SNOWFLAKE GEOMETRY IN CAT(0) GROUPS

NOEL BRADY AND MAX FORESTER

ABSTRACT. We construct CAT(0) groups containing subgroups whose Dehn functions are given by x^s, for a dense set of numbers $s \in [2, \infty)$. This significantly expands the known geometric behavior of subgroups of CAT(0) groups.

1. INTRODUCTION

In this paper we are concerned with the geometry of subgroups of CAT(0) groups. Such subgroups need not be CAT(0) themselves, and in fact this realm includes many instances of exotic or unusual group-theoretic behavior. For example, there exist finitely presented subgroups of CAT(0) groups having unsolvable membership and conjugacy problems [BBMS00, Bri13], or possessing infinitely many conjugacy classes of finite subgroups (implicitly in [FM91], explicitly in [LN03, BCD08]). Many well known examples of groups having interesting homological properties, such as the Stallings–Bieri groups or other Bestvina–Brady kernels, arise naturally as subgroups of CAT(0) groups.

Regarding the geometry, specifically, of subgroups of CAT(0) groups, only a few types of unusual behavior have been observed to date. One common feature to all subgroups of CAT(0) groups is that cyclic subgroups are always undistorted. This constraint immediately rules out a great many examples from being embeddable into CAT(0) groups. The examples found so far with interesting geometric properties have been constructed using fairly sophisticated techniques. Such examples include the hydra groups of [DR13], which are CAT(0) groups possessing free subgroups with extreme distortion; subgroups of CAT(0) groups having exponentional or polynomial Dehn functions [BBMS97, Bri13], [BRS07, ABD$^+$13, BGL11]; and subgroups of CAT(0) groups having different homological and ordinary Dehn functions [ABDY13].

Our goal in the present paper is to construct CAT(0) groups containing subgroups that exhibit a wide range of isoperimetric behavior. Our main theorem is the following.

Theorem A. *Let m, n be positive integers such that $\alpha = n \log_m(1+\sqrt{2}) \geqslant 1$. Then there exists a 6–dimensional CAT(0) group G which contains a finitely presented subgroup S whose Dehn function is given by $\delta_S(x) = x^{2\alpha}$.*

Varying m and n, the resulting isoperimetric exponents 2α form a dense set in the interval $[2,\infty)$. Thus, the *isoperimetric spectrum* of exponents arising from subgroups of CAT(0) groups resembles, in its coarse structure, that of finitely presented groups generally [BB00].

CONTENTS

The group S in the main theorem is constructed with reference to certain parameters; it is denoted $S_{T,n}$ where T is a finite tree and n is an integer. It will embed into a CAT(0) group denoted $G_{T,n}$.

Comparison with prior constructions of snowflake groups. The construction of the snowflake groups $S_{T,n}$ builds on the methods used in [BBFS09] to construct groups with specified power functions as Dehn functions. The latter groups do not embed into any CAT(0) group, for the simple reason that they contain distorted cyclic subgroups. The same is true for the groups constructed in [SBR02], and many other groups having interesting Dehn functions. Indeed, it is the presence of precisely distorted cyclic subgroups that enables the computation of the Dehn functions in [BBFS09] in the first place.

The basic structure of the groups in [BBFS09] is that of a graph of groups, with vertex groups of cohomological dimension 2, and infinite cyclic edge groups. The lower bound for the Dehn function was easy to establish, using asphericity of the presentation 2–complex. The distortion of the cyclic edge groups was then precisely determined, and this distortion estimate led to a matching upper bound for the Dehn function of the ambient group.

If one hopes to embed examples into CAT(0) groups, one cannot have distorted cyclic subgroups. The groups $S_{T,n}$ we construct here are fundamental groups of graphs of groups in which the edge groups are *free groups* of rank 2. Moreover, the vertex groups are groups of cohomological dimension 3, and therefore $S_{T,n}$ does not admit an aspherical presentation. For this reason, establishing the lower bound for the Dehn function of $S_{T,n}$ requires some effort.

The change to free edge groups introduces some major challenges. In [BBFS09], given any element of an edge group, conjugation by the appropriate stable letter resulted in a word r times longer, for a uniform factor r. In the Cayley 2–complex, each edge space had a well defined long side and short side. A key geometric property was that no geodesic could pass from the short side to the long side of any edge space, and back again. It thus became possible to determine the large scale behavior of geodesics, which in turn led to very explicit distance estimates in the group.

In the groups $S_{T,n}$ defined here, the edge groups have monodromy modeled on a hyperbolic free group automorphism. Under this automorphism, some words get longer and some get shorter. It is no longer possible to constrain the behavior of geodesics relative to the Bass-Serre tree for the graph of groups structure as in [BBFS09]. This local stretching and compressing behavior of the automorphism is a phenomenon that was similarly faced in the work of Bridson and Groves in [BG10]. To accommodate this behavior, we rely on a coarser and more robust approach. We begin with the framework used in [BB00], and develop additional techniques for handling the lack of local control over geodesics. These ideas are described in more detail in the subsection "Edge group distortion" near the end of this Introduction.

The embedding trick. The method we use for embedding $S_{T,n}$ into a CAT(0) group is based on a twisting phenomenon for graphs of groups. Suppose $\alpha\colon G_e \to G_v$ is the inclusion map from an edge group to a vertex group, where e is incident to v. If one replaces α by $\varphi \circ \alpha$, where φ is an automorphism of $\alpha(G_e)$, then typically one obtains a very different fundamental group. However, if φ is the restriction of an *inner automorphism* of G_v, then the fundamental group remains unchanged.

The graph of groups structure of the group $S_{T,n}$ incorporates in an essential way the "twisting automorphisms" φ just described; properties of these automorphisms influence strongly the geometry of $S_{T,n}$. The presence of twisting accounts for the lack of non-positive curvature and the possibility of larger-than-quadratic isoperimetric behavior.

In order to embed $S_{T,n}$ into a larger group $G_{T,n}$, we give the latter the structure of a graph of groups that is very similar to that of $S_{T,n}$. It has the same underlying graph, somewhat larger edge and vertex groups, and inclusion maps which make use of twisting automorphisms closely related to those in $S_{T,n}$. This is to ensure the existence of a morphism of graphs of groups which induces a homomorphism

$S_{T,n} \to G_{T,n}$. Injectivity of this homomorphism is established using a criterion due to Bass [Bas93]; see Lemma 5.1.

The basic trick now is that the vertex group in $G_{T,n}$ has been enlarged specifically to arrange that the twisting automorphisms become *inner*. Thus the twisting can be undone, without changing the fundamental group. With no twisting, it becomes a simple matter to put a CAT(0) structure on $G_{T,n}$, provided the enlarged vertex group is already CAT(0).

Bieri doubles and the embedding trick. Earlier we referred to examples of subgroups of CAT(0) groups having exponentional or polynomial Dehn functions. The examples from [BBMS97, BRS07] made use of an embedding theorem from [BBMS97]. The authors embed certain doubles of groups into the direct product of a related group and a free group; the latter product may then have a CAT(0) structure, while the embedded subgroup has interesting geometric behavior. The embedded subgroup is called a *Bieri double*.

It turns out that some instances of this method can be viewed as instances of the embedding trick discussed above, resulting in an alternative way of looking at certain Bieri doubles.

Here is an example, which includes the [BRS07] examples of subgroups of CAT(0) groups having polynomial Dehn functions. Let N be any group with an automorphism φ, and consider the double of $N \rtimes_{\varphi} \mathbb{Z}$ along N: $(N \rtimes \langle t_1\rangle) *_N (N \rtimes \langle t_2\rangle)$. There is a homomorphism

$$(N \rtimes \langle t_1\rangle) *_N (N \rtimes \langle t_2\rangle) \;\to\; (N \rtimes \langle t\rangle) \times \langle u, v\rangle$$

which sends N to N, t_1 to tu, and t_2 to tv. The arguments of [BBMS97] show that this map is an embedding.

An alternative viewpoint is to note that the double $(N \rtimes \langle t_1\rangle) *_N (N \rtimes \langle t_2\rangle)$ has a graph of groups decomposition $\mathscr{A}$ with underlying graph a figure-eight, with all edge and vertex groups N, and with monodromy along each loop given by φ. The elements t_1 and t_2 are the two stable letters. The larger group $(N \rtimes \langle t\rangle) \times \langle u, v\rangle$ has a similar decomposition with the same underlying graph, with vertex and edge groups $N \rtimes \langle t\rangle$, with monodromy the identity. In this case u and v are the two stable letters.

This description of the larger group is the "untwisted" presentation for it. If we change the monodromy to be by the automorphism $\varphi \times \mathrm{id}$ of $N \rtimes \langle t\rangle$, then the new graph of groups $\mathscr{B}$ still has fundamental group $(N \rtimes \langle t\rangle) \times \langle u, v\rangle$, since $\varphi \times \mathrm{id}$ is conjugation by t in the vertex group $N \rtimes \langle t\rangle$. The two stable letters for this new decomposition are tu and tv. There is an obvious morphism of graphs of groups $\mathscr{A} \to \mathscr{B}$ since $\varphi \times \mathrm{id}$ restricts to φ on N. The induced map on fundamental groups is exactly the embedding displayed above.

An overview of the construction. Recall that our basic aim is to construct a pair of finitely presented groups S and G such that G is CAT(0), S embeds into G, and S has a specified Dehn function. The starting data needed for our constructions is as follows.

First choose a palindromic monotone automorphism $\varphi\colon \langle x, y\rangle \to \langle x, y\rangle$ (see Section 2 for these notions). Let λ be the Perron-Frobenius eigenvalue of the transition matrix of φ. Also let T be a finite tree with valence at most 3, and let $m = |T| + 1$ (here, $|T|$ denotes the number of vertices of T). Choose a positive integer n such that $\alpha = n\log_m(\lambda) \geqslant 1$. Based on these choices, we define finitely presented groups $S_{T,n}$ and $G_{T,n}$.

These groups are multiple HNN extensions of groups V_T and W_T respectively, with m stable letters. In Section 3 we define these vertex groups; they are, themselves, fundamental groups of graphs of groups with underlying graph T. The group V_T, in particular, is required to have some very special properties, notably the *balancing property*. The structure needed to achieve this property also dictates how W_T should be constructed.

Briefly, V_T has vertex groups isomorphic to $F_2 \times F_2 \times F_2$, and each such group has three designated *peripheral subgroups* which are used as edge groups; these edge groups are isomorphic to F_2. They are the "antidiagonal" copies of F_2 in each factor $F_2 \times F_2$ in each vertex group.

The group W_T is built in a similar way, with the free-by-cyclic group $G = F_2 \rtimes_\varphi \mathbb{Z}$ used in place of F_2; thus the vertex groups are $G \times G \times G$, and the edge groups are isomorphic to G. The groups G possess 2–dimensional CAT(0) structures, by work of Tom Brady [Bra95].

Given the CAT(0) structure on W_T, the ambient group $G_{T,n}$ can also be made CAT(0), as in the discussion of the Embedding Trick above; see Section 4. When building this CAT(0) structure, we fix the free group automorphism φ to be a particular one which is palindromic. Tom Brady's CAT(0) structure has a symmetry that respects the palindromic nature of φ. In the case of the other free-by-cyclic groups arising in [Bra95], we believe that analogous symmetries exist (in the palindromic case), but establishing this would take us too far afield. It is for this reason that the exponents in the main theorem involve $1 + \sqrt{2}$ (which is λ for this choice of φ). Apart from the CAT(0) statement for $G_{T,n}$, the rest of the paper does not require any particular choice of φ, and we compute the Dehn functions of $S_{T,n}$ in terms of a general (monotone, palindromic) φ.

Corridor schemes and σ–corridors. There are several places where we make use of corridor arguments in van Kampen diagrams. There are many different types of corridors under consideration simultaneously, and these different types are codified as *corridor schemes*. In particular, the group V_T possesses a large number of distinct corridor schemes, which are parametrized by maximal segments σ in a tree $\widehat{T}$. Each such segment determines a corridor scheme, whose corridors are called *σ–corridors*.

These corridors provide the primary means by which we establish the properties of V_T that are needed, such as the *balancing property* (Proposition 6.7). This latter property is somewhat awkward to explain, but it concerns the distribution of generators in words representing the trivial element. The property has the most force when the generators all lie in peripheral subgroups of V_T. It implies, for instance, that an element of a peripheral subgroup cannot be expressed efficiently using only generators from *other* peripheral subgroups; see Remarks 6.8. The balancing property plays a key role in several parts of the paper, most notably in the distortion bound of Proposition 9.7, and also in the area bound for V_T given in Proposition 10.3.

Least-area diagrams. In sections 7 and 8 we define families of van Kampen diagrams in order to establish the lower bound for the Dehn function of $S_{T,n}$. We use basic building blocks called *canonical diagrams*, which are defined for every palindromic word in $\langle x, y\rangle$. These are diagrams over the vertex group V_T. We then build *snowflake diagrams* over $S_{T,n}$, using canonical diagrams joined along strips dual to the stable letters of $S_{T,n}$.

To establish the lower bound, we show that all of these diagrams minimize area relative to their boundary words. Proving this requires a detailed study of σ–corridors and their intersection properties for various σ.

Edge group distortion. The heart of the computation of the Dehn function of $S_{T,n}$ lies in Proposition 9.7, in which we establish the distortion of the edge groups inside $S_{T,n}$. This argument requires first some properties of *folded corridors*, analogous to those studied in [BG10] in the context of free-by-cyclic groups. These properties are established in the first half of Section 9, culminating in Lemma 9.5.

Next we establish the bound on edge group distortion. The proof is an inductive proof based on Britton's Lemma. It falls into two cases, which require very different treatments. In the first case, the proof is based on a method from [BB00]. It is only thanks to the balancing property of V_T that this argument can be carried out.

The second case is where the folded corridors and Lemma 9.5 come in. The inductive framework is based on the nested structure of r_j–corridors in a putative van Kampen diagram. These corridors may appear in two possible orientations, *forwards* and *backwards*. Forward-facing corridors present few difficulties and can be handled using the method above based on the balancing property. If there is a backwards-facing r_j–corridor, then it may introduce undesirable geometric effects that threaten to spoil the inductive calculation. The argument in this case is to show that when this occurs, there will be forward facing corridors just behind the first one, and perfectly matching segments along these corridors, along which any metric distortion introduced by the first corridor is exactly undone.

Computing the Dehn function. Establishing the upper bound for the Dehn function of $S_{T,n}$ proceeds along similar lines as in [BBFS09]. One important step in this argument is a statement about area in the vertex group. This occurs in Proposition 10.3 in the present paper. Due to the more complicated structure of V_T (being constructed from free groups), the argument is considerably more involved than the corresponding result in [BBFS09]. It requires many of the tools developed here, such as the balancing property and corridor schemes.

Acknowledgments. Noel Brady acknowledges support from the NSF and from NSF award DMS-0906962. Max Forester acknowledges support from NSF award DMS-1105765.

2. Preliminaries

In this section we review some basic definitions and properties concerning Dehn functions, van Kampen diagrams, and words in the free group.

Dehn functions. Let $G = \langle A \mid R \rangle$ be a finitely presented group and w a word in the generators $A^{\pm 1}$ representing the trivial element of G. We define the *area* of w to be

$$\mathrm{Area}(w) \ = \ \min\Big\{N \in \mathbb{N} \;\Big|\; \exists \text{ equality } w = \prod_{i=1}^{N} u_j r_j u_j^{-1} \text{ freely, where } r_j \in R^{\pm 1}\Big\}.$$

The *Dehn function* $\delta(x)$ of the finite presentation $\langle A \mid R \rangle$ is given by

$$\delta(x) \ = \ \max\Big\{\mathrm{Area}(w) \;\Big|\; w \in \ker(F(A) \to G), |w| \leqslant x\Big\}$$

where $|w|$ denotes the length of the word w.

There is an equivalence relation on functions $f\colon \mathbb{N} \to \mathbb{N}$ defined as follows. First, we say that $f \preccurlyeq g$ is there is a constant $C > 0$ such that

$$f(x) \ \leqslant \ Cg(Cx) + Cx$$

for all $x \in \mathbb{N}$. If $f \preccurlyeq g$ and $g \preccurlyeq f$ then we say that f and g are *equivalent*, denoted $f \simeq g$. It is not difficult to show that two finite presentations of the same group define equivalent Dehn functions; we therefore speak of "the" Dehn function of G, which is well defined up to equivalence.

Remark 2.1. In order to show that $f \preccurlyeq g$ for non-decreasing functions f and g, it is sufficient to prove that $f(n_i) \leqslant g(n_i)$ for an unbounded sequence of positive integers $\{n_i\}$ such that the ratios n_{i+1}/n_i are bounded. For, if $n_{i+1} \leqslant Cn_i$ for all i and x is any integer between n_i and n_{i+1}, say, we have $f(x) \leqslant f(n_{i+1}) \leqslant g(n_{i+1}) \leqslant g(Cn_i) \leqslant g(Cx)$. Therefore $f(x) \leqslant Cg(x)$ for all x.

Let X be a 2–dimensional cell complex. We call X a *presentation 2–complex* if it has one 0–cell, and every 2–cell is attached by a map $f\colon S^1 \to X^{(1)}$ which is *regular* in the following sense: there is a cell structure for S^1 such that the restriction of f to each edge maps monotonically over a 1–cell of X.

The presentation 2–complex of the presentation $\langle A \mid R\rangle$ has oriented 1–cells labeled by the generators in A, and a 2–cell for each relator r in R, attached via a map $S^1 \to X^{(1)}$ which traverses edges sequentially, following the word r.

Given a presentation 2–complex X, one then has the notion of a *van Kampen diagram* over X. Briefly, a van Kampen diagram for the word w is a contractible, planar 2–complex with edges labeled by generators, with each 2–cell boundary word equal to a relator, with outer boundary word w. The *area* of the diagram is the number of its 2–cells. It is a standard fact that $\mathrm{Area}(w)$ as defined above can be interpreted as the minimal area of a van Kampen diagram over X for w. See [Bri02] for details on this interpretation of $\mathrm{Area}(w)$. We refer to [Bri02] for background on Dehn functions generally, and also to [BH99] for background on CAT(0) spaces.

Words and automorphisms. A word $w(x,y)$ is *palindromic* if $w(x^{-1},y^{-1}) = w(x,y)^{-1}$ as words in the free group. An automorphism $\varphi\colon \langle x,y\rangle \to \langle x,y\rangle$ is called *palindromic* if it takes palindromic words to palindromic words. Note that φ will be palindromic if the two words $\varphi(x)$, $\varphi(y)$ are palindromic. If φ is palindromic, so is φ^n for any $n \geqslant 1$.

A word $w(x,y)$ is *positive* if it does not contain occurrences of x^{-1} or y^{-1}. It is *negative* if it does not contain x or y. It is *monotone* if it is positive or negative. Note that monotone words are reduced, and if w is monotone then so is w^{-1}. If an automorphism φ takes x and y to monotone words of the same kind, then it takes all monotone words to monotone words. We call φ *monotone* if it has this property. The same will then be true of φ^n for any $n \geqslant 1$. (The inverse φ^{-1} will typically *not* be monotone.)

3. Group-theoretic constructions

In this section we begin by constructing groups V_T and W_T, which will then serve as vertex groups in graph of groups decompositions defining the snowflake group $S_{T,n}$ and the CAT(0) group $G_{T,n}$.

Our constructions make use of a free group automorphism $\varphi\colon \langle x,y\rangle \to \langle x,y\rangle$ which is both palindromic and monotone. For concreteness, we define φ to be the automorphism given by $\varphi(x) = xyx$, $\varphi(y) = x$. In most of what follows, only the palindromic and monotone properties of φ are relevant. However, in Section 4 we define the CAT(0) structure for $G_{T,n}$ based on the knowledge that φ is the automorphism just defined. This section is the only place where explicit knowledge of φ is used. In particular, all results concerning the groups $S_{T,n}$ are valid for any palindromic, monotone φ.

Let λ be the exponential growth rate of φ. That is, λ is the Perron-Frobenius eigenvalue of the transition matrix $M_\varphi = \begin{pmatrix} |\varphi(x)|_x & |\varphi(y)|_x \\ |\varphi(x)|_y & |\varphi(y)|_y \end{pmatrix}$. From the beginning, we will fix an integer $n \geqslant 1$ and work with the automorphism φ^n.

The groups F and G. In this paper F will always denote the free group of rank two, $\langle x, y\rangle$. We define G to be the free-by-cyclic group $F \rtimes_{\varphi^n} \langle t\rangle$. That is,

$$G \;=\; \langle\, x, y, t \mid txt^{-1} = \varphi^n(x),\; tyt^{-1} = \varphi^n(y)\,\rangle.$$

One verifies easily that because φ^n is palindromic, there is an involution $\tau_G\colon\, G \to G$ defined by $\tau_G(x) = \overline{x}$, $\tau_G(y) = \overline{y}$, and $\tau_G(t) = t$ (bar denotes inverse). Similarly, define the involution $\tau_F\colon\, F \to F$ by $\tau_F(x) = \overline{x}$, $\tau_F(y) = \overline{y}$. We may refer to either involution simply as τ.

The groups we construct will contain many copies of F and G. The different copies will be indexed along with their generators: $F_i = \langle x_i, y_i\rangle$ and $G_i = \langle x_i, y_i\rangle \rtimes_{\varphi^n} \langle t_i\rangle$.

The groups V and W. These groups are defined as follows:

$$\begin{aligned} V \;&=\; F_0 \times F_1 \times F_2, \\ W \;&=\; G_0 \times G_1 \times G_2. \end{aligned}$$

Thus V has generators $x_0, y_0, x_1, y_1, x_2, y_2$ and W contains V along with the additional generators t_0, t_1, t_2.

Define the following subgroups (indices mod 3):

$$\begin{aligned} A_i \;&=\; \langle \overline{x}_i x_{i+1}, \overline{y}_i y_{i+1}\rangle \;<\; F_i \times F_{i+1} \;<\; V, \\ B_i \;&=\; \langle \overline{x}_i x_{i+1}, \overline{y}_i y_{i+1}, t_i t_{i+1}\rangle \;<\; G_i \times G_{i+1} \;<\; W. \end{aligned}$$

There are injective homomorphisms $F \to F_i \times F_{i+1}$ and $G \to G_i \times G_{i+1}$ given by $\tau_F \times \mathrm{id}$ and $\tau_G \times \mathrm{id}$ respectively, taking x to $\overline{x}_i x_{i+1}$, y to $\overline{y}_i y_{i+1}$, and t to $t_i t_{i+1}$. The subgroups A_i and B_i are the images of these homomorphisms, and therefore are isomorphic to F and G, with the generators listed above corresponding to the standard generators x, y, t. We name these generators as follows:

$$a_i = \overline{x}_i x_{i+1}, \;\; b_i = \overline{y}_i y_{i+1}, \;\; c_i = t_i t_{i+1}.$$

Thus, $A_i = \langle a_i, b_i\rangle$ and $B_i = \langle a_i, b_i\rangle \rtimes_{\varphi^n} \langle c_i\rangle$. These subgroups will be called *peripheral subgroups.*

The groups V_T and W_T. These groups will be obtained by amalgamating copies of V (respectively, W) together along peripheral subgroups.

Let T be a finite tree of valence at most 3. Choose a copy of V for each vertex of T, and assign distinct peripheral subgroups of V to each of the outgoing edges at that vertex. Then, for each edge in T, amalgamate the associated peripheral subgroups of the two copies of V via the isomorphism τ (relative to their standard generating sets).[1] The resulting group is V_T. In the case where T is a single vertex, V_T is just V.

The group W_T is defined by the same procedure with vertex groups W instead of V. Again, if T has one vertex, then $W_T = W$.

In order to have consistent notation, assign non-overlapping triples of indices $(0,1,2)$, $(3,4,5)$, etc. to the vertices of T, and use these in place of $(0,1,2)$ in the definitions of V and W above. For example, in the case of W_T, if a vertex has triple $(3,4,5)$, then the vertex group is $G_3 \times G_4 \times G_5$ with peripheral subgroups labeled B_3, B_4, and B_5. The subgroup B_5 has standard generators $a_5 = \overline{x}_5 x_3$, $b_5 = \overline{y}_5 y_3$, and $c_5 = t_5 t_3$.

With this notation, if an edge e of T is assigned the peripheral subgroups B_i and B_j in its neighboring vertex groups, then amalgamating along e adds the relations $a_i = \overline{a}_j$, $b_i = \overline{b}_j$, and $c_i = c_j$.

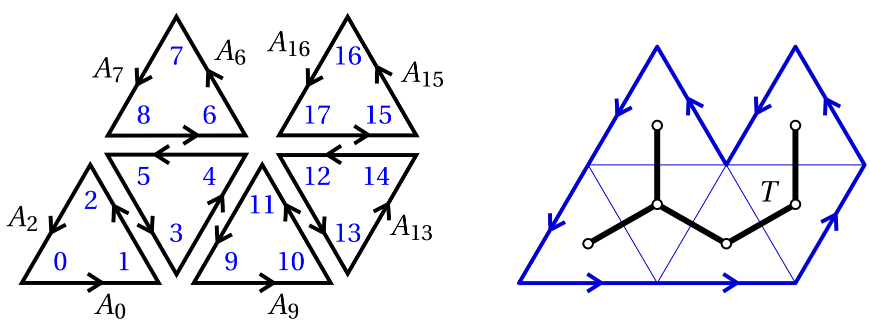

FIGURE 1. The diagram D on the left, and the triangulated $(|T|+2)$–gon on the right, with dual graph T. Each triangle with corners labeled i, j, k represents the vertex group $F_i \times F_j \times F_k$ or $G_i \times G_j \times G_k$. The side pairings depict the amalgamations $A_1 =_\tau A_5$, $A_3 =_\tau A_{11}$, $A_4 =_\tau A_8$, $A_{10} =_\tau A_{12}$, and $A_{14} =_\tau A_{17}$ (in the case of V_T).

It may be helpful to consider the diagram D as shown in Figure 1. It has a triangle for each vertex of T, with corners corresponding to the factors F_i or G_i of the vertex group there. The edges correspond to peripheral subgroups, and the edge-pairings between triangles correspond to amalgamations. The triangles assemble into a triangulated $(|T|+2)$–gon with dual graph T. (Here, $|T|$ denotes the number of vertices of T.)

[1]There is no need to specify a direction because τ is an involution.

The orientations on edges indicate the standard generating sets of the peripheral subgroups, relative to the indexing of the groups F_i or G_i. The orientation-reversing nature of the side pairings reflects the fact that the amalgamations are performed using τ.

The *peripheral subgroups* of V_T (or of W_T) are defined to be the remaining peripheral subgroups of the vertex groups that were not assigned to edges of T. In terms of the diagram D, these are the peripheral subgroups corresponding to the edges forming the boundary $(|T|+2)$–gon.

Next we re-index the peripheral subgroups. Note that the boundary edges along D are coherently oriented. Start with one and let v_0 be the index of the corresponding peripheral subgroup. Following the orientation, let v_1 be the index of the next edge along ∂D. Repeat in this way and define the indices $v_0 \ldots, v_m$, allowing us to refer to the peripheral subgroups as $A_{v_0}, \ldots, A_{v_m}$ (or $B_{v_0}, \ldots, B_{v_m}$). Note, $m = |T| + 1$.

The groups $S_{T,n}$ and $G_{T,n}$. Fix a tree T as above and let $m = |T| + 1$. The group $S_{T,n}$ is defined to be a multiple HNN extension over V_T with stable letters $r_1, \ldots, r_m$, where r_i conjugates the peripheral subgroup A_{v_0} to A_{v_i} via the automorphism φ^n. That is,

$$\begin{aligned} S_{T,n} = \langle\, V_T, r_1, \ldots, r_m \mid r_i a_{v_0} r_i^{-1} &= \varphi^n(a_{v_i}), \\ r_i b_{v_0} r_i^{-1} &= \varphi^n(b_{v_i}) \text{ for each } i \,\rangle. \end{aligned} \tag{3.1}$$

Thus $S_{T,n}$ is the fundamental group of a graph of groups whose underlying graph is the m–rose (having one vertex and m loops). The vertex group is V_T and the edge groups are all F.

We define $G_{T,n}$ in a similar manner, but *without twisting*. It is a multiple HNN extension over W_T with stable letters $s_1, \ldots, s_m$, where s_i conjugates B_{v_0} to B_{v_i} via the identity map:

$$\begin{aligned} G_{T,n} = \langle\, W_T, s_1, \ldots, s_m \mid s_i a_{v_0} s_i^{-1} &= a_{v_i}, \\ s_i b_{v_0} s_i^{-1} = b_{v_i},\ s_i c_{v_0} s_i^{-1} &= c_{v_i} \text{ for each } i \,\rangle. \end{aligned} \tag{3.2}$$

Again, $G_{T,n}$ is the fundamental group of a graph of groups over the m–rose. The vertex group is W_T and the edge groups are all G.

4. The CAT(0) structure

In this section we build the CAT(0) structure for the group $G_{T,n}$.

The space Z. Recall that we have chosen a specific monotone palindromic automorphism $\varphi\colon F \to F$ given by $\varphi(x) = xyx$, $\varphi(y) = x$. Let

$$G_0 \;=\; \langle x, y\rangle \rtimes_\varphi \langle t_0\rangle \;=\; \langle\, x, y, t_0 \mid t_0 x t_0{}^{-1} = xyx,\ t_0 y t_0{}^{-1} = x \,\rangle. \tag{4.1}$$

T. Brady [Bra95] has constructed a piecewise Euclidean, locally CAT(0) 2–complex Z_0 with fundamental group G_0. This 2–complex has two vertices, four edges, and two 2–cells, consisting of a Euclidean octagon and quadrilateral as shown in Figure 2. The angle $\alpha \in (0,\pi)$ is a free parameter, and the rest of the geometry (up to scaling) is then determined. It is easy to check that both vertices satisfy the link condition, making Z_0 locally CAT(0).

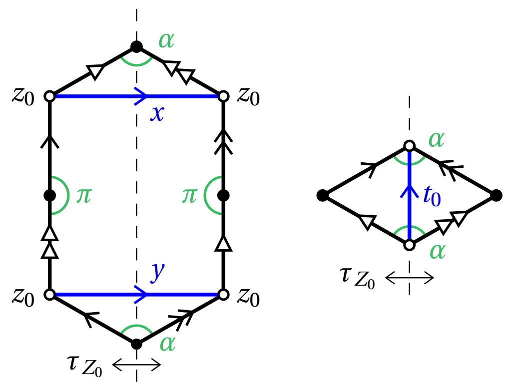

FIGURE 2. The piecewise Euclidean 2–complex Z_0 and its isometric involution τ_{Z_0}.

The figure also shows three arcs crossing the interiors of the 2–cells, representing the elements x, y, and t_0 in $\pi_1(Z_0, z_0)$; we leave it to the reader to verify that $\pi_1(Z_0, z_0)$ indeed has the presentation (4.1) relative to these generators.

Reflection of each 2–cell across the vertical dotted lines in Figure 2 respects the edge identifications, and induces an isometric involution $\tau_{Z_0}\colon Z_0 \to Z_0$. The induced homomorphism $\tau_{G_0}\colon G_0 \to G_0$ is given by $\tau_{G_0}(x) = \overline{x}$, $\tau_{G_0}(y) = \overline{y}$, and $\tau_{G_0}(t_0) = t_0$.

Let $t = (t_0)^n$ and note that the index-n subgroup $\langle x, y, t\rangle \lhd G_0$ is the group G defined earlier. The corresponding covering space Z of Z_0 has a locally CAT(0) structure made from n octagons and n quadrilaterals. The involution τ_{Z_0} lifts to an isometric involution $\tau_Z\colon Z \to Z$, with induced homomorphism τ_G. The case $n = 3$ is shown in Figure 3.

To summarize, we now have a locally CAT(0) space Z with fundamental group G, and an isometric involution $\tau_Z\colon Z \to Z$ whose induced homomorphism is given by the involution τ_G.

The space K_T. We shall need the following "gluing with a tube" result.

Proposition 4.2 ([BH99], II.11.13)**.** *Let X and A be locally CAT(0) metric spaces. If A is compact and $\varphi, \psi\colon A \to X$ are locally isometric immersions, then the quotient of $X \sqcup (A \times [0,1])$ by the equivalence relation generated by $(a,0) \sim \varphi(a)$; $(a,1) \sim \psi(a)$, $\forall a \in A$ is locally CAT(0).*

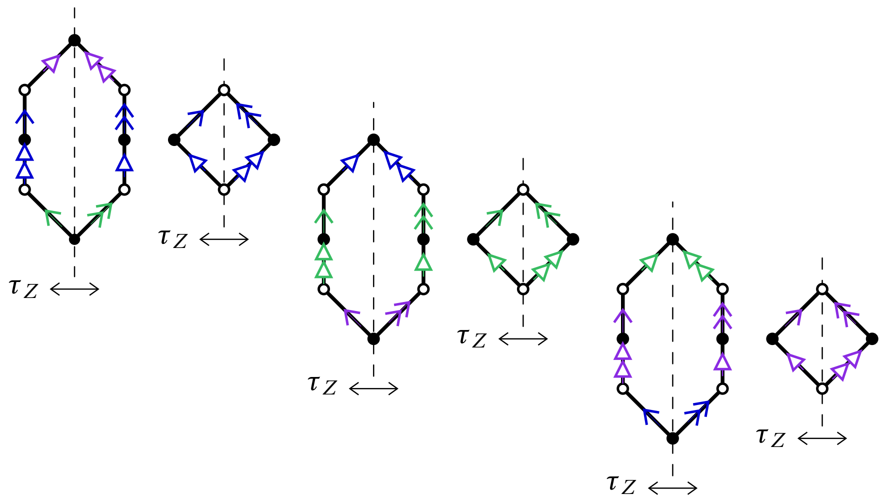

FIGURE 3. The locally CAT(0) 2–complex Z and its involution τ_Z, when $n = 3$.

Let Z_i and Z_j be copies of Z with fundamental groups G_i and G_j repsectively. Let $Z_i \times Z_j$ be given the product metric. Define the map $f_{i,j}\colon Z \to Z_i \times Z_j$ by $f_{i,j}(p) = (\tau_z(p), p)$. Note that the induced homomorphism $(f_{i,j})_*\colon G \to G_i \times G_j$ is given by $x \mapsto \overline{x}_i x_j$, $y \mapsto \overline{y}_i y_j$, $t \mapsto t_i t_j$. Metrically, $f_{i,j}$ behaves as follows:

$$\begin{aligned} d(f_{i,j}(p), f_{i,j}(q)) &= d((\tau_Z(p), p), (\tau_Z(q), q)) \\ &= (d(\tau_Z(p), \tau_Z(q))^2 + d(p,q)^2)^{1/2} \\ &= (2d(p,q)^2)^{1/2} = \sqrt{2}d(p,q). \end{aligned}$$

Hence $f_{i,j}$ is an isometric embedding of the scaled metric space $(\sqrt{2})Z$ into $Z_i \times Z_j$.

Now we define the locally CAT(0) space K_T with fundamental group W_T. Let

$$K = Z_0 \times Z_1 \times Z_2$$

where each Z_i is a copy of Z. Thus, K is locally CAT(0) and has fundamental group $W = G_0 \times G_1 \times G_2$. Fix a basepoint $v \in Z$ and let $v_i \in Z_i$ be the corresponding points. Each product space $Z_i \times Z_{i+1}$ isometrically embeds into K using the basepoint v_{i-1} as the missing coordinate (indices mod 3). Define the *peripheral subspace* L_i to be the image of the map

$$Z \xrightarrow{f_{i,i+1}} Z_i \times Z_{i+1} \hookrightarrow K. \tag{4.3}$$

Note that L_i has fundamental group $B_i < W$ and the induced map $\pi_1(Z, v) \to W$ is the standard inclusion map $x \mapsto a_i$, $y \mapsto b_i$, $t \mapsto c_i$ with image B_i.

The space K_T is formed from copies of K in the same way that W_T is built from copies of W. Take a copy of K for each vertex of T, with all indices re-named to agree with the triple of indices assigned to that vertex. Whenever B_i and B_j were amalgamated in W_T, glue the ends of a *tube* $(\sqrt{2})Z \times [0,1]$ to the peripheral subspaces L_i and L_j, using the isometric embedding (4.3) from $(\sqrt{2})Z \times \{0\}$ to the copy of K containing L_i, and using a similar isometric embedding

$$Z \xrightarrow{\tau_Z} Z \xrightarrow{f_{j,j+1}} Z_j \times Z_{j+1} \hookrightarrow K$$

from $(\sqrt{2})Z \times \{1\}$ to the copy of K containing L_j. The involution τ_Z is being used to obtain the correct identification between the subgroups B_i and B_j. The resulting space K_T has fundamental group W_T, and is locally CAT(0) by Proposition 4.2. In particular, W_T is CAT(0).

Remark 4.4. The reasoning above also shows that V_T is CAT(0). One simply re-defines Z to be the space $S^1 \vee S^1$ with any path metric (which will be locally CAT(0)). There is an obvious isometric involution τ_Z which reverses the direction of each loop in Z, and induces $\tau_F \colon F \to F$. The rest is entirely similar.

The space $K_{T,n}$. Inside K_T there are peripheral subspaces $L_{v_0}, \dots, L_{v_m}$. For each $j = 1, \dots, m$ glue the ends of a tube $(\sqrt{2})Z \times [0,1]$ to L_{v_0} and L_{v_j} using the isometric embeddings (4.3) from $(\sqrt{2})Z \times \{0\}$ and $(\sqrt{2})Z \times \{1\}$ to the appropriate copies of K in K_T. The resulting space $K_{T,n}$ is the total space of a graph of spaces corresponding to the description (3.2) of $G_{T,n}$ as the fundamental group of a graph of groups. In particular, $K_{T,n}$ has fundamental group $G_{T,n}$. It is locally CAT(0) by Proposition 4.2. Thus we have proved:

Theorem 4.5. *$G_{T,n}$ is CAT(0).* □

5. Embedding results

In this section we define the embedding $S_{T,n} \to G_{T,n}$ and prove that it is injective. The map is defined step by step, following the constructions defining $S_{T,n}$ and $G_{T,n}$. In several places we use the following lemma to establish injectivity. It is a special case of a basic result of Bass [Bas93], reformulated slightly.

Lemma 5.1 (Injectivity for graphs of groups)**.** *Suppose $\mathscr{A}$ and $\mathscr{B}$ are graphs of groups such that the underlying graph $\Gamma_{\mathscr{A}}$ of $\mathscr{A}$ is a subgraph of the underlying graph of $\mathscr{B}$. Let A and B be their respective fundamental groups. Suppose that there are injective homomorphisms $\psi_e \colon A_e \to B_e$ and $\psi_v \colon A_v \to B_v$ between edge and vertex groups, for all edges e and vertices v in $\Gamma_{\mathscr{A}}$, which are compatible with the edge-inclusion maps.*

(That is, whenever e has initial vertex v, the diagram

$$\begin{tikzcd} A_e \arrow[r] \arrow[d, "\psi_e"'] & A_v \arrow[d, "\psi_v"] \\ B_e \arrow[r] & B_v \end{tikzcd}$$

commutes.)

If $\psi_e(A_e) = \psi_v(A_v) \cap B_e$ whenever e has initial vertex v, then the induced homomorphism $\psi\colon A \to B$ is injective.

Remark 5.2. Given the initial assumptions, it is always true that $\psi_e(A_e) \subset (\psi_v(A_v) \cap B_e)$. In practice one only needs to verify that $\psi_e(A_e)$ contains $\psi_v(A_v) \cap B_e$.

Proof. The homomorphisms ψ_e, ψ_v combine to give a morphism of graphs of groups in the sense of Bass [Bas93]. According to Proposition 2.7 of [Bas93], $\psi\colon A \to B$ will be injective if, whenever e has initial vertex v, the function $A_v/A_e \to B_v/B_e$ induced by ψ_v is injective.

To prove that the latter statement holds, suppose that the cosets $\psi_v(a_1)B_e$ and $\psi_v(a_2)B_e$ are equal for some $a_1, a_2 \in A_v$. Then $\psi_v(a_1a_2^{-1}) \in B_e$, and hence (by the main assumption) $\psi_v(a_1a_2^{-1}) \in \psi_e(A_e)$. Since ψ_e is injective (and agrees with ψ_v), $a_1a_2^{-1} \in A_e$, and therefore $a_1A_e = a_2A_e$. □

Lemma 5.3. *Let $\iota\colon V \to W$ be inclusion. Then $\iota(A_i) = \iota(V) \cap B_i$ for each i.*

Proof. Without loss of generality let $i = 0$. Note that $\iota(A_0)$ and $\iota(V) \cap B_0$ are both contained in the subgroup $G_0 \times G_1$, so it suffices to show that $\iota(A_0) = \iota(F_0 \times F_1) \cap B_0$ in $G_0 \times G_1$. One direction, $\iota(A_0) \subset \iota(F_0 \times F_1) \cap B_0$, is obvious.

For the other direction, consider an element of $\iota(F_0 \times F_1) \cap B_0$. It can be expressed as a word $w(\overline{x}_0x_1, \overline{y}_0y_1, t_0t_1)$, which equals $w(\overline{x}_0, \overline{y}_0, t_0)w(x_1, y_1, t_1)$ in $G_0 \times G_1$. Being in $\iota(F_0 \times F_1)$ it also has an expression of the form $u(\overline{x}_0, \overline{y}_0)v(x_1, y_1)$ where u and v are reduced words in the free group. Projecting onto the second factor of $G_0 \times G_1$, one obtains the equation in G:

$$w(x, y, t) = v(x, y).$$

Considering G as an HNN extension with vertex group F and stable letter t, the right hand side is a word in normal form (that is, a word of length 1 consisting of an element of F), and therefore gives the (unique) normal form representative for the element w. Similarly, considering $\{\overline{x}_0, \overline{y}_0\}$ as a basis for F_0, projecting onto the first factor gives the equation in G

$$w(x, y, t) = u(x, y)$$

and hence $u(x, y)$ is also the normal form for w. We conclude that u and v represent the same element of F. Since both words are reduced, they are equal as words and so $u(\overline{x}_0, \overline{y}_0)v(x_1, y_1) = u(\overline{x}_0, \overline{y}_0)u(x_1, y_1) = u(\overline{x}_0x_1, \overline{y}_0y_1) \in \iota(A_0)$. □

Proposition 5.4. *The inclusion maps $F_i \hookrightarrow G_i$ induce an injective homomorphism $V_T \to W_T$.*

Henceforth we will regard V_T as a subgroup of W_T.

Proof. We will use Lemma 5.1 since V_T and W_T are both fundamental groups of graphs of groups with underlying graph T.

To elaborate on the graph of groups structures of V_T and W_T, fix an orientation of each edge of T and use these to specify the edge-inclusion maps as follows. For V_T, each edge group is F and the two neighboring vertex groups are isomorphic to V. For the *initial* vertex, the inclusion map $F \to F_i \times F_{i+1}$ is $\tau_F \times \mathrm{id}$, and for the *terminal* vertex, the inclusion map $F \to F_j \times F_{j+1}$ is $\mathrm{id} \times \tau_F$. In the case of W_T, the same convention is used: inclusion maps $G \to G_i \times G_{i+1}$ are $\tau_G \times \mathrm{id}$ for initial vertices and $\mathrm{id} \times \tau_G$ for terminal vertices.

The inclusion maps $F_i \to G_i$ induce inclusions between corresponding vertex groups $V \to W$. The compatibility diagrams become

$$\begin{array}{ccccc} F & \xrightarrow{\tau_F \times \mathrm{id}} & F_i \times F_{i+1} & \longrightarrow & V \\ \downarrow & & & & \downarrow \\ G & \xrightarrow{\tau_G \times \mathrm{id}} & G_i \times G_{i+1} & \longrightarrow & W \end{array} \quad \text{or} \quad \begin{array}{ccccc} F & \xrightarrow{\mathrm{id} \times \tau_F} & F_i \times F_{i+1} & \longrightarrow & V \\ \downarrow & & & & \downarrow \\ G & \xrightarrow{\mathrm{id} \times \tau_G} & G_i \times G_{i+1} & \longrightarrow & W \end{array}$$

and these clearly commute (all unnamed maps are inclusion). Thus there is an induced homomorphism $V_T \to W_T$. The last condition needed by Lemma 5.1 is provided by Lemma 5.3, and so we conclude from 5.1 that $V_T \to W_T$ is injective. $\square$

Change of coordinates in $G_{T,n}$. We plan to use Lemma 5.1 to embed $S_{T,n}$ into $G_{T,n}$, but first we must modify the graph of groups structure of $G_{T,n}$. The modification amounts to a change in the choice of stable letters in the multiple HNN extension description of $G_{T,n}$. Alternatively, it can be seen as an application of Tietze transformations.

Indeed, one can start with the presentation (3.2) defining $G_{T,n}$, add new generators $u_1, \ldots, u_m$ and relations $u_i = c_{v_i} s_i$, replace occurrences of s_i with $c_{v_i}^{-1} u_i$, and delete the generators s_i. The relation $s_i a_{v_0} s_i^{-1} = a_{v_i}$ becomes $c_{v_i}^{-1} u_i a_{v_0} u_i^{-1} c_{v_i} = a_{v_i}$, or equivalently, $u_i a_{v_0} u_i^{-1} = c_{v_i} a_{v_i} c_{v_i}^{-1} = \varphi^n(a_{v_i})$. Similarly, the relation $s_i b_{v_0} s_i^{-1} = b_{v_i}$ becomes $u_i b_{v_0} u_i^{-1} = \varphi^n(b_{v_i})$ and the relation $s_i c_{v_0} s_i^{-1} = c_{v_i}$ becomes $u_i c_{v_0} u_i^{-1} = c_{v_i}$. Thus one obtains the new presentation

$$\begin{aligned} G_{T,n} = \langle W_T, u_1, \ldots, u_m \mid u_i a_{v_0} u_i^{-1} &= \varphi^n(a_{v_i}), \\ u_i b_{v_0} u_i^{-1} &= \varphi^n(b_{v_i}),\ u_i c_{v_0} u_i^{-1} = c_{v_i} \text{ for each } i \rangle. \end{aligned} \tag{5.5}$$

This is evidently the presentation arising from a new description of $G_{T,n}$ as a multiple HNN extension of W_T with stable letters $u_1, \ldots, u_m$, where u_i conjugates B_{v_0} to B_{v_i} via $\varphi^n \times \mathrm{id}$.

Theorem 5.6. *The homomorphism $S_{T,n} \to G_{T,n}$ induced by the inclusion $V_T \hookrightarrow W_T$ and the assigment $r_i \mapsto u_i$ is injective.*

Proof. First, given the presentations (3.1) and (5.5), it is clear that the given assignment defines a homomorphism $S_{T,n} \to G_{T,n}$. Furthermore, this is the homomorphism induced by the injective maps on vertex and edge groups: $V_T \hookrightarrow W_T$ in the case of the vertex, and $F \hookrightarrow G$ for each edge of the m–rose. The corresponding compatibility diagrams are

$$\begin{tikzcd} F \arrow[r, "\mathrm{id}"] \arrow[d] & A_{v_0} \arrow[r] & V_T \arrow[d] \\ G \arrow[r, "\mathrm{id}"] & B_{v_0} \arrow[r] & W_T \end{tikzcd} \quad \text{and} \quad \begin{tikzcd} F \arrow[r, "\varphi^n"] \arrow[d] & A_{v_i} \arrow[r] & V_T \arrow[d] \\ G \arrow[r, "\varphi^n \times \mathrm{id}"] & B_{v_i} \arrow[r] & W_T \end{tikzcd}$$

where A_j and B_j are canonically identified with F and G via their standard generating sets, and the unnamed maps are inclusion. These diagrams commute.

We have all of the initial hypotheses of Lemma 5.1 satisfied. It remains to verify that $A_{v_i} = V_T \cap B_{v_i}$ in W_T for $i = 0, \dots, m$. Consider the vertex of T whose triple of indices includes v_i. Let V and W be the vertex groups at that vertex (for the graph of groups decompositions of V_T and W_T). Then A_{v_i} and $V_T \cap B_{v_i}$ are both contained in W. Moreover $W \cap V_T = V$, and so it suffices to show that $A_{v_i} = V \cap B_{v_i}$. This holds by Lemma 5.3. Hence, by Lemma 5.1, the map $S_{T,n} \to G_{T,n}$ is injective. □

6. Corridor schemes and the balancing property of V_T

In this section we develop two key tools which will play an essential role throughout the rest of the paper. These tools are specific to the groups V_T, and they are the primary means by which we establish the various properties of V_T that are needed. The first of these is the notion of σ–corridors in van Kampen diagrams over V_T. The second is the *balancing property* of V_T, given in Proposition 6.7.

In order to discuss σ–corridors we first define *corridor schemes*. We will make use of several corridor schemes in this paper, in addition to σ–corridors. In this section we also discuss the standard generating set for V_T, and various notions of length associated with this generating set.

The 2–complex X_T. In order to discuss area in V_T we will work with a specific 2–complex X_T with fundamental group V_T.

The group V has a presentation with generators x_i, y_i, a_i, b_i for $i = 0, 1, 2$ and eighteen relations (see also Figure 5):

$$\begin{aligned} &a_i = \overline{x}_i x_{i+1}, \;\; a_i = x_{i+1}\overline{x}_i, \;\; b_i = \overline{y}_i y_{i+1}, \;\; b_i = y_{i+1}\overline{y}_i, \\ &x_i y_{i+1} = y_{i+1} x_i, \;\; x_{i+1} y_i = y_i x_{i+1} \;\; (i = 0, 1, 2 \mod 3) \end{aligned} \tag{6.1}$$

We define X to be the presentation 2–complex for this presentation of V. Thus X has one 0–cell, twelve labeled, oriented 1–cells, twelve triangular 2–cells, and six quadrilateral 2–cells.

For each i, the subcomplex $Y_i \subset X$ consisting of the two 1–cells labeled a_i and b_i is called a *peripheral subspace*. It is homeomorphic to $S^1 \vee S^1$ and has fundamental group $A_i \subset V$.

The 2–complex X_T is formed from copies of X in the same way that V_T is built from copies of V. Take a copy of X for each vertex of T, with edge labels re-indexed according to the triple of indices assigned to that vertex. Whenever A_i and A_j were amalgamated in V_T, glue the peripheral subspaces Y_i and Y_j via a cellular homeomorphism which induces τ between A_i and A_j. The resulting space X_T, with fundamental group V_T, has a natural cell structure. The 1–cells are labeled by the generators x_i, y_i, a_i, and b_i, where in some cases, a 1–cell labeled a_i or b_i is also labeled a_j or b_j in the opposite direction. The 2–cells are the same as those of the copies of X, with the same triangular and quadrilateral boundary relations.

Area. In order to simplify the area calculations to follow, we declare each triangular cell of X_T to have area 1, and each quadrilateral cell to have area 2. (Think of it as being made of two triangles.)

Corridor schemes. Let Z be any presentation 2–complex. A *corridor scheme* for Z is a subset $\mathscr{S}$ of the set of edges of Z such that every 2–cell of Z has either zero or two occurrences of edges of $\mathscr{S}$ in its boundary. Given such an $\mathscr{S}$, one can then define *corridors* in any van Kampen diagram over Z. Call the 2–cells having two $\mathscr{S}$–edges in their boundaries *corridor cells*. Given a van Kampen diagram Δ, two corridor cells in Δ are called *neighbors* if they meet along an $\mathscr{S}$–edge in their boundaries. A corridor cell has zero, one, or two neighbors. A *corridor* is a minimal collection C of corridor cells in Δ such that if $c \in C$ then all neighbors of c are also in C. Every corridor cell is contained in a corridor.

Corridors come in two types: those in which every corridor cell has neighbors along both of its $\mathscr{S}$–edges, called *annulus type*, and the others, called *band type*. Each band type corridor joins an $\mathscr{S}$–edge on the boundary of Δ to another $\mathscr{S}$–edge on the boundary of Δ, and contains no other $\mathscr{S}$–edges on the boundary of Δ. An annulus type corridor may have 2–cells meeting the boundary of Δ, but the $\mathscr{S}$–edges of such 2–cells will not be on the boundary.

If C is a corridor in Δ, then the subset formed by taking the union of the interiors of its 2–cells along with the interiors of its $\mathscr{S}$–edges is an open set, homeomorphic to a tubular neighborhood of a properly embedded connected 1–dimensional submanifold of Δ. The 1–manifold meets each corridor cell in an arc joining the two $\mathscr{S}$–edges of the cell. Thus an annulus type corridor contains an embedded open annulus, and a band type corridor contains an embedded open band $[0,1]\times(0,1)$ meeting the boundary of the diagram in its boundary $\{0,1\}\times(0,1)$.

Corridors have two key properties. First, two corridors in a diagram Δ will never have 2–cells or $\mathscr{S}$–edges in common. In particular, corridors cannot *cross*. Second, every $\mathscr{S}$–edge appearing on the boundary of Δ is part of a band type corridor, unless that edge is not in any 2–cell of Δ. In particular, given an $\mathscr{S}$–edge in the boundary of Δ, if there is a 2–cell containing that edge, then one can pass from neighbor to neighbor in the corridor, until one arrives at a second, uniquely determined, $\mathscr{S}$–edge on the boundary of Δ. Also, in the boundary, this pair of $\mathscr{S}$–edges cannot be linked with another such pair, because corridors do not cross.

Orientable corridor schemes. A corridor scheme $\mathscr{S}$ is *orientable* if there is a choice of orientations of the edges of $\mathscr{S}$ such that in each corridor cell, the two $\mathscr{S}$–edges have oppposite orientations relative to the boundary of the cell. It follows that in any corridor, the transverse orientations of the $\mathscr{S}$–edges along the corridor all agree.

Remark 6.2. A corridor scheme defines a 1–dimensional $\mathbb{Z}_2$–valued cellular cocycle in Z. (If Z happens to be a simplicial complex, then every 1–cocycle is a corridor scheme.) An orientable corridor scheme defines a $\mathbb{Z}$–valued 1–cocycle in Z. See Gersten [Ger98] for a thorough study of corridors from the cohomological point of view.

σ–corridors. Recall that D was a diagram made of triangles, with dual graph T, which may be regarded as being embedded as a subspace of a triangulated $(|T|+2)$–gon. Let $\widehat{T}$ be a tree obtained from T by joining the midpoint of each boundary edge to the vertex in the neighboring triangle. Then $\widehat{T}$ has $m+1$ leaves, corresponding to the peripheral subgroups of V_T, and interior vertices all of valence three, which are the original vertices of T. Denote the leaves of $\widehat{T}$ by $v_{v_0}, \dots, v_{v_m}$, so that v_{v_i} corresponds to the peripheral subgroup A_{v_i}.

Let σ be a maximal segment in $\widehat{T}$. Note that σ is uniquely determined by its endpoints; choosing σ amounts to choosing a pair of peripheral subgroups of V_T. For each such σ we will define a corridor scheme $\mathscr{S}_\sigma$ for X_T.

In the $(|T|+2)$–gon, σ starts on a boundary edge, passes through a sequence of triangles, and ends on a boundary edge. Its intersection with each of these triangles is an arc joining two sides. It separates one corner of the triangle from the other two. If i is the index of this corner, put the edges of X_T labeled by x_i and y_i into $\mathscr{S}_\sigma$. Also, if σ passes through the side of a triangle associated with the subgroup A_j, put the edges labeled by a_j and b_j into $\mathscr{S}_\sigma$. Do this for each triangle that intersects σ to obtain $\mathscr{S}_\sigma$. The fact that some edges of X_T have two labels is not a problem; either both labels or neither label will be chosen for inclusion in $\mathscr{S}_\sigma$. See Figure 4.

One verifies easily that $\mathscr{S}_\sigma$ is a corridor scheme, by examining the relations (6.1) for each triangle of D. See also Figure 5. (Because of the two-label phenomenon, it is important here that σ is *maximal*.) Corridors defined by this scheme will be called *σ–corridors.*

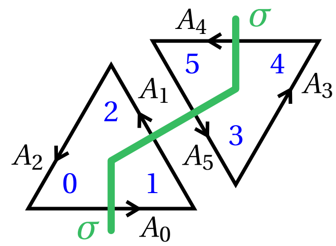

FIGURE 4. This portion of σ passing through D contributes the following edges of X_T to $\mathscr{S}_\sigma$: a_0, b_0, x_1, y_1, a_1 $(= a_5)$, b_1 $(= b_5)$, x_5, y_5, a_4, b_4.

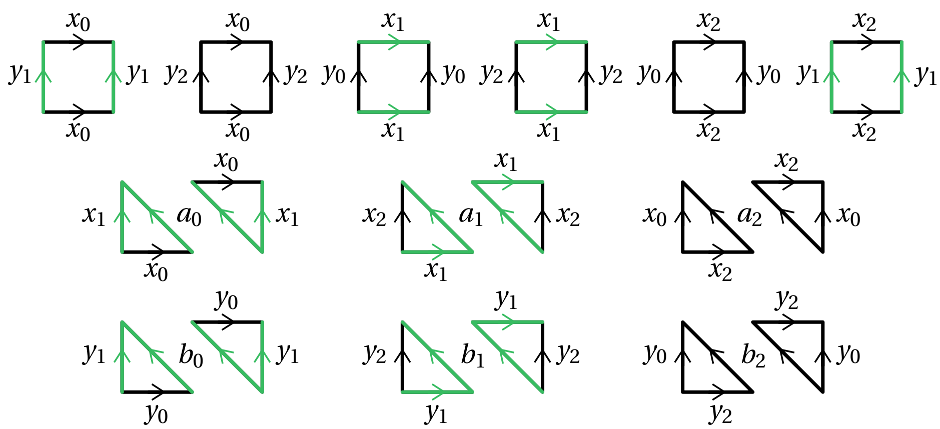

FIGURE 5. Part of the corridor scheme $\mathscr{S}_\sigma$, in green, for the segment σ from Figure 4. The 2–cells from the copy of X with triple $(0,1,2)$ are shown. There is an analogous collection of corridor cells for every triangle that σ meets.

Remark 6.3. Looking closely at the corridor scheme $\mathscr{S}_\sigma$, two additional properties become evident. First, the $\mathscr{S}_\sigma$–edges appearing in a single corridor are all labeled x_i or a_i for various indices i, or they are all labeled y_i or b_i. (That is, $\mathscr{S}_\sigma$ is the disjoint union of two smaller corridor schemes.)

Second, the corridor scheme $\mathscr{S}_\sigma$ is orientable. Referring to Figure 5, we can give positive orientations (relative to the labeling) to the edges labeled x_0, y_0, x_1, y_1, a_0, and b_0, and negative orientations to those labeled a_1 and b_1. This set of choices, or its opposite, can be imposed on any copy of X in X_T that contains edges of $\mathscr{S}_\sigma$. If two copies are adjacent, meaning that they intersect in a subspace Y_i, then the orientations on each copy can be made to agree on Y_i, by reversing the choices on one side if necessary. Now recall that the copies of X containing edges of $\mathscr{S}_\sigma$ all lie along σ. Starting with the copy of X at one end, one may propagate these choices consistently over all of $\mathscr{S}_\sigma$.

Orientability implies that if an $\mathscr{S}_\sigma$–edge label appears more than once along a corridor, then it is oriented the same way across the corridor in each occurrence. Furthermore, if a band type corridor joins two $\mathscr{S}_\sigma$–edges in the boundary which carry the same label, then those labels have opposite orientations relative to the boundary of the diagram.

Standard generators for V_T. Recall that V_T contains many free subgroups $A_i = \langle a_i, b_i\rangle$ which were the peripheral subgroups of the vertex groups V. Some of these subgroups were assigned to edges of T and amalgamated together; these subgroups of V_T will be called the *internal* subgroups, and their generators the *internal generators*. Recall that every A_i that is not an internal subgroup is called a peripheral subgroup of V_T.

The *standard generating set* for V_T will be the union of the generators of the vertex groups (all the generators x_i and y_i) and the generators a_i, b_i of the peripheral subgroups. The internal generators are *not* included.

Definition 6.4. If w is a word let $|w|$ denote the length of w. We define some additional lengths for a word w in the standard generators of V_T:

- $|w|_x$ is the number of occurrences of letters $x_i^{\pm1}$ (for any i) in w
- $|w|_y$ is the number of occurrences of letters $y_i^{\pm1}$ (for any i) in w
- for each i, $|w|_i$ is the number of occurrences of letters $a_i^{\pm1}, b_i^{\pm1}$ in w

Clearly, $|w| = |w|_x + |w|_y + \sum_i |w|_i$.

We use similar notation to count occurrences of $x^{\pm1}$ and $y^{\pm1}$ in words representing elements of $\langle x, y\rangle$.

Definition 6.5. We also define *weighted* word lengths $\|w\|$ similar to the lengths above, where letters are counted with real-valued weights.

Recall that φ has transition matrix $M_\varphi = \begin{pmatrix} |\varphi(x)|_x & |\varphi(y)|_x \\ |\varphi(x)|_y & |\varphi(y)|_y \end{pmatrix}$ with Perron-Frobenius eigenvalue $\lambda > 1$. Let $\vec{d}$ be a left eigenvector for λ (so that $\vec{d}M_\varphi = \lambda\vec{d}$) with positive entries d_1 and d_2.

To define the weighted word lengths, we assign the weight d_1 to the letters x_i and a_i, and we assign d_2 to each y_i and b_i. Thus,

- $\|w\|_x = d_1|w|_x$
- $\|w\|_y = d_2|w|_y$

The weighted length functions are needed for the sake of Lemma 6.6 below. Up to scaling, this is the only choice of weights for which the conclusion of the lemma holds.

Lemma 6.6. *Suppose w is a word in the free group $\langle x, y\rangle$ and v is the reduced word representing $\varphi(w)$. Let $\|\cdot\|$ denote the weighted word length which assigns weight d_1 to $x^{\pm1}$ and weight d_2 to $y^{\pm1}$, where $\vec{d}M_\varphi = \lambda\vec{d}$. Then $\|v\| \leqslant \lambda\|w\|$.*

Proof. This is a simple calculation:

$$\begin{aligned}
\|\varphi(w)\| &= d_1|\varphi(w)|_x + d_2|\varphi(w)|_y \\
&= d_1|\varphi(x)|_x|w|_x + d_1|\varphi(y)|_x|w|_y + d_2|\varphi(x)|_y|w|_x + d_2|\varphi(y)|_y|w|_y \\
&= \lambda d_1|w|_x + \lambda d_2|w|_y \\
&= \lambda\|w\|.
\end{aligned}$$

Now, $\|v\| \leqslant \|\varphi(w)\| = \lambda\|w\|$. ☐

The balancing property. A fundamental property of V_T and its standard generating set, the *balancing property*, is given in the next proposition.

Proposition 6.7. *Suppose w and $z(a_{v_0}, b_{v_0})$ represent the same element of $A_{v_0} \subset V_T$, where w is a word in the standard generators of V_T and $z(a_{v_0}, b_{v_0})$ is reduced. Then for each $i = 1, \ldots, m$ there is an inequality*

$$|z| \leqslant |w|_{v_i} + |w|_{v_0} + |w|_x + |w|_y.$$

Remarks 6.8. (1) The proposition says that an element of a peripheral subgroup cannot be expressed efficiently using generators from *other* peripheral subgroups. For instance, if w contains only generators from $A_{v_1}, \ldots, A_{v_m}$, then $|w|_{v_0} = |w|_x = |w|_y = 0$ and $|w|_{v_i} \geqslant |z|$ for every i, whence $|w| \geqslant m|z|$. An example of such a word w is given in (7.1) below, where z is the initial subword $w(a_{v_0}, b_{v_0})$ and w is the inverse of the remaining expression (see also Figure 7).

(2) There is nothing special about v_0. By re-indexing the peripheral subgroups, there is a corresponding statement that holds for each peripheral subgroup of V_T.

Proof of Proposition 6.7. Let σ be the maximal segment in $\widehat{T}$ with endpoints v_{v_0} and v_{v_i}. The corridor scheme $\mathscr{S}_\sigma$ contains exactly four edges whose labels are peripheral generators of V_T; these generators are a_{v_0}, b_{v_0}, a_{v_i}, and b_{v_i}. Every other standard generator occurring as the label of an edge in $\mathscr{S}_\sigma$ is of the form x_j or y_j.

We may assume without loss of generality that w is reduced. We may further assume that the word zw^{-1} is cyclically reduced, since cancellation of letters between z and w^{-1} does not change the status of the inequality.

Let Δ be a reduced van Kampen diagram over X_T with boundary labeled by zw^{-1}. We may assume that Δ is topologically a disk. Every edge on z is an $\mathscr{S}_\sigma$–edge, and is joined by a σ–corridor to another $\mathscr{S}_\sigma$–edge on the boundary of Δ. If this latter edge is not in z then it contributes 1 to the right hand side of the inequality, since it is labeled by a standard generator.

We claim that no σ–corridor can join two edges of z. Then, since σ–corridors never have $\mathscr{S}_\sigma$–edges in common, there will be at least $|z|$ $\mathscr{S}_\sigma$–edges along w^{-1}, which establishes the result.

If a σ–corridor joins two edges of z, then since corridors do not cross, there is an innermost such corridor. The $\mathscr{S}_\sigma$–edges that it joins must be adjacent edges of z, by the innermost property. Suppose (without loss of generality) the label on one of the edges is a_{v_0}. By Remark 6.3 the label on the other edge must then be $a_{v_0}^{-1}$, but this contradicts the assumption that z is reduced. □

Remark 6.9. Proposition 6.7 remains true if *weighted* word lengths are used throughout:

$$\|z\| \leqslant \|w\|_{v_i} + \|w\|_{v_0} + \|w\|_x + \|w\|_y$$

for each $i = 1, \ldots, m$. Recall that the proof entailed finding corridors joining letters of z to letters of w. The letters occurring at the ends of such a corridor will have the same weights, by Remark 6.3. Therefore, each contribution to the left hand side of the inequality has a matching contribution on the right hand side.

7. Canonical diagrams

In this section we construct a large family of van Kampen diagrams over X_T called canonical diagrams. These will be used in the construction of snowflake diagrams in the next section. We also develop properties of σ–corridors in order to show that canonical diagrams and snowflake diagrams minimize area relative to their boundaries.

Canonical diagrams over X_T. Let $w(x,y)$ be a palindromic word in the free group. In V_T, for each i, one has the relation $w(a_i,b_i) = w(\overline{x}_i,\overline{y}_i)w(x_{i+1},y_{i+1})$ where "$i+1$" is interpreted appropriately. Since w is palindromic, this relation is identical to the relation $w(a_i,b_i) = w(x_i,y_i)^{-1}w(x_{i+1},y_{i+1})$. It bounds a triangular van Kampen diagram over X_T of area $|w|^2$; see Figure 6. Assembling three such diagrams in cyclic

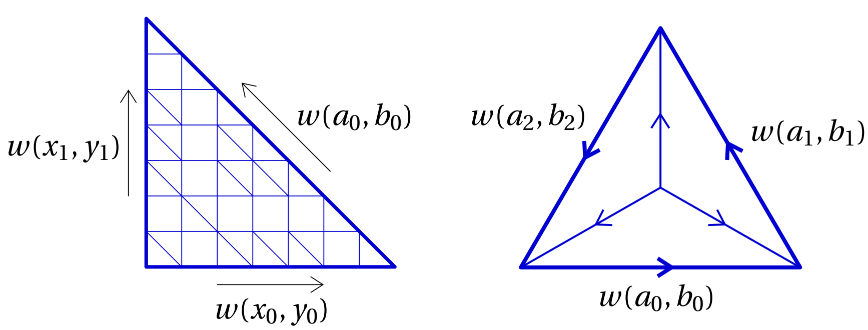

FIGURE 6. Assembling diagrams over X_T

fashion, one obtains, for each vertex group V with index triple (i,j,k), a diagram of area $3|w|^2$ with boundary word $w(a_i,b_i)w(a_j,b_j)w(a_k,b_k)$. Finally, taking diagrams of the latter kind, one for each vertex of T, and assembling them according to T (just like the triangles in D) one obtains a van Kampen diagram over X_T of area $3|T||w|^2$ with boundary word

$$w(a_{v_0}, b_{v_0})\, w(a_{v_1}, b_{v_1}) \cdots w(a_{v_m}, b_{v_m}). \tag{7.1}$$

See Figure 7. In assembling this diagram, we are using the fact that $w(a_i, b_i) = w(a_j, b_j)^{-1}$

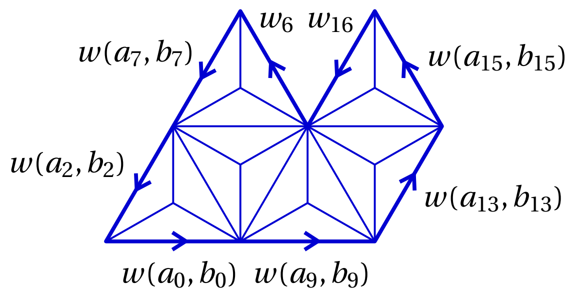

FIGURE 7. The canonical diagram of w with boundary word (7.1), where $w_i = w(a_i, b_i)$.

whenever A_i and A_j were amalgamated, which also relies on the palindromic property of w.

This van Kampen diagram will be called the *canonical diagram of w*, and it is defined for every palindromic word. If w is reduced then the canonical diagram is also reduced.

Our remaining objective in this section is to show that canonical diagrams (and their "doubled" variants) minimize area. To this end, we need to establish some additional properties of σ–corridors.

Lemma 7.2. *Let Δ be a van Kampen diagram over X_T and suppose that C is a σ–corridor and C' is a σ'–corridor in Δ. If C and C' have intersection of positive area, then $C \cap C'$ contains one of the following:*

(1) *a quadrilateral relator*
(2) *two neighboring triangular relators with a common $\mathcal{S}_\sigma \cap \mathcal{S}_{\sigma'}$–edge labeled a_i or b_i*
(3) *a triangular relator with a side labeled by a_i or b_i, which is a $\mathcal{S}_\sigma \cap \mathcal{S}_{\sigma'}$–edge in the boundary of Δ.*

We will refer to the quadrilateral relator in (1) and the union of the two neighboring triangular relators in (2) as *crossing squares* for $C \cap C'$. The triangular relators in (3) will be called *crossing triangles.* Note that crossing squares have area 2.

We shall see that canonical diagrams are completely filled by crossing squares and triangles for various pairs of corridors, and that these crossing regions must be present in any diagram with the same boundary.

Proof. If case (1) does not occur, then $C \cap C'$ contains a triangular relator. Note that $\mathscr{S}_\sigma$ has the property that a corridor cell is triangular if and only if one of its boundary $\mathscr{S}_\sigma$–edges is labeled a_i or b_i; see Figure 5. The same is true of $\mathscr{S}_{\sigma'}$. Thus, the side of the triangular relator labeled a_i or b_i is an $\mathscr{S}_\sigma \cap \mathscr{S}_{\sigma'}$–edge. If this edge is in the boundary of Δ then case (3) occurs. Otherwise, the neighboring 2–cell across that edge is a corridor cell for both $\mathscr{S}_\sigma$ and $\mathscr{S}_{\sigma'}$ and case (2) occurs. □

Lemma 7.3. *If σ and σ' are maximal segments in $\widehat{T}$ with no edges in common, then $\mathscr{S}_\sigma \cap \mathscr{S}_{\sigma'}$ is empty and no 2–cell is a corridor cell for both $\mathscr{S}_\sigma$ and $\mathscr{S}_{\sigma'}$.*

Proof. Because $\widehat{T}$ has valence at most 3, σ and σ' must actually be disjoint. Thus, they never pass through the same triangle of D, which shows that $\mathscr{S}_\sigma \cap \mathscr{S}_{\sigma'}$ is empty. It follows immediately that no triangular 2–cell can be a corridor cell for both $\mathscr{S}_\sigma$ and $\mathscr{S}_{\sigma'}$. The same is true for quadrilateral 2–cells because each such 2–cell has all of its side labels coming from a single triangle in D. □

Definition 7.4. For each edge e in $\widehat{T}$ choose maximal segments σ_e, σ'_e in $\widehat{T}$ whose intersection is exactly e. If e is an interior edge, we also require that the endpoints of σ_e and σ'_e are linked in the boundary of the $(|T|+2)$–gon; see Figure 8.

If C and C' are σ_e– and σ'_e–corridors respectively, a crossing square or crossing triangle for $C \cap C'$ will be called an *e–crossing square* or an *e–crossing triangle* (or *e–crossing region* in either case). Figure 8 shows the location of e–crossing regions in a canonical diagram Δ.

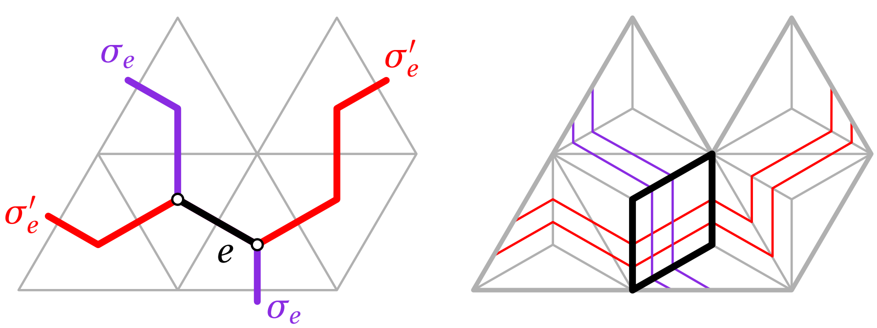

FIGURE 8. An edge e in $\widehat{T}$ and the segments σ_e, σ'_e. The e–crossing squares in Δ fill a quadrilateral region of area $2|w|^2$ as shown. If e were a peripheral edge, the e–crossing squares and triangles would fill a triangular region next to the boundary of Δ, of area $|w|^2$.

Lemma 7.5. *If e and f are distinct edges of $\widehat{T}$ then e–crossing regions and f–crossing regions have no 2–cells in common.*

Proof. It suffices to show that no 2–cell is a corridor cell simultaneously for all four corridor schemes $\mathscr{S}_{\sigma_e}, \mathscr{S}_{\sigma'_e}, \mathscr{S}_{\sigma_f}, \mathscr{S}_{\sigma'_f}$.

If e and f are separated by a third edge g, then at least one of σ_e, σ'_e and one of σ_f, σ'_f does not contain g. Hence, these two segments are disjoint and Lemma 7.3 applies.

Otherwise, e and f have a common vertex v. Consider the triangle in D centered at v. There are three ways that a segment σ can pass through the triangle, and the four segments must use all three of these. Of the eighteen 2–cells associated with this triangle, one can verify easily that each 2–cell is a corridor cell for exactly two of the three possible schemes. It follows that no 2–cell associated with this triangle can be a corridor cell for all four corridor schemes. No other triangle in D can meet all four segments, so the same is true for the other 2–cells of X_T. □

Definition 7.6. A van Kampen diagram over X is called *least-area* if it has the smallest area of all van Kampen diagrams over X having the same boundary word.

Proposition 7.7. *Let $w(x, y)$ be a reduced palindromic word. The canonical diagram of w is least-area.*

Proof. Let Δ be the canonical diagram of w and let Δ' be an arbitrary van Kampen diagram with the same boundary word. Let $\ell = |w|$.

First we claim that for any choice of σ, say with endpoints v_i and v_j, there are exactly ℓ band type σ–corridors in Δ', each joining a letter in $w(a_i, b_i)$ with a letter in $w(a_j, b_j)$. Certainly, these two subwords of (7.1) contain the only occurrences of $\mathscr{S}_\sigma$–edges in the boundary of Δ', so the number of such corridors can only be ℓ. Also, no such corridor can join two letters of the same subword $w(a_i, b_i)$ or $w(a_j, b_j)$; using Remark 6.3 as in the proof of Proposition 6.7 one finds that w must then fail to be reduced.

Now let us identify crossing squares and triangles in Δ'. If e is an internal edge of $\widehat{T}$ then every σ_e–corridor crosses every σ'_e–corridor, by the linking requirement on σ_e and σ'_e (cf. Figure 8). Thus there are exactly ℓ^2 e-crossing squares for such e.

If e is a peripheral edge of $\widehat{T}$ incident to v_i, say, then some pairs of σ_e– and σ'_e–corridors cross and some do not. There is one corridor of each type (σ_e or σ'_e) emanating from each letter in the subword $w(a_i, b_i)$ on the boundary of Δ', and this accounts for all σ_e– and σ'_e–corridors. For each letter in $w(a_i, b_i)$ the two corridors emanating there will contain an e–crossing triangle; there are ℓ such corridor pairs. Of the remaining corridor pairs, half of them definitely cross (because their endpoints on the boundary are linked), yielding e–crossing squares. There are at least $\ell(\ell-1)/2$ of these. In total we have identified e-crossing regions of total area $2\ell^2$ when e is an internal edge, and of total area ℓ^2 when e is a peripheral edge. By Lemma 7.5 we conclude that $\mathrm{Area}(\Delta') \geqslant \mathrm{Area}(\Delta)$. □

Doubled canonical diagrams. For any palindromic word w, take the canonical diagrams of w and of w^{-1} and join them along their boundary subwords labeled $w(a_{v_0}, b_{v_0})$ and $w(a_{v_0}, b_{v_0})^{-1}$ to form a new diagram, called the *doubled canonical diagram of* w. Its boundary word is given by

$$w(a_{v_1}, b_{v_1}) \cdots w(a_{v_m}, b_{v_m}) w(a_{v_1}, b_{v_1})^{-1} \cdots w(a_{v_m}, b_{v_m})^{-1}. \tag{7.8}$$

If w is reduced, then so is its doubled canonical diagram.

Proposition 7.9. *Let $w(x, y)$ be a reduced palindromic word. The doubled canonical diagram of w is least-area.*

Proof. Let $D\Delta$ be the doubled canonical diagram of w and let Δ' be an arbitrary van Kampen diagram with the same boundary word.

Let σ be a maximal segment in $\widehat{T}$ that does not contain v_{v_0}. The $\mathscr{S}_\sigma$–edges on the boundary of Δ' comprise four subwords $w(a_i, b_i)$, $w(a_j, b_j)$, $w(a_i, b_i)^{-1}$, $w(a_j, b_j)^{-1}$, arranged in this cyclic ordering. We claim that the σ–corridors joining letters in these subwords must in fact join all the letters of $w(a_i, b_i)$ to those of $w(a_j, b_j)$, and similarly with $w(a_i, b_i)^{-1}$ and $w(a_j, b_j)^{-1}$.

First, as before, no σ–corridor joins two letters of the same subword, because w is reduced. Next, no corridor runs between $w(a_i, b_i)$ and $w(a_i, b_i)^{-1}$ (or $w(a_j, b_j)$ and $w(a_j, b_j)^{-1}$) because then there is no room for the remaining corridors to be disjoint. Thus, the σ–corridors must be arranged as in Figure 9. It is evident that if any corridor

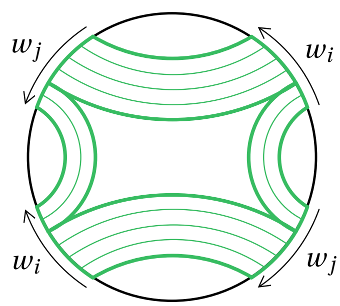

FIGURE 9. A configuration of σ–corridors.

joins $w(a_i, b_i)$ to $w(a_j, b_j)^{-1}$, then there is such a corridor joining the *first* letter of $w(a_i, b_i)$ to the *last* letter of $w(a_j, b_j)^{-1}$. If the first letter is, say, a_i (in the orientation of the boundary of Δ'), then the corridor joins it to a_j^{-1}. On the other hand, in the *canonical* diagram of w, there is a corridor joining a_i in the boundary to the last letter of $w(a_j, b_j)$, which is a_j (because w is palindromic). The existence of both corridors, even in different diagrams, contradicts the orientability of $\mathscr{S}_\sigma$ established in Remark 6.3. Therefore all corridors join $w(a_i, b_i)$ to $w(a_j, b_j)$ or $w(a_i, b_i)^{-1}$ to $w(a_j, b_j)^{-1}$, as claimed.

If σ is a maximal segment with endpoints v_{v_0} and v_i then the only $\mathscr{S}_\sigma$–edges on the boundary are the two subwords $w(a_i, b_i)$ and $w(a_i, b_i)^{-1}$, and all σ–corridors run between them.

The rest of the proof now proceeds without difficulty just like Proposition 7.7. We have complete knowledge of which pairs of edges in the boundary of Δ' are joined by corridors of various kinds, and these pairings are in agreement with those of $D\Delta$. One easily finds the requisite numbers of e–crossing regions for each e and concludes that $\mathrm{Area}(\Delta') \geqslant \mathrm{Area}(D\Delta)$. □

8. SNOWFLAKE DIAGRAMS

Snowflake diagrams, defined below, will be used to establish the lower bound in the proof of Theorem 10.14.

The 2–complex $Y_{T,n}$. Recall that $S_{T,n}$ was defined via the relative presentation (3.1). Starting with X_T, adjoin 1–cells and 2–cells according to this relative presentation to obtain the 2–complex $Y_{T,n}$ with fundamental group $S_{T,n}$. There will be m new 1–cells labeled $r_1, \ldots, r_m$ and $2m$ new 2–cells with boundary words given by the relators of (3.1).

r_i–corridors. For each $i = 1, \ldots, m$ there is an orientable corridor scheme consisting of the single edge labeled r_i. It has two corridor cells which we think of as being rectangular, with sides labeled by the words

$$\varphi^n(a_{v_i})^{-1},\ r_i,\ a_{v_0},\ r_i^{-1} \tag{8.1}$$

and

$$\varphi^n(b_{v_i})^{-1},\ r_i,\ b_{v_0},\ r_i^{-1}. \tag{8.2}$$

The sides labeled by a_{v_0} or b_{v_0} will be called the *short sides* and the sides labeled by $\varphi^n(a_{v_i})^{-1}$ or $\varphi^n(b_{v_i})^{-1}$ the *long sides* of the corridor cells. The corridors for this scheme are called *r_i–corridors*.

Note that in any r_i–corridor in a *reduced* van Kampen diagram, the short sides of the corridor cells join up to form a single arc in the boundary of the corridor, labeled by a reduced word w in the generators a_{v_0}, b_{v_0}. If this word happens to be monotone, then the long sides of the corridor cells also assemble to form a monotone (and reduced) word $\varphi^n(w)(a_{v_i}, b_{v_i})$.

Snowflake diagrams. Let $w(x, y)$ be a *monotone* palindromic word. We will define van Kampen diagrams over $Y_{T,n}$ based on w and an integer d (the *depth*) denoted $\Delta(w, d)$. To begin, we define $\Delta(w, 0)$ to be the doubled canonical diagram of w.

Next, to define $\Delta(w, d)$ for $d > 0$, start with the diagram $\Delta(\varphi^n(w), d-1)$ (noting that $\varphi^n(w)$ is also monotone and palindromic, by our assumptions on φ). Its boundary word will have subwords of the form $\varphi^n(w)(a_{v_i}, b_{v_i})^{\pm 1}$ for each $i = 1, \ldots, m$. Alongside each subword $\varphi^n(w)(a_{v_i}, b_{v_i})^{\varepsilon}$ adjoin a rectangular strip made of $|w|$ 2–cells whose four sides are labeled by the words

$$\varphi^n(w)(a_{v_i}, b_{v_i})^{-\varepsilon},\ r_i,\ w(a_{v_0}, b_{v_0})^{\varepsilon},\ r_i^{-1}.$$

Then, adjoin a copy of the canonical diagram of $w^{-\varepsilon}$, which contains a side labeled $w(a_{v_0}, b_{v_0})^{-\varepsilon}$.

Doing this for each subword as described, one obtains $\Delta(w,d)$. See Figure 10. Note that the boundary of $\Delta(w,d)$ contains many copies of the subwords $w(a_{v_i}, b_{v_i})^{\pm 1}$ for each $i = 1,\ldots,m$. In fact, it is easy to verify by induction on d that the boundary word is made entirely of copies of these words, together with occurrences of the letters $r_i^{\pm 1}$.

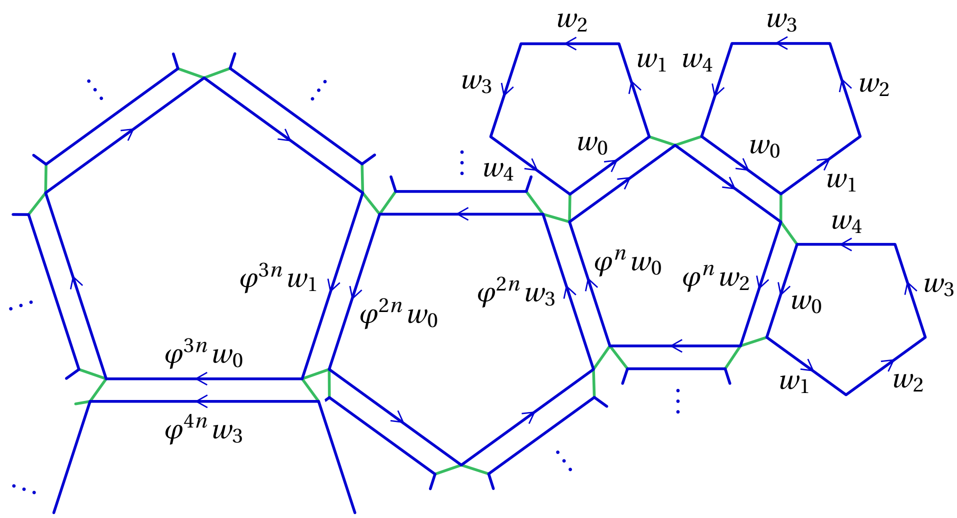

FIGURE 10. Part of the snowflake diagram $\Delta(w,d)$, with $m = 4$ (and $d \geqslant 4$). The word $\varphi^{in}(w)(a_{v_j}, b_{v_j})$ is abbreviated as $\varphi^{in} w_j$. The pentagonal regions are canonical diagrams and the strips between them are r_j–corridors for various j.

Note that $\Delta(w,d)$ contains a sub-diagram $\Delta(\varphi^{in}(w), d-i)$ for each i between 0 and d. In particular, it contains a copy of $\Delta(\varphi^{dn}(w), 0)$, which is a doubled canonical diagram of area $6|T||\varphi^{dn}(w)|^2$.

Proposition 8.3. *Let $w(x,y)$ be a monotone palindromic word. For each $d \geqslant 0$ the diagram $\Delta(w,d)$ is least-area.*

Proof. First note that $\Delta(w,d)$ is reduced, since it is made of reduced sub-diagrams, separated by reduced r_i–corridors, which have no 2–cells in common with the sub-diagrams. (The sub-diagrams being reduced depends on the monotonicity of w, which implies that the words $\varphi^{in}(w)$ are reduced.)

Now suppose that $d = 0$ and let Δ' be any reduced diagram over $Y_{T,n}$ with the same boundary as $\Delta(w,0)$. We claim that there are no r_i–corridors for any i. If there were, they would be of annulus type, and the short side of the corridor would be labeled by a cyclically reduced word $v(a_{v_0}, b_{v_0})$ representing the trivial element. Since A_{v_0} is free on a_{v_0}, b_{v_0}, no such corridors can exist. Therefore Δ' is actually a diagram over X_T, and Proposition 7.9 says that $\mathrm{Area}(\Delta') \geqslant \mathrm{Area}(\Delta(w,0))$.

Proceeding by induction on d, suppose that $d \geqslant 1$ and Δ' is a reduced diagram over $Y_{T,n}$ of smallest area, with the same boundary as $\Delta(w,d)$. As before, there can be no r_i–corridors of annulus type. There will be band type r_i–corridors joining occurrences of $r_i^{\pm 1}$ on the boundary. Note that r_i– and r_j–corridors cannot cross for any i, j (no 2–cell is a corridor cell for both corridor schemes). Hence the r_i–edges on the boundary must be paired by corridors in the same way as in $\Delta(w,d)$.

Consider an outermost r_i–corridor. Its complement in Δ' is two sub-diagrams, one of which is a diagram over X_T with boundary word

$$\widehat{w}(a_{v_0}, b_{v_0})^{\varepsilon}\, w(a_{v_1}, b_{v_1})^{\varepsilon} \cdots w(a_{v_m}, b_{v_m})^{\varepsilon}$$

for some word $\widehat{w}$ and some $\varepsilon = \pm 1$. Here, $\widehat{w}(a_{v_0}, b_{v_0})$ is the word along the short side of the corridor. Recall that in V_T, the word $w(a_{v_1}, b_{v_1})^{\varepsilon} \cdots w(a_{v_m}, b_{v_m})^{\varepsilon}$ represents the element $w(a_{v_0}, b_{v_0})^{-\varepsilon}$, which moreover is in the free subgroup A_{v_0}. Since $\widehat{w}(a_{v_0}, b_{v_0})$ is reduced, it must equal $w(a_{v_0}, b_{v_0})$. It follows that the corridor, considered as a sub-diagram, is identical to the corresponding corridor in $\Delta(w,d)$. Also, the part of Δ' on the short side of the corridor is a diagram over X_T with the same boundary as the canonical diagram of w^{ε}. By Proposition 7.7 its area agrees with that of the canonical diagram.

Taking all the outermost r–corridors and the sub-diagrams that they separate from the central region in Δ', we have found that these have total area equal to the corresponding regions in $\Delta(w,d)$. Moreover, if we delete these regions, the resulting boundary word is the boundary word of the corresponding sub-diagram $\Delta(\varphi^n(w), d-1)$ of $\Delta(w,d)$. By induction, the central portion of Δ' has area equal to that of $\Delta(\varphi^n(w), d-1)$ and we are done. □

9. Folded corridors and subgroup distortion

The main result of this section is Proposition 9.7 (and its variant Corollary 9.14) which bounds the distortion of the edge group A_{v_0} in $S_{T,n}$. After discussing some preliminaries, we proceed to study *folded corridors*, culminating in Lemma 9.5. This lemma plays an important role in the proof of Proposition 9.7, which occupies the rest of the section.

Define the *standard generating set* for $S_{T,n}$ to be the standard generating set for V_T together with the generators $r_1, \dots, r_m$. Recall that the former generators include all generators x_i and y_i, and the peripheral generators a_{v_i}, b_{v_i} $(i = 0, \dots, m)$.

For $g \in A_{v_i}$ let $|g|_{A_{v_i}}$ denote the length of the *reduced* word in the basis a_{v_i}, b_{v_i} representing g. Similarly, let $\|g\|_{A_{v_i}}$ be the *weighted* word length of the reduced representative (cf. Definition 6.5). Recall that the letters $a_i^{\pm 1}$ and $x_i^{\pm 1}$ have weight $1+\sqrt{2}$ and the letters $b_i^{\pm 1}$ and $y_i^{\pm 1}$ have weight 1. Let the letters $r_i^{\pm 1}$ also be given weight 1.

Now we assign lengths to the edges of $Y_{T,n}$, and correspondingly to the edges in any van Kampen diagram over $Y_{T,n}$, as follows. Edges labeled by x_i or a_i are given length $1+\sqrt{2}$, and all other edges (those labeled y_i, b_i, or r_i) are given length 1. In this section, lengths of paths in a van Kampen diagram will always be meant with respect to these edge lengths.

With this convention, the length of a path in the 1–skeleton will agree with the weighted length of the word labeling it.

Folded corridors. Recall that each r_i–corridor cell has a short side and a long side. In any r_i–corridor, the embedded open annulus or open band inside it separates all of the short sides of the corridor cells from the long sides.

The boundary of an r_i–corridor is a 1–complex containing zero or two r_i–edges. The *partial boundary* is defined to be the boundary with the interiors of the r_i–edges removed.

If C is an r_i–corridor in a reduced diagram, then the short sides of its cells join to form a component of the partial boundary which is labeled by a reduced word in the generators a_{v_0}, b_{v_0}. If C were of annulus type, then we would have a cyclically reduced word in the free group A_{v_0} representing the trivial element. Hence, C must be of band type.

Following [BG10], a band type r_i–corridor is called *folded* if it is reduced and every component of its partial boundary is labeled by a reduced word in the generators of $S_{T,n}$. We have noted already that the short sides of corridor cells form a single such component, which we now call the *bottom* of the corridor. Any other component is labeled by a reduced word in the generators a_{v_i}, b_{v_i}. Again, such a component cannot be a loop (since A_{v_i} is free) and hence there is only one other component, which we call the *top* of the corridor. If $w(a_{v_0}, b_{v_0})$ is the reduced word along the bottom, then the top is labeled by the reduced word in a_{v_i}, b_{v_i} representing $\varphi^n(w)(a_{v_i}, b_{v_i})$.

Remark 9.1. Given any reduced word $w(a_{v_0}, b_{v_0})$, one can build a folded corridor with bottom labeled by w. Start by joining corridor cells end to end along r_i–edges to form a corridor with bottom side labeled by w. Then, the long sides of the corridor cells form an arc labeled by a possibly *unreduced* word representing $\varphi^n(w)(a_{v_i}, b_{v_i})$. By successively folding together adjacent pairs of edges along the top (with matching labels), one eventually obtains a folded corridor. Each folding operation corresponds to a free reduction in the word labeling the top side of the corridor. The final word along the top of the folded corridor is uniquely determined (being the reduced form of $\varphi^n(w)(a_{v_i}, b_{v_i})$) but the internal structure will depend on the particular sequence of folds chosen.

Let C be a folded r_i–corridor. Define $S \subset C$ to be the smallest subcomplex containing all the r_i–edges, and all the open 1–cells which lie in the interior of C (informally, the *seams* in C). Note that S contains exactly those edges of C that are not in the top or bottom.

Lemma 9.2. *Let S_0 be a connected component of S. Then*

(1) *S_0 is a tree;*
(2) *S_0 contains exactly one vertex in the top of C;*
(3) *every valence-one vertex in S_0 lies in the top or bottom of C.*

Conclusions (2) and (3) imply that S_0 contains at least one vertex in the bottom of C.

Proof. First note that every 2–cell of C meets the bottom in exactly one edge. Conclusion (1) follows immediately since a loop in S_0 would separate a 2–cell from the bottom. For the same reason, S_0 cannot contain two or more vertices of the top.

A second observation is that the boundary of every 2–cell is labeled by a cyclically reduced word (namely, (8.1) or (8.2)). Hence no two adjacent edges of the same cell can be folded together. Therefore S_0 cannot have a valence-one vertex in the interior of C, whence (3).

It remains to show that S_0 contains a vertex of the top. If not, then it is separated from the top by a 2–cell which then must meet the bottom in a disconnected set, contradicting the initial observation above. □

Let p be a vertex in the top of C and q a vertex in the bottom. We say that *p is above q* if both vertices are in the same connected component of S. We have the following "bounded cancellation" lemma, which is a restatement of Lemma 1.2.4 of [BG10]:

Lemma 9.3. *There is a constant $K_0 = K_0(\varphi^n)$ such that if p is a vertex in the top of a folded corridor and q_1, q_2 are vertices that are both below p, then the sub-segment $[q_1, q_2]$ of the bottom has at most K_0 edges.* □

For vertices p in the top and q in the bottom, we say that *p is nearly above q* if there is a vertex p' above q such that p and p' are in the boundary of a common 2–cell. It is clear that every vertex in the top is nearly above some vertex on the bottom.

Lemma 9.4. *There is a constant $K_1 = K_1(\varphi^n)$ such that if C is a folded corridor and p is nearly above q in C, then there is a path in the 1–skeleton of C from p to q, containing no bottom edges, of length at most K_1.*

Proof. Let L be the maximum of the boundary lengths of the two r_i–corridor cells. Let p' be a vertex above q such that p and p' are in the boundary of a common 2–cell e^2, and let S_0 be the component of S containing p' and q. Let q' be the unique bottom vertex of S_0 which is in the boundary of e^2. There are paths $[p, p']$ in the top and $[p', q']$ in S_0. These have total length at most L, since they are in the boundary of e^2.

Next, any two adjacent bottom vertices in S_0 are in the boundary of a 2–cell, and hence are joined by a path in S_0 of length at most L. It follows that q' and q are joined by a path $[q', q]$ in S_0 of length at most K_0L, with K_0 given by Lemma 9.3. Now the path $[p, p']\cdot[p', q']\cdot[q', q]$ has length at most $K_1 = (K_0 + 1)L$. □

Lemma 9.5. *There is a constant $K_2 = K_2(\varphi^n)$ such that if C is a folded corridor and $[p_1, p_2]$, $[q_1, q_2]$ are sub-segments of the top and bottom, respectively, with p_j nearly above q_j for $j = 1,2$, and $u(a_{v_i}, b_{v_i})$ is the word labeling $[p_1, p_2]$, then*

$$\|\varphi^{-n}(u)\|_{A_{v_i}} - K_2 \leqslant |[q_1, q_2]| \leqslant \|\varphi^{-n}(u)\|_{A_{v_i}} + K_2.$$

The essential point is that the length of $[q_1, q_2]$ is determined, up to an *additive* error, by the word u labeling $[p_1, p_2]$. (It is certainly not determined by the length of $[p_1, p_2]$ alone.)

Proof. Let $w(a_{v_0}, b_{v_0})$ be the reduced word labeling $[q_1, q_2]$. Let $\widehat{S}$ be the union of S and the top of C. For $j = 1,2$ let $[q_j, p_j]$ be a shortest path in $\widehat{S}$ from q_j to p_j. Its first edge is an r_i–edge labeled r_i^{-1}, and so the label on $[q_j, p_j]$ has the form $r_i^{-1}\cdot v_j(a_{v_i}, b_{v_i})$ for some reduced word v_j. This word has weighted length less than K_1 by Lemma 9.4.

The four segments form a relation in $S_{T,n}$, namely

$$\begin{aligned} w(a_{v_0}, b_{v_0}) &= r_i^{-1} v_1(a_{v_i}, b_{v_i}) u(a_{v_i}, b_{v_i}) v_2(a_{v_i}, b_{v_i})^{-1} r_i \\ &= \varphi^{-n}(v_1)(a_{v_0}, b_{v_0})\, \varphi^{-n}(u)(a_{v_0}, b_{v_0})\, \varphi^{-n}(v_2)(a_{v_0}, b_{v_0}). \end{aligned} \tag{9.6}$$

This is an equality of elements of the free subgroup A_{v_0}. Now define

$$K_2 = 2\max\left\{\|\varphi^{-n}(v)\|_{A_{v_0}} \mid v \text{ is a word in } a_{v_0}, b_{v_0} \text{ of weighted length } < K_1\right\}.$$

Observe that the left hand side of (9.6) has reduced weighted length $|[q_1, q_2]|$ and the right hand side has reduced weighted length within K_2 of $\|\varphi^{-n}(u)(a_{v_0}, b_{v_0})\|_{A_{v_0}}$, which is equal to $\|\varphi^{-n}(u)\|_{A_{v_i}}$. □

We turn now to the main result of this section, the bound on edge group distortion. Recall from the Introduction that the proof is an inductive proof based on Britton's Lemma. It falls into two cases, requiring very different methods.

In the first case, the proof is based on a method from [BB00]. It is the balancing property of V_T that allows us to carry out this argument. The crucial moment occurs in (9.8) and the choice of index j', and the subsequent reasoning.

The second case makes use of folded corridors and Lemma 9.5. The overvall induction argument is based on the nested structure of r_j–corridors in a van Kampen diagram. If these corridors are always oriented in the correct direction, then the argument based on the balancing property would suffice. If there exists a backwards-facing r_j–corridor, then the inductive process will inevitably land in Case II.

The backwards-facing corridor may then introduce geometric effects that adversely affect the inductive calculation. When this occurs, we prove that there will be correctly oriented corridors just behind the first one, and perfectly matching segments along these corridors, along which any metric distortion introduced by the first corridor is exactly undone. This occurs in (9.9), using Lemma 9.5. This argument also depends crucially on σ–corridors.

Proposition 9.7 (Edge group distortion)**.** *Given T and n there is a constant K_3 such that if w is a word in the standard generators of $S_{T,n}$ representing an element $g \in A_{v_0}$ then*

$$\|g\|_{A_{v_0}} \leqslant K_3 \|w\|^{\alpha}$$

where $\alpha = n\log_m(\lambda)$.

Proof. Let $K_3 = \max\{1, 3K_2/2, \Lambda^n\}$ where K_2 is given by Lemma 9.5 and Λ is the maximum stretch factor for φ^{-1} with respect to the weighted word length $\|\cdot\|$.

The universal cover $\widetilde{Y}_{T,n}$ is the total space of a tree of spaces, with vertex spaces equal to copies of the universal cover $\widetilde{X}_T$. Every 1–cell of $\widetilde{Y}_{T,n}$ is either contained in a vertex space, or is labeled $r_j^{\pm 1}$ for some j and has endpoints in two neighboring vertex spaces.

We argue by induction on the number of occurrences of $r_j^{\pm 1}$ in w (for all j). We may assume that w is a shortest word in the generators of $S_{T,n}$ representing g. This word describes a labeled geodesic in the 1–skeleton of $\widetilde{Y}_{T,n}$ with endpoints in the same vertex space. Using the tree structure, one finds a decomposition of w as $w_1 \cdots w_k$ where each w_i satisfies one of the following:

(1) $w_i = r_j u_i r_j^{-1}$ for some j and w_i represents an element of A_{v_j}. Let $v_i(a_{v_j}, b_{v_j})$ be the reduced word representing w_i in this case.
(2) $w_i = r_j^{-1} u_i r_j$ for some j and w_i represents an element of A_{v_0}. Let $v_i(a_{v_0}, b_{v_0})$ be the reduced word representing w_i in this case.
(3) w_i is a word in the generators a_{v_j}, b_{v_j} for some j.
(4) w_i is a word in the generators $\{x_j, y_j\}$ (allowing all j).

Let $v_i = w_i$ in cases (3) and (4), and define $v = v_1 \cdots v_k$. The proof now splits into cases, based on the structure of the decomposition $w = w_1 \cdots w_k$.

- Case IA: all subwords w_i are of types (3) or (4). This is simply the base case of the induction.
- Case IB: either there is a subword w_i of type (1), or there are two or more subwords of type (2).
- Case II: exactly one subword w_i is of type (2) and all others are of types (3) or (4).

Proof in Cases IA and IB. Define the sets

$$\begin{aligned} I_j &= \{i \mid v_i \text{ is a word in the generators } a_{v_j}, b_{v_j}\}, \\ I_{xy} &= \{i \mid v_i \text{ is a word in the generators } \{x_j, y_j\}\}. \end{aligned}$$

Note that $\|v\|_{v_j} = \sum_{i\in I_j}\|v_i\|$ and $\|v\|_x + \|v\|_y = \sum_{i\in I_{xy}}\|v_i\|$. We begin by establishing two claims.

Claim 1: if $i \in I_j$ and $j \neq 0$ then $\|v_i\| \leqslant K_3(m\|w_i\|)^\alpha$. If w_i satisfies (3) then the claim is trivial: $\|v_i\| = \|w_i\| \leqslant K_3(m\|w_i\|)^\alpha$ since $K_3, m, \alpha \geqslant 1$. Otherwise, w_i satisfies (1) and $w_i = r_j u_i r_j^{-1}$ where u_i represents an element of A_{v_0}. Let z_i be the reduced word in a_{v_0}, b_{v_0} equal to u_i in A_{v_0}, and note that $\varphi^n(z_i)(a_{v_j}, b_{v_j}) = v_i$ in A_{v_j}. Since v_i is reduced we have

$$\|v_i\| \leqslant \|\varphi^n(z_i)(a_{v_j}, b_{v_j})\| = \|\varphi^n(z_i)\| \leqslant \lambda^n\|z_i\| = \lambda^n\|u_i\|_{A_{v_0}} = m^\alpha\|u_i\|_{A_{v_0}}.$$

Here, the second inequality follows from Lemma 6.6. The final quantity is at most

$$K_3 m^\alpha(\|w_i\| - 2)^\alpha \leqslant K_3(m\|w_i\|)^\alpha$$

by the induction hypothesis.

Claim 2: if $i \in I_0 \cup I_{xy}$ then $\|v_i\| \leqslant K_3\|w_i\|^\alpha$. As in Claim 1, if w_i satisfies (3) or (4) then Claim 2 is trivially true. The remaining case is when w_i satisfies (2). Now the claim is an instance of the induction hypothesis: since v_i is reduced we have $\|v_i\| = \|w_i\|_{A_{v_0}} \leqslant K_3\|w_i\|^\alpha$. Note that the induction hypothesis applies precisely because we are not in Case II, and w_i contains fewer occurrences of the letters $r_j^{\pm 1}$ than w.

Among the indices $1, \ldots, m$, choose j' to minimize the sum $\sum_{i\in I_{j'}}\|w_i\|$. Thus we have

$$m \sum_{i\in I_{j'}} \|w_i\| \leqslant \sum_{i\in I_1\cup\cdots\cup I_m} \|w_i\|. \tag{9.8}$$

Observe that $v = v_1 \cdots v_k$ is a word in the standard generators of V_T representing the element g. Applying Proposition 6.7 (and Remark 6.9) to this word yields the inequality

$$\begin{aligned} \|g\|_{A_{v_0}} &\leqslant \|v\|_{v_0} + \|v\|_x + \|v\|_y + \|v\|_{v_{j'}} \\ &= \sum_{i\in I_0}\|v_i\| + \sum_{i\in I_{xy}}\|v_i\| + \sum_{i\in I_{j'}}\|v_i\|. \end{aligned}$$

Then we have

$$\begin{aligned}\|g\|_{A_{v_0}} &\leqslant \sum_{i\in I_0\cup I_{xy}} K_3\|w_i\|^\alpha + \sum_{i\in I_{j'}} K_3\,(m\|w_i\|)^\alpha \\ &\leqslant \sum_{i\in I_0\cup I_{xy}} K_3\|w_i\|^\alpha + K_3\left(\sum_{i\in I_{j'}} m\|w_i\|\right)^\alpha\end{aligned}$$

by Claims 1 and 2, and

$$\begin{aligned}\|g\|_{A_{v_0}} &\leqslant \sum_{i\in I_0\cup I_{xy}} K_3\|w_i\|^\alpha + K_3\left(\sum_{i\in(I_1\cup\cdots\cup I_m)} \|w_i\|\right)^\alpha \\ &\leqslant K_3\left(\sum_i \|w_i\|\right)^\alpha \\ &= K_3\|w\|^\alpha\end{aligned}$$

By (9.8).

Proof in Case II. In this case we write w as $w_L\hat{w}w_R$ where $\hat{w}$ is the subword of type (2) and w_L, w_R are products of subwords of types (3) and (4). Thus $\hat{w} = r_{j'}^{-1}ur_{j'}$ for some index $j' \neq 0$, where u represents an element of $A_{v_{j'}}$. Let $\hat{v}$ be the reduced word in a_{v_0}, b_{v_0} representing $\hat{w}$, and let $\hat{u}$ be the reduced word in $a_{v_{j'}}, b_{v_{j'}}$ representing u.

We proceed by decomposing the word u using the tree of spaces structure of $\widetilde{Y}_{T,n}$. First, choose an index $j'' \neq j'$ from the set $\{1,\ldots,m\}$. Then write $u = \hat{w}_1\cdots\hat{w}_{k'}$ where each $\hat{w}_i$ satisfies one of the following:

(5) $\hat{w}_i = r_{j'}\hat{u}_ir_{j'}^{-1}$ or $\hat{w}_i = r_{j''}\hat{u}_ir_{j''}^{-1}$. Let $\hat{v}_i$ be the reduced word in $a_{v_{j'}}, b_{v_{j'}}$ or $a_{v_{j''}}, b_{v_{j''}}$ representing $\hat{w}_i$.

(6) $\hat{w}_i = r_j\hat{u}_ir_j^{-1}$ for some $j \neq j', j''$. Let $\hat{v}_i(a_{v_j}, b_{v_j})$ be the reduced word representing $\hat{w}_i$.

(7) $\hat{w}_i = r_j^{-1}\hat{u}_ir_j$ for some j. Let $\hat{v}_i(a_{v_0}, b_{v_0})$ be the reduced word representing $\hat{w}_i$.

(8) $\hat{w}_i$ is a word in the standard generators of V_T. Let $\hat{v}_i = \hat{w}_i$ in this case.

There is an equality $\hat{u} = \hat{v}_1\cdots\hat{v}_{k'}$ in V_T. Let Δ be a van Kampen diagram over X_T with boundary word $\hat{v}_1\cdots\hat{v}_{k'}\hat{u}^{-1}$. The arcs in the boundary of Δ labeled by the words $\hat{v}_i$ or $\hat{u}^{-1}$ are called the *sides* of Δ. The side labeled by $\hat{v}_i$ is declared to be of type (5), (6), (7), or (8) accordingly as $\hat{w}_i$ is of one of these types. The remaining side will be regarded as being oriented in the opposite direction, so that its label reads $\hat{u}$.

We enlarge Δ to a diagram Δ' by adjoining folded r–corridors, as follows. First, for each subword $\hat{w}_i$ of type (5), build a folded corridor C_i whose top is labeled by $\hat{v}_i$ and whose bottom is labeled by the reduced word in a_{v_0}, b_{v_0} representing $\hat{u}_i$ in A_{v_0}. Adjoin this corridor to Δ along the side labeled by $\hat{v}_i$. Next, build a folded $r_{j'}$–corridor

C_0 whose top is labeled by $\hat{u}$ and whose bottom is labeled by $\hat{v}$. Adjoin it to Δ along the side labeled by $\hat{u}$. The resulting diagram is Δ'.

Now let $\sigma \subset \widehat{T}$ be the maximal segment whose endpoints correspond to the peripheral subgroups $A_{v_{j'}}$ and $A_{v_{j''}}$. The corridor scheme $\mathscr{S}_\sigma$ defines σ–corridors in Δ. Every edge in the side labeled $\hat{u}$ has a σ–corridor emanating from it, landing on a side of type (5) or (8). (They cannot land on the other sides, because their edges are not members of $\mathscr{S}_\sigma$. They cannot land on the side labeled $\hat{u}$, because $\hat{u}$ is reduced; cf. the proof of Proposition 6.7.) Decompose $\hat{u}$ as $z_1 \cdots z_\ell$ such that

- for each z_i, the σ–corridors emanating from the edges labeled by z_i either land on a single side of type (5), or on a union of sides of type (8)
- each z_i is maximal with respect to the preceding property.

Let $p_0, \dots, p_\ell$ be the vertices along the side labeled $\hat{u}$ such that for each i, the arc labeled by z_i has endpoints p_{i-1} and p_i. These vertices lie along the top of the folded $r_{j'}$–corridor C_0. Choose vertices $q_0, \dots, q_\ell$ along the bottom of C_0 such that p_i is nearly above q_i for each i and q_0, q_ℓ are the endpoints of the bottom. Let $[p_{i-1}, p_i]$ and $[q_{i-1}, q_i]$ denote the segments along the top and bottom, respectively, with the indicated endpoints.

Next define the index sets

$$\begin{aligned} I_{(5)} &= \{ i \mid \text{the } \sigma\text{–corridors emanating from } z_i \text{ land on a side of type (5)} \}, \\ I_{(8)} &= \{ i \mid \text{the } \sigma\text{–corridors emanating from } z_i \text{ land on sides of type (8)} \} \end{aligned}$$

so that $I_{(5)} \cup I_{(8)} = \{1, \dots, \ell\}$. For each $i \in I_{(5)}$ let j_i be the index such that the σ–corridors emanating from z_i land on the side labeled by $\hat{v}_{j_i}$.

If C and C' are two σ–corridors emanating from $\hat{u}$ and landing on the same side $\hat{v}_j$, then every corridor emanating from $\hat{u}$ between C and C' must also land on $\hat{v}_j$, since corridors do not cross. Moreover, if C and C' are adjacent in $\hat{u}$ then they will land on adjacent edges of $\hat{v}_j$; otherwise, any σ–corridor emanating from $\hat{v}_j$ between C and C' is forced to land on $\hat{v}_j$, contradicting that $\hat{v}_j$ is reduced.

These remarks imply that for every $i \in I_{(5)}$, the corridors emanating from z_i land on a connected subarc α_i of the side labeled $\hat{v}_{j_i}$, and in fact the label along α_i is the word $z_i(a_{v_{j'}}, b_{v_{j'}})$ or $z_i(a_{v_{j''}}, b_{v_{j''}})$, by Remark 6.3. Note that α_i lies along the top of the r–corridor C_{j_i}. Let β_i be a subarc of the bottom of C_{j_i} such that the endpoints of α_i are nearly above those of β_i. Since $[p_{i-1}, p_i]$ and α_i are labeled by the same word, we have

$$\big| |[q_{i-1}, q_i]| - |\beta_i| \big| \leqslant 2K_2 \tag{9.9}$$

by Lemma 9.5. Therefore,

$$|[q_{i-1}, q_i]| \leqslant \|\hat{u}_{j_i}\|_{A_{v_0}} + 2K_2.$$

Applying the induction hypothesis to $\hat{u}_{j_i}$ we obtain

$$\begin{aligned} |[q_{i-1},q_i]| &\leqslant K_3\|\hat{u}_{j_i}\|^{\alpha} + 2K_2 \\ &= K_3(\|\hat{w}_{j_i}\|-2)^{\alpha} + 2K_2. \end{aligned} \tag{9.10}$$

Next, if $i \in I_{(8)}$, we have

$$\begin{aligned} |[q_{i-1},q_i]| &\leqslant \|\varphi^{-n}(z_i)\|_{A_{v_{j'}}} + K_2 \\ &\leqslant \Lambda^n\|z_i\| + K_2 \\ &\leqslant K_3\|z_i\| + K_2 \end{aligned} \tag{9.11}$$

by Lemma 9.5. Now define the disjoint sets

$$\begin{aligned} J_{(5)} &= \{j_i \mid i \in I_{(5)}\}, \\ J_{(8)} &= \{j \mid \hat{w}_j \text{ is of type (8)}\} \end{aligned}$$

and note that $|J_{(5)}| = |I_{(5)}|$. Summing the inequalities (9.10) and (9.11) over all $i \in I_{(5)} \cup I_{(8)}$ we obtain

$$\|\hat{v}\| = \sum_{i=1}^{\ell}|q_{i-1},q_i| \leqslant K_3\left(\sum_{i\in I_{(5)}}(\|\hat{w}_{j_i}\|-2)^{\alpha}\right) + K_3\left(\sum_{i\in I_{(8)}}\|z_i\|\right) + \big(2|I_{(5)}|+|I_{(8)}|\big)K_2.$$

By considering σ–corridors we have $\sum_{i\in I_{(8)}}\|z_i\| \leqslant \sum_{j\in J_{(8)}}\|\hat{w}_j\|$, and therefore

$$\|\hat{v}\| \leqslant K_3\left(\sum_{j\in J_{(5)}}(\|\hat{w}_j\|-2)^{\alpha}\right) + K_3\left(\sum_{j\in J_{(8)}}\|\hat{w}_j\|\right) + \big(2|I_{(5)}|+|I_{(8)}|\big)K_2. \tag{9.12}$$

Now observe that no two adjacent indices can both be in $I_{(8)}$, by the maximality property of the words z_i. It follows that $|I_{(8)}| \leqslant |I_{(5)}|+1$, and since $K_3 \geqslant 3K_2/2$, we have

$$\begin{aligned} \big(2|I_{(5)}|+|I_{(8)}|\big)K_2 &\leqslant \big(3|I_{(5)}|+1\big)K_2 \\ &\leqslant K_3\big(2|I_{(5)}|+2\big). \end{aligned}$$

Combining this with (9.12) we obtain

$$\begin{aligned} \|\hat{v}\| &\leqslant K_3\left(2 + \sum_{j\in J_{(5)}}\big((\|\hat{w}_j\|-2)^{\alpha}+2\big)\right) + K_3\left(\sum_{j\in J_{(8)}}\|\hat{w}_j\|\right) \\ &\leqslant K_3\left(2^{\alpha} + \sum_{j\in J_{(5)}}\|\hat{w}_j\|^{\alpha} + \sum_{j\in J_{(8)}}\|\hat{w}_j\|^{\alpha}\right) \\ &\leqslant K_3\left(2+\sum_{j=1}^{k'}\|\hat{w}_j\|\right)^{\alpha} = K_3\|\hat{w}\|^{\alpha}. \end{aligned} \tag{9.13}$$

Finally, consider the words w_L and w_R, and note that $v = w_L\hat{v}w_R$. Choose any index $j \in \{1,\dots,m\}$ and apply Proposition 6.7/Remark 6.9 to obtain

$$\|g\|_{A_{v_0}} \leqslant \|v\|_{v_j} + \|v\|_{v_0} + \|v\|_x + \|v\|_y \leqslant \|v\| = \|w_L\| + \|\hat{v}\| + \|w_R\|.$$

Combining this with (9.13) yields

$$\|g\|_{A_{v_0}} \leqslant \|w_L\| + K_3\|\hat{w}\|^{\alpha} + \|w_R\| \leqslant K_3\|w_L\hat{w}w_R\|^{\alpha} = K_3\|w\|^{\alpha}.$$

This completes the proof in Case II, and the proof of the proposition. □

Corollary 9.14. *Given T and n there is a constant K_4 such that if w is a word in the standard generators of $S_{T,n}$ representing an element $g \in A_{v_0}$ then*

$$|g|_{A_{v_0}} \leqslant K_4|w|^{\alpha}$$

where $\alpha = n\log_m(\lambda)$.

Proof. One simply enlarges the constant K_3 to K_4, to account for the maximum scaling factor between $|\cdot|$ and $\|\cdot\|$. □

10. The Dehn function of $S_{T,n}$

Before proceeding we need to establish some additional properties of σ–corridors. First, we identify the subgroups of V_T generated by the labels along the sides of σ–corridors. Fixing σ, this subgroup will be denoted S_σ (the *side subgroup*). Looking at the corridor cells in Figure 5, it is clear that S_σ is generated by the subgroups F_j for various j; namely, whenever σ passes through a triangle in the $(|T|+2)$–gon with corners i, j, k, if σ separates i from j and k, then F_j and F_k are in S_σ.

Lemma 10.1. *The side subgroup S_σ is the free product of the factors $F_j \times F_k$, with one factor for each triangle in the $(|T|+2)$–gon through which σ passes.*

Proof. First, there is a surjective homomorphism from the free product onto S_σ, induced by inclusion of the subgroups $F_j \times F_k$. Injectivity then follows from Lemma 5.1, once we observe that in each vertex group $F_i \times F_j \times F_k$ of V_T, the subgroup $F_j \times F_k$ has trivial intersection with the edge groups $A_i \subset F_i \times F_j$ and $A_k \subset F_k \times F_i$. □

Next we need some observations about corridors in diagrams over the subgroup $F_0 \times F_1 \subset V$. Let $X_{01} \subset X$ be the subcomplex whose 2–cells have boundary edges labeled by the elements $x_0, y_0, x_1, y_1, a_0, b_0$. It has two quadrilateral cells and four triangular cells (see Figure 5). Its fundamental group is $F_0 \times F_1$. Define two orientable corridor schemes $\mathscr{S}_0, \mathscr{S}_1$ over X_{01} as follows: $\mathscr{S}_0$ contains the edges labeled x_0, y_0, a_0, b_0 and $\mathscr{S}_1$ contains the edges labeled x_1, y_1, a_0, b_0. Note that the *side subgroup* associated with $\mathscr{S}_0$ is F_1, and the side subgroup of $\mathscr{S}_1$ is F_0.

Remarks 10.2. Let Δ be a reduced diagram over X_{01}.

(1) Let α be a path in the 1–skeleton of Δ along the side of an $\mathscr{S}_0$–corridor, and let β be a path along the side of an $\mathscr{S}_1$–corridor. Then α and β intersect in at most one point. If there were two points in the intersection, then the labels along α and β

between these points give non-trivial elements of F_1 and F_0 respectively, representing the same element of $F_0 \times F_1$.

(2) Every 2–cell of Δ is contained in both an $\mathscr{S}_0$– and an $\mathscr{S}_1$–corridor. This is immediate by examining the corridor cells for $\mathscr{S}_0$ and $\mathscr{S}_1$.

(3) There are no $\mathscr{S}_0$– or $\mathscr{S}_1$–corridors of annulus type (because the side subgroups are free and Δ is reduced).

Proposition 10.3 (Area in V_T)**.** *Given T there is a constant K_5 with the following property. Suppose w and $z(a_{v_\ell}, b_{v_\ell})$ represent the same element of $A_{v_\ell} \subset V_T$, where z is reduced and w decomposes as $w_1 \cdots w_k$ where each w_i is a reduced word in the generators x_j, y_j for some j, or the peripheral generators a_{v_j}, b_{v_j} for some j. Let $\mathscr{I}$ be the set of pairs of indices (i, j) such that either $i \neq j$ or $i = j$ and w_i is not a word in peripheral generators of V_T. Then*

$$\mathrm{Area}(wz^{-1}) \leqslant K_5 \sum_{(i,j)\in\mathscr{I}} |w_i||w_j|.$$

Remark 10.4. Let $\mathscr{P}$ be the set of indices i such that w_i is a word in peripheral generators of V_T. Then

$$\sum_{(i,j)\in\mathscr{I}} |w_i||w_j| + \sum_{i\in\mathscr{P}} |w_i|^2 = |w|^2.$$

Proof of Proposition 10.3. Recall from Remark 4.4 that V_T is CAT(0), and therefore its Dehn function is quadratic. Thus there is a constant $C \geqslant 1$ (independent of z, w) such that $\mathrm{Area}(wz^{-1}) \leqslant C(|w| + |z|)^2$. Also, $|z| \leqslant |w|$ by Proposition 6.7, so $\mathrm{Area}(wz^{-1}) \leqslant 4C|w|^2$.

The proof now falls into two cases: the generic case and a special case which is somewhat more difficult. The latter case is when there is an index i' such that $w_{i'}$ is a word in the generators a_{v_ℓ}, b_{v_ℓ} (the same as z) and $|w_{i'}| > (1/2)|w|$.

Consider first the generic case (that is, whenever w_i is a word in the generators a_{v_ℓ}, b_{v_ℓ} we have $|w_i| \leqslant (1/2)|w|$). We claim that $|w_i| \leqslant (1/2)|w|$ for *every* $i \in \mathscr{P}$. To see this, suppose w_i is a word in a_{v_j}, b_{v_j} with $j \neq \ell$ and let $v_{j'}$ be a peripheral index not equal to v_j or v_ℓ. Let $\hat{w}$ be the complement of w_i in the cyclic word wz^{-1}. Apply Proposition 6.7 to obtain

$$|w_i| \leqslant |\hat{w}|_{v_{j'}} + |\hat{w}|_{v_j} + |\hat{w}|_x + |\hat{w}|_y.$$

The right hand side counts no letters of z because $j, j' \neq \ell$. Thus, $|w_i| \leqslant (1/2)|w|$.

This claim implies (see Figure 11) that

$$\sum_{i\in\mathscr{P}} |w_i|^2 \leqslant \sum_{i\neq j} |w_i||w_j|. \tag{10.5}$$

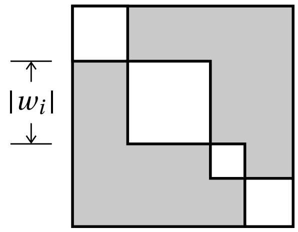

FIGURE 11. The area of the squares does not exceed the shaded area, if the side lengths are each at most half the total side length.

From Remark 10.4 we deduce that

$$\sum_{i\in\mathscr{P}} |w_i|^2 \leqslant (1/2)|w|^2$$

and therefore

$$\sum_{(i,j)\in\mathscr{I}} |w_i||w_j| \geqslant (1/2)|w|^2.$$

Finally, we have

$$\text{Area}(wz^{-1}) \leqslant 4C|w|^2 \leqslant 8C \sum_{(i,j)\in\mathscr{I}} |w_i||w_j|$$

so we are done by taking $K_5 \geqslant 8C$.

Next consider the special case: $w_{i'}$ is a word in a_{v_ℓ}, b_{v_ℓ} and $|w_{i'}| > (1/2)|w|$. Write $w = w_L w_{i'} w_R$, so that $|w_L| + |w_R| < |w_{i'}|$. Note that it will suffice for us to prove that $\text{Area}(wz^{-1}) \leqslant K_5|w_{i'}|(|w_L| + |w_R|)$.

Let σ and σ' be maximal segments in $\widehat{T}$ with common endpoint v_{v_ℓ}, which diverge immediately. That is, their intersection consists of the single edge from the leaf v_{v_ℓ} to its parent vertex in T. This edge lies inside a triangle in the $(|T|+2)$–gon. Suppose without loss of generality that this triangle has index triple $(0,1,2)$ and that $A_{v_\ell} = A_0 \subset F_0 \times F_1$. The vertex group in V_T corresponding to this triangle is $F_0 \times F_1 \times F_2$; denote this subgroup by V. For concreteness, suppose that σ separates corner 0 from corners 1 and 2, and σ' separates corner 1 from corners 2 and 0.

We can express the $(|T|+2)$–gon as a union of two smaller sub-diagrams whose intersection is the $(0,1,2)$ triangle. Note that σ and σ' each lie wholly inside one of these sub-diagrams. Thus V_T has an expression as $A *_V B$ where A is the fundamental group of the sub-diagram containing σ and B is the fundamental group of the sub-diagram containing σ'. See Figure 12. In particular, $S_\sigma \subset A$, $S_{\sigma'} \subset B$, and $A \cap B = V$.

Let Δ be a reduced van Kampen diagram over X_T with boundary wz^{-1}. Think of the boundary as being two arcs with the same endpoints, labeled by w and z respectively. Every edge along $w_{i'}$ has a σ–corridor emanating from it. Since $w_{i'}$ and z are reduced words in a_{v_ℓ}, b_{v_ℓ}, we can argue (as usual) that no σ–corridor has both ends on $w_{i'}$ or on z, and that the σ–corridors emanating from $w_{i'}$ and landing on z all land on a

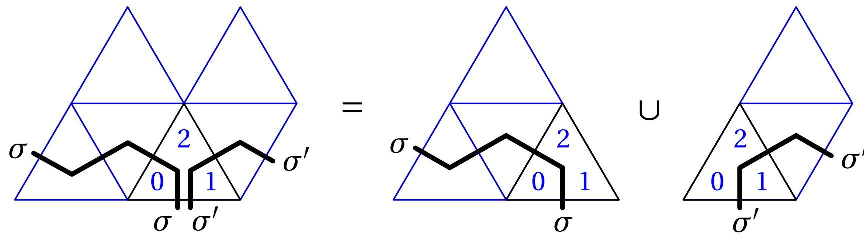

FIGURE 12. Expressing V_T as $A *_V B$ with $S_\sigma \subset A$ and $S_{\sigma'} \subset B$.

connected subarc of the arc labeled z. The σ–corridors not landing on z must land on w_L or w_R, and one finds that these comprise at most 1/3 of the corridors emanating from $w_{i'}$ (because $|w_{i'}| > |w_L| + |w_R|$).

Let p_1, p_2 be the initial and final endpoints of the maximal segment along $w_{i'}$ whose σ–corridors land on z. Let q_1, q_2 be the analogous points along z, so that p_i is joined to q_i by the side of a σ–corridor ($i = 1,2$). Note that the subsegments $[p_1, p_2]$ and $[q_1, q_2]$ are labeled by the same word in a_{v_ℓ}, b_{v_ℓ}, by Remark 6.3.

In a similar fashion, there are σ'–corridors emanating from $w_{i'}$, at least 2/3 of which land on a connected subsegment of z. Define p'_1, p'_2, q'_1, q'_2 analogously to p_1, p_2, q_1, and q_2.

Now let $[p_i, q_i]$ denote the side of the σ–corridor joining p_i to q_i ($i = 1,2$). Define $[p'_i, q'_i]$ analogously. Note that the label along $[p_1, q_1]$ represents an element of $S_\sigma \subset A$. However, p_1 and q_1 are also joined by the path $[p_1, p'_1] \cdot [p'_1, q'_1] \cdot [q'_1, q_1]$ where the first and third segments run along $w_{i'}$ and z respectively. This path represents an element of B. Therefore, $[p_1, q_1]$ represents an element of $A \cap B = V$. By Lemma 10.1 we have $S_\sigma \cap V = F_1 \times F_2$, and so $[p_1, q_1]$ in fact represents an element of this latter subgroup. By the same argument, $[p_2, q_2]$ also represents an element of $F_1 \times F_2$.

We need to introduce a little more notation. Let α be the path along the boundary of Δ from q_1 to p_1 which contains the segment labeled w_L. Similarly, let β be the path in the boundary from p_2 to q_2 which contains w_R. Let p_0, p_3 be the initial and final endpoints of the segment labeled $w_{i'}$. Let q_0, q_3 be the endpoints of the segment labeled z. Note that $|[p_0, p_1]| \leqslant |w_L|$ since the σ–corridors from $[p_0, p_1]$ land on w_L. Similarly, $|[p_2, p_3]| \leqslant |w_R|$. Because $|z| \leqslant |w|$, we have

$$\begin{aligned} |[q_0, q_1]| + |[q_2, q_3]| &\leqslant |w_L| + |[p_0, p_1]| + |[p_2, p_3]| + |w_R| \\ &\leqslant 2(|w_L| + |w_R|) \end{aligned}$$

and therefore

$$|\alpha| + |\beta| \leqslant 4(|w_L| + |w_R|). \tag{10.6}$$

At this point we discard the diagram Δ and start over with its boundary loop. We have seen that α and β represent elements g_α and g_β of the subgroup $F_1 \times F_2 \subset V$.

Attach a segment $[q_1, p_1]$ to the points q_1 and p_1, and label it by a reduced word representing g_α, of the form $u(x_1, y_1)\,v(x_2, y_2)$. Similarly, attach a segment $[q_2, p_2]$ labeled by a reduced word $u'(x_1, y_1)\,v'(x_2, y_2)$ representing g_β^{-1} to the points q_2 and p_2. Note that the loop $[p_1, p_2]\cdot[p_2, q_2]\cdot[q_2, q_1]\cdot[q_1, p_1]$ is labeled by generators of V, and the only occurrences of generators of F_2 are in the reduced words v and v'. It follows that $v = v'$.

Now define o_1 to be the point along $[q_1, p_1]$ such that $[q_1, o_1]$ is labeled by $u(x_1, y_1)$ and $[o_1, p_1]$ is labeled by $v(x_2, y_2)$. Similarly let o_2 be the point on $[q_2, p_2]$ such that $[q_2, o_2]$ is labeled by $u'(x_1, y_1)$ and $[o_2, p_2]$ is labeled by $v(x_2, y_2)$. Attach one more segment $[o_1, o_2]$ to the points o_1 and o_2, labeled by the same word in a_0, b_0 $(= a_{v_\ell}, b_{v_\ell})$ as $[p_1, p_2]$ and $[q_1, q_2]$.

We proceed now to fill the original boundary loop with a van Kampen diagram, in four parts. Fill the loop $\alpha\cdot[p_1, q_1]$ with a least-area van Kampen diagram Δ_α over X_T. Fill $\beta\cdot[q_2, p_2]$ with a least-area diagram Δ_β over X_T. To estimate the areas of Δ_α and Δ_β, choose a segment $\sigma'' \subset \widehat{T}$ which passes through the $(0,1,2)$ triangle and separates corner 2 from corners 0 and 1. Note that $x_1, y_1 \in \mathscr{S}_{\sigma'} - \mathscr{S}_{\sigma''}$ and $x_2, y_2 \in \mathscr{S}_{\sigma''} - \mathscr{S}_{\sigma'}$. Every edge along $[p_1, q_1]$ has either a σ'–corridor or a σ''–corridor in Δ_α emanating from it, and not both. Since the words $u(x_1, y_1)$, $v(x_2, y_2)$ are reduced, these corridors can only land on α. It follows that

$$|[p_1, q_1]| \leqslant 2|\alpha|. \tag{10.7}$$

By a similar argument using Δ_β, we have

$$|[p_2, q_2]| \leqslant 2|\beta|. \tag{10.8}$$

From (10.7) and (10.8) we deduce that

$$\mathrm{Area}(\Delta_\alpha) \leqslant C(3|\alpha|)^2 \text{ and } \mathrm{Area}(\Delta_\beta) \leqslant C(3|\beta|)^2.$$

Two more loops remain to be filled. Let $\hat{w}(a_0, b_0)$ be the word labeling $[p_1, p_2]$ and $[o_1, o_2]$. The loop $[p_1, p_2]\cdot[p_2, o_2]\cdot[o_2, o_1]\cdot[o_1, p_1]$ is labeled by the commutator $[\hat{w}(a_0, b_0), v(x_2, y_2)]$. Adjoin a triangular relator to each edge of $[p_1, p_2]$ and $[o_1, o_2]$ to obtain paths labeled by the word $\hat{w}(\overline{x}_0 x_1, \overline{y}_0 y_1)$. Now these paths and the paths $[p_1, o_1]$, $[p_2, o_2]$ can be filled using commutator relators. In this way the loop $[p_1, p_2]\cdot[p_2, o_2]\cdot[o_2, o_1]\cdot[o_1, p_1]$ bounds a diagram Δ_{012} over X_T of area $2|[p_1, p_2]| + 2(2|[p_1, p_2]|)|v|$.

Finally consider the loop $[o_1, o_2]\cdot[o_2, q_2]\cdot[q_2, q_1]\cdot[q_1, o_1]$. It is labeled entirely by generators of $F_0 \times F_1$ and represents the trivial element, so it bounds a least-area diagram Δ_{01} over X_{01}. Its boundary is labeled by the word

$$\hat{w}(a_0, b_0)\,u'(x_1, y_1)^{-1}\,\hat{w}(a_0, b_0)^{-1}\,u(x_1, y_1)$$

with each of the four subwords being reduced. Consider corridors in Δ_{01} for the corridor schemes $\mathscr{S}_0$ and $\mathscr{S}_1$. The $\mathscr{S}_0$–corridors can only land on the sides $[q_1, q_2]$ and

$[o_1, o_2]$, and since these sides are labeled by reduced words, each such corridor has one end on each side. Thus, each edge of $[q_1, q_2]$ is joined by an $\mathscr{S}_0$–corridor to the edge of $[o_1, o_2]$ corresponding to the same letter of $\hat{w}$.

Consider the arrangement of the $\mathscr{S}_1$–corridors in Δ_{01}. Every edge in the boundary has an $\mathscr{S}_1$–corridor emanating from it. There are no corridors of annulus type (see Remarks 10.2), and corridors cannot join two edges in the same side of the boundary. If there is a corridor joining $[q_1, o_1]$ to $[q_2, o_2]$ then we must have

$$2|[p_1, p_2]| \leqslant |[q_1, o_1]| + |[q_2, o_2]| \tag{10.9}$$

and we will be satisfied for the moment. Assume now that no such corridor is present. Then, the most general arrangement is shown in Figure 13.

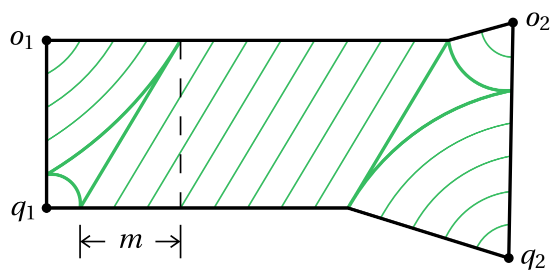

FIGURE 13. A generic configuration of $\mathscr{S}_1$–corridors in Δ_{01}.

Notice that corridors running between $[o_1, o_2]$ and $[q_1, q_2]$ will land on edges that are offset by a fixed amount $m \leqslant \min\{|[q_1, o_1]|, |[q_2, o_2]|\}$. In particular an $\mathscr{S}_0$–corridor can be crossed by at most $m+1$ $\mathscr{S}_1$–corridors of this type. There are $|[q_1, o_1]| + |[q_2, o_2]|$ $\mathscr{S}_1$–corridors not of this type, so an $\mathscr{S}_0$–corridor can cross no more than $2(|[q_1, o_1]| + |[q_2, o_2]|)$ $\mathscr{S}_1$–corridors overall.

By Remarks 10.2, an $\mathscr{S}_0$–corridor and an $\mathscr{S}_1$–corridor will intersect if and only if their endpoints are linked (or equal) in the boundary of Δ_{01}, and when this occurs, the intersection will have area 1 or 2. Moreover, every 2–cell of Δ_{01} is in the intersection of such a pair. Since every $\mathscr{S}_0$–corridor is crossed by at most $2(|[q_1, o_1]| + |[q_2, o_2]|)$ $\mathscr{S}_1$–corridors, we conclude that

$$\operatorname{Area}(\Delta_{01}) \leqslant 4(|[q_1, o_1]| + |[q_2, o_2]|)|[o_1, o_2]|.$$

To finish, first suppose we are in the special case where (10.9) holds. Then, the total perimeter is at most $12(|w_L| + |w_R|)$ by (10.7), (10.8), and (10.6). Then

$$\operatorname{Area}(wz^{-1}) \leqslant C(12(|w_L| + |w_R|))^2 \leqslant 144C|w_{i'}|(|w_L| + |w_R|)$$

so we require $K_5 \geqslant 144C$ to cover this case.

Otherwise, we use our estimates for the areas of the four diagrams Δ_α, Δ_β, Δ_{012}, and Δ_{01}. Since $|[p_1,p_2]| = |[o_1,o_2]| = |[q_1,q_2]| \leqslant |w_{i'}|$, these estimates yield

$$\mathrm{Area}(wz^{-1}) \leqslant 9C|\alpha|^2 + 9C|\beta|^2 + (2|w_{i'}| + 4|w_{i'}||v|) + 4(|u|+|u'|)|w_{i'}|.$$

Using (10.7) and (10.8) this reduces to

$$\mathrm{Area}(wz^{-1}) \leqslant 9C(|\alpha|+|\beta|)^2 + |w_{i'}|(2+8(|\alpha|+|\beta|)).$$

Note that $(|\alpha|+|\beta|) \leqslant 4|w_{i'}|$ by (10.6), and a further application of (10.6) yields the following:

$$\begin{aligned}\mathrm{Area}(wz^{-1}) &\leqslant 36C|w_{i'}|(|\alpha|+|\beta|) + |w_{i'}|(2+8(|\alpha|+|\beta|))\\ &\leqslant (36C+10)|w_{i'}|(|\alpha|+|\beta|)\\ &\leqslant (144C+40)|w_{i'}|(|w_L|+|w_R|).\end{aligned}$$

Taking $K_5 \geqslant 144C+40$ completes the proof. □

Proposition 10.10 (Area in $S_{T,n}$)**.** *Given T and n there is a constant K_6 with the following property. Suppose w and $z(a_{v_j}, b_{v_j})$ represent the same element of $A_{v_j} \subset S_{T,n}$, where w is a word in the standard generators of $S_{T,n}$ and z is reduced. Then* $\mathrm{Area}(wz^{-1}) \leqslant K_6|w|^{2\alpha}$.

We follow the proof of Proposition 5.5 in [BBFS09].

Proof. The proof is by induction on $|w|$. Let $M = \max\{|\varphi^n(x)|, |\varphi^n(y)|, |\varphi^{-n}(x)|, |\varphi^{-n}(y)|\}$. Let $K_6 = M^2K_4^2K_5$, where K_4 is given by Corollary 9.14 and K_5 is given by Proposition 10.3. Write w as $w_1 \dots w_k$ where each w_i either is a word in the standard generators of V_T, or $w_i = r_j^{\pm1}u_ir_j^{\mp1}$ for some j. Let I_r be the set of indices for which the latter case occurs, and note that w_i represents an element of a peripheral subgroup of V_T. Let v_i be the reduced word in the generators of that subgroup representing w_i. For $i \notin I_r$ let $v_i = w_i$, and define $v = v_1 \cdots v_k$.

By Proposition 10.3 we have $\mathrm{Area}(vz^{-1}) \leqslant K_5\sum_{(i,j)\in\mathscr{I}}|v_i||v_j|$. For $i \in I_r$ we have either $|v_i| \leqslant K_4|w_i|^\alpha$ (if $v_i \in A_{v_0}$) or $|v_i| \leqslant MK_4|u_i|^\alpha$ (if $u_i \in A_{v_0}$), by Corollary 9.14. In any case (including $i \notin I_r$) we have $|v_i| \leqslant MK_4|w_i|^\alpha$. Therefore,

$$\mathrm{Area}(vz^{-1}) \leqslant K_6 \sum_{(i,j)\in\mathscr{I}} |w_i|^\alpha|w_j|^\alpha. \tag{10.11}$$

Next note that

$$\mathrm{Area}(wv^{-1}) \leqslant \sum_{i\in I_r} \mathrm{Area}(w_iv_i^{-1})$$

since $w_i = v_i$ for $i \notin I_r$. For each term $\mathrm{Area}(w_iv_i^{-1})$, recall that $w_i = r_j^{\pm1}u_ir_j^{\mp1}$. Let z_i be the reduced word in peripheral generators representing the same peripheral element as u_i. Apply the induction hypothesis to u_i to obtain

$$\mathrm{Area}(u_iz_i^{-1}) \leqslant K_6|u_i|^{2\alpha} = K_6(|w_i|-2)^{2\alpha}. \tag{10.12}$$

One of the words z_i, v_i is an element of A_{v_0}, so there is a folded r_j–corridor with boundary word $r_j^{\pm 1} z_i r_j^{\mp 1} v_i^{-1}$, where one of the boundary arcs labeled z_i or v_i is the bottom and the other is the top. The area of this corridor is the length of the bottom, which is at most $M|v_i|$. As noted above, $|v_i| \leqslant MK_4|w_i|^{\alpha}$, and so

$$\mathrm{Area}(r_j^{\pm 1} z_i r_j^{\mp 1} v_i^{-1}) \leqslant MK_4|w_i|^{\alpha}.$$

Together with (10.12) we obtain

$$\begin{aligned} \mathrm{Area}(w_i v_i^{-1}) &\leqslant K_6((|w_i|-2)^{2\alpha} + |w_i|^{\alpha}) \\ &\leqslant K_6|w_i|^{2\alpha} \end{aligned} \tag{10.13}$$

for $i \in I_r$. The last inequality above holds exactly as in [BBFS09, Proposition 5.5]: for numbers $x \geqslant 0$ one has $(x+2)^{2\alpha} \geqslant x^{\alpha}(x+2)^{\alpha} + 2^{\alpha}(x+2)^{\alpha} \geqslant x^{2\alpha} + (x+2)^{\alpha}$.

Lastly, add together (10.11) and (10.13) for each $i \in I_r$ to obtain

$$\begin{aligned} \mathrm{Area}(wz^{-1}) &\leqslant K_6 \sum_{(i,j)\in\mathscr{I}} |w_i|^{\alpha}|w_j|^{\alpha} + K_6 \sum_{i\in I_r} |w_i|^{2\alpha} \\ &\leqslant K_6 \sum_{i\neq j} |w_i|^{\alpha}|w_j|^{\alpha} + K_6 \sum_{i} |w_i|^{2\alpha}. \end{aligned}$$

The latter quantity is $K_6|w|^{2\alpha}$, as desired. □

Theorem 10.14. *Given T and n let $m = |T|+1$, let $\lambda > 1$ be the Perron-Frobenius eigenvalue of φ, and let $\alpha = n\log_m(\lambda)$. If $\alpha \geqslant 1$ then the Dehn function of $S_{T,n}$ is given by $\delta(x) = x^{2\alpha}$.*

Proof. First we establish the lower bound $\delta(x) \succcurlyeq x^{2\alpha}$. Let $w(x,y)$ be a monotone palindromic word (eg. x) and consider the snowflake diagrams $\Delta(w,i)$ for $i \geqslant 1$. Let n_i be the boundary length of $\Delta(w,i)$. The boundary word has $4m^{i-1}$ occurrences of the letters r_j in it, and two adjacent such letters are never separated by more than $m|w|$ letters from V_T. Thus we have

$$4m^{i-1} \leqslant n_i \leqslant 4(m|w|+1)m^{i-1}. \tag{10.15}$$

From the second of these inequalities we obtain

$$\begin{aligned} (n_i)^{\alpha} &\leqslant \left(\frac{4(m|w|+1)}{m}\right)^{\alpha} (m^{\alpha})^i \\ &= \left(\frac{4(m|w|+1)}{m}\right)^{\alpha} \lambda^{ni} \end{aligned}$$

and so

$$\left(\frac{m}{4(m|w|+1)}\right)^{\alpha} (n_i)^{\alpha} \leqslant \lambda^{ni}.$$

Next, $\Delta(w,i)$ has area at least $6|T||\varphi^{in}(w)|^2$, which is the area of the doubled canonical diagram at its center. There is a constant C such that $|\varphi^k(v)| \geqslant C\lambda^k|v|$ for every nontrivial word v. Thus,

$$\begin{aligned}\operatorname{Area}(\Delta(w,i)) &\geqslant 6|T|C^2|w|^2\left(\lambda^{ni}\right)^2\\ &\geqslant 6|T|C^2|w|^2\left(\frac{m}{4(m|w|+1)}\right)^{2\alpha}(n_i)^{2\alpha}.\end{aligned}$$

Taking $D = 6|T|C^2|w|^2\left(\frac{m}{4(m|w|+1)}\right)^{2\alpha}$, we have shown that $\delta_{Y_{T,n}}(n_i) \geqslant D(n_i)^{2\alpha}$ for each i, because $\Delta(w,i)$ is a least-area diagram over $Y_{T,n}$ with boundary length n_i. By (10.15) the ratios n_{i+1}/n_i are bounded, and so we conclude by Remark 2.1 that $\delta(x) \succcurlyeq x^{2\alpha}$.

The upper bound follows immediately from Proposition 10.10: taking z to be the empty word, $\operatorname{Area}(w) \leqslant K_6|w|^{2\alpha}$ for every word w representing the trivial element of $S_{T,n}$. Thus $\delta(x) \preccurlyeq x^{2\alpha}$. □

References

[ABD+13] Aaron Abrams, Noel Brady, Pallavi Dani, Moon Duchin, and Robert Young, *Pushing fillings in right-angled Artin groups*, J. Lond. Math. Soc. (2) **87** (2013), no. 3, 663–688. MR 3073670

[ABDY13] Aaron Abrams, Noel Brady, Pallavi Dani, and Robert Young, *Homological and homotopical Dehn functions are different*, Proc. Natl. Acad. Sci. USA **110** (2013), no. 48, 19206–19212. MR 3153947

[Bas93] Hyman Bass, *Covering theory for graphs of groups*, J. Pure Appl. Algebra **89** (1993), no. 1-2, 3–47. MR 1239551 (94j:20028)

[BB00] N. Brady and M. R. Bridson, *There is only one gap in the isoperimetric spectrum*, Geom. Funct. Anal. **10** (2000), no. 5, 1053–1070. MR 1800063 (2001j:20046)

[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 2469516 (2010a:20090)

[BBMS97] Gilbert Baumslag, Martin R. Bridson, Charles F. Miller, III, and Hamish Short, *Finitely presented subgroups of automatic groups and their isoperimetric functions*, J. London Math. Soc. (2) **56** (1997), no. 2, 292–304. MR 1489138 (98j:20034)

[BBMS00] ______, *Fibre products, non-positive curvature, and decision problems*, Comment. Math. Helv. **75** (2000), no. 3, 457–477. MR 1793798 (2001k:20091)

[BCD08] Noel Brady, Matt Clay, and Pallavi Dani, *Morse theory and conjugacy classes of finite subgroups*, Geom. Dedicata **135** (2008), 15–22. MR 2413325 (2009d:20103)

[BG10] Martin R. Bridson and Daniel Groves, *The quadratic isoperimetric inequality for mapping tori of free group automorphisms*, Mem. Amer. Math. Soc. **203** (2010), no. 955, xii+152. MR 2590896 (2011g:20058)

[BGL11] Noel Brady, Dan P. Guralnik, and Sang Rae Lee, *Dehn functions and finiteness properties of subgroups of perturbed right-angled Artin groups*, preprint, `http://arxiv.org/abs/1102.5551`, 2011.

[BH99] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999. MR 1744486 (2000k:53038)

[Bra95] Thomas Brady, *Complexes of nonpositive curvature for extensions of F_2 by* $\mathbf{Z}$, Topology Appl. **63** (1995), no. 3, 267–275. MR 1334311 (96e:20056)

[Bri02] Martin 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 1967746 (2004g:20056)

[Bri13] ______, *On the subgroups of right-angled Artin groups and mapping class groups*, Math. Res. Lett. **20** (2013), no. 2, 203–212. MR 3151642

[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 2281936 (2009j:20053)

[DR13] Will Dison and Timothy R. Riley, *Hydra groups*, Comment. Math. Helv. **88** (2013), no. 3, 507–540. MR 3093501

[FM91] Mark Feighn and Geoffrey Mess, *Conjugacy classes of finite subgroups of Kleinian groups*, Amer. J. Math. **113** (1991), no. 1, 179–188. MR 1087807 (92a:57042)

[Ger98] S. M. Gersten, *Cohomological lower bounds for isoperimetric functions on groups*, Topology **37** (1998), no. 5, 1031–1072. MR 1650363 (2000c:20063)

[LN03] Ian J. Leary and Brita E. A. Nucinkis, *Some groups of type VF*, Invent. Math. **151** (2003), no. 1, 135–165. MR 1943744 (2003k:20086)

[SBR02] Mark V. Sapir, Jean-Camille Birget, and Eliyahu Rips, *Isoperimetric and isodiametric functions of groups*, Ann. of Math. (2) **156** (2002), no. 2, 345–466. MR 1933723 (2005b:20077a)

MATHEMATICS DEPARTMENT, UNIVERSITY OF OKLAHOMA, NORMAN, OK 73019, USA
E-mail address: `nbrady@math.ou.edu, forester@math.ou.edu`