k)$, and so we can obtain useful information by specializing

to spheres whose mass is of this order. We therefore introduce a new parameter α and write $\text{div}_{\rho,\alpha}^k(r)$ for $\text{div}_{\rho}^k(\alpha r^k, r)$. Then, formally, $\text{Div}^k(X)$ is the two-parameter family of functions:

$$\text{Div}^k(X) := \{\text{div}_{\rho,\alpha}^k(r)\}_{\alpha,\rho}, \quad (k > 0)$$

where $\alpha > 0$ and $0 \leq \rho \leq \bar{\rho}$.

In the case $k = 0$, we are filling pairs of points (0-cycles) by paths (1-chains). The 0-mass of a cycle does not restrict its diameter, so instead we require the 0-cycle to lie on the boundary of the deleted ball. Set

$$\text{div}_{\rho}^0(r) := \sup_{x,y \in S_r} \inf_P |P|,$$

where the sup is over pairs of points on S_r and the inf is over ρr -avoidant paths P with endpoints x and y .

In this case we get a one-parameter family of functions of one variable:

$$\text{Div}^0(X) := \{\text{div}_{\rho}^0(r)\}_{\rho},$$

where $0 \leq \rho \leq \bar{\rho}$.

Let $F = \{f_{\rho,\alpha}\}$ and $F' = \{f'_{\rho,\alpha}\}$ be two-parameter families of functions $\mathbb{R}^+ \rightarrow \mathbb{R}^+$, indexed over $\alpha > 0$ and $0 < \rho \leq \bar{\rho}$. Then we write $F \preceq F'$ if there exist thresholds $0 < \rho_0 \leq \bar{\rho}$, $\alpha_0 \geq 0$, and constants $L, M > 1$ such that for all $\rho \leq \rho_0$ and all $\alpha \geq \alpha_0$,

$$f_{\rho,\alpha} \preceq f'_{L\rho, M\alpha}.$$

Here the relation \preceq is the one defined in the previous section for functions $\mathbb{R}^+ \rightarrow \mathbb{R}^+$.




FIGURE 1. Comparison of two-parameter families of functions: for each coordinate position in the rectangle on the left, there is a corresponding position in the rectangle on the right. We say $F \preceq F'$ if the functions in those positions satisfy $f \preceq f'$ over the whole rectangle.

From this description it is clear that \preceq is a partial order. Finally, $F \asymp F'$ if and only if $F \preceq F'$ and $F \succeq F'$.

Proposition 2.3 (Quasi-isometry invariance). *Let X and Y be cell complexes with finitely many isometry types of cells. If X is quasi-isometric to Y then $\text{Div}^k(X) \asymp \text{Div}^k(Y)$ for $0 \leq k \leq \min\{\text{divdim } X, \text{divdim } Y\}$.*

Quasi-isometry invariance allows us to write $\text{Div}^k(G)$ for the equivalence class of two-parameter families $\{\text{Div}^k(X)\}$ where X has a geometric G -action.

Proof. If $k = 0$ then this is a result of Gersten [Ger94b]. Fix $k > 0$ in the indicated range.

Some technical lemmas from [AWP99] imply the following: if cell complexes X and Y are k -acyclic at infinity and have finitely many isometry types of cells, and if X and Y are quasi-isometric, then there are quasi-isometries $\varphi : X \rightarrow Y$ and $\bar{\varphi} : Y \rightarrow X$ that are quasi-inverses of each other and that are cellular and C -Lipschitz on the $(k+1)$ -skeleton for some $C \geq 1$. Fix such maps $\varphi, \bar{\varphi}$ and constant C . It is furthermore possible to choose a constant M , dependent on C and k , such that (i) the mass of any push-forward $\varphi_{\#}(\sigma)$ or $\bar{\varphi}_{\#}(\sigma)$ is at most $M \cdot \text{mass}(\sigma)$ for any Lipschitz k - or $(k+1)$ -chain in X or Y , and (ii) every Lipschitz k -chain a in X is homotopic to $\bar{\varphi}_{\#}\varphi_{\#}(a)$ by a Lipschitz homotopy of mass at most $M \cdot \text{mass}(a)$.

Let $0 < \bar{\rho}(X), \bar{\rho}(Y) \leq 1$ be the maximal values for which $\text{div}_{\rho, \alpha}^k(X)$ and $\text{div}_{\rho, \alpha}^k(Y)$ are defined. Let $\rho_0 = \min\{\bar{\rho}(X), C^{-2}\bar{\rho}(Y)\}$. Let $\alpha_0 = 0$, let $L = C^2$, and let M be as described in the previous paragraph. Now fix $0 < \rho \leq \rho_0$ and $\alpha \geq \alpha_0$. We will show that

$$\text{div}_{\rho, \alpha}^k(X) \preceq \text{div}_{L\rho, M\alpha}^k(Y),$$

from which we conclude $\text{Div}^k(X) \preceq \text{Div}^k(Y)$. A symmetric argument shows the other inequality, giving the desired equivalence.

Specifically, let $r > 0$ be given and let a be an r -avoidant Lipschitz k -cycle in X with mass $\leq \alpha r^k$. It suffices to show that a can be filled by a ρr -avoidant Lipschitz $(k+1)$ -cycle b that has mass at most $A \text{div}_{L\rho, M\alpha}^k(Y)(r/C)$ for some constant $A > 0$.

Note that the pushforward $a' = \varphi_{\#}(a)$ is a Lipschitz k -cycle in Y ; it is r/C -avoidant and has mass at most $M\alpha r^k$. Therefore for any $0 < \rho' < \bar{\rho}(Y)$ there exists a filling b' of $\varphi_{\#}(a)$ (that is, b' is a Lipschitz $(k+1)$ -chain) that is $(\rho'r/C)$ -avoidant and that satisfies

$$\text{mass}(b') \leq \text{div}_{\rho', M\alpha}^k(r/C).$$

Choose $\rho' = L\rho$, so that b' is $C\rho r$ -avoidant in Y .

Now consider $b'' = \bar{\varphi}_{\#}(b')$. This is a Lipschitz $(k+1)$ -chain in X that is $\rho'r/C^2 = \rho r$ -avoidant and that has mass at most $M \text{mass}(b')$. However b'' is not quite a filling of a ; we know only that its boundary a'' is bounded distance from a . Since $a'' = \bar{\varphi}_{\#}\varphi_{\#}(a)$ there is a (Lipschitz) homotopy between a'' and a with mass at most $M \cdot \text{mass}(a)$; thus we have

$$\text{div}_{\rho, \alpha}^k(X)(r) \leq M \text{div}_{L\rho, M\alpha}^k(Y)(r/C) + M\alpha r^k.$$

Now since X and Y are at least $(k+2)$ -dimensional, there exist avoidant k -cycles in each space with filling volume on the order of their mass. (In fact, dimension $k+1$ suffices; if the dimension of X were just k , then every k -cycle would have filling volume 0.) Thus we can ignore the last term of the above inequality and we have $\text{div}_{\rho, \alpha}^k(X) \preceq \text{div}_{L\rho, M\alpha}^k(Y)$, as desired. \square

For a function $h : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, we say that the family F has order $h(r)$ and write $F \asymp h(r)$ if $F \asymp \{h(r)\}$, that is, F is equivalent to the family that contains the same function $h(r)$ for every value of the parameters. Then $F \preceq h(r)$ and $F \succeq h(r)$ can be defined similarly.




FIGURE 2.

Remark 2.4 (Remarks on comparison and equivalence).

- (1) Note that the statement $h(r) \preceq \text{Div}^k(X)$ is equivalent to the statement that there exist ρ'_0 and α'_0 such that $h(r) \preceq \text{div}_{\rho,\alpha}(r)$ for all $\rho \leq \rho'_0$ and $\alpha \geq \alpha'_0$. (Here $\rho'_0 = \rho_0 L$ and $\alpha'_0 = \alpha_0 M$ in the definition of \preceq .)

Similarly, note that if ρ or α is decreased, then the value of $\text{div}_{\rho,\alpha}^k(r)$ decreases. Thus in order to establish that $\text{Div}^k(X) \preceq h(r)$, it suffices to show that there exist $\rho_0 \leq 1$ and $\alpha_0 \geq 0$ such that $\text{div}_{\rho_0,\alpha}(r) \preceq h(r)$ for all $\alpha \geq \alpha_0$.

Taken together, these give sufficient criteria to establish that $\text{Div}(X) \asymp h(r)$, as shown in Figure 2. However, for a particular X there is no guarantee that $\text{Div}(X) \asymp h(r)$ for any h .

- (2) Every family F satisfies $F \succeq r$, since all sublinear functions are equivalent to the linear function r . In particular $\text{Div}^0(G) \succeq r$ for any group G . The following (true) statement is slightly stronger in two ways: if G is a finitely generated infinite group, then for every ρ we have $\text{div}_\rho^0(r) \geq 2r$ for all $r > 0$.
- (3) For CAT(0) spaces with extendable geodesics, one sees only one \asymp class of functions in $\text{Div}^0(X)$. On the other hand, for $k \geq 1$ the family $\text{Div}^k(X)$ often contains functions from multiple \asymp classes, as we will see in the example of orthoplex groups in §5.

For groups that are direct products, it is easy to see that Div^0 is exactly linear; we will show this below (Lemma 7.2).

2.3. Example: Euclidean space. Constructing avoidant fillings is sometimes difficult; removing a ball of radius r from a space breaks its symmetry, making it harder to apply methods from group theory. One method of constructing avoidant fillings is to first construct a filling in the entire space, then modify that filling to be avoidant. In this paper, we modify fillings using maps $X \rightarrow X - B_r$; we call these pushing maps.

Our constructions generally follow the following outline: given an avoidant k -cycle a in X , we will find a filling b (typically not avoidant) of a and a pushing map $X \rightarrow X - B_r$, where B_r is the ball of radius r . This map generally has singularities, and we use techniques from geometric measure theory to move b off these singularities.

The basic example to consider is \mathbb{R}^d , where one has the “pushing” map given by radial projection to $\mathbb{R}^d - B_r$, namely

$$(2.1) \quad \pi_r(v) = \begin{cases} r \frac{v}{\|v\|}, & v \in B_r \\ v, & v \notin B_r. \end{cases}$$

This map is undefined at 0, but if the filling has dimension $< d$, it can be perturbed to miss the origin, and the Federer-Fleming Deformation Theorem (Theorem 2.1) can be used to control the volume. The theorem allows us to approximate singular k -chains in X by cellular k -chains in X , and if the k -skeleton of X misses the singularity, then so will the approximation.

We will prove that filling an avoidant cycle by an avoidant chain is roughly as hard as filling a cycle by a chain. Specifically, we will show the following.

Proposition 2.5 (Euclidean bounds). *Let d be a positive integer and let $1 \leq k \leq \text{divdim}(\mathbb{R}^d) = d - 2$. There is a constant c depending only on the dimension d such that if $r, l \geq 0$ and a is an r -avoidant k -cycle in \mathbb{R}^d of mass at most l , then there is an r -avoidant $(k + 1)$ -chain b such that $\partial b = a$ and*

$$\text{mass } b \leq cl^{\frac{k+1}{k}}.$$

Further, there is a constant c' depending only on d and an r -avoidant k -cycle a with mass l such that for every chain b with $\partial b = a$,

$$\text{mass } b \geq c'l^{\frac{k+1}{k}}.$$

This implies, in particular, that for $0 \leq k \leq d - 2$,

$$\text{Div}^k(\mathbb{R}^d) \asymp r^{k+1}.$$

As \mathbb{R}^d is a model space for \mathbb{Z}^d , Proposition 2.3 gives

$$\text{Div}^k(\mathbb{Z}^d) \asymp r^{k+1}.$$

Proof. Wenger’s work (Prop. 2.2, though cf. Federer and Fleming [FF60] in the Euclidean case) implies that there exists an $m > 0$ independent of a and there exists a chain b such that $\partial b = a$, $\text{mass } b \leq ml^{(k+1)/k}$, and $\text{supp } b$ is contained in a $ml^{1/k}$ -neighborhood of $\text{supp } a$. We will modify this to find an avoidant cycle, using different arguments when $l \leq r^k$ and when $l \geq r^k$.

First, set $c_0 = (2m)^{-k}$ and note that if $l \leq c_0 r^k$ then b is $r/2$ -avoidant. In this case $b' = (\pi_r)_\#(b)$ fills a , is r -avoidant, and satisfies

$$\text{mass } b' \leq 2^k \text{mass } b \leq 2^k ml^{(k+1)/k},$$

so the conclusion of the lemma holds.

We can thus assume that $l \geq c_0 r^k$. We will show the proposition when $r = 1$, and then use scaling to prove the general case. Let a be a 1-avoidant Lipschitz k -cycle of volume $l \geq c_0$. There is a $(k + 1)$ -chain b (which is not necessarily avoidant) filling a such that $\text{mass } b \leq ml^{\frac{k+1}{k}}$. We will approximate b by a cellular chain, then “push” it out of the 1-sphere.

Let τ be the grid of cubes of side length $\frac{1}{2d}$ translated so that the center of one of the cubes lies at the origin. Let $P(b)$, $Q(b)$, and $R(b)$ be as in Theorem 2.1, so that $P(b)$ is a chain in $\tau^{(k+1)}$ approximating b , and

$$\partial R(b) = \partial b - \partial P(b) = a - \partial P(b).$$

Each cell of τ has diameter at most $1/2$, so the smallest subcomplex of τ containing the support of a is $1/2$ -avoidant. It follows that $R(b)$ is $1/2$ -avoidant. Since any k -cell of τ is $1/4d$ -avoidant, so is $P(b)$. Thus $b' := R(b) + P(b)$ is $1/4d$ -avoidant filling of a . Further, there is a constant c_1 , which comes from Theorem 2.1 and depends only on d , such that

$$\text{mass } b' \leq c_1(\text{mass } a + \text{mass } b) \leq c_1(l + ml^{\frac{k+1}{k}}) \leq c_2l^{\frac{k+1}{k}}$$

for some c_2 , where the last bound uses the lower bound on l .

Pushing b' forward under the map π_1 defined in (2.1), we get a chain $b'' := (\pi_1)_\#(b')$. This is a 1-avoidant filling of a . Furthermore, since b' is $1/4d$ -avoidant, π_1 is $4d$ -Lipschitz on the support of b' , so there is a constant c such that

$$\text{mass } b'' \leq (4d)^{k+1} \text{mass } b' \leq c \text{mass } a^{\frac{k+1}{k}}$$

as desired.

Now, consider the case that $r \neq 1$. Let $s_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be the homothety $v \mapsto tv$. If γ is a k -chain, then $\text{mass } s_t(\gamma) = t^k \text{mass } \gamma$.

If a is an r -avoidant Lipschitz k -cycle of mass l , then $a_1 = s_{r^{-1}}_\#(a)$ is 1-avoidant, and $\text{mass } a_1 = r^{-k}l \geq c_0$. By the argument above, there is a 1-avoidant $(k+1)$ -chain b_1 filling a_1 such that

$$\text{mass } b_1 \leq c(\text{mass } a_1)^{\frac{k+1}{k}} = cr^{-(k+1)}l^{\frac{k+1}{k}}.$$

Rescaling this by letting $b = (s_r)_\#(b_1)$, we find that b is an r -avoidant filling of a and $\text{mass } b \leq cl^{\frac{k+1}{k}}$.

For the second statement, simply take a to be a round sphere far from B_r . Then the estimate is just the isoperimetric theorem for \mathbb{R}^d . \square

Thus the best avoidant fillings have roughly the same volume as the most efficient (not necessarily avoidant) fillings.

Much of the rest of this paper will be dedicated to generalizing this technique to right-angled Artin groups. These groups act on a complex X which consists of a union of flats, and as with \mathbb{R}^d , we will construct avoidant fillings by using a pushing map to modify non-avoidant fillings. The pushing map is singular in the sense that it cannot be defined continuously on all of X , but as with \mathbb{R}^d , we will delete small neighborhoods of the singularities, enabling us to define the pushing map continuously on a subset of X . Because X generally has more complicated topology than \mathbb{R}^d , the pushing map has more singularities, and these singularities lead to larger bounds on Div^k .

3. BACKGROUND ON RIGHT-ANGLED ARTIN GROUPS

In this section we will introduce some of the key background on RAAGs. We refer the reader to [Cha07] for a more complete treatment.

A right-angled Artin group is a finitely generated group given by a presentation in which all the relators are commutators of generators. A RAAG can be described by a graph which keeps track of which pairs of generators commute, and the structure of this graph affects the geometry of the group and its subgroups. Let Γ be a finite graph with no loops or multiple edges, and with vertices labeled a_1, \dots, a_k . The right-angled Artin group based on Γ is the group

$$A_\Gamma := \langle a_1, \dots, a_k \mid R \rangle,$$

where R is the set of relators

$$R = \{[a_i, a_j] \mid a_i \text{ and } a_j \text{ are connected by an edge of } \Gamma\}.$$

We call A_Γ the *defining graph* of A_Γ . Let L be the flag complex of Γ ; that is, the simplicial complex with the same vertex set as Γ , and in which a set S of vertices spans a simplex if and only if every pair of vertices of S is connected by an edge of Γ .

The group A_Γ acts freely on a CAT(0) cube-complex X_Γ , defined as follows. Let Y be a subcomplex of the torus $(S^1)^k$, where the circle S^1 is given a cell structure with one 0-cell and one 1-cell, and each S^1 factor corresponds to a vertex a_i . Thus Y is a cube complex with one vertex. A d -cell σ of $(S^1)^k$ is contained in Y if and only if the vertices corresponding to the S^1 factors of σ span a simplex in L . Then X_Γ is the universal cover of Y .

Since Y has one vertex, all the vertices of X_Γ are in the same A_Γ -orbit. We pick one of the vertices of X_Γ as a basepoint, which we denote e , and identify the element $a \in A_\Gamma$ with the vertex $a \cdot e$ of X_Γ . We will often refer to elements of A_Γ and vertices of X_Γ interchangeably. Similarly, the edges of Y correspond to the generators of A_Γ , and we say that an edge of X_Γ is *labeled* by its corresponding generator.

Since each cell in Y is part of a torus, each cell of X_Γ is part of a flat. Typically, each cell is part of infinitely many flats, but we can use the group structure to pick a canonical one. If σ is a d -cube in X , its edges are labeled by d different generators; if S is this set of labels, and if $v \in A_\Gamma$ is a vertex of σ , then the elements of S commute and generate an undistorted copy of \mathbb{Z}^d , which we denote A_S . This subgroup spans a d -dimensional flat through the origin, and the translate $v \cdot A_S$ spans a flat containing σ , which we call the *standard flat* containing σ and denote by F_σ .

The link of a vertex of X_Γ , which we denote $S(L)$, is the union of the links of all the standard flats. It has two vertices for each generator v of A_Γ , one corresponding to v and one to v^{-1} . We will denote the vertex in the v direction by $\hat{v} = +\hat{v}$ and the one in the v^{-1} direction by $-\hat{v}$. Furthermore, if v_1, \dots, v_d are the vertices of a simplex Δ of L , that simplex corresponds to a d -dimensional standard flat containing x . The link of this flat is an *orthoplex* (i.e., the boundary of a cross-polytope; in the case that $d = 3$, it is an octahedron), so Δ corresponds to 2^d simplices in $S(L)$, each with vertices $\pm\hat{v}_1, \dots, \pm\hat{v}_d$. If L has m vertices, labeled v_1, \dots, v_m , then $S(L)$ contains 2^m copies of L with vertices $\pm v_1, \dots, \pm v_m$; we call these *signed copies* of L .

We will use two metrics on X_Γ . The first metric on X_Γ , with respect to which it is CAT(0), is the Euclidean (or ℓ^2) metric on each cube, extended as a length metric to X_Γ . (That is, the distance between two points is the infimal length of a path connecting them, where length is measured piecewise within each cube.)

The second metric restricts instead to the ℓ^1 metric on each cube, and is extended as a length metric from the cubes to the whole space. This has the property that its restriction to the one-skeleton of X_Γ is the word metric on a Cayley graph for A_Γ . We will mainly use the ℓ^1 metric to define balls and spheres in X_Γ which coincide with balls and spheres in A_Γ . The notation $B_r(x) := \{y \in X : d_{\ell^1}(x, y) < r\}$ will denote the open ℓ^1 ball and $\bar{B}_r(x)$, $S_r(x)$ will denote the closed ball and sphere, respectively, so that $B_r(x) \sqcup S_r(x) = \bar{B}_r(x)$. When there is no center specified for

a ball, it is taken to be centered at the basepoint e . Note that all words whose reduced spellings have length r are vertices in S_r .

As an illustration, \mathbb{Z}^3 is a RAAG, and the corresponding X_Γ is \mathbb{R}^3 , with the structure of a cube complex. The sphere of radius r in the ℓ^1 metric is a Euclidean octahedron with equilateral triangle faces. All vertices corresponding to group elements of length r in the word metric lie on this sphere.

Recall that RAAGs themselves, being CAT(0) groups, have at worst Euclidean Dehn functions (Proposition 2.2). To find bigger Dehn functions, one must look at subgroups such as those defined in the next section.

We will study divergence functions for RAAGs below, so we remark that the divergence dimension can be read off of the defining graph. Brady and Meier showed that the group A_Γ is k -acyclic at infinity if and only if the link $S(L)$ is k -acyclic. In fact, their construction shows that k -acyclicity of the link is equivalent to $(1, k)$ -acyclicity at infinity of the group (and therefore (ρ, k) -acyclicity at infinity for all $0 < \rho \leq 1$). Furthermore, it is clear that the dimension of X_Γ is the size of the largest clique in the defining graph Γ . Thus, $\text{divdim}(A_\Gamma)$ is the largest k such that $S(L)$ is k -acyclic and there exists a clique of $k + 2$ vertices in Γ .

3.1. Bestvina-Brady groups. Let $h : A_\Gamma \rightarrow \mathbb{Z}$ be the homomorphism which sends each generator to 1; we call h the *height function* of A_Γ . Let $H_\Gamma := \ker h$; a group H_Γ constructed in this fashion is called a Bestvina-Brady group. These subgroups were studied by Bestvina and Brady in [BB97], and provide a fertile source of examples of groups satisfying some finiteness properties but not others. Brady [BRS07] showed that there are graphs Γ such that H_Γ has a quartic (l^4) Dehn function, and Dison [Dis08] recently showed that this is the largest Dehn function possible, that is, the Dehn function of any Bestvina-Brady group is at most l^4 . We will generalize Dison's result to higher-order Dehn functions in Section 5 below.

Abusing notation slightly, let $h : X_\Gamma \rightarrow \mathbb{R}$ also denote the usual height map defined by linear extension of the homomorphism above; it is a Morse function on X_Γ , in the sense of [BB97]. Let $Z_\Gamma = h^{-1}(0)$ be the zero level set. The action of A_Γ on X_Γ restricts to a geometric action of H_Γ on Z_Γ . The topology of Z_Γ is closely related to Γ ; indeed, if L is the flag complex corresponding to Γ , then Z_Γ contains infinitely many scaled copies of L and is homotopy equivalent to a wedge of infinitely many copies of L [BB97]. In particular, if L is k -connected, then Z_Γ is also k -connected, so H_Γ is type F_k .

3.2. Tools for RAAGs. We introduce several basic tools: the *orthant* associated to a cube in the complex $X = X_\Gamma$, a related *scaling map* on X , and an *absolute value map* on X .

Fix attention on a particular d -cube σ in X and let v be its closest vertex to the origin. The vertices of the standard flat F_σ correspond to a coset vA_S where S is the set of labels on edges of σ . For each $a_i \in S$ let γ_i be the geodesic ray in F_σ that starts at v , traverses the edge of σ labeled a_i in time one, and continues at this speed inside F_σ , so that $\gamma_i(n) = va_i^{\pm n}$, with the sign fixed once and for all depending on whether v or va_i is closer to the origin. We take Orth_σ to be the flat orthant within vA_S spanned by the rays γ_i , so that the cube σ itself is contained in Orth_σ , and v is its extreme point. Note that if τ is a face of σ , then $\text{Orth}_\tau \subset \text{Orth}_\sigma$ as an orthant with appropriate codimension. In particular, if τ is an edge, then Orth_τ is a ray starting at one endpoint of τ and pointing “away” from e .

Next we define a *scaling map* $s_r : S(L) \rightarrow X_\Gamma$. The sphere S_1 (the unit sphere in the ℓ^1 metric) is homeomorphic to $S(L)$; we associate points of $S(L)$ with points of S_1 . Because of the abundance of flats in X_Γ , these correspond to canonically extendable directions in X_Γ , as follows. If $x \in S_1$, then x is in some maximal cube σ corresponding to commuting generators; we define $\gamma_x : [0, \infty) \rightarrow X_\Gamma$ to be the unique geodesic ray in F_σ that is based at the identity, goes through x , and is parametrized by arc length in the ℓ^1 metric. For instance, if x is a vertex corresponding to a generator a , then γ_x is a standard ray along edges labeled a , so that $\gamma_x(n) = a^n$ for $n = 0, 1, 2, \dots$. Note that the map $x \mapsto \gamma_x$ is continuous. The scaling map is defined by $s_r(x) = \gamma_x(r)$.

Finally we define the *absolute value* map. Given an element $g \in A_\Gamma$, let

$$w = a_{i_1}^{\pm 1} \dots a_{i_r}^{\pm 1}$$

be a geodesic word representing g . Then the absolute value of g is given by

$$\text{abs}(g) = a_{i_1} \dots a_{i_r}.$$

We claim this is well-defined; indeed, if w and w' are two geodesic words representing g , then w can be transformed to w' by a process of switching adjacent commuting letters. If $a_{i_j}^{\pm 1}$ and $a_{i_{j+1}}^{\pm 1}$ commute, then so do a_{i_j} and $a_{i_{j+1}}$, so the choice of geodesic spelling does not affect $\text{abs}(g)$.

The absolute value map preserves adjacencies. If g_1 and g_2 are adjacent in the Cayley graph of A_Γ , then since A_Γ has no relations of odd length, we may assume that $|g_1| + 1 = |g_2|$. Let a be a generator such that $g_2 = g_1 a^{\pm 1}$. If w is a geodesic word representing g_1 , then $wa^{\pm 1}$ is a geodesic word representing g_2 , so $\text{abs}(g_2) = \text{abs}(g_1)a$.

Like the height function h , abs can be extended to X_Γ by extending linearly over each face. This extension is 1-Lipschitz and satisfies the property that $h(\text{abs}(x)) = |x| = |\text{abs}(x)|$ for all $x \in X_\Gamma$; in other words, abs maps the r -sphere S_r into the coset $h^{-1}(r)$ of H_Γ . Furthermore, abs is idempotent; if $h(x) = |x|$, then $\text{abs}(x) = x$.

4. PUSHING MAPS

Throughout this section, let Γ be the defining graph of a RAAG, let $X = X_\Gamma$, $A = A_\Gamma$, and $Z = Z_\Gamma$.

In §2.3, we constructed avoidant fillings in \mathbb{R}^d by using a pushing map $\mathbb{R}^d - \{0\} \rightarrow \mathbb{R}^d - B_r$; in this section, we will construct pushing maps for general RAAGs. Here the pushing maps become more complicated, because typically there is branching at the vertices. This has two implications: first, there are many more singularities to work around, and second, one needs to be careful in computing the amount by which the map distorts volumes. In contrast to the situation in \mathbb{R}^d , where avoidant fillings are roughly the same size as ordinary fillings, avoidant fillings in RAAGs (when they exist) may be much larger.

We will define two pushing maps: heuristically, one pushes radially to $X - B_r$ from the basepoint, and the other pushes along the height gradient to the 0-level set Z . The constructions are similar; in both cases we delete neighborhoods of certain vertices, put a cell structure on the remaining space, and then define a map cell by cell to the target space. The Lipschitz constants of the maps produce upper bounds on the volume expansion of fillings.

As the theorems are formulated, the pushing maps are not cellular, but it is easy to modify the cell structure of X to make the pushing maps cellular. Alternatively,

if one has applied a pushing map to a filling, one can modify the pushed filling to be cellular. Either way, the same volume expansion can be observed by counting cells rather than measuring the Lipschitz constant. Thus these pushing maps, and the results that use them in the following sections, are compatible with the admissible maps definitions found in the Dehn function literature.

Given $X = X_\Gamma$, we define modified spaces X_r and Y , which have ℓ^1 -neighborhoods of some vertices removed. Let

$$X_r = X - \bigcup_{v \in B_r} B_{1/4}(v) \quad \text{and} \quad Y = X - \bigcup_{v \notin Z} B_{1/4}(v).$$

Endow X_r and Y with their length metrics from the ℓ^1 metric on X . In a moment we will describe cell structures on the spaces X_r and Y . Recall that B_t denotes the open ball, so that a modified space $X - B_t(v)$ includes the boundary $S_t(v)$.

Theorem 4.1 (Radial pushing map). *For $r > 0$ there are Lipschitz retractions*

$$\mathcal{P}_r : X_r \rightarrow X - B_r$$

which are $O(r)$ -Lipschitz. That is, there is a constant $c = c(\Gamma)$ such that \mathcal{P}_r is $(cr + c)$ -Lipschitz for each r .

Proof. We first give X_r the structure of a polyhedral complex. If σ is a cell of X that intersects X_r , then $\sigma \cap X_r$ is a cube or a truncated cube, and we let $\sigma \cap X_r$ be a cell of X_r . We call such a cell an *original face* of X_r . The boundary of an original face consists of other original faces (cubes and truncated cubes) as well as some simplices arising from the truncation; these simplices are also cells of X_r and we call them *link faces*. Note that every link face is a simplex in a translate of $S_{1/4}$, which we identify with $S(L)$. If τ is an original face of X_r , so that $\tau = \sigma \cap X_r$ for some cube σ of X , then we define $\text{Orth}_\tau = \text{Orth}_\sigma$.

The pushing map \mathcal{P}_r is the identity on $X - B_r$. To define it on $X_r \cap B_r$, we will first define a map on the original edges, then extend it linearly to the link faces, and finally extend inductively to the remaining original faces. We will require that $\mathcal{P}_r(\tau \cap B_r) \subset \text{Orth}_\tau \cap S_r$ for every original face τ of X_r . Note that $\text{Orth}_\tau \cap S_r$ is the intersection of Orth_τ with a hyperplane.

We first consider the original edges of X_r . If τ is an original edge, it is part of a ray Orth_τ (in X) traveling away from the origin; we push all its points along Orth_τ until they hit S_r , setting $\mathcal{P}_r(\tau) = \text{Orth}_\tau \cap S_r$ (so that the image of each such edge is a single point on S_r). To be explicit, suppose τ comes from an edge with endpoints v and va , with $|v|, |va| \leq r$. Then for a point $x \in \tau$,

$$\mathcal{P}_r(x) = \begin{cases} va^{r-|v|}, & |va| > |v|; \\ va^{|v|-r}, & |va| < |v|. \end{cases}$$

Then the image of each original edge is a point in S_r and the images of adjacent edges are separated by distance $\leq r$. If two edges of X are adjacent, points on the corresponding original edges in X_r are separated by distance at least $1/4$, so the map is Lipschitz on the edges with constant $\leq r$.

If σ is a link face, and τ is an original face which contains σ , then \mathcal{P}_r sends the vertices of σ to points in $\text{Orth}_\tau \cap S_r$. We can extend \mathcal{P}_r linearly to σ : every point x of σ is a unique convex combination of its vertices, so we define the image $\mathcal{P}_r(x)$ to be the same convex combination of the images of the vertices. This is clearly independent of our choice of τ , so the map is well-defined. Since incident edges have images no more than $O(r)$ apart, this extension is also $O(r)$ -Lipschitz.




FIGURE 3. This figure shows a portion of X with X_r (a union of truncated squares) shaded. The vertical and horizontal edges are *original edges* and the thick diagonal lines are *link edges*. The map \mathcal{P}_r sends each original edge to a point on S_r , and sends each link edge to a line segment in S_r (vertices perturbed to avoid overlaps). The boundary of each octagonal cell is sent to a loop of length $\asymp r$ with zero area.

It remains to define \mathcal{P}_r on the rest of the original faces. We will proceed inductively on the dimension of the faces, using edges as the base case. Recall that if $K \subset \mathbb{R}^d$ is a convex subset of Euclidean space and $f : S^m \rightarrow K$ is a Lipschitz map, we can extend f to the ball D^{m+1} via

$$g(t, \theta) = f(x_0) + t(f(\theta) - f(x_0)),$$

where x_0 is a basepoint on S^m and $t \in [0, 1]$, $\theta \in S^m$ are polar coordinates for D^{m+1} . This extension has Lipschitz constant bounded by a multiple of $\text{Lip}(f)$.

Each d -dimensional original face τ of X_r is of the form $\tau = \sigma \cap X_r$ for some d -cube σ of X . We may assume by induction that \mathcal{P}_r is already defined on the boundary of τ . Define \mathcal{P}_r to be the identity on $\tau - B_r$. The remaining part, $\tau' = \tau \cap B_r$, is isometric to one of finitely many polyhedra, so it is bilipschitz equivalent to a ball D^d , with uniformly bounded Lipschitz constant.

Since the boundary of τ' is mapped to a convex subset of a flat, namely $\text{Orth}_\sigma \cap S_r$, and the Lipschitz constant of $\mathcal{P}_r|_{\partial\tau'}$ is $O(r)$, we can extend \mathcal{P}_r to a Lipschitz map sending τ' to $\text{Orth}_\sigma \cap S_r$ which is again $O(r)$ -Lipschitz, with the constant enlarged by a factor depending on the dimension. \square

A similar pushing map on RAAGs can be used to find bounds on higher-order fillings in Bestvina-Brady groups. We proceed slightly differently here; instead of defining a map on the original edges and extending it to the rest of the space, we start with a map on the link faces and extend it.

Recall that the vertices of $S(L)$ are labeled $\pm\hat{a}_1, \dots, \pm\hat{a}_d$, where the a_i are generators of A . Let $\text{pos} : S(L) \rightarrow S(L)$ be the simplicial map satisfying $\text{pos}(\pm\hat{a}_i) = \hat{a}_i$. The image of this map is the copy of L inside $S(L)$ whose vertices all have positive signs; we denote this *positive link* by $S(L)^+$. Similarly define $\text{neg} : S(L) \rightarrow S(L)$ so that $\text{neg}(\pm\hat{a}_i) = -\hat{a}_i$ and define the *negative link* $S(L)^-$ to be its image.

Theorem 4.2 (Height-pushing map). *There is an H -equivariant retraction*

$$\mathcal{Q} : Y \rightarrow Z,$$

such that the Lipschitz constant of \mathcal{Q} grows linearly with distance from Z . That is, there is a uniform constant $c = c(\Gamma)$ such that the restriction of \mathcal{Q} to $h^{-1}([-t, t])$ is $(ct + c)$ -Lipschitz.

Proof. Here, instead of mapping an original face τ into $\text{Orth}_\tau \cap S_r$, we map it to $F_\tau \cap Z$, where F_τ is the standard flat containing τ .

Like X_r , the space Y inherits the structure of a polyhedral complex from X . The cells of Y are either original faces or link faces; the union of the link faces is a union of translates of $S_{1/4}$, which we identify with $S(L)$. Thus a link vertex is denoted by $v \cdot (\pm\hat{a})$ for some $v \in A$ and some generator a of A .

We construct a map on the link faces, then extend. Each copy of $S(L)$ in the boundary of Y is essentially the set of directions at some vertex of X . We construct a map $v \cdot S(L) \rightarrow Z$ by flipping each direction, if necessary, to point towards Z , and then pushing along standard rays in X . That is, if $v \in A - H$ and $x \in S(L)$ we define

$$\mathcal{Q}(v \cdot x) = \begin{cases} v \cdot s_{|h(v)|}(\text{neg}(x)) & \text{if } h(v) > 0 \\ v \cdot s_{|h(v)|}(\text{pos}(x)) & \text{if } h(v) < 0. \end{cases}$$

It is easy to check that the image of this map lies in Z . Furthermore, the Lipschitz constant of this map restricted to $v \cdot S(L)$ is $4|h(v)|$, and if σ is a link face of Y contained in an original face τ , then $\mathcal{Q}(\sigma) \subset F_\tau \cap Z$.

We have defined the map on link faces of Y , and we extend to the rest of Y by the same inductive procedure as in Theorem 4.1, obtaining an $O(t)$ Lipschitz constant on $h^{-1}([-t, t])$. \square

Remark 4.3 (Signed copies of L). A fact that will be useful in the sequel is that if L' is a signed copy of L in $S(L)$ (that is, an isomorphic copy of L with some of the vertices of L replaced by their negatives), then $\mathcal{Q}(L')$ is a scaled copy of either $\text{pos}(L)$ or $\text{neg}(L)$; recall that Bestvina and Brady proved that Z is a union of infinitely many scaled copies of L [BB97].

In §5–6, we will start with arbitrary $(k+1)$ -dimensional fillings, and approximate them by fillings in the $(k+1)$ -skeleton. In order to use the pushing maps to construct fillings that are either r -avoidant or in Z , we need to extend the maps \mathcal{P}_r and \mathcal{Q} to the $(k+1)$ -skeleta of the deleted balls. The connectivity hypotheses in the next two lemmas will be satisfied in the applications.

Lemma 4.4 (Extended radial pushing map). *If $S(L)$ is k -connected, the radial pushing map \mathcal{P}_r can be extended to a map*

$$\mathcal{P}_r : X_r \cup X^{(k+1)} \rightarrow X - B_r$$

which is $O(r)$ -Lipschitz.

Proof. This requires extending \mathcal{P}_r to the $(k+1)$ -skeleta of the removed balls. Note that $B_{1/4}$ is the cone over $S(L)$, so it is a simplicial complex in a natural way. We will define a retraction

$$\rho : B_{1/4}^{(k+1)} \cup S(L) \rightarrow S(L)$$

and use it to extend \mathcal{P}_r . We construct ρ by extending the map $id_{S(L)}$ to $B_{1/4}^{(k+1)}$. Since $S(L)$ is k -connected, there is no obstruction to constructing such an extension, and since $S(L)$ is a finite simplicial complex (and thus a compact Lipschitz neighborhood retract) one can choose it to be Lipschitz.

Then if v is a vertex in B_r , we can extend \mathcal{P}_r to $v \cdot B_{1/4}^{(k+1)}$ by letting $\mathcal{P}_r(v \cdot x) = \mathcal{P}_r(v \cdot \rho(x))$. This extension may increase the Lipschitz constant, but we still have $\text{Lip}(\mathcal{P}_r) = O(r)$. \square

Lemma 4.5 (Extended height-pushing map). *If L is k -connected, the height-pushing map \mathcal{Q} can be extended to a map*

$$\mathcal{Q} : Y \cup X^{(k+1)} \rightarrow Z$$

which is $O(r)$ -Lipschitz.

Proof. As before, it suffices to extend \mathcal{Q} over the $(k+1)$ -skeleta of the removed balls. Consider $S(L)$ as a subcomplex of $B_{1/4}$. We can extend $\text{pos} : S(L) \rightarrow S(L)^+$ and $\text{neg} : S(L) \rightarrow S(L)^-$ to maps pos' and neg' defined on $B_{1/4}^{(k+1)} \cup S(L)$; since $S(L)^\pm$ is homeomorphic to L , which is k -connected, there is no obstruction to constructing this extension and the extensions can be chosen to be Lipschitz.

Then we can extend \mathcal{Q} by letting

$$\mathcal{Q}(v \cdot x) = \begin{cases} v \cdot s_{|h(v)|}(\text{neg}'(x)) & \text{if } h(v) > 0 \\ v \cdot s_{|h(v)|}(\text{pos}'(x)) & \text{if } h(v) < 0. \end{cases}$$

for all $x \in B_{1/4}^{(k+1)}$ and $v \in A - H$. Again this extension may increase the Lipschitz constant, but we still have $\text{Lip}(\mathcal{Q}|_{h^{-1}([-t, t])}) \leq ct + c$. \square

We will gently abuse notation so that if α is a chain or cycle then we will write the push-forward map $\mathcal{Q}_\#(\alpha)$ as simply $\mathcal{Q}(\alpha)$, and similarly for \mathcal{P}_r .

5. DEHN FUNCTIONS IN BESTVINA-BRADY GROUPS

The kernel H_Γ of the height map acts geometrically on the zero level set Z , so the Dehn function δ_H measures the difficulty of filling in Z .

When H_Γ is of type F_{k+1} , the results of [BB97] imply that L is k -connected. This means that we have a height-pushing map $\mathcal{Q} : Y \cup X^{(k+1)} \rightarrow Z$ as defined in the previous section (Lemma 4.5). The fact that it is $O(t)$ -Lipschitz on the part of X up to height t immediately yields bounds on higher Dehn functions; we will see below that these bounds turn out to be sharp.

Theorem 5.1 (Dehn function bound for kernels). *If $H = H_\Gamma$ is a Bestvina-Brady group and H is type F_{k+1} , then*

$$\delta_H^k(l) \preceq l^{2(k+1)/k}.$$

The proof is straightforward: push the CAT(0) filling (Proposition 2.2) into the zero level set, and observe that its volume can't have increased too much while pushing.

Proof. Let a be a Lipschitz k -cycle in Z of mass at most l , and let $t_a = l^{1/k}$. By Federer-Fleming approximation (Theorem 2.1), we may assume that a is supported in $Z^{(k)} = X^{(k+1)} \cap Z$.

We know that \mathcal{Q} has Lipschitz constant $ct + c$ on heights up to t . Since X is CAT(0), there is a constant $m > 0$ and a chain $b \in C_{k+1}^{\text{Lip}}(X)$ such that $\text{mass } b \leq mt_a^{k+1}$ and b is supported in a mt_a -neighborhood of $\text{supp } a$. In particular, the height is bounded: $h(\text{supp } b) \subset [-mt_a, mt_a]$. Approximating again, we may assume that b is supported in $X^{(k+1)}$. Then $b' = \mathcal{Q}(b)$ is a $(k+1)$ -chain in Z whose boundary is a , and

$$\delta_H^k(a) \preceq \text{mass } b' \preceq (cmt_a + c)^{k+1} \cdot \text{mass } b \preceq t_a^{2k+2} = l^{\frac{2(k+1)}{k}}.$$

□

This recovers a theorem of Dison [Dis08] in the case $k = 1$.

Brady [BRS07] constructed examples of Bestvina-Brady groups with quartic (l^4) Dehn functions, showing that this upper bound is sharp when $k = 1$. We next generalize these examples to find Bestvina-Brady groups with large higher-order Dehn functions, showing that the upper bound in the previous theorem is sharp for all k .

Definition 5.2 (Orthoplex groups). Recall that a k -dimensional orthoplex (also known as a cross-polytope) is the join of $k+1$ zero-spheres. The standard k -orthoplex is the orthoplex in \mathbb{R}^{k+1} whose extreme points are $\pm e_i$ for the standard basis vectors $\{e_i\}_{i=0}^k$. For $k \geq 0$, we call a graph Γ a k -orthoplex graph, and its associated group A_Γ a k -orthoplex group, if the flag complex L on Γ has the following properties:

- L is a $(k+1)$ -complex that is a triangulation of a $(k+1)$ -dimensional ball;
- the boundary of L is isomorphic to a k -dimensional orthoplex;
- there exists a top-dimensional simplex in L whose closure is contained in the interior of L . We call this a *strictly interior* simplex.

The boundary is isomorphic as a complex to the standard orthoplex, so it has $2(k+1)$ vertices that we label by a_i, b_i for $0 \leq i \leq k$, where a_i corresponds to e_1 and b_i to $-e_i$.

For example, a path with at least three edges is a 0-orthoplex graph.

If A_Γ is a k -orthoplex group, then the flag complex L is $(k+1)$ -connected. It follows that the associated Bestvina-Brady group H_Γ is of type F_{k+1} , and Z is $(k+1)$ -connected [BB97].

Theorem 5.3 (Kernels of orthoplex groups have hard-to-fill spheres). *If A_Γ is a k -orthoplex group, then*

$$\delta_H^k(l) \succeq l^{2(k+1)/k}.$$




FIGURE 4. The figure on the left is a 1-orthoplex graph. It is shown in [BRS07] that H_Γ has quartic Dehn function for this Γ . On the right is a 2-orthoplex graph. Note that a copy of the 1-orthoplex graph appears on the “equator” of the 2-orthoplex example; similarly, the 1-orthoplex graph contains an “equatorial” path of length three, which is a 0-orthoplex graph. These are not the simplest examples of k -orthoplex graphs for $k = 1, 2$, but they are the simplest symmetric examples.

Proof. We write A , X , Z , and H as usual, with $h : Z \rightarrow \mathbb{R}$ the height function. Let γ_i , $0 \leq i \leq k$ be the bi-infinite geodesic along edges of X such that $\gamma_i(0) = e$,

$$\gamma_i|_{\mathbb{R}^+} = a_i b_i a_i b_i \dots,$$

and

$$\gamma_i|_{\mathbb{R}^-} = b_i a_i b_i a_i \dots;$$

that is, $\gamma_i(-1) = b_i$, $\gamma_i(-2) = b_i a_i$, etc. The idea of this proof is that the γ_i span a flat that is not collapsed very much by being pushed down to Z , so that we can get quantitative control on the filling volume in Z for spheres coming from that flat.

In this proof we will adopt the notation that $x = (x_0, \dots, x_k) \in \mathbb{Z}^{k+1}$. Let $F : \mathbb{Z}^{k+1} \rightarrow A$ be given by

$$F(x) = \prod_{i=0}^k \gamma_i(x_i);$$

note that $\gamma_i(n)$ commutes with $\gamma_j(m)$ for all $i \neq j$ and all $n, m \in \mathbb{Z}$ and that the image of F lies in the non-abelian subgroup $\langle a_i, b_i \rangle$. The image of F forms the set of vertices of a (nonstandard) $(k+1)$ -dimensional flat \bar{F} . This flat \bar{F} is entirely at non-negative height, with a unique vertex, $F(0, \dots, 0)$, at height zero. Since $h(F(x)) = \sum |x_i|$, if $r > 0$, then the part of \bar{F} at height r is an orthoplex and is homeomorphic to S^k . Define a k -sphere Σ_r to be a translate of this sphere back to height zero:

$$\Sigma_r := [(a_0)^{-r} \bar{F}] \cap Z.$$

We will regard Σ_r as a Lipschitz k -cycle and show that it is difficult to fill in Z .

Recall that Z is a $(k+1)$ -connected $(k+1)$ -dimensional complex. If a is a cellular k -cycle in Z , it has a unique cellular filling b_0 , and this filling has minimal mass among Lipschitz fillings. If $b = \sum b_i \sigma_i$ is a Lipschitz $(k+1)$ -chain filling a , where $b_i \in \mathbb{Z}$ and σ_i are maps from the $(k+1)$ -simplex to Z , then it may have larger mass than b_0 , because the different simplices making up b may partially cancel.




FIGURE 5. The $k = 1$ case. The vertical and horizontal rays in the figure on the left fit together to form the geodesic γ_1 , and the other two rays fit together to form γ_0 . The four quadrants glue together to form the plane \bar{F} (a nonstandard 2-flat). In red we see the orthoplex $\bar{F} \cap h^{-1}(r)$ whose translate is a hard-to-fill sphere in Z . It has a unique efficient filling in \bar{F} (of area $O(r^2)$), given by coning to the origin. The figure on the right depicts the pushing of that filling to Z , with signed copies of L shown. That filling has area $O(r^4)$.

This cancellation, however, is the only way that b may differ from b_0 . Since the boundary of b is in the k -skeleton of Z , the degree with which it covers any $(k + 1)$ -cell is well defined, and this degree must equal the corresponding coefficient of b_0 . In particular, the mass of the parts of b which do not cancel provide a lower bound for $\|b_0\|$. Thus, if one of the σ_i is disjoint from all the others then $|b_i|$ mass σ_i is a lower bound on the mass of any filling. We will use this general fact to show that Σ_r is hard to fill.

We first use the height-pushing map to find a chain in Z that fills Σ_r ; then we show that this chain is large. To construct the chain in Z note that the cycle Σ_r comes from the sphere of radius r in the flat \bar{F} , so it has an obvious filling T_r from the ball it bounds in \bar{F} . Formally, we write

$$T_r := (a_0)^{-r} \bar{F} \cap h^{-1}([-r, 0]),$$

considered as a $(k + 1)$ -chain.

Before we can apply the height-pushing map, we must perturb T_r so that it misses all vertices of X of nonzero height. In the perturbation T'_r , neighborhoods of the vertices of nonzero height are replaced with copies of L , as follows. Recall that the link $S(L)$ consists of signed copies of the simplices of L (Remark 4.3); for each simplex of L with vertices v_0, \dots, v_d , there are 2^{d+1} simplices in $S(L)$, with vertices $\pm v_0, \dots, \pm v_d$. For any vertex $v \in T_r$, the link $v \cdot S_{1/4} \cap T_r$ is an orthoplex: it is a join of $k + 1$ zero-spheres, and the i th zero-sphere is labeled by a_i and b_i with some signs, as in Figure 5. So in $S(L)$, there exists a copy of L with this orthoplex as its boundary (in fact there are many, each specified by a choice of signs on the interior vertices of L). Perform a surgery at each vertex v of T_r , replacing the

1/4-neighborhood of v in T_r with a copy of L in this way. The modified filling lives in Y , and we call it T'_r .

Now, we push the perturbed filling T'_r into Z . The result, $T''_r = \mathcal{Q}(T'_r)$, is a chain that fills Σ_r in Z . By the remark above, it suffices to find a lower bound on the size of the uncanceled pieces of this filling.

To get such a bound, we only need to consider images under \mathcal{Q} of link faces of T'_r , since original faces of Y are sent to lower-dimensional pieces. All of these link faces occur as part of a copy of L . Recall that \mathcal{Q} takes each copy of L in its domain to a scaled copy (with some orientation) of $S(L)^+$ or $S(L)^-$ in Z (Remark 4.3); copies with positive height go to $S(L)^-$ and copies with negative height go to $S(L)^+$. We can thus write, with $x = (x_0, \dots, x_k)$,

$$(5.1) \quad T''_r = \sum_{j=0}^r \sum_{\sum |x_i|=j} (-1)^j (a_0)^{-r} F(x) s_{r-j}(\lambda),$$

where λ is the fundamental class of $S(L)^+$. We are trying to estimate $\|T''_r\|$, but some of the terms in (5.1) may cancel. We will obtain a lower bound on $\|T''_r\|$ by showing that many of the scaled simplices making up the sum are disjoint from all other cells of T''_r .

First, note that if σ and σ' are two different k -simplices of $S(L)^+$, and $g, g' \in A$ and $t, t' > 0$, then $gs_t(\sigma)$ and $g's_{t'}(\sigma')$ intersect in at most a $(k-1)$ -dimensional set. We thus only need to consider the case of overlap between scaled copies $gs_t(\sigma)$ and $g's_{t'}(\sigma)$ of the same σ . Let σ be a strictly interior $(k+1)$ -simplex in L , with vertex set $S = \{g_0, \dots, g_k\}$; note that S is disjoint from $\{a_i, b_i\}$. Let $A_S = \langle g_0, \dots, g_k \rangle$. We claim that the scaled copies of σ in (5.1) are disjoint, and so none of them is canceled in the sum. If v and v' are vertices of scaled copies of σ , based at $x \in \bar{F}$ and $x' \in \bar{F}$ respectively, then

$$v = (a_0)^{-r} x \prod_i g_i^{t_i}$$

$$v' = (a_0)^{-r} x' \prod_i g_i^{t'_i}$$

for some vertex x of \bar{F} and some $t_i \in \mathbb{Z}^+$. Note that $x \in \langle a_i, b_i \rangle$ and $\prod_i g_i^{t_i} \in A_S$. Since $A_S \cap \langle a_i, b_i \rangle = \{0\}$, if $v = v'$, then $x = x'$, so no two distinct scaled copies of σ intersect.

Consequently, these terms do not cancel in T''_r , and as we mentioned previously, the mass of any filling of Σ_r is bounded below by the mass of these terms. Thus

$$\begin{aligned} \delta_X^k(\Sigma_r) &\geq \sum_{j=0}^r \sum_{\sum |x_i|=j} \text{mass}(s_{r-j}(\sigma)) \\ &\geq \sum_{j=0}^r \sum_{\sum |x_i|=j} \text{mass}(\sigma)(r-j)^{k+1} \geq cr^{2k+2} \end{aligned}$$

for some $c > 0$ depending only on k . Since $\text{mass}(\Sigma_r) \asymp r^k$ and r was arbitrary, we have $\delta_H(l) \succeq l^{\frac{2k+2}{k}}$ as desired. \square

6. HIGHER DIVERGENCE IN RAAGS

There are still many subgroups of right-angled Artin groups for which the Dehn function is unknown, and the higher divergence functions have a similar level of difficulty. For the groups themselves, however, we can get sharp bounds on the possible rates of divergence.

Theorem 6.1 (Higher divergence in RAAGs). *For $0 \leq k \leq \text{divdim}(A_\Gamma)$,*

$$r^{k+1} \preceq \text{Div}^k(A_\Gamma) \preceq r^{2k+2}.$$

The upper and lower bounds are sharp: for every k there are examples of right-angled Artin groups realizing these bounds.

In the next section, we will give a sharper result in the case $k = 0$.

In §2.3 we saw that the general lower bounds are realized by free abelian groups. We divide the rest of the theorem into several pieces: the general lower bounds, the general upper bounds, and a construction of groups whose divergence realizes the general upper bounds.

Proposition 6.2 (RAAG lower bounds). *For $k \leq \text{divdim}(A_\Gamma)$,*

$$r^{k+1} \preceq \text{Div}^k(A_\Gamma).$$

Proof. The Dehn function of A_Γ , evaluated at αr^k , is a lower bound for $\text{div}_\rho^k(\alpha r^k, r)$. This is because there are cycles in A_Γ of mass at most αr^k whose most efficient fillings have mass arbitrarily close to $\delta_A^k(\alpha r^k)$, and these can be translated to be r -avoidant. Since Γ has a clique of $k + 1$ vertices, A_Γ retracts onto a subgroup $\mathbb{Z}^{k+1} \subset A_\Gamma$. Thus

$$\delta_{A_\Gamma}^k(l) \succeq \delta_{\mathbb{Z}^{k+1}}^k(l) \succeq l^{\frac{k+1}{k}}.$$

□

Proposition 6.3 (RAAG upper bounds). *For $k \leq \text{divdim}(A_\Gamma)$, we have*

$$\text{Div}^k(A_\Gamma) \preceq r^{2k+2}.$$

Proof. By Remark 2.4, it suffices to show that there is a c such that for all sufficiently large r ,

$$\text{div}_1^k(l, r) \leq cl^{\frac{k+1}{k}} r^{k+1}.$$

Let $a \in Z_k^{\text{cell}}(X_\Gamma)$ be an r -avoidant k -cycle and let $l = \|a\|$. Since X_Γ is CAT(0), there is a $(k + 1)$ -chain $b \in C_{k+1}^{\text{cell}}(X_\Gamma)$ such that $\partial b = a$ and $\|b\|_1 \preceq l^{\frac{k+1}{k}}$. We will consider b as a Lipschitz chain and push it out of B_r .

Let \mathcal{P}_r be the map constructed in Lemma 4.4; this map is $O(r)$ -Lipschitz. The image $\mathcal{P}_r(b)$ is an r -avoidant filling of a , and there is a c such that

$$\text{mass}(\mathcal{P}_r(b)) \leq (\text{Lip } \mathcal{P}_r)^{k+1} \text{mass } b \leq cr^{k+1} l^{\frac{k+1}{k}},$$

as desired. □

The final step of Theorem 6.1 is to construct groups that have the stated divergences. The key to this construction is to use the connection between divergence functions and the Dehn functions of Bestvina-Brady groups; large portions of S_r can be embedded the level sets corresponding to Bestvina-Brady groups, so avoidant fillings can be converted into fillings in Bestvina-Brady groups.

Theorem 6.4 (Sharpness of upper bounds). *If A_Γ is a k -orthoplex group, then*

$$\text{Div}^k(A_\Gamma) \asymp r^{2k+2}.$$

Proof. By Remark 2.4, it suffices to show that there is a $c_k > 0$ such that for all $0 < \rho \leq 1$,

$$\text{div}_\rho^k(c_k r^k, r) \succeq r^{2k+2}.$$

Recall that when A_Γ is a $(k+1)$ -dimensional orthoplex group, we constructed cycles in Z_Γ which have volume r^k and filling volume r^{k+1} by defining a flat \bar{F} and considering the intersections

$$\Sigma_r := ((a_0)^{-r} \bar{F}) \cap Z.$$

These are cycles in Z of volume $c_k r^k$ for some $c_k > 0$ with filling volume $\asymp r^{2k+2}$. Let

$$\Sigma'_r := (a_0)^r \Sigma_r = \bar{F} \cap h^{-1}(r).$$

This is an r -avoidant cycle of volume $c_k r^k$ and we will show that every ρr -avoidant filling of Σ'_r has volume $\succeq r^{2k+2}$.

We first define a retraction $\pi_t : X_\Gamma - B_t \rightarrow S_t$ for all t . If $x \in X_\Gamma - B_t$, there is a unique CAT(0) geodesic from x to e . Let $\pi_t(x)$ be the intersection of S_t with this geodesic. (As before, we will also write π_t for the induced map $(\pi_t)_\#$ on chains and cycles.) This is clearly the identity on S_t , and since B_t is convex, it is 1-Lipschitz (distance-nonincreasing). Furthermore, if $x \in \bar{F}$, then since \bar{F} is a flat, the geodesic from x to e is a straight line in \bar{F} , and $\pi_t(\Sigma'_r) = \pi_t(\Sigma'_t)$ for all $t \leq r$.

Consider the map $\text{abs} \circ \pi_t$. We claim that this is a 1-Lipschitz retraction from $X_\Gamma - B_t$ to $S_t \cap h^{-1}(t)$. If $x \in S_t \cap h^{-1}(t)$, then $\pi_t(x) = x$, and $\text{abs}(x) = x$, so this map is a retraction, and since abs and π_t are each 1-Lipschitz, the composition is as well. Furthermore, we have

$$\text{abs} \circ \pi_t(\Sigma'_r) = \Sigma'_t$$

for all $t \leq r$.

Let b be a ρr -avoidant $(k+1)$ -chain whose boundary is Σ'_r . Then $b' = \text{abs} \circ \pi_{\rho r}(b)$ is a chain in $S_t \cap h^{-1}(\rho r)$ whose boundary is $\partial b' = \Sigma'_{\rho r}$. Its translate $(a_0)^{-\rho r} b'$ is a chain in Z whose boundary is $\Sigma_{\rho r}$, so for all $0 < \rho \leq 1$,

$$\text{mass } b' \geq \delta_H^k(\Sigma_{\rho r}) \succeq r^{2k+2}.$$

Since $\text{abs} \circ \pi_t$ is 1-Lipschitz, $\text{mass } b \geq \text{mass } b'$, so

$$\text{div}_\rho^k(c_k r^k, r) \succeq r^{2k+2}$$

as desired. □

7. A REFINED RESULT FOR DIVERGENCE OF GEODESICS

The case $k = 0$ gives a quantitative measure of how fast geodesics spread apart. Here, the answer is completely determined by whether or not the group is a direct product. Note that A_Γ is a direct product if and only if the vertices can be partitioned into two nonempty subsets A and B such that each vertex of A is joined to each vertex of B by an edge of Γ . Equivalently, A_Γ is a direct product if and only if the complement of Γ is not connected.

The following result also appears in [BC10].

Theorem 7.1 (Divergence of geodesics in RAAGs). *For a right-angled Artin group A_Γ , Div^0 exists if and only if the defining graph Γ is connected. In this case, $\text{Div}^0(A_\Gamma) \asymp r$ if and only if A_Γ is a nontrivial direct product, and $\text{Div}^0(A_\Gamma) \asymp r^2$ otherwise.*

Note that if Γ is not connected, then A_Γ has infinitely many ends, so Div^0 is not defined. We have already established (Theorem 6.1 with $k = 0$) that $r \preceq \text{Div}^0(A_\Gamma) \preceq r^2$. We proceed by considering the presence of a product structure.

Lemma 7.2 (Linear if direct product). *If $A_\Gamma = H \times K$ is a direct product of nontrivial factors, then $\text{Div}^0(A_\Gamma) \asymp r$.*

Proof. A linear-length path is easily constructed by pushing outside the ball in one factor first, then modifying the other factor as desired before adjusting the first factor. Writing elements of A_Γ as ordered pairs, let (h_1, k_1) and (h_2, k_2) be elements of A_Γ of length r , that is, $|h_i|_H + |k_i|_K = r$. There exists a $u \in H$ such that $|u|_H \leq r$ and $|h_1 u|_H = r$. Similarly there exists an element $v \in K$ such that $|v|_K \leq r$ and $|k_2 v|_K \geq r$. Now since the length of an element of A_Γ is simply the sum of the lengths of its components in H and K , the vertices representing the elements

$$(h_1, k_1), (h_1 u, k_1), (h_1 u, e), (h_1 u, k_2 v), (e, k_2 v), (h_2, k_2 v), (h_2, k_2)$$

lie on or outside the ball of radius r centered at the identity. Moreover, successive elements of the sequence can be connected by r -avoidant paths, each of which has length at most r . Thus any two vertices on the ball of radius r at the identity can be connected by a r -avoidant path of length at most $6r$, and the lemma follows. \square

In fact the proof uses little about RAAGs, and the same result holds for all direct products $G = H \times K$ where H and K each have the property that every point lies on a geodesic ray based at e .

Lemma 7.3 (Quadratic if not direct product). *If A_Γ is not a direct product, then $\text{Div}^0(A_\Gamma) \asymp r^2$.*

Proof. We only need to show that $\text{Div}^0 \succeq r^2$. As A_Γ is not a direct product, the complement Γ^c of Γ is connected. Choose a closed path in Γ^c that visits each vertex (possibly with repetitions) and let $w = a_1 a_2 \dots a_n$ be the word made up of the generators (i.e., vertices) encountered along this path. Introduce the symbol a_{n+1} as another name for a_1 . Note that for $1 \leq i \leq n$, the vertices a_i and a_{i+1} are not connected by an edge in Γ , so the corresponding generators do not commute. As a consequence, w is a geodesic word, as is any nontrivial power w^k ; let η be the unique bi-infinite geodesic going through e and all of these powers, so that the letters of η cycle through the a_i . (Note this is a geodesic with respect to either the word metric on A_Γ or the CAT(0) metric on X .)

To prove the lemma, it is enough to show that for each r , any ρr -avoidant path connecting $\eta(r)$ and $\eta(-r)$ has length at least on the order of r^2 .

Let γ be a ρr -avoidant path from $\eta(r)$ to $\eta(-r)$ in $X^{(1)}$. Let η_0 be the segment of η between the same two endpoints, labeled by the word w^{2r} . The union of γ and η_0 is a loop in X labeled by a word representing the identity in A_Γ (see Figure 6). So there exists a van Kampen diagram Δ whose boundary cycle $\partial\Delta$ is labeled by this word. We refer to [Bri02] for background on van Kampen diagrams.

There is a combinatorial map $\Delta \rightarrow X$ such that $\partial\Delta$ maps to $\gamma \cup \eta_0$. In the terminology of [Bri02], $\partial\Delta$ consists of a thick part, consisting of edges which lie in the boundary of a 2-cell of Δ , and a thin part. (In the informal “lollipop” language used to talk about these diagrams, the thick part is the candy and the thin part is the sticks.) Let σ be an edge of η_0 in the interior of $B_{\rho r}$. Here, $0 < \rho < 1$ is the Gersten parameter, so that γ lies outside the ball $B_{\rho r}$. Then σ is in the image of a thick edge of Δ . (The thin parts must be traveled in both directions to form a loop; since η is a geodesic, the only way that an edge of η_0 can come from the thin part is if it coincides with some edge of γ , but this contradicts the choice of γ .)

Therefore, σ is contained in the boundary of some 2-cell of Δ , with boundary label $xyx^{-1}y^{-1}$. If y is the label on σ , then σ is one end of a y -corridor. A corridor that starts at $\partial\Delta$ cannot intersect itself in a 2-cell, so its other end is an edge of $\partial\Delta$ with label y , but orientation opposite to that of σ . Since the map from Δ to X preserves orientations of edge labels, and all the edges of η have the same orientation as σ , the other end of the corridor must be an edge of γ .




FIGURE 6. Corridors in Δ . For each of the corridors from η to γ (vertical in this picture), every cell has a corridor that crosses it with both ends on γ .

Thus each edge of η_0 whose vertices lie in $B_{\rho r}$ bounds a corridor whose other end is an edge of γ , as shown on the left side of Figure 6.) An x -corridor intersects a y -corridor in a 2-cell if and only if x and y commute. Since a_i does not commute with a_{i+1} for any i , no two of these corridors intersect.

Since all the relations in the presentation are commutation relations, a y -corridor has boundary label of the form $yvy^{-1}v^{-1}$ for some word v , which we will refer to as the lateral boundary word of the corridor. We now show that by performing some surgeries on the Δ which do not change $\partial\Delta$, we may assume that the lateral boundary words of the a_i -corridors emanating from η are geodesic words (that is, minimal representatives of the group elements that they represent).

Suppose the boundary word of an a_i corridor is not geodesic. Since every word can be reduced to a geodesic by shuffling neighboring commuting pairs (see [HM95]), there has to be a sub-segment of the form xux^{-1} where x is a generator (or its inverse), u is a word, and x commutes with all the generators appearing in u . Then we can perform the *tennis-ball move* shown in Figure 7.




FIGURE 7. Here, the long edges are labeled by a word u and the square faces are x, a_i commutators. In a *tennis-ball move*, the shaded disk is replaced with the unshaded disk, noting that the boundary words $(xux^{-1}a_ixu^{-1}x^{-1}a_i^{-1})$ are equal.

Note that $\partial\Delta$ remains unchanged at the end of such a move, and that the length of the a_i -corridor is reduced. After performing enough of these moves so that all the boundaries of corridors emanating from η are geodesic words we get a new van Kampen diagram, which we also call Δ , with map from Δ to X such that $\partial\Delta$ maps to $\gamma \cup \eta_0$ (in fact it is the same map we had before.)

We now restrict our attention to a_1 -corridors emanating from η . (There may be j such that $a_j = a_1$ as generators, but we ignore the a_j -corridors.) A 2-cell in such a corridor has boundary label $a_1xa_1^{-1}x^{-1}$ for some x , and is therefore part of an x -corridor or annulus. We claim that this is in fact an x -corridor which intersects the a_1 -corridor in exactly one 2-cell. (In particular it is not an annulus.)

If the intersection contains more than one 2-cell, then Δ contains the picture shown in Figure 8, where u is a geodesic word whose individual letters commute with a_1 and v is a word (not necessarily geodesic) whose individual letters commute with x . Then the words u and v represent the same group element. Note that for right-angled Artin groups, since the only relators are commutators, the set of letters appearing in a geodesic representative of a word, called the *support* of the word, is well-defined (this follows from [HM95]). Since u was geodesic, we have that $\text{Support}(u) \subseteq \text{Support}(v)$, so that x commutes with the individual letters in u , and hence commutes with u . This contradicts the fact that the boundary word of the a_1 corridor was geodesic, and proves the claim.

Now $x = a_i$ for some i , and two x -corridors can't intersect in a 2-cell (since that would require the existence of a 2-cell with boundary label $xxx^{-1}x^{-1}$). Thus the x -corridor is trapped in the region bounded by two a_i -corridors, as shown on the right in Figure 6, and its two ends are edges of γ .




FIGURE 8. Each crossing corridor intersects the vertical corridor only once.

To summarize, each 2-cell along each of the a_1 -corridors emanating from η_0 is part of a corridor whose ends are edges of γ . No two of these edges coincide. Thus the total area of the a_i -corridors is a lower bound for the length of γ .

To obtain a lower bound on this area, note that for each $0 \leq i \leq \lfloor r/n \rfloor$, there is an a_1 -corridor whose η -end has vertices whose distances from the origin are in and $in + 1$ respectively. Since the other end of this corridor is an edge along γ , and γ is nr -avoidant, the length of the corridor is at least $\rho r - in$. So the total area of the a_1 -corridors is at least $\sum_{i=0}^{\lfloor r/n \rfloor} (\rho r - in)$, which is on the order of r^2 . \square

8. FILLING LOOPS IN THE MAPPING CLASS GROUP

In this section we use pushing maps to obtain results on Div^1 for the mapping class group of a surface. Mapping class groups have quadratic Dehn functions because they are automatic, hence combable (see [Mos94, ECH⁺92]). But how flexible is this filling? That is, if we insist that a filling avoid a large ball, must its area be much worse than quadratic? We will show in this section that, just as for right-angled Artin groups, loops can be avoidantly filled using area at most polynomial of degree four. Heuristically, one might expect this: considering that there are many pairs of points which require quadratic-length paths to connect avoidantly ([DR09]), one can seek loops with two such “diameters,” for which the filling area should be on the order of $r^2 \cdot r^2 = r^4$.

For a topological surface $S = S_{g,b}$ (where g is the genus and b is the number of punctures/boundary components), we write $\text{Mod}(S)$ for its group of orientation-preserving diffeomorphisms up to isotopy, or *mapping class group*. $\text{Mod}(S)$ can be realized as a quotient of a RAAG whose generators are Dehn twists (which commute if the corresponding curves are disjoint). Thus, in a two-complex for such a presentation, a filling has two types of 2-cells, one type coming from the RAAG and one from the other relators of the mapping class group. We can deal with the cells coming from commuting relators using the ideas we have already developed: push radially by post-multiplying with a high power of a generator, which translates along a standard ray in an orthant. For the other types of relators, we can do something similar if we can find a common commutator for all of the labels appearing

in the loop. To make use of this, we would like a presentation in which every relator has “small support” (in the surface), so that there are other generators disjoint from those in the relator.

The Gervais presentation of the mapping class group (see [Ger01]) fits the bill: this is a finite presentation with Dehn twist generators for which every relator is supported on a small subsurface (at most a three-times punctured torus).

All the generators are Dehn twists and all relations are of three kinds: commuting, braid, and so-called *star* relations. Commuting relations arise from twists around disjoint curves. The braid relation, which one obtains whenever two curves intersect once, is of the form $aba = bab$. Each star relation, which occurs when a collection of seven curves is in a particular topological configuration, is of the form $(abcd)^3 = xyz$. If two of the twisting curves for a, b, c, d are isotopic (say a and b), then there is a degenerate star relation of the form $(a^2cd)^3 = xy$.

For this presentation of $\text{Mod}(S)$, let X be the universal cover of the presentation 2-complex, so that its 1-skeleton is the Cayley graph. Then all of the 2-cells are squares, hexagons, 14-gons, and 15-gons corresponding to the relators described above.

In this complex, fillings cannot be perturbed off of the vertices, and we use a different method to construct avoidant fillings. We will take an efficient filling and push out the 2-cells of the filling in a controlled way to get a collection of disjoint avoidant cells. These are then patched together avoidantly using only commuting relations. The possible directions in which 2-cells are pushed, as well as the possible commuting relations used for the patchwork, are determined by a particular abelian subgroup of $\text{Mod}(S)$ generated by Dehn twists around the set of curves described in the following lemma.

Lemma 8.1 (Common commutators in the Gervais presentation). *Let S be an orientable surface of genus $g \geq 5$ and any $b \geq 0$. There is a set \mathcal{H} of g mutually disjoint curves with the following properties:*

- (1) *The Dehn twist around each curve in \mathcal{H} is a Gervais generator.*
- (2) *For every relation in the Gervais presentation, at least one of the Dehn twists about curves from \mathcal{H} commutes with every Dehn twist appearing in the relation.*

Proof. Gervais gives his presentation in terms of three types of curves: α curves (separating out the topology), β curves (one around each handle, dual to some of the α 's), and γ curves (derived from pairs of α 's). Let \mathcal{H} be the set of β curves, which are all disjoint. Each star relation is supported on a three-times punctured torus, and the relation involves three α curves, one β curve, and three γ curves. One easily verifies that the maximum number of other β curves intersecting any of these support curves is three, if dual to the α 's. Thus if there are at least five β curves in total, one of them must be disjoint from the support curves. Clearly, in this case, the support of any commuting or braid relation also misses some β curve. The total number of β curves is g , the genus. \square

Let h_1, \dots, h_g denote the Dehn twists about the curves in \mathcal{H} , and let H denote the subgroup $\langle h_1, \dots, h_g \rangle$. The group H is abelian, since the curves in \mathcal{H} are disjoint. Abelian subgroups of $\text{Mod}(S)$ are undistorted ([FLM01]); that is, for the group H , if $d(w, v)$ is the word metric with respect to the Gervais presentation, then there is

a $c > 1$ such that for all $t_i \in \mathbb{Z}$,

$$(8.1) \quad \sum |t_i| \leq c \cdot d(\prod h_i^{t_i}, e) + c.$$

This allows us to push points outside the r -ball. For instance, if $d(v, e) \leq r$, then $d(vh_i^{2rc+c}, e) \geq r$ by the triangle inequality. Equation (8.1) also lets us find avoidant loops and disks in $\text{Mod}(S)$ by looking at avoidant loops and disks in H and its cosets.

Lemma 8.2 (Avoidance in cosets). *There is a $D > 0$ depending only on S and H such that if $H' \subset H$ is generated by a subset of the generators of H , then:*

- (1) *Assume that H' has rank at least 2. Let $v \in \text{Mod}(S)$ and let $x_1, x_2 \in v \cdot H'$. If $r = \min d(x_i, e) > D$, there is a path γ connecting x_1 and x_2 in $v \cdot H'$ which is r/D -avoidant and has length at most $D \cdot d_{H'}(x_1, x_2)$, where $d_{H'}$ is the word metric with respect to the generators of H' .*
- (2) *Let $v \in \text{Mod}(S)$, let $1 \leq k \leq \text{divdim}(H')$, and let a be a k -cycle in $v \cdot H'$. If $r = d(\text{supp } a, e) > D$, there is a $(k+1)$ -chain b in $v \cdot H'$ such that $\partial b = a$, b is r/D -avoidant, and $\text{mass } b \leq (\text{mass } a)^{(k+1)/k}$.*

Proof. First we prove (1) by showing that the statement holds with $D = 5c$, where c is the distortion constant from (8.1) above. Consider $d(v \cdot H', e)$. If this is larger than r/D , we can take γ to be a geodesic in $v \cdot H'$ which connects x_1 and x_2 . Otherwise, let $y \in v \cdot H'$ be a point such that $d(y, e) = d(v \cdot H', e)$. Then

$$d_{H'}(y, x_i) \geq d(y, x_i) \geq d(x_i, e) - d(y, e) \geq r - \frac{r}{D} > \frac{4r}{5}.$$

Let

$$B := \{h \in H' \mid d_{H'}(h, y) \leq \frac{4r}{5}\}.$$

Both x_1 and x_2 are outside of B , and since Div^0 is linear for H' (Lemma 7.2), there is a γ which has the required length and connects x_1 and x_2 outside of B . We claim that this γ is avoidant in $\text{Mod}(S)$. This follows from (8.1); if $h \in H' - B$, then

$$d(h, e) \geq d(h, y) - d(y, e) \geq \frac{d_{H'}(h, y)}{c} - 1 - \frac{r}{D} \geq \frac{r}{D},$$

as desired.

The proof of (2) is similar, except we use the Euclidean Dehn function (Proposition 2.5) to construct a filling instead of linearity of Div^0 . \square

Thus we can construct avoidant fillings of cycles that live in cosets of H ; we will use these to build avoidant fillings of arbitrary cycles.

Theorem 8.3 (Filling loops at infinity in the mapping class group). *If S has genus at least 5 and any number of punctures, then $\text{Div}^1(\text{Mod}(S)) \preceq r^4$.*

In this proof we will work with cellular maps rather than Lipschitz maps. That is, we start with a loop in the 1-skeleton of X and produce an avoidant map from a van Kampen diagram into X whose boundary maps to the loop.

Proof. Begin with an avoidant loop in X of length l . Since the Dehn function is quadratic, there exists a (not necessarily avoidant) filling Δ with area $\preceq l^2$. We use Δ as a combinatorial model for an avoidant filling of the same loop. The new filling is obtained by making the following replacements, which are described more precisely later.

- (1) Each 2-cell in the filling is pushed to an avoidant copy of itself.
- (2) Each edge is replaced by a (possibly degenerate) avoidant strip of squares of length $\preceq r$. An edge adjacent to two 2-cells is replaced by a strip connecting the two pushed copies of the cells. An edge adjacent to a single 2-cell is necessarily part of the boundary loop, and is replaced by a strip connecting the edge to the pushed copy of the cell.
- (3) The result of steps (1) and (2) is topologically a punctured disk, with one boundary component equal to the original loop, and an additional boundary component corresponding to each vertex in Δ . Each boundary component of the latter kind is filled by an avoidant disk.

Pushing 2-cells. We replace each 2-cell σ with an avoidant pushed cell σ' . If σ is already r -avoidant, we let $\sigma' = \sigma$. Otherwise, σ is partially contained in the ball of radius r . It corresponds to a relation in the Gervais presentation, and we choose a common commuter h_σ for the generators in this relation, which is guaranteed by Lemma 8.1.

Let $R = (2r + 30)c + c$. If the vertices of σ are v_1, \dots, v_k , then $v_1 h_\sigma^R, \dots, v_k h_\sigma^R$ are the vertices of a copy of σ , since h_σ commutes with the edge labels of σ . Denote this copy by σ' . Since σ is partially contained in the ball of radius r , it is entirely contained in the ball of radius $r + 15$ (relators have length at most 15). Then σ' is r -avoidant by (8.1).

Connecting pushed cells with strips. Consider an edge e with vertices v and w , labeled by a Gervais generator f . Let $H_f \subset H$ be the subgroup of H generated by the generators of H which commute with f ; this has rank at least $g - 1$, where g is the genus of S .

First, we find vertices v_1 and v_2 corresponding to v in the pushed filling. If e is adjacent to two 2-cells σ_1 and σ_2 , let v_1, v_2 be the vertices of σ'_1 and σ'_2 which correspond to v . Otherwise, e is a boundary edge of Δ . If it is adjacent to a face σ ; let $v_1 = v$ and let v_2 be the vertex in σ' corresponding to v . Otherwise, e is used twice in the boundary of Δ ; we can let $v_1 = v_2 = v$. In any case, v_1 and v_2 are r -avoidant and are both contained in $v \cdot H_f$, and $f_{H_f}(v_1, v_2) \preceq r$. By Lemma 8.2, there is a r/D -avoidant path γ in $v \cdot H_f$ which connects v_1 to v_2 ; we can interpret this as a word representing $v_1^{-1} v_2$ whose letters all commute with f . Then there is a strip built out of squares (that is, RAAG relations) whose boundary label is the commutator $[f, \gamma]$; this strip is $(r/D - 1)$ -avoidant and has length $\preceq r$.




FIGURE 9. The cells σ'_i are obtained by pushing the σ_i out of the ball of radius r . The path γ is r/D -avoidant, and the letters in the corresponding word all commute with f .

Filling in the holes. The partial filling constructed above has one boundary component for each vertex of Δ . Each boundary component is a polygon with sides coming from parts of boundaries of strips as shown in Figure 10. The number of sides of the polygon associated to v is the number of edges incident to v . Each




FIGURE 10. The edges in bold in this figure are part of the boundary loop of Δ . Each strip is r/D -avoidant and has length $\leq r$.

vertex is r -avoidant, and each side is an $(r/D - 1)$ -avoidant curve of length $\leq r$. Indeed, each vertex is distance at most R away from v , so any two vertices are distance $\leq 2R$ apart. The entire polygon is contained in $v \cdot H$.

To fill these polygons, we first subdivide an n -gon into $n - 2$ triangles by adding additional r/D -avoidant curves between the vertices; these exist by Lemma 8.2. The resulting triangles are $(r/D - 1)$ -avoidant and have length $\leq r$. By Lemma 8.2, they can be filled by $(r/D^2 - 1)$ -avoidant disks of area $\leq r^2$. Let $\rho = 1/D^2$, so that when r is sufficiently large, $r/D^2 - 1 \geq \rho r$.

The union of the pushed cells, strips and filled triangles above is a ρr -avoidant filling of the boundary loop; it remains to estimate the area of this filling. Since the area (the number of 2-cells) of Δ is $\leq l^2$, the number of vertices and edges in Δ is $\leq l^2$ as well. Each strip introduced in this construction has length (and area) $\leq r$. There is a constant M such that the area of any of the triangle fillings above is at most Mr^2 . Moreover, the number of triangles is certainly bounded above by twice the number of edges in Δ . So the total area of the new filling is $\leq l^2 + l^2(2r) + l^2(Mr^2)$. Specializing to the case that l is a constant multiple of r^2 , we see that the area of the ρr -avoidant filling constructed above is $\leq r^4$.

The above argument shows that $\text{div}_\rho(\alpha r, r) \leq r^4$ for all α . This establishes that $\text{Div}^1(\text{Mod}(S)) \leq r^4$. \square

REFERENCES

- [AWP99] J. M. Alonso, X. Wang, and S. J. Pride, *Higher-dimensional isoperimetric (or Dehn) functions of groups*, J. Group Theory **2** (1999), no. 1, 81–112. MR MR1670329 (2000e:20113)
- [BB97] Mladen Bestvina and Noel Brady, *Morse theory and finiteness properties of groups*, Invent. Math. **129** (1997), no. 3, 445–470. MR MR1465330 (98i:20039)
- [BBFS09] Noel Brady, Martin R. Bridson, Max Forester, and Krishnan Shankar, *Snowflake groups, Perron-Frobenius eigenvalues and isoperimetric spectra*, Geom. Topol. **13** (2009), no. 1, 141–187. MR MR2469516
- [BC10] Jason Behrstock and Ruth Charney, *Divergence and quasimorphisms of right-angled artin groups*, Preprint (2010).
- [BF98] Noel Brady and Benson Farb, *Filling-invariants at infinity for manifolds of nonpositive curvature*, Trans. Amer. Math. Soc. **350** (1998), no. 8, 3393–3405. MR MR1608281 (99c:53039)

- [BM01] Noel Brady and John Meier, *Connectivity at infinity for right angled Artin groups*, Trans. Amer. Math. Soc. **353** (2001), no. 1, 117–132. MR MR1675166 (2001b:20068)
- [Bri02] M. R. Bridson, *The geometry of the word problem*, Invitations to geometry and topology, Oxf. Grad. Texts Math., vol. 7, Oxford Univ. Press, Oxford, 2002, pp. 29–91. MR MR1967746 (2004g:20056)
- [BRS07] Noel Brady, Tim Riley, and Hamish Short, *The geometry of the word problem for finitely generated groups*, Advanced Courses in Mathematics. CRM Barcelona, Birkhäuser Verlag, Basel, 2007, Papers from the Advanced Course held in Barcelona, July 5–15, 2005. MR MR2281936 (2009j:20053)
- [Cha07] Ruth Charney, *An introduction to right-angled Artin groups*, Geom. Dedicata **125** (2007), 141–158. MR MR2322545 (2008f:20076)
- [Dis08] Will Dison, *An isoperimetric function for Bestvina-Brady groups*, Bull. Lond. Math. Soc. **40** (2008), no. 3, 384–394. MR MR2418794 (2009e:20091)
- [DMS09] Cornelia Druțu, Shahar Mozes, and Mark Sapir, *Divergence in lattices in semisimple Lie groups and graphs of groups*, TAMS, to appear (2009).
- [DR09] Moon Duchin and Kasra Rafi, *Divergence of geodesics in Teichmüller space and the mapping class group*, Geom. Funct. Anal. **19** (2009), no. 3, 722–742.
- [ECH⁺92] D. B. A. Epstein, J. W. Cannon, D. F. Holt, S. V. F. Levy, M. S. Paterson, and W. P. Thurston, *Word processing in groups*, Jones and Bartlett Publishers, Boston, MA, 1992. MR MR1161694 (93i:20036)
- [FF60] H. Federer and W. H. Fleming, *Normal and integral currents*, Ann. of Math. (2) **72** (1960), 458–520. MR MR0123260 (23 #A588)
- [FLM01] Benson Farb, Alexander Lubotzky, and Yair Minsky, *Rank-1 phenomena for mapping class groups*, Duke Math. J. **106** (2001), no. 3, 581–597. MR MR1813237 (2001k:20076)
- [Ger94a] S. M. Gersten, *Divergence in 3-manifold groups*, Geom. Funct. Anal. **4** (1994), no. 6, 633–647. MR MR1302334 (95h:57013)
- [Ger94b] ———, *Quadratic divergence of geodesics in CAT(0) spaces*, Geom. Funct. Anal. **4** (1994), no. 1, 37–51. MR MR1254309 (95h:53057)
- [Ger01] Sylvain Gervais, *A finite presentation of the mapping class group of a punctured surface*, Topology **40** (2001), no. 4, 703–725. MR MR1851559 (2002m:57025)
- [Groa] Chad Groft, *Generalized Dehn functions I*.
- [Grob] ———, *Generalized Dehn functions II*.
- [Gro83] Mikhail Gromov, *Filling Riemannian manifolds*, J. Differential Geom. **18** (1983), no. 1, 1–147. MR MR697984 (85h:53029)
- [Gro87] M. Gromov, *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263. MR MR919829 (89e:20070)
- [Gro93] ———, *Asymptotic invariants of infinite groups*, Geometric group theory, Vol. 2 (Sussex, 1991), London Math. Soc. Lecture Note Ser., vol. 182, Cambridge Univ. Press, Cambridge, 1993, pp. 1–295. MR MR1253544 (95m:20041)
- [Hin05] Mohamad A. Hindawi, *On the filling invariants at infinity of Hadamard manifolds*, Geom. Dedicata **116** (2005), 67–85. MR MR2195442 (2006i:53056)
- [HM95] Susan Hermiller and John Meier, *Algorithms and geometry for graph products of groups*, J. Algebra **171** (1995), no. 1, 230–257. MR MR1314099 (96a:20052)
- [KL98] M. Kapovich and B. Leeb, *3-manifold groups and nonpositive curvature*, Geom. Funct. Anal. **8** (1998), no. 5, 841–852. MR MR1650098 (2000a:57040)
- [Leu00] E. Leuzinger, *Corank and asymptotic filling-invariants for symmetric spaces*, Geom. Funct. Anal. **10** (2000), no. 4, 863–873. MR MR1791143 (2001k:53074)
- [Mac08] Natasha Macura, *CAT(0) spaces with polynomial divergence of geodesics*, Preprint (2008).
- [Mos94] Lee Mosher, *Mapping class groups are automatic*, Math. Res. Lett. **1** (1994), no. 2, 249–255. MR MR1266763 (95a:57023)
- [Wen06] Stefan Wenger, *Filling invariants at infinity and the Euclidean rank of Hadamard spaces*, Int. Math. Res. Not. (2006), Art. ID 83090, 33. MR MR2250014 (2007m:53041)
- [Wen08] ———, *A short proof of Gromov’s filling inequality*, Proc. Amer. Math. Soc. **136** (2008), no. 8, 2937–2941. MR MR2399061 (2009a:53072)
- [Youa] R. Young, *Filling inequalities for nilpotent groups*.
- [Youb] R. Young, *Homological and homotopical higher-order filling functions*, Groups, Geometry, and Dynamics.