

DISTORTION OF SURFACE GROUPS IN CAT(0) FREE-BY-CYCLIC GROUPS

JOSH BARNARD AND NOEL BRADY¹

ABSTRACT. Given a non-positively curved 2-complex with a circle-valued Morse function satisfying some extra combinatorial conditions, we describe how to locally isometrically embed this in a larger non-positively curved 2-complex with free-by-cyclic fundamental group.

This embedding procedure is used to produce examples of CAT(0) free-by-cyclic groups that contain closed hyperbolic surface subgroups with polynomial distortion of arbitrary degree. We also produce examples of CAT(0) hyperbolic free-by-cyclic groups that contain closed hyperbolic surface subgroups that are exponentially distorted.

CONTENTS

Introduction	1
1. Morse functions on 2-complexes	3
2. Quasi-convex surface subgroups	8
3. Polynomially distorted surface subgroups	10
4. Exponentially distorted surface subgroups	14
5. Questions	22
References	22

INTRODUCTION

We introduce a technique for constructing non-positively curved 2-complexes with free-by-cyclic fundamental groups. As an application, we have a series of existence results concerning closed hyperbolic surface subgroups in free-by-cyclic groups.

It is an open question as to whether every hyperbolic free-by-cyclic group contains a closed hyperbolic surface subgroup. Theorem 1 below can be viewed as a positive answer to the much more basic question: do there exist hyperbolic free-by-cyclic groups that contain closed hyperbolic surface subgroups?

Theorem 1 (Quasi-convex surface subgroups). *There exist hyperbolic free-by-cyclic groups containing quasi-convex closed hyperbolic surface subgroups.*

The following two theorems show that the distortion of closed hyperbolic surface subgroups in CAT(0) free-by-cyclic groups can be as varied as the

Date: September 20, 2005.

¹N. Brady was partially supported by NSF grant no. DMS-0124344. He also thanks the Centre de Recerca Matemàtica for their support and hospitality.

distortion of free subgroups. We produce examples where the distortion is polynomial of arbitrary degree, or exponential. In fact, the proofs given here rely on the fact that there are $CAT(0)$ free-by-cyclic groups whose free kernels have the appropriate distortion function.

Theorem 2 (Polynomially distorted surface subgroups). *For each integer $n \geq 2$, there exists a $CAT(0)$ free-by-cyclic group containing the fundamental group of a closed hyperbolic surface that has polynomial distortion of degree n .*

In the exponential distortion case, we show that with quite a bit of extra care, one can require that the ambient free-by-cyclic group be hyperbolic.

Theorem 3 (Exponentially distorted surface subgroups). *There exist hyperbolic free-by-cyclic groups that contain exponentially distorted closed surface subgroups.*

The basic examples of hyperbolic groups that contain exponentially distorted, closed, hyperbolic surface subgroups are the fundamental groups of closed, hyperbolic 3-manifolds that fiber over the circle. Theorem 3 shows that one can find hyperbolic groups of cohomological dimension 2 containing exponentially distorted closed surface subgroups, and that, in particular, one can make the ambient hyperbolic group be free-by-cyclic.

In all three theorems, the groups in question are fundamental groups of non-positively curved squared 2-complexes.

The paper is organized as follows. In Theorem 1.8 of Section 1 we establish a criterion for when a non-positively curved 2-complex can be locally isometrically embedded in a larger non-positively curved 2-complex with free-by-cyclic fundamental group. The local isometric embedding ensures that the fundamental group of the original 2-complex embeds into the free-by-cyclic group. This is the basic way that we shall produce closed surface groups inside free-by-cyclic groups, but the technical difficulties in each section increase as we require more conditions either on the distortion of the surface group, or on the curvature of the ambient free-by-cyclic group, or both.

In Section 2 we apply this technique to a squared structure on a genus 2 surface to obtain a linearly distorted, or quasi-convex, subgroup of a hyperbolic free-by-cyclic group. The hyperbolicity is ensured by a careful application of the techniques used to prove Theorem 1.8 in Section 1. In Section 3 we take a $CAT(0)$ amalgamation of (i) a free-by-cyclic group with polynomially distorted free subgroup and (ii) a closed surface group, amalgamating along the free subgroup. This produces a closed surface group with polynomial distortion inside a 2-dimensional $CAT(0)$ group. We then embed this into a $CAT(0)$ free-by-cyclic group by the methods of Theorem 1.8. In Section 4 we proceed as in Section 3, but this time starting with the amalgam of a surface group and a *hyperbolic* free-by-cyclic group, and proceed carefully as in Section 2 in order to end with an ambient hyperbolic free-by-cyclic group. Finally, in Section 5 we pose questions that are related to these ideas and constructions.

Noel Brady thanks Martin Bridson and Jens Harlander for conversations about $CAT(0)$ LOT groups and free-by-cyclic constructions.

1. MORSE FUNCTIONS ON 2-COMPLEXES

In this section we recall a Morse theory way of recognizing that a 2-complex is aspherical and has free-by-cyclic fundamental group. We then describe a combinatorial condition on non-positively curved 2-complexes (the existence of *f-link filling graphs*) that ensures that they locally isometrically embed into non-positively curved 2-complexes with free-by-cyclic fundamental groups. The material in this section is key to understanding the constructions in the subsequent sections of the paper.

The first two definitions below are from [1].

Definition 1.1 (Affine cell complexes). A finite-dimensional cell-complex X is said to be an *affine cell-complex* if it is equipped with the following structure. An integer $m \geq \dim(X)$ is given, and for each cell e of X we are given a convex polyhedral cell $C_e \subset \mathbb{R}^m$ and a characteristic function $\chi_e : C_e \rightarrow e$ such that the restriction of χ_e to any face of C_e is a characteristic function of another cell, possibly precomposed by a partial affine homeomorphism (=restriction of an affine homeomorphism) of \mathbb{R}^m .

Definition 1.2 (Morse function). A map $f : X \rightarrow \mathbb{R}$ defined on an affine cell complex X is a *Morse function* if

- for every cell e of X , $f\chi_e : C_e \rightarrow \mathbb{R}$ extends to an affine map $\mathbb{R}^m \rightarrow \mathbb{R}$, and $f\chi_e$ is constant only when $\dim e = 0$, and
- the image of the 0-skeleton is discrete in \mathbb{R} .

Definition 1.3 (Circle-valued Morse function). A circle-valued Morse function on an affine cell complex X is a cellular map $f : X \rightarrow S^1$, with the property that f lifts to a Morse function between universal covers.

By composing with a covering map $S^1 \rightarrow S^1$ if necessary, we will assume that all circle-valued Morse functions map to a circle with a single vertex.

Definition 1.4 (Ascending and descending links). Suppose X is an affine complex and $f : X \rightarrow S^1$ is a circle-valued Morse function. Choose an orientation of S^1 , which lifts to one of \mathbb{R} , and lift f to a map of universal covers $\tilde{f} : \tilde{X} \rightarrow \mathbb{R}$. Let $v \in X^{(0)}$, and note that the link of v in X is naturally isomorphic to the link of any lift \tilde{v} of v in \tilde{X} . We say that a cell $\tilde{e} \subset \tilde{X}$ contributes to the *ascending* (respectively *descending*) link of \tilde{v} if $\tilde{v} \in \tilde{e}$ and if $\tilde{f}|_{\tilde{e}}$ achieves its minimum (respectively maximum) value at \tilde{v} . The ascending (respectively descending) link of v is then defined to be the subset of $Lk(v, X)$ naturally identified with the ascending (respectively descending) link of \tilde{v} . Note that in the case X is a 2-complex, all the ascending, descending and entire links will be graphs.

Henceforth, all cell complexes will be 2-complexes. The next proposition gives a local way of telling that a 2-complex is aspherical and has free-by-cyclic fundamental group.

Proposition 1.5 (Free-by-cyclic). *If $f : X \rightarrow S^1$ is a circle-valued Morse function on the 2-complex X all of whose ascending and descending links are trees, then X is aspherical, and $\pi_1(X)$ is free-by-cyclic.*

Remark 1.6. This is basically a result of J. Howie [5]. In [5] it is proved that if all the ascending links (or all the descending links) of a given LOT 2-complex are trees, then the 2-complex is aspherical. Also, if both ascending and descending links are trees, then the LOT 2-complex has free-by-cyclic fundamental group.

The current proposition says that if all of both the ascending and the descending links of a general affine 2-complex are trees, then the 2-complex is aspherical and has free-by-cyclic fundamental group.

Note that if just the ascending links are trees and the descending links are not, then the aspherical 2-complex need not have free-by-cyclic fundamental group. As a simple example, take the group with five generators t and a_i ($i = 0, \dots, 3$) and four conjugation relations $a_i a_{i+1} a_i^{-1} = t$ ($i \bmod 4$). In any map to \mathbb{Z} , all five generators must have the same image, so we can assume this common image is a generator of \mathbb{Z} . The corresponding circle-valued Morse function takes the a_i and the t loops once around a target circle. Its ascending link is a tree, but the descending link is the disjoint union of a single vertex and a circle. The kernel of the map to \mathbb{Z} is not finitely generated.

Proof of Proposition 1.5. Look at generic f point preimages in X . Since X has an affine structure and $f : X \rightarrow S^1$ is affine, the point preimages are affine graphs in X . We may then view X as the total space of a graph of spaces: the underlying space is the target circle, the edge space is the generic point preimage graph, and the vertex space is the S^1 -vertex preimage. The maps from an edge space to the adjacent vertex spaces collapse either the ascending or the descending subgraphs of the corresponding link. These are trees by hypothesis, so these maps are homotopy equivalences. It follows that X is homotopy equivalent to a graph bundle over S^1 , and therefore is aspherical and has free-by-cyclic fundamental group. \square

The next definition is a little technical, so we preface it with a discussion of the intuition behind it. We are given a circle-valued Morse function on a locally CAT(0) 2-complex X , with the property that all ascending and descending links are forests. We would like to be able to attach 1-cells and 2-cells to X to obtain a bigger locally CAT(0) 2-complex in which X locally isometrically embeds. Furthermore, we want to extend the Morse function over the resulting 2-complex, so that all ascending and descending links are now trees. The description of the new 2-complex will be as a graph of spaces, whose underlying graph is a circle with one vertex, and one 1-cell. The vertex space is the original 2-complex X , and the edge space is a graph Γ . That is, the new 2-complex is obtained from X by attaching a “handle” of the form $\Gamma \times [0, 1]$. Condition (1) in Definition 1.7 is there to ensure that the geometry works out, namely, that the new 2-complex is locally CAT(0), and that the inclusion map of X is a local isometric embedding. Condition (2) in Definition 1.7 is there to ensure that all the ascending and descending forests get connected up precisely into trees. This intuition will be made precise in the proof of Theorem 1.8.

Definition 1.7 (*f*-link filling graphs). Suppose $f : X \rightarrow S^1$ is a circle-valued Morse function on a locally CAT(0) 2-complex X . An *f*-link filling graph is a (not necessarily connected) graph Γ satisfying the following:

- (1) there are two combinatorial immersions $\varphi_1, \varphi_2 : \Gamma \rightarrow X^{(1)}$ that are local isometric embeddings,
- (2) the maps $(f \circ \varphi_i)_* : \pi_1(\Gamma) \rightarrow \pi_1(S^1)$ agree, and
 - (a) for every vertex $v \in X$ the induced map on the direct sum of reduced homology groups

$$\bigoplus_{\substack{w \in \Gamma^{(0)} \\ \varphi_1(w)=v}} \tilde{H}_0(LK_{\uparrow}(w, \Gamma); \mathbb{Z}) \rightarrow \tilde{H}_0(Lk_{\uparrow}(v, X); \mathbb{Z})$$

is an isomorphism, and

- (b) for every vertex $v \in X$ the induced map on the direct sum of reduced homology groups

$$\bigoplus_{\substack{w \in \Gamma^{(0)} \\ \varphi_2(w)=v}} \tilde{H}_0(LK_{\downarrow}(w, \Gamma); \mathbb{Z}) \rightarrow \tilde{H}_0(Lk_{\downarrow}(v, X); \mathbb{Z})$$

is an isomorphism.

Note that in part 2(a) the ascending links of Γ are with respect to the Morse function $f \circ \varphi_1$, and in part 2(b) the descending links are with respect to the Morse function $f \circ \varphi_2$. We often refer explicitly to the immersions and write that $\varphi_1, \varphi_2 : \Gamma \rightarrow X$ is f -link filling.

As indicated before the definition, the main application we have in mind for an f -link filling graph Γ is when the Morse function $f : X \rightarrow S^1$ has forests as its ascending and descending links. Then the new 2-complex obtained from X by attaching $\Gamma \times [0, 1]$ via the map φ_1 at the $\Gamma \times 0$ end, and the map φ_2 at the $\Gamma \times 1$ end, will admit a circle-valued Morse function with trees as ascending and descending links. Here are the details.

Theorem 1.8 (Free-by-cyclic completions). *Let X be a $CAT(0)$ 2-complex X with a circle-valued Morse function $f : X \rightarrow S^1$ whose ascending and descending links are all forests. If $\varphi_1, \varphi_2 : \Gamma \rightarrow X$ is an f -link filling graph, then the 2-complex, $C(X)$, obtained from X by adjoining $\Gamma \times [0, 1]$ via φ_1 on $\Gamma \times 0$, and φ_2 on $\Gamma \times 1$ satisfies:*

- (1) *the complex $C(X)$ is locally $CAT(0)$ and the inclusion $X \rightarrow C(X)$ is a local isometric embedding,*
- (2) *the Morse function $f : X \rightarrow S^1$ extends to a Morse function $\hat{f} : C(X) \rightarrow S^1$ all of whose ascending and descending links are trees.*

In particular, $\pi_1(X)$ is a subgroup of $\pi_1(C(X))$, and $\pi_1(C(X))$ is a $CAT(0)$, free-by-cyclic group.

Proof of Theorem 1.8. The 2-complex $C(X)$ is obtained from the non-positively curved 2-complex X by attaching $\Gamma \times [0, 1]$ via the local isometric embeddings φ_i . Here $\Gamma \times [0, 1]$ is given the product metric, and the complex $C(X)$ gets the induced length metric. Note that we have just attached 1-cells (in the form of $\{w\} \times [0, 1]$ where w is a vertex of Γ) and 2-cells (in the form of $e \times [0, 1]$ where e is an edge of Γ) to X . Thus, $C(X)$ and X have the same vertex set.

Proof of part (1). We show that $C(X)$ is non-positively curved, and that the inclusion $X \subset C(X)$ is a local isometric embedding. This follows from

the standard theory [4], but we give a proof here, as the local picture that is developed below is used in the proof of part (2).

For non-positive curvature, we must verify that $Lk(v, C(X))$ is ‘large’ for every vertex $v \in C(X)$; i.e., we check that all cycles in $Lk(v, C(X))$ have length at least 2π . From the previous paragraph, we know that $v \in X$, and the non-positive curvature of X ensures that $Lk(v, X)$ is large. Thus we have to describe the effect on $Lk(v, X)$ of attaching $\Gamma \times [0, 1]$ via the maps φ_i . Let us consider the φ_1 attaching map; the analysis of φ_2 is completely analogous. If $w \in \Gamma^{(0)}$ then $Lk((w, 0), \Gamma \times [0, 1])$ is simply an orthogonal cone on $Lk(w, \Gamma)$. The cone point corresponds to the edge $w \times [0, 1]$ at $(w, 0)$, and the adjective ‘orthogonal’ just means that the lengths of the edges are all $\pi/2$. This latter fact holds since we’ve taken the product metric on $\Gamma \times [0, 1]$. Since φ_1 is a local isometric embedding we know that the φ_1 -induced image of $Lk(w, \Gamma)$ in $Lk(v, X)$, denoted by $(\varphi_1)_*(Lk(w, \Gamma))$, is convex. Specifically, $(\varphi_1)_*(Lk(w, \Gamma))$ is a discrete set of points (in 1-1 correspondence with $Lk(w, \Gamma)$) that is π -separated.

Thus, the local effect of attaching a neighborhood of $(w, 0)$ in $\Gamma \times [0, 1]$ via φ_1 is to attach an orthogonal cone on $Lk(w, \Gamma)$ to $Lk(v, X)$ by identifying $Lk(w, \Gamma)$ with the π -separated subset $(\varphi_1)_*(Lk(w, \Gamma))$ of $Lk(v, X)$. This new graph is still large, and $Lk(v, X)$ is a convex subgraph (in the CAT(1) sense). We continue to attach orthogonal cones, one for every $w \in \Gamma$ with $\varphi_1(w) = v$ and one for every $w \in \Gamma$ with $\varphi_2(w) = v$. Thus the final result, $Lk(v, C(X))$, is large, and $Lk(v, X)$ is a convex subgraph. This shows that $C(X)$ is non-positively curved, and that $X \subset C(X)$ is a local isometric embedding. Note that we have needed only the hypothesis that the φ_i are local isometric embeddings for this part.

Proof of part (2). The requirement that $(f \circ \varphi_1)_* = (f \circ \varphi_2)_*$ ensures that the Morse function $f : X \rightarrow S^1$ extends to a Morse function $\tilde{f} : C(X) \rightarrow S^1$.

We can explicitly construct this extension as follows. For any edge $e \subset \Gamma$, we let $N_1(e)$ denote the (signed) number of times that $(f \circ \varphi_1)(e)$ winds around S^1 , and similarly define $N_2(e)$. Choose a vertex $w_0 \in \Gamma$ and an integer t_0 so that

$$t_0 > \sum_{e \in \Gamma} |N_2(e) - N_1(e)|.$$

We map both end points of the new 1-cell $\{w_0\} \times [0, 1]$ to the vertex of the circle and send the $[0, 1]$ interval around the circle t_0 times in the positive direction. Now suppose w_1 is a vertex of Γ adjacent to w_0 via the edge e_1 . We map the new 1-cell $\{w_1\} \times [0, 1]$ around the circle monotonically $t_1 = t_0 + (N_2(e_1) - N_1(e))$ times. Note that t_1 is always positive by choice of t_0 , and so this 1-cell gets mapped positively around the target circle. Continuing in this fashion, suppose we have some sequence of edges e_1, \dots, e_n with e_i joining w_{i-1} to w_i and with $w_n = w_0$. We need to check that $t_0 = t_n$. Set $\gamma \in \pi_1(\Gamma)$ to be that element represented by the concatenation of the edges e_1, \dots, e_n , and note that

$$\begin{aligned} t_n - t_0 &= \sum_{i=1}^n N_2(e_i) - N_1(e_i) = \sum_{i=1}^n N_2(e_i) - \sum_{i=1}^n N_1(e_i) \\ &= (f \circ \varphi_2)_*(\gamma) - (f \circ \varphi_1)_*(\gamma) = 0. \end{aligned}$$

It follows that we may extend consistently across the edges of $C(X)$. Note also that this construction implies that the boundaries of the 2-cells $\{e\} \times [0, 1]$ are mapped null-homotopically into S^1 , so that we may extend “linearly” across these new 2-cells, completing the extension of f to \widehat{f} .

We have to argue that the new ascending and descending links are trees. We give the argument for \uparrow -links. The argument for \downarrow -links is analogous: just interchange the roles of φ_1 and φ_2 in the argument below, and use hypothesis 2(b) in place of hypothesis 2(a).

Let $v \in C^{(0)}(X)$. Since \widehat{f} is an extension of f , we have that $Lk_{\uparrow}(v, X) \subset Lk_{\uparrow}(v, C(X))$. Here the first \uparrow -link is with respect to f , and the second \uparrow -link is with respect to \widehat{f} . By the definition of the extension, any lift of \widehat{f} to the universal covers will attain a maximum on an edge $\{w\} \times [0, 1]$ at the $(w, 1)$ endpoint, and will attain a minimum at the $(w, 0)$ endpoint. Thus there are no contributions to $Lk_{\uparrow}(v, C(X))$ by the $(\varphi_2)_*$ -attached orthogonal cones. Let $w \in \Gamma$ be such that $\varphi_1(w) = v$. In the proof of part (1) we saw that this contributed an orthogonal cone on $(\varphi_1)_*(Lk(w, \Gamma))$ to $Lk(v, C(X))$. Of this, the contribution to $Lk_{\uparrow}(v, C(X))$ is simply the cone on the subset $(\varphi_1)_*(Lk_{\uparrow}(w, \Gamma))$. Here we think of φ_1 as defining a Morse function on Γ by composition with f .

To complete the proof, we claim that by coning off each of the sets $(\varphi_1)_*(Lk_{\uparrow}(w, \Gamma))$ as w ranges over all those vertices of Γ that are in $\varphi_1^{-1}(v)$, the components of the forest $Lk_{\uparrow}(v, X)$ are connected to produce a tree. For this we attach one $C(Lk_{\uparrow}(w, \Gamma))$ at a time to $Lk_{\uparrow}(v, X)$ and apply the following lemma. The precise argument is continued after the lemma.

Lemma 1.9. *If $f : P \rightarrow Q$ is a map of spaces so that $f_* : \tilde{H}_0(P) \rightarrow \tilde{H}_0(Q)$ is injective, and $\tilde{H}_1(Q) = 0$, then*

- (1) $\tilde{H}_1(C(P) \cup_f Q) = 0$, and
- (2) $\tilde{H}_0(C(P) \cup_f Q) = \tilde{H}_0(Q) / f_*(\tilde{H}_0(P))$

where $C(P)$ denotes the cone on P .

Proof of Lemma 1.9. This is just reduced Mayer-Vietoris. We have

$$\begin{aligned} \rightarrow \tilde{H}_1(C(P)) \oplus \tilde{H}_1(Q) &\rightarrow \tilde{H}_1(C(P) \cup_f Q) \rightarrow \tilde{H}_0(C(P) \cap Q) \rightarrow \\ &\tilde{H}_0(C(P)) \oplus \tilde{H}_0(Q) \rightarrow \tilde{H}_0(C(P) \cup_f Q) \rightarrow 0 \end{aligned}$$

Since $C(P)$ is contractible, all its reduced homologies vanish. By hypothesis, the reduced 1st homology of Q is trivial. Thus the first term on the left is trivial. The $\tilde{H}_0(C(P) \cap Q)$ term is just a copy of $\tilde{H}_0(P)$ and so injects (via f_*) into $\tilde{H}_0(Q)$ by hypothesis. Since this map injects, its kernel is trivial, and so we conclude that $\tilde{H}_1(C(P) \cup_f Q)$ is trivial. This proves (1).

The sequence now becomes

$$0 \rightarrow \tilde{H}_0(P) \rightarrow \tilde{H}_0(Q) \rightarrow \tilde{H}_0(C(P) \cup_f Q) \rightarrow 0$$

which gives us part (2). \square

Let $\{w_1, \dots, w_n\}$ be the set of vertices of Γ that are mapped to v by φ_1 . Start with $P = Lk_{\uparrow}(w_1, \Gamma)$ and $Q = Lk_{\uparrow}(v, X)$. Use Lemma 1.9 to determine the effect of attaching a cone on $Lk_{\uparrow}(w_1, \Gamma)$ to $Q = Lk_{\uparrow}(v, X)$. Continue to attach cones on the $Lk(w_i, \Gamma)$, using Lemma 1.9 at each stage

to determine that the resulting graphs all have trivial first homology, and that their reduced 0-homology groups are quotients of $\tilde{H}_0(Lk_\uparrow(v, X))$ by sums $\oplus_{1 \leq j \leq i} f_*(\tilde{H}_0(Lk_\uparrow(w_j, \Gamma)))$. The isomorphism hypothesis 2(a) in the definition of f -link filling ensures that we can continue applying Lemma 1.9 until the last cone has been attached, at which point the reduced 0-homology group will be trivial. This concludes the proof of Theorem 1.8. \square

2. QUASI-CONVEX SURFACE SUBGROUPS

Proof of Theorem 1. We start by taking the 2-complex X of Theorem 1.8 to be a closed hyperbolic surface. For example we can take X to be the surface of genus 2 obtained from the top 3 squares of Figure 1 by performing the indicated edge-pair identifications. The circle-valued Morse function is defined to take each oriented edge once around the target circle, and is extended linearly over the 2-cells of the surface X .




FIGURE 1. Hyperbolic free-by- \mathbb{Z} group with linearly distorted genus 2 surface subgroup.

We now perform the completion outlined in Theorem 1.8 so that the universal cover of the resulting complex $C(X)$ does not contain any isometrically embedded flat planes. This will ensure that the resulting group is hyperbolic.

This is done by taking the graph Γ of Theorem 1.8 to be the disjoint union of two copies of the wedge of two circles. The maps φ_i of Theorem 1.8 are defined as follows: φ_1 maps the first bouquet of two circles to the bouquet in $X^{(1)}$ labeled by the ordered pair (x, y) , and the second bouquet to the bouquet in $X^{(1)}$ labeled by (a, b) , while the map φ_2 takes the first bouquet to the bouquet in $X^{(1)}$ labeled by (z, x) and takes the second bouquet to the bouquet in $X^{(1)}$ labeled by (b, z) . This can be easily read off from the bottom 4 squares in Figure 1; identifying the vertical s and t edges gives the complex $\Gamma \times [0, 1]$, and the labels on the top and bottom bouquets define the maps φ_i . One can either check that the graph Γ and the two immersions φ_i satisfy the f -link connectivity condition of Theorem 1.8, or simply check directly that the ascending and descending links in the final 2-complex are both trees.

Now suppose the universal cover of $C(X)$ contains an isometrically embedded flat plane P . If $w \in \sigma \subset P$ is a vertex in a 2-cell of P we have the following inclusion of links

$$Lk(w, \sigma) \subset Lk(w, P) \subset Lk(w, \widetilde{C(X)}).$$

This means that P must be tiled by 2-cells σ with the property that for each vertex $w \in \sigma$ the segment $Lk(w, \sigma)$ lies in an embedded circuit of length 4 in $Lk(w, \widetilde{C(X)})$. This means that σ is a lift of a 2-cell of $C(X)$ with the same property; each corner of the 2-cell is contained in an embedded circuit of length 4 in $Lk(v, C(X))$. Here v is the vertex of $C(X)$. Note that $Lk(v, C(X))$ is just obtained from $Lk(v, X)$, which is a circle of length 12, by attaching 4 cones with cone vertices $\{s^\pm, t^\pm\}$, and so it is not hard to analyze. It is shown in Figure 2.




FIGURE 2. Link of the sole vertex $v \in C(X)$.

We can determine which 2-cells of $C(X)$ do not have the property above by determining which edges of $Lk(v, C(X))$ are not contained in embedded circuits of length 4. The only such edges are the edge from y^+ to c^- and the edge from z^- to c^- . Thus the plane P cannot contain any lifts of either the $bxc^{-1}y^{-1}$ 2-cell or of the $cza^{-1}z^{-1}$ 2-cell.

Now consider edge labels of the remaining five 2-cells:

$$ayb^{-1}x^{-1}, \quad sxs^{-1}z^{-1}, \quad sys^{-1}x^{-1}, \quad tbt^{-1}z^{-1}, \quad tat^{-1}b^{-1}.$$

Note that the edge labels a , y , and z each occur exactly twice among these five 2-cells. In particular, if the plane P contains a lift of the y -edge, then it must contain lifts of the $ayb^{-1}x^{-1}$ and the $sys^{-1}z^{-1}$ 2-cells. This means that it contains a lift of the a -edge, and therefore must also contain a lift of the $tat^{-1}b^{-1}$ 2-cell. We conclude that if the plane contains a lift of one of these three 2-cells it has to contain lifts of all three. Consider the vertex w at the

terminal endpoint of the y -edge. The only circuit of combinatorial length 4 in $Lk(v, C(X))$ that contains y^+ is $y^+s^+x^-b^+$. But the x^-b^+ edge implies that P must contain a lift of the illegal 2-cell $bxc^{-1}y^{-1}$, a contradiction. Thus, we have increased the list of illegal 2-cells types from two to five.

Finally, if P contains a lift of one of the last two 2-cells, $sxs^{-1}z^{-1}$ and $tbt^{-1}z^{-1}$, then it contains the other (since they share a z -edge), and so P contains a lift of an x -edge. Therefore, P must contain a lift of one of the five illegal 2-cells. This final contradiction implies that no such plane P exists in the universal cover of $C(X)$, and hence that $\pi_1(C(X))$ is hyperbolic. \square

The graph Γ and the maps φ_i were found by inspection, by analyzing the components of the ascending and descending links of the Morse function on the genus 2 surface, and seeing how they were arranged in the link of the surface. The bouquets and maps were chosen specifically in order to connect up components of the ascending and descending links, while being careful to preserve non-positive curvature and hyperbolicity. One could have chosen different hyperbolic completions $C(X)$ of the same squared surface X , or one could start with a different squared structure on a closed hyperbolic surface of possibly different genus, and find suitable f -filling graphs Γ . There are many possibilities.

3. POLYNOMIALLY DISTORTED SURFACE SUBGROUPS

In this section we prove Theorem 2. Here is an outline of the construction.

- (1) We begin with a CAT(0) squared complex with free-by-cyclic fundamental group, where the free kernel has polynomial distortion of degree n , say. In [3] there are examples of such groups where n is arbitrary.
- (2) Let K be the double cover of this complex, corresponding to the kernel of a map to \mathbb{Z}_2 . We glue K to a suitable surface group to form X , and we define a Morse function on X so that the ascending and descending links are forests.
- (3) We then find an f -link filling graph and apply Theorem 1.8 above. It will follow from the construction that the distortion of the surface group in $\pi_1(C(X))$ will be polynomial of degree n .

The explicit examples of [3] with which we begin are encoded by labeled oriented graphs (LOGs), so we begin by recalling the definition of LOG presentations.

Definition 3.1 (LOGs and LOTs). A *labeled, oriented graph*, or LOG, consists of a finite, directed graph with distinct labels on all the vertices, and oriented edge labels taken from the set of vertex labels. An LOG defines a finite presentation as follows. The set of generators is in 1–1 correspondence with the set of vertex labels, and the set of relators is in 1–1 correspondence with the set of edges, so that an edge labeled a oriented from vertex u to vertex v corresponds to a conjugation relation $ava^{-1} = u$. In case the graph is a tree we call it a *labeled, oriented tree* or LOT and call the corresponding finite presentation an LOT presentation.

Note that the presentation 2-complex of an LOG group is a squared complex. In the examples that follow, the squared complexes will be non-positively curved.

Proof of Theorem 2.

(1): The free-by-cyclic group, and the complex K .

We begin with CAT(0) groups of the form $F_n \rtimes_{\varphi} \mathbb{Z}$ where the F_n subgroup has polynomial distortion of degree n . We describe specific examples as follows (from [3]). Let $\varphi : F_n \rightarrow F_n$ be the automorphism that is defined on the free basis $\{a_1, \dots, a_n\}$ as follows:

$$\varphi(a_j) = a_j a_{j-1} \quad (j \text{ odd}), \quad \varphi(a_j) = a_{j-1} a_j \quad (j \text{ even})$$

where a_0 denotes the identity element 1.




FIGURE 3. LOG presentation of $F_n \rtimes_{\varphi} \mathbb{Z}$, and link of vertex in associated squared 2-complex.

Now consider the group $F_n \rtimes_{\varphi} \mathbb{Z}$ where $\mathbb{Z} = \langle \tau \rangle$. Define α_i by the following

$$\alpha_i a_i = \tau \quad (\text{for } i \text{ odd}), \quad \text{and } \alpha_i = a_i \tau \quad (\text{for } i \text{ even}).$$

Then $F_n \rtimes_{\varphi} \mathbb{Z}$ is isomorphic to the group with the LOG presentation shown in Figure 3. The graph has $n + 1$ vertices labeled τ and α_i for $1 \leq i \leq n$. The vertex α_n is isolated, there is a loop labeled α_1 at τ , and directed edges α_i from α_{i-1} to τ for $2 \leq i \leq n$. Figure 3 also shows the link of the single vertex in the presentation 2-complex. This link is large, and so the 2-complex is non-positively curved. We define a circle-valued Morse function on this complex by sending each generator around the circle once in the positive direction and extending linearly over the 2-cells. This labels the vertices with + or - as shown in Figure 3.

In [3] it is shown that the distortion of the free group F_n in $F_n \rtimes_{\varphi} \mathbb{Z}$ is polynomial of degree n . Define K to be the double cover of this square complex corresponding to the kernel of the homomorphism $F_n \rtimes_{\varphi} \mathbb{Z} \rightarrow \mathbb{Z}_2$. Thus K is a non-positively curved squared complex with fundamental group $F_n \rtimes_{\varphi^2} \mathbb{Z}$, and the F_n subgroup also has polynomial distortion of degree n in $\pi_1(K)$. Note that K has two vertices, v_1 and v_2 say, and the graph, Θ , consisting of directed edges $\tau, \alpha_1, \dots, \alpha_n$ from v_1 to v_2 will correspond to the F_n subgroup. We let $\hat{\tau}, \hat{\alpha}_1, \dots, \hat{\alpha}_n$ denote the corresponding edges of K directed from v_2 to v_1 .

(2): The hyperbolic surface Σ , and the complex $X = \Sigma \cup_{\Theta} K$.




FIGURE 4. The non-positively curved, squared surface Σ_n .

Note that the $2(n + 1)$ squares in Figure 4 glue together to form a non-positively curved squared structure on a closed hyperbolic surface Σ_n . All the topmost $2(n + 1)$ vertices are identified to one vertex, called v_1 , and all the bottommost $2(n + 1)$ vertices are identified to a second vertex, called v_2 . The α_i edges and the τ edge form a copy of the graph Θ in the surface Σ_n . The squared complex X is formed by gluing the surface to the complex K by identifying the two copies of the graph Θ .




FIGURE 5. Ascending and descending links of v_1 in X .

It is easy to check that X is non-positively curved, and that the complex K locally isometrically embeds into X . The key point to note in checking this is that the links $Lk(v_i, \Sigma_n)$ are both circles of combinatorial length $4n + 4$. Moreover if p, q are distinct vertices in $Lk(v_i, K)$ that are also in $Lk(v_i, \Sigma_n)$, then their combinatorial distance in $Lk(v_i, \Sigma_n)$ is at least 4 (actually is a positive integral multiple of 4). Thus when we identify corresponding vertices in $Lk(v_i, \Sigma_n)$ and $Lk(v_i, K)$ to get $Lk(v_i, X)$, the resulting links are still large, and so X is non-positively curved. Moreover, $Lk(v_i, K)$ will be convex subcomplexes, in the CAT(1) sense, of $Lk(v_i, X)$ and so $K \rightarrow X$ is a local isometric embedding.

Note that the surface does not locally isometrically embed in X , but that $\pi_1(X)$ is the free product with amalgamation of the surface group with $\pi_1(K)$, amalgamated along the $F_n = \pi_1(\Theta)$ subgroup. Therefore, the surface group will have polynomial distortion of degree n in $\pi_1(X)$.

The circle-valued Morse function on K is extended to all of X by mapping the directed edges of the squares in Figure 4 once around the target circle positively, and extending linearly over the 2-cells.




FIGURE 6. Ascending and descending links of v_2 in X .

Figures 5 and 6 show the ascending ($-$) and descending ($+$) links of the two vertices of X . We see that $Lk_-(v_1, X)$ and $Lk_+(v_2, X)$ are both trees, but that the remaining two links are forests with $(n + 2)$ components each.




FIGURE 7. The graph Γ in the application of Theorem 1.8.

(3): The completion of X to $C(X)$.

We now choose an f -link filling graph Γ and attach $\Gamma \times [0, 1]$ as in Theorem 1.8 to obtain the complex $C(X)$ with $CAT(0)$ free-by-cyclic fundamental group into which X locally isometrically embeds. Thus $\pi_1(X)$ will be undistorted in $\pi_1(C(X))$, and so the surface subgroup of $\pi_1(X)$ will still

have polynomial distortion of degree n in the CAT(0) free-by-cyclic group $\pi_1(C(X))$.

We claim that the graph Γ shown in Figure 7 works for any choice of i, j, k and m with $i \neq k, j + 1$. We define a map $\varphi : \Gamma \rightarrow X^{(1)}$ according to the labels and orientations on the edges in Figure 7, and set $\varphi_1 = \varphi_2 = \varphi$ in constructing $C(X)$. It is easy to check that φ is a local isometric embedding. We now show that the resulting ascending and descending links are trees.

We let t_1, \dots, t_4 denote the edges of $\Gamma \times [0, 1]$ reading from left to right in Figure 7. Each ascending and descending link in Figures 5 and 6 receives two new vertices, labeled and attached as follows:

To $Lk_{\uparrow}(v_1, X)$ we add a new vertex t_1^- , joined by an edge to x_{i-1}^- , and a new vertex t_2^- , joined to x_j^- . To $Lk_{\downarrow}(v_1, X)$ we add a new vertex t_1^+ , joined to x_j^+ , and a new vertex t_2^+ , joined to $x_{i-1}^+, \widehat{\alpha}_m^+$, and all the β^+ except β_i^+ .

To $Lk_{\uparrow}(v_2, X)$ we add a new vertex t_4^- , joined to y_k^- , and a new vertex t_3^- , joined to $y_i^-, \widehat{\alpha}_m^-$, and all the β^- except β_i^- . To $Lk_{\downarrow}(v_2, X)$ we add a new vertex t_3^+ , joined to y_k^+ , and a new vertex t_4^+ , joined to y_i^+ .

It is easy to check that the new ascending and descending links are all trees. □

4. EXPONENTIALLY DISTORTED SURFACE SUBGROUPS

In this section we prove Theorem 3. We do this by constructing a non-positively curved squared complex with hyperbolic, free-by-cyclic fundamental group, that contains an exponentially distorted closed surface group. The outline of the construction is similar to that of the previous section.

- (1) We begin with a non-positively curved squared complex with hyperbolic $F_7 \rtimes \mathbb{Z}$ fundamental group. Note that the fiber F_7 is exponentially distorted in $F_7 \rtimes \mathbb{Z}$.
- (2) As in the previous example, we next glue a double cover K of this complex to a suitably chosen closed surface to obtain X . Note that the surface group injects into $\pi_1(X)$ by the standard theory of free products with amalgamation. Moreover, we choose the gluing, essentially identifying the F_7 fiber with a highly convex F_7 in the surface group, so that the distortion of the surface group in $\pi_1(X)$ is exponential.
- (3) Finally, we apply Theorem 1.8 to obtain a non-positively curved 2-complex $C(X)$ so that $X \rightarrow C(X)$ is a local isometric embedding. As in the proof of Theorem 1, the completion is done with care to ensure that $\pi_1(C(X))$ is hyperbolic. In our construction $\pi_1(C(X))$ is isomorphic to $F_{65} \rtimes \mathbb{Z}$.

Here are the details.

Proof of Theorem 3.

(1): The free-by-cyclic group, and the complex K .

We start with the LOT group corresponding to the labeled oriented tree on the left hand side of Figure 8. The link of the vertex in the corresponding presentation 2-complex is shown on the right side of Figure 8. From this we




FIGURE 8. LOT presentation and link of vertex in presentation 2-complex.

see that the LOT group is CAT(0). As in the proof of Theorem 1, if P is an isometrically embedded flat plane in the universal cover of the presentation 2-complex, and if $w \in \sigma \subset P$ is a vertex of a square σ in P , then the segment $Lk(w, \sigma)$ is contained in the embedded length 4 circuit, $Lk(w, P)$, inside the link of w in the universal cover. Again, links in the universal cover are isomorphic to the link of the single vertex v in the LOT complex. The only circuits of length 4 in this link contain a_0^- and a_0^+ as opposite vertices. Thus each 2-cell in P contains at least two adjacent corners of the form Xa_0^{\pm} , where $X \in \{a_1^-, a_2^-, a_3^-, a_4^-, a_5^+, a_6^-, a_7^-\}$, and thus two corners of the form \overline{XY} , where $\overline{X} \in \{a_1^+, a_2^+, a_3^+, a_4^+, a_5^-, a_6^+, a_7^+\}$. Such corners correspond to segments not contained in any embedded circuits of length four in the ambient link. Therefore, the universal cover of the presentation 2-complex does not contain any flat planes, and so the LOT group is hyperbolic. Essentially, the largest pieces of flat planes that exists in the universal cover are strips of width 2 squares which are centered on translates of the a_0 -axis.

There is a circle-valued Morse function on the presentation 2-complex of this group, defined by mapping each directed edge once around the target circle. The ascending and descending links are trees; they are the sublinks of the link in Figure 8 spanned by the negative and by the positive vertices respectively. Proposition 1.5 implies that the LOT group is free-by-cyclic, and an examination of the rank of the vertex preimage shows that it is in fact $F_7 \rtimes \mathbb{Z}$. Finally, since the ambient group is hyperbolic, the F_7 subgroup is exponentially distorted in the LOT group.

The LOT group can be mapped onto \mathbb{Z}_2 by sending each a_i to the non-trivial element of \mathbb{Z}_2 . The complex K is the double cover of the presentation 2-complex corresponding to the kernel of this epimorphism. Its fundamental group is also hyperbolic of the form $F_7 \rtimes \mathbb{Z}$ where the F_7 subgroup is exponentially distorted. The 1-skeleton of K has two vertices, v_1 and v_2 , and the subgraph $\Theta \subset K^{(1)}$ consisting of directed edges a_0, \dots, a_7 from v_1 to v_2 has fundamental group isomorphic to this F_7 . We let $\hat{a}_0, \dots, \hat{a}_7$ denote the corresponding edges directed from v_2 to v_1 .

(2): The hyperbolic surface Σ , and the complex $X = \Sigma \cup_{\Theta} K$.




FIGURE 9. The non-positively curved, squared surface Σ .

Identifying the a_0 edge pairs, the x_i edge pairs, and the y_i edge pairs in Figure 9 produces a closed hyperbolic surface Σ with two vertices. Let v_1 be the vertex where the y_i loops are based, and v_2 be the vertex where the x_i loops are based. The surface Σ contains a copy of the graph Θ . Identifying the surface Σ with the complex K along the copies of Θ gives a non-positively curved squared complex $X = \Sigma \cup_{\Theta} K$ with hyperbolic fundamental group. To verify this, it is sufficient to check that $Lk(v_1, X)$ and $Lk(v_2, X)$ are both large, and then to check that \tilde{X} does not contain any isometrically embedded flat planes. To this end, note that because the links $Lk(v_i, K)$ are large, any short loop in $Lk(v_i, X)$ must contain a segment of $Lk(v_i, \Sigma)$. But each $Lk(v_i, \Sigma)$ is a circle with combinatorial length 32 with the distance between any two a_i^{\pm} at least four (see Figure 10). Thus the links $Lk(v_i, X)$ are large. Moreover, the only circuits of length four come from those in K . As the original presentation 2-complex of the LOT group (which is double covered by K) has no flat planes in its universal cover, it follows that $\pi_1(X)$ is hyperbolic.




FIGURE 10. Pieces of the surface links $Lk(v_1, \Sigma)$ and $Lk(v_2, \Sigma)$.

Note that the surface group $\pi_1(\Sigma)$ is a subgroup of the hyperbolic group $\pi_1(X)$, since $\pi_1(X)$ is obtained from $\pi_1(\Sigma)$ via free product with amalgamation. Note also that this surface subgroup is exponentially distorted in $\pi_1(X)$, since the amalgamating free subgroup $\pi_1(\Theta)$ is quasi-convex in $\pi_1(\Sigma)$ and is exponentially distorted in $\pi_1(K)$ and hence in $\pi_1(X)$. All that remains is to apply Theorem 1.8, choosing Γ carefully so that $\pi_1(C(X))$ is hyperbolic.

(3): The completion of X to $C(X)$.

Figure 11, without edge labels and orientations, shows $\Gamma \times [0, 1]$, with $\Gamma \times 0$ at the bottom and $\Gamma \times 1$ at the top. The maps φ_1 and φ_2 are defined as indicated by the edge labels and orientations on $\Gamma \times 0$ and $\Gamma \times 1$, respectively. The complex $C(X)$ is the union of X with $\Gamma \times [0, 1]$ glued together using the maps φ_1 and φ_2 . Hereinafter, all subscripts are to be interpreted modulo 8. There are several conditions that must be satisfied by the various unspecified

subscripts involved in the definition of the maps φ_i . For the present, we note that there are precisely four choices of triple (z, p, q) that work: $(3, 6, 7)$, $(4, 7, 3)$, $(6, 2, 1)$, and $(7, 2, 6)$. The conditions imposed on the choices of subscripts for the x and y edges are more easily met. In particular, the construction will work as long as j, k, n, m are distinct and different both from z and from $z - 1$.




FIGURE 11. The complex $\Gamma \times [0, 1]$.

One may now check that these maps φ_1 and φ_2 are f -link filling. It is also possible to check directly that the conclusions of Theorem 1.8 hold for $C(X)$. We proceed as follows: describe the extension of the Morse function, show that the new ascending and descending links are all trees, argue that the complex $C(X)$ has large links, and that the subcomplex X locally isometrically embeds, and finally argue that $\pi_1(C(X))$ is hyperbolic.

Morse function extension. Here is an explicit description of a Morse function $\hat{f} : C(X) \rightarrow S^1$ which extends f . We need to extend the function over the 2-cells of $\Gamma \times [0, 1]$. This will be by the usual “linear” extension, once we have defined the Morse function on the 1-cells of $\Gamma \times [0, 1]$. We have already defined the function on the 1-cells in $\Gamma \times 0$ and $\Gamma \times 1$ since these are identified with 1-cells of X , and f is already defined on X . So we only have to define the function \hat{f} on the six 1-cells of the form $w \times [0, 1]$, where w is a vertex of Γ . Label these edges as $r_1, r_2, s_1, s_2, t_1, t_2$, as indicated in Figure 11, and orient each edge from the 0 end to the 1 end. The map \hat{f} sends each t_i edge once around the target circle, sends s_2 and each r_i three times around the target circle, and sends s_1 five times around the target circle. Thus from the perspective of the Morse function \hat{f} , the 2-cells of $\Gamma \times [0, 1]$ are parallelograms and trapezoids, as shown in Figure 12. The unlabeled edges

in the center trapezoid of Figure 12 range over $\{b_{z+1}, \dots, b_7, \widehat{a}_q, b_0, \dots, b_{z-1}\}$ for the upper edge, and over $\{b_z, \dots, b_7, \widehat{a}_p, b_0, \dots, b_{z-2}\}$ for the lower edge.




FIGURE 12. The Morse function \widehat{f} on the new 2-cells of $C(X)$.

Ascending and descending links. The ascending and descending links with respect to \widehat{f} are described as follows. Start with the ascending and descending links of v_1 and v_2 in the 2-complex X . These are drawn in Figure 13 for v_1 and in Figure 14 for v_2 .




FIGURE 13. Descending and ascending links of v_1 in X .

To form the corresponding ascending and descending links in $C(X)$, we add three new vertices to each of these links, labeled and attached as follows:

In $Lk_{\uparrow}(v_1, C(X))$, the three new vertices are: r_1^- , which is joined by a new edge to y_m^- ; r_2^- , which is joined by an edge to y_n^- ; and s_2^- , which is joined by an edge to a_z^- . In $Lk_{\downarrow}(v_1, C(X))$, the three new vertices are: s_1^+ , which is joined to both \widehat{a}_q^+ and all the b_i^+ ; t_1^+ , which is joined to y_k^+ ; and t_2^+ , which is joined to y_j^+ .

In $Lk_{\uparrow}(v_2, C(X))$, the new vertices are: t_2^- , joined to x_k^- ; t_1^- , joined to x_j^- ; and s_1^- , joined to both \widehat{a}_p^- and all the b_i^- . In $Lk_{\downarrow}(v_2, C(X))$, the new vertices are: r_2^+ , joined to x_m^+ ; r_1^+ , joined to x_n^+ ; and s_2^+ , joined to a_z^+ .

Clearly all ascending and descending links in $C(X)$ are trees.




FIGURE 14. Ascending and descending links of v_2 in X .

Large links and local isometric embedding. Now we verify that the links of the vertices v_i in $C(X)$ are large. We have already seen that the links of the v_i in X are large. From the proof of Theorem 1.8, we know that we only have to check that the maps φ_i are local isometric embeddings. To check that φ_1 is a local isometric embedding at v_1 we need to verify that the following two subsets of $Lk(v_1, X)$ are π -separated:

$$\{y_m^-, y_n^+\}, \{y_m^+, y_n^-, b_{z-1}^+\}, \{a_z^-, b_0^+, \dots, b_{z-2}^+, b_z^+, \dots, b_7^+, \widehat{a}_p^+\},$$

and to check that φ_1 is a local isometric embedding at v_2 we need to verify that the following subsets of $Lk(v_2, X)$ are π -separated:

$$\{b_0^-, \dots, b_7^-, \widehat{a}_p^-\}, \{a_z^+, x_j^-, x_k^+\}, \{x_j^+, x_k^-\}.$$

Likewise, to see that φ_2 is a local isometric embedding at v_1 we need to verify that the following subset of $Lk(v_1, X)$ is π -separated:

$$\{b_0^+, \dots, b_7^+, \widehat{a}_q^+\}, \{a_z^-, y_j^-, y_k^+\}, \{y_j^+, y_k^-\},$$

and to check that φ_2 is a local isometric embedding at v_2 we need to verify that the following two subsets of $Lk(v_2, X)$ are π -separated:

$$\{x_n^+, x_m^-\}, \{x_n^-, x_m^+, b_z^-\}, \{a_z^+, b_0^-, \dots, b_{z-1}^-, b_{z+1}^-, \dots, b_7^-, \widehat{a}_q^-\}.$$

As long as we choose $z \neq j \neq k \neq z - 1$ and $z \neq n \neq m \neq z - 1$, and choose p and q so that neither (\widehat{a}_p^+, a_z^-) nor (\widehat{a}_q^-, a_z^+) are adjacent pairs in the appropriate $Lk(v_i, K)$, these sets will be π -separated in the links both of K and of Σ . It follows that they are π -separated also in the appropriate links of X .

The fundamental group of $C(X)$ is hyperbolic. It remains only to show that $\pi_1(C(X))$ is hyperbolic. Equivalently, we have to check that the universal cover of $C(X)$ does not contain any isometrically embedded flat planes. For this it is helpful to have a more complete description of the links in $C(X)$. Each $Lk(v_i, C(X))$ is the union of $Lk(v_i, K)$, $Lk(v_i, \Sigma)$, and edges coming from the corners of squares in $\Gamma \times [0, 1]$. Figure 8 shows the links in K (replacing a_i^+ with \widehat{a}_i^+ in the link of v_1 , while a_i^- becomes \widehat{a}_i^- in the link of v_2). Figure 10 shows $Lk(v_i, \Sigma)$, while the new portions coming from $\Gamma \times [0, 1]$ are shown in Figure 15, where the unlabeled vertices are, on the left, all the b_i^+ except b_{z-1}^+ , and on the right, all the b_i^- except b_z^- .




FIGURE 15. The portions of $Lk(v_1, C(X))$ and $Lk(v_2, C(X))$ contributed to by $\Gamma \times [0, 1]$.

We first show that no squares lifted from Σ can be contained in any flat plane P . Suppose σ is a square in P covering a square in Σ . The square in Σ covered by σ has either two x_i sides or two y_i sides. Consider first the case that this square has two x_i sides, and let w be the corner of σ shared by the lifts of these two sides in σ . Thus $Lk(w, P)$ is a circuit of length four with vertices labeled $x_i^+ x_i^- AB$ for some labels A and B . We will show that there are no choices of A and B that can complete this circuit. Figure 16 shows the possible choices for A and B , as they depend on i . We need to show that none of the vertices on the far right in Figure 16 is adjacent to x_i^+ in $Lk(v_2, C(X))$. From Figures 10 and 15 we see that x_i^+ is not adjacent to a_i^+ , b_z^- , or any other x_j^+ . We can rule out a_z^+ by choosing $j \neq z - 1$.

The case that σ contains two sides labeled y_i is similar. If the purported length four circuit in P is of the form $y_i^- y_i^+ AB$, one finds that possibilities for B are other y_i^- , b_{z-1}^+ , or a_{i+1}^- (none of which is adjacent to y_i^-), or possibly a_z^- , in the case $i = k$. Thus by choosing $k \neq z$, we ensure there are




FIGURE 16. Finding a circuit around a vertex in P .

no length four circuits of this form. This rules out the presence of squares lifted from Σ in any purported flat plane P .

Now consider the possibility that P contains squares lifted from K and from $\Gamma \times [0, 1]$. In particular, there is some edge of P along which a square from K meets a square from $\Gamma \times [0, 1]$. This edge must be a lift either of a_z , \hat{a}_p , or \hat{a}_q . In any case, this edge joins lifts of v_1 and of v_2 . Let v be the lift of v_1 on this edge. Then $Lk(v, P)$ is a circuit of length four in $Lk(v_1, C(X))$, with vertices labeled, say $XYZW$, where $X \in \{a_z^-, \hat{a}_p^+, \hat{a}_q^+\}$. Note that there must be some other edge incident to v lying between a square lifted from K and a square lifted from $\Gamma \times [0, 1]$. In particular, one of Y , Z , or W is also labeled a_z^-, \hat{a}_p^+ , or \hat{a}_q^- . If we choose z , p , and q so that these vertices all have combinatorial distance at least three in $Lk(v_1, C(X))$, it will follow that P cannot contain squares lifted from K and from $\Gamma \times [0, 1]$. There are precisely four choices for the triple (z, p, q) satisfying this, and all previously imposed, conditions. They are $(3, 6, 7)$, $(4, 7, 3)$, $(6, 2, 1)$, and $(7, 2, 6)$.

Because $\pi_1(K)$ is hyperbolic, no flat plane can consist entirely of squares lifted from K . Thus with any of the four allowed choices for (z, p, q) , we find that any flat plane P consists entirely of squares lifted from $\Gamma \times [0, 1]$. To see that such a plane cannot exist, we first note that it cannot contain any of the four squares of $\Gamma \times [0, 1]$ with x or y sides, as each side of this type occurs exactly once in $\Gamma \times [0, 1]$, so that it is not possible to extend P across such an edge. It follows then that no square with an r side can be contained in P , as such a square must be adjacent across its r edge to a square with an x or y edge. Thus P consists entirely of squares from the “middle” section of $\Gamma \times [0, 1]$. Fix a point u in the interior of such a square, and consider a lift in P of the “vertical” segment $u \times [0, 1] \subset \Gamma \times [0, 1]$. Because P consists entirely of squares from $\Gamma \times [0, 1]$, this segment lifts to a bi-infinite line in P . On the other hand, moving in the positive direction along this line, one successively crosses successive b_i edges before finally intersecting \hat{a}_q , at which point the line cannot be extended in $\Gamma \times [0, 1]$. It follows that no such plane P exists, so $\pi(C(X))$ is hyperbolic. \square

Remark 4.1. We mentioned in the introductory paragraph to this section that the ambient free-by-cyclic group in this example is of the form $F_{65} \rtimes \mathbb{Z}$. Here is one way to see this. From Proposition 1.5 we know that the rank of the free kernel is equal to the rank of the preimage of the basepoint of




FIGURE 17. Circle vertex preimages in the 2-cells of $\Gamma \times [0, 1]$ (cf. Figure 12).

the target circle. This preimage is a graph with 12 vertices: v_1, v_2 , two points dividing the edges r_1, r_2, s_2 into thirds, and four points dividing s_1 into fifths. This graph inherits 14 edges from the subcomplex K (diagonals of squares), 16 edges from the surface Σ (diagonals of squares), and a total of 46 from the 2-cells of $\Gamma \times [0, 1]$ (see Figure 17). This makes a total of 76 edges, and so an Euler number of $12 - 76 = -64 = 1 - 65$. Therefore, the graph has rank 65.

5. QUESTIONS

We have seen how to construct hyperbolic free-by-cyclic groups containing exponentially distorted surface subgroups. Do these groups also contain quasi-convex surface subgroups?

Question 5.1. Does every $\text{CAT}(0)$ free-by-cyclic group that contains a non-linearly distorted closed hyperbolic surface subgroup, also contain a quasi-convex closed hyperbolic surface subgroup?

The following related question is also open.

Question 5.2. Does every hyperbolic free-by-cyclic group contain a closed hyperbolic surface subgroup?

This latter question is a special case of the question (due to Gromov) of whether all one-ended hyperbolic groups contain closed hyperbolic surface subgroups.

Finally, we note that all the free-by-cyclic groups which we construct in this paper are $\text{CAT}(0)$. It is tempting to think that Theorem 1.8 could be used to construct $\text{CAT}(0)$ structures for all hyperbolic free-by-cyclic groups. However, John Crisp and Noel Brady [2] have found infinitely many hyperbolic free-by-cyclic groups that are not $\text{CAT}(0)$ in dimension 2. It is not known if these groups are $\text{CAT}(0)$. The following question still remains open.

Question 5.3. Is every hyperbolic free-by-cyclic group $\text{CAT}(0)$?

REFERENCES

- [1] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.*, 129(3):445–470, 1997.

- [2] Noel Brady and John Crisp. Cat(0) and cat(-1) dimensions of torsion free hyperbolic groups. *Preprint*, 2005.
- [3] Noel Brady, Max Forester, and Krishnan Shankar. Dehn functions of subgroups of cat(0) groups. *Preprint*, 2005.
- [4] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [5] James Howie. On the asphericity of ribbon disc complements. *Trans. Amer. Math. Soc.*, 289(1):281–302, 1985.

DEPT. OF MATHEMATICS, UNIVERSITY OF OKLAHOMA, NORMAN, OK 73019
E-mail address: `jbarnard@math.ou.edu`

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