CAT(-1) STRUCTURES FOR FREE-BY-FREE GROUPS

Noel Brady and Andy Miller

Abstract. We give a procedure for constructing CAT(-1) free-by-free groups which uses tree valued Morse theory and branched coverings of 2-complexes.

Introduction

There are two main paradigms in modern geometric group theory. The more classically inspired of the two involves studying lattices in the isometry group of some interesting geometric space. The spaces one typically considers are CAT(0) and CAT(-1) spaces, which are respectively generalizations of simply connected manifolds of nonpositive and negative sectional curvatures. If a group acts cocompactly and discretely by isometries on some CAT(0) space, then one can deduce a lot of topological and algebraic information about the group. For example, the group is FP, and the centralizers of elements are also FP.

The second paradigm was first propounded by M. Gromov in [Gr1] and involves defining interesting classes of groups by placing restrictions on their large scale (or *coarse*) geometry. In [Gr2] Gromov introduced the class of hyperbolic groups, by placing restrictions on the large scale geometry of geodesic triangles in these groups.

These two paradigms give us two different perspectives on negative curvature. The coarse viewpoint of hyperbolicity has the advantage of being a quasi-isometry invariant. It also interacts well with other quasi-isometry invariant definitions such as isoperimetric inequalities. If a group has a sub-quadratic isoperimetric inequality, then it has a linear isoperimetric inequality and is in fact hyperbolic – see [Gr2], [Bo], [P], and [Ol]. This enabled Bestvina and Feighn [BF] to prove a combination theorem for hyperbolic groups. The CAT(-1) viewpoint is useful for constructing concrete examples of hyperbolic groups as the fundamental groups of finite nonpositively curved piecewise hyperbolic complexes. One advantage of working with such complexes is that there are only finitely many local conditions to check for nonpositive curvature.

The first author received partial support from the NSF grant DMS-9704417.

It turns out that a group which acts cocompactly and discretely by isometries on a CAT(-1) space is hyperbolic. It is natural to ask (and Gromov did in [Gr2]) if hyperbolic groups are CAT(-1). This question remains unanswered even when one restricts to special subclasses of hyperbolic groups, such as free-by-cyclic hyperbolic groups or free-by-free hyperbolic groups. In fact, it is not obvious how to construct any examples of CAT(-1) free-by-cyclic or CAT(-1) free-by-free groups. The existence of hyperbolic free-by-cyclic groups and hyperbolic free-by-free groups follows from the results in [BF] and [BFH]. The first examples of CAT(-1) free-by-cyclic groups were constructed in [McW]. The first examples of CAT(-1) free-by-free groups are given in sections 3 and 4 of this paper.

Trying to produce CAT(-1) free-by-free groups directly by glueing two CAT(-1) free-by-cyclic groups together appears to be quite difficult. To see why, recall that the fiber subgroup in a hyperbolic free-by-cyclic group is exponentially distorted. This means that given a free-by-cyclic group which is the fundamental group of a nonpositively curved piecewise hyperbolic complex, the fiber subgroup can not be carried by a convex subcomplex. Therefore, it is not clear how one might try to glue two such free-by-cyclic examples together to get a free-by-free one.

This paper offers a procedure for constructing many examples of CAT(-1) free-by-free groups (and free-by-cyclic groups). Here are the main results.

Theorem A. *For each $n \geq 2$ there exist short exact sequences of the form*

$$1 \to K \to G \to F_n \to 1$$

where G is the fundamental group of a compact, nonpositively curved piecewise hyperbolic 2-complex, and K is finitely generated but not finitely presented.

Theorem B. *There exist CAT(-1) free-by-free groups.*

Note that the groups G in Theorem A are not coherent. Dani Wise [W] gave the first examples of non-coherent CAT(-1) groups. Some of his examples were of the form

$$1 \to K \to G \to \mathbb{Z} \to 1$$

where K is finitely generated but not finitely presented. Theorem A gives examples with higher rank free quotients.

Here is how the construction works. Both theorems assert the existence of short exact sequences with a CAT(-1) middle group, and a free quotient. In Theorem A the kernel is finitely generated but not finitely presented, and in Theorem B the kernel is free of finite rank. We shall realize these sequences by constructing compact, nonpositively curved piecewise hyperbolic 2-complexes and maps to finite rank graphs. The 2-complexes for Theorem A are constructed first of all, and the 2-complexes for Theorem B are taken to be special subcomplexes of these.

We shall use tree valued Morse theory to determine the structure of the kernel groups. The finiteness properties of the kernel will follow from connectivity properties of *e-links* (which are generalizations of ascending/descending links [BB] to the tree valued setting). In the case of Theorem A we will have connected e-links which have nontrivial H_1, and in the case of Theorem B we will require that the e-links be trees. We shall see in section 2 that connectivity of the e-links ensures finite generation of the kernels, that nontrivial 1-cycles in the e-links produce a non-finitely presented kernel, and that contractible 1-complex e-links give us a free kernel.

The construction of the 2-complexes themselves involves a 2-step process. First, we construct some model 2-complexes and maps to finite graphs with rank at least two which satisfy all the e-link conditions we want, but which are not nonpositively curved. Next we introduce nonpositive curvature by pre-composing the maps with finite branched covering maps. It is well known that one can make 2-complexes nonpositively curved by taking appropriate branched coverings [Ger]. However, in our case we have to be careful to preserve connectivity of the e-links, which makes constructing an appropriate branched cover that bit harder. The problem of preserving connectivity of e-links while constructing a branched cover was also a concern in [Bra].

This paper is organized as follows. In the first section we review the basic concepts of nonpositive curvature and hyperbolic groups. In section 2 we define tree valued Morse functions on simplicial complexes, and prove some results which relate the connectivity properties of point pre-images to the connectivity properties of the e-links. In the case that the Morse function is a lift to the universal covers of a map from a nonpositively curved complex to a finite rank graph, these connectivity properties are translated into finiteness properties of the kernel of the corresponding exact sequence in π_1. In section 3 we show that it is possible to construct finite branched coverings while preserving the connectivity properties of the e-links for a broad class of model 2-complexes. Then we give the proofs of Theorem A and Theorem B. In section 4 we give some explicit examples of the groups in the main theorems, and point out some of the difficulties we have encountered in constructing CAT(-1) free-by-free groups.

§1 Nonpositive curvature and hyperbolic groups

In this section we discuss various CAT notions of nonpositive curvature for metric spaces due to Alexandroff. One way of constructing these CAT(0) and CAT(-1) metric spaces is by constructing simply connected nonpositively curved piecewise euclidean and piecewise hyperbolic complexes. We also define hyperbolic groups, and give a precise statement of the result quoted in the introduction that a CAT(-1) group is hyperbolic. This material can be found in [GdH] and in the forthcoming book [BH]. Although we only need the results about piecewise hyperbolic 2-complexes for the rest of the paper we state them in more generality here. The reader is invited to consider higher dimensional analogs of the construction in section 3. Here is a particularly challenging example: can these ideas be used to produce an example of a CAT(-1) surface-by-free group? Lee Mosher [Mo] has shown that there exist hyperbolic groups which are (closed) surface-by-free.

We begin by recalling the notions of a geodesic metric space, and of geodesic triangles in a metric space. Let x, y be two points in a metric space X. A *geodesic* from x to y is an isometric embedding $g : [0, d(x,y)] \to X$ such that $g(0) = x$ and $g(d(x,y)) = y$. If no ambiguity is likely (for example in the case of unique geodesics) we denote the image of g by $[xy]$. A metric space X is said to be a *geodesic metric space* if any two points can be joined by a geodesic. A *geodesic triangle* $\triangle xyz$ in a metric space X consists of three points x, y and z in X together with three geodesic segments connecting the points.

Given a geodesic triangle $\triangle xyz$ in a metric space X, a *comparison triangle* for $\triangle xyz$ in a model space M^2_χ is a geodesic triangle $\triangle x'y'z' \subset M^2_\chi$, such that

$$d_X(x,y) \,=\, d(x',y')\,, \quad d_X(x,z) \,=\, d(x',z')\,, \text{ and } d_X(z,y) \,=\, d(z',y')$$

where d_X denotes distance in the metric space X and d denotes distance in M^2_χ. Typically, one takes the model space to be one of the classical 2-dimensional geometries: $M^2_0 = \mathbb{E}^2$, $M^2_{-1} = \mathbb{H}^2$, or $M^2_1 = S^2$. We say that the geodesic triangle $\triangle xyz$ satisfies a CAT(χ) inequality ($\chi = 0, \pm 1$) if for all points $p \in [yz]$ and comparison points $p' \in [y'z']$ (that is $d_X(y,p) = d(y',p')$) we have

$$d_X(x,p) \;\leq\; d(x',p')\,.$$

For $\chi = 0$ or 1, a complete, locally compact, geodesic metric space X is said to be a CAT(χ) space if every geodesic triangle in X satisfies the CAT(χ) inequality. Also X is said to be a CAT(1) space if every geodesic triangle in X of perimeter at most 2π satisfies the CAT(1) inequality. Note that triangles in these spaces are at least as thin as triangles in the corresponding model geometries.

The CAT(0) and CAT(-1) spaces are generalizations of Hadamard manifolds from riemannian geometry. It is of great interest to geometric group theorists to study discrete cocompact groups of isometries of such spaces. These groups are termed CAT(0) and CAT(-1) groups respectively. A practical way of producing such groups is to take the fundamental group of a nonpositively curved piecewise euclidean or piecewise hyperbolic cell complex. We will use the latter method to produce our examples for Theorem A and Theorem B. First, we give some definitions which are relevant to these constructions.

Let M^n_χ denote the n-dimensional model geometries; euclidean for $\chi = 0$, hyperbolic for $\chi = -1$, and spherical for $\chi = 1$. A *convex cell* in the model space M^n_χ is a compact intersection of a finite number of half spaces. Let v be a vertex of a convex n-cell e in M^n_χ. By the *link of v in e* (denoted by $Lk(v,e)$) we mean the set of unit tangent vectors at v which point into e. It is a spherical $(n-1)$-cell.

A *piecewise* M_χ cell complex is a regular cell complex which is obtained from a disjoint collection of convex M_χ cells by identifying faces via isometries. Let v be a vertex in a piecewise M_χ cell complex K. The *link of v in K* is denoted by $Lk(v,K)$, and is defined as the union

$$\bigcup\{Lk(v,e) \mid \; e \text{ a cell of } K \text{ which contains } V\}\,.$$

It inherits a natural piecewise spherical structure from the M_χ structure on K. We use this piecewise spherical structure to define local nonpositive curvature on K as follows. A piecewise euclidean or piecewise hyperbolic cell complex K is said to be *nonpositively curved* if the link of every vertex of K is CAT(1).

The first result of this section gives a concrete way of producing CAT(0) and CAT(-1) metric spaces.

Theorem 1.1. (see [GdH] ch 10, Th 7 and 15) *Let K be a nonpositively curved piecewise euclidean (resp. piecewise hyperbolic) cell complex. Then the universal cover of K is a CAT(0) (resp. CAT(-1)) metric space.*

The examples in this paper will all be piecewise hyperbolic 2-complexes. In this case, the links of vertices are just finite graphs, and so the CAT(1) condition reduces to the following: the lengths of loops in the links of vertices of K must be at least 2π. In the case the cell complex is composed entirely of small equilateral hyperbolic triangles (angles are close to $\pi/3$) the nonpositive curvature condition becomes a simple combinatorial one. Loops in the links of vertices must have combinatorial length at least seven. This fact will be used in sections 3 and 4.

Finally, we state the relationship between CAT(-1) groups and hyperbolic groups in the sense of Gromov. Recall that a geodesic metric space X is said to be *δ-hyperbolic* if there exists a constant $\delta \geq 0$ such that every geodesic triangle in X is δ-thin: each edge is contained in the δ-neighborhood of the union of the other two edges. A finitely generated group is said to be *hyperbolic* (in the sense of Gromov) if it's Cayley graph with respect to some finite generating set is a δ-hyperbolic metric space.

Theorem 1.2. (see [GdH] ch 1, Th 29 and ch 10, Rk 2) *Suppose G is a CAT(-1) group. Then G is hyperbolic in sense of Gromov.*

Combining the results 1.1 and 1.2, we see that the fundamental group of a compact, nonpositively curved, piecewise hyperbolic cell complex is Gromov hyperbolic.

§2 Tree valued Morse functions

We begin with the definition of a $\triangle$*-complex*. This is similar to the notion in [Ha], except that we do not keep track of orientations on the simplices. Basically, a $\triangle$*-complex* is a cell complex obtained from a disjoint collection of simplices of various dimensions, by glueing certain faces together by linear homeomorphisms. Here is a more precise definition.

Definition 2.1. A $\triangle$*-complex* is a cell complex with the following additional structure.

(1) For each cell $c \subset X$ we are given a standard simplex $\triangle_c \subset \mathbb{R}^n$ and a characteristic map $f_c : \triangle_c \to c$. An *admissible* characteristic map for c is any map obtained from f_c by pre-composing with a (partial) linear homeomorphism of $\mathbb{R}^n$.

(2) If $f_c : \triangle_c \to X$ is the characteristic map for a cell c, then the restriction of f_c to any face of $\triangle_c$ is an admissible characteristic map for another cell. Here we are assuming that for $m < n$ we identify $\mathbb{R}^m$ with the first m coordinates of $\mathbb{R}^n$.

Let $\tau \subset \sigma$ be a face of a standard simplex σ in $\mathbb{R}^n$. Recall that the *link of* τ *in* σ is defined to be the face of σ which is spanned by all the vertices of σ which are not in τ. We shall modify this definition to include links of cells in a $\triangle$-complex. Since a $\triangle$-complex is not necessarily a simplicial complex, we will have to be a bit careful.

Definition 2.2. Let σ be a standard simplex in $\mathbb{R}^n$, with vertices $\{v_0, \dots, v_s\}$, and let τ be the face of σ which is spanned by the vertices $\{v_0, \dots, v_t\}$. The *link of* τ *in* σ is defined as follows. For each $i \geq t+1$ set

$$w_i \;=\; \frac{2}{3}\frac{(v_0 + \cdots + v_t)}{t+1} \;+\; \frac{v_i}{3}\,.$$

Thus, w_i divides the line segment from the barycentre of τ to the vertex v_i into a 1/3: 2/3 ratio. Now define

$$Lk(\tau, \sigma) \;=\; \sum_{i=t+1}^{s} t_i w_i \qquad \text{where} \qquad \sum_{i=t+1}^{s} t_i = 1\,.$$

Note that this link can be identified with the usual link (defined in previous paragraph) by radial projection from the barycentre of the face τ.

Now, suppose that c is a cell in a $\triangle$-complex X. Let Σ_c be the collection of cells b of X such that the attaching map $f_b : \triangle_b \to X$ restricts to a face $\tau \subset \triangle_b$ to give an admissible characteristic map for the cell c. The *link of* c *in* X is denoted by $Lk(c, X)$, and is defined to be the $\triangle$-complex obtained from the disjoint union

$$\bigcup \{\, Lk(\tau, \triangle_b) \mid b \in \Sigma_c \,\}$$

by identifying faces whenever their f_b images are the same.

Note that this definition makes sense since the glueing maps of X are linear homeomorphisms. Note also that in the case X is a simplicial complex, $Lk(c, X)$ can be identified with the usual simplicial link, by radial projection from the barycentre of c.

Now we are in a position to describe the setup for tree-valued Morse theory.

Definition 2.3. A *ϕ-equivariant tree-valued Morse function* $f : \tilde{X} \to T$ consists of the following data.

(1) A simplicial complex $\tilde{X}$ and a free, properly discontinuous, cocompact, cellular G action on $\tilde{X}$.
(2) A tree T with a free, properly discontinuous, cocompact, cellular F_n action.
(3) An epimorphism $\phi : G \to F_n$
(4) A ϕ-equivariant, simplicial map $f : \tilde{X} \to T$.

Since f is simplicial, we see that cells of $\tilde{X}$ will either map to a vertex or an edge in T.

Definition 2.4. Let $f : \tilde{X} \to T$ be a simplicial map from a simplicial complex $\tilde{X}$ to a tree T. A cell $c \subset \tilde{X}$ is said to be *horizontal* if $f(e)$ is a vertex of T.

Now we define certain important subcomplexes of the link of a horizontal cell of $\tilde{X}$.

Definition 2.5. Suppose that $f : \tilde{X} \to T$ is given as in 2.4 above. Let $c \subset \tilde{X}$ be a horizontal cell, and let $e \subset T$ be an (oriented) edge such that $i(e) = f(c)$. The *e-link of c in $\tilde{X}$* is denoted by $Lk_e(c, \tilde{X})$, and is defined to be the full subcomplex of $Lk(c, \tilde{X})$ which contains the vertices which map to $\tau(e)$ under f. Note that in the case where the quotient of T consists of just a single vertex and one oriented edge labeled a, and $v \in \tilde{X}$ is a vertex, then

$$Lk_a(v, \tilde{X}) \;=\; Lk_{\uparrow}(v, \tilde{X}) \qquad \text{and} \qquad Lk_{\overline{a}}(v, \tilde{X}) \;=\; Lk_{\downarrow}(v, \tilde{X})\,.$$

where $\overline{a}$ denotes the edge a with the opposite orientation.

Our examples of equivariant Morse functions arise as lifts to universal covers of *linear maps* $l : X \to \Gamma$ from a finite nonpositively curved piecewise euclidean or piecewise hyperbolic Δ-complex to a finite, oriented graph Γ. A *linear map* between Δ-complexes K_1 and K_2 is map $l : K_1 \to K_2$ which sends cells to cells, and is linear on each cell; ie. for each cell $e \subset K_1$ there is a linear map λ such that

$$l \circ f_e \;=\; f_{l(e)} \circ \lambda$$

where f_e and $f_{l(e)}$ are the characteristic maps for the cells e and $l(e)$ respectively. Note that a linear map between simplicial complexes is just a simplicial map.

Now, definitions 2.4 and 2.5 can be extended to cells in the quotient complex X in this case. A simplex $\sigma \subset X$ is said to be *horizontal* if it lifts to a horizontal simplex in $\tilde{X}$. Equivalently, $\sigma \subset X$ is said to be horizontal if $l(\sigma)$ is a vertex in Γ. We can talk about e-links in the quotient as follows. Note that the covering map gives an isomorphism between the link of a simplex $\sigma \subset \tilde{X}$ and the link of its image in X. This allows us to define e-links of horizontal cells in X. It is also possible to define e-links directly using the map $l : X \to \Gamma$.

In the next definition we give sufficient conditions which control the connectivity properties of point preimages under the Morse function f. We let $|\sigma|$ denote the dimension of the simplex σ.

Definition 2.6. We say that $f : \tilde{X} \to T$ satisfies *condition* (C_n) if

$$Lk_e(\sigma, \tilde{X}) \text{ are all } (n - |\sigma|)\text{-connected for every horizontal simplex } \sigma \subset \tilde{X}\,.$$

We say that $f : \tilde{X} \to T$ satisfies *condition* (HC_n) if

$$Lk_e(\sigma, \tilde{X}) \text{ are all homologically } (n - |\sigma|)\text{-connected for every horizontal simplex } \sigma \subset \tilde{X}\,.$$

We adopt the convention that "(-1)-connected" means "nonempty", and that the $(-k)$-connected conditions are vacuously true for all $k \geq 2$.

The next two results relate the conditions (C_n) and (HC_n) to the connectivity properties of $f^{-1}(t_0)$. They are both *local to global* results, translating local conditions on e-links into (global) connectivity properties of f point preimages. Once they are established, we will then be able to read off the finiteness properties of $Ker(\phi)$ in Corollary 2.9.

The first result below only uses the facts that f is simplicial, and that $\tilde{X}$ is sufficiently highly connected. It can be viewed as a positive it local to global result, which uses the local connectivity properties, (HC_n) and (C_n), to conclude (global) connectivity properties for the point preimages of f.

Theorem 2.7. *Let $f : \tilde{X} \to T$ be a simplicial map from a contractible simplicial complex $\tilde{X}$ to a tree T, and let $t_0 \in T$ be a vertex. Then the following are true.*

(1) *If condition (HC_n) holds, then $f^{-1}(t_0)$ is homologically n-connected.*
(2) *If condition (C_1) holds, then $f^{-1}(t_0)$ is 1-connected.*

Proof. The proof proceeds as follows. Pick a basepoint $t_0 \in T$, and let $B_r(t_0)$ and $S_r(t_0)$ denote the ball and sphere of radius r (where $r \geq 0$ is an integer) about t_0 in T respectively. Note that $S_r(t_0)$ consists of a discrete set of vertices. We shall determine how the topology of the preimages $f^{-1}(t_0)$ change as $r \to \infty$. This information together with the contractibility of $f^{-1}(B_\infty(t_0)) = f^{-1}(T) = \tilde{X}$ will enable us to deduce the connectivity of $f^{-1}(t_0)$.

The main step is to determine the change in topology in passing from $f^{-1}(B_{r-1}(t_0))$ to $f^{-1}(B_r(t_0))$. We shall build up $f^{-1}(B_r(t_0))$ from $f^{-1}(B_{r-1}(t_0))$ by attaching simplices which map to $B_r(t_0)$ but not to $B_{r-1}(t_0)$ in an order determined by the dimension of (necessarily horizontal) simplices in the subcomplex $f^{-1}(S_r(t_0))$. Here are the details.

(1) Begin by attaching the simplices of $v * Lk_e(v, \tilde{X})$ where $v \in f^{-1}(S_r(t_0))$ and $e \subset T$ is an edge from $f(v)$ to a point in $S_{r-1}(t_0)$. Note that the net effect of attaching these simplices is to cone off copies of $Lk_e(v, \tilde{X}) \subset f^{-1}(B_{r-1}(t_0))$ for vertices $v \in f^{-1}(S_r(t_0))$. Call the resulting complex Y_0.
(2) Now, for an integer $k \geq 0$, suppose that we have obtained a complex Y_k by attaching the simplices of $\sigma * Lk_e(\sigma, \tilde{X})$ for all simplices $\sigma \subset f^{-1}(S_r(t_0))$ of dimension at most k, and where $e \subset T$ connects $f(\sigma)$ to $S_{r-1}(t_0)$. Now, attach all the simplices of $\sigma * Lk_e(\sigma, \tilde{X})$ for each $(k+1)$-simplex $\sigma \subset f^{-1}(S_r(t_0))$ to obtain a complex Y_{k+1}. Again, $e \subset T$ connects $f(\sigma)$ to $S_{r-1}(t_0)$. Note that this is equivalent to coning off copies of $\partial\sigma * Lk_e(\sigma, \tilde{X}) \subset Y_k$ where the cone points are the barycenters of the $(k+1)$-simplices $\sigma \subset f^{-1}(S_r(t_0))$.

We finish with $Y_m = f^{-1}(B_r(t_0))$ where m is the maximal dimension of simplices $\sigma \subset f^{-1}(S_r(t_0))$. Figure 1 shows the effect of attaching the simplices of $\sigma * Lk_e(\sigma, \tilde{X})$ for a horizontal 1-simplex σ. In this figure, $Lk_e(\sigma, \tilde{X})$ consists of three points.

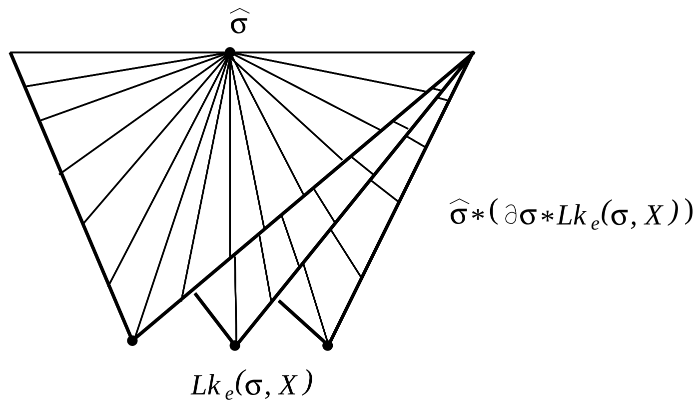

FIGURE 1. The cone structure of $\widehat{\sigma} * \partial\sigma * Lk_e(\sigma, \tilde{X})$ for a horizontal 1-simplex σ.

From the procedure above, we see that $f^{-1}(B_r(t_0))$ is obtained from $f^{-1}(B_{r-1}(t_0))$ by coning off copies of

$$\partial\sigma * Lk_e(\sigma, \tilde{X})$$

for all simplices $\sigma \subset \tilde{X}$ such that $f(\sigma) \in S_r(t_0)$. Suppose that condition (HC_n) holds. Then each $Lk_e(\sigma, \tilde{X})$ is homologically $(n - |\sigma|)$-connected, and so each $\partial\sigma * Lk_e(\sigma, \tilde{X})$ is homologically n-connected, since it is homotopy equivalent to the $|\sigma|$-times iterated suspension of $Lk_e(\sigma, \tilde{X})$. Thus, the Mayer-Vietoris sequence implies that the inclusion map

$$f^{-1}(B_{r-1}(t_0)) \;\rightarrow\; f^{-1}(B_r(t_0))$$

induces an isomorphism in homology up through dimension n and an epimorphism in $(n+1)$st homology. Similarly, if condition (C_1) holds, then the Van Kampen theorem and Mayer-Vietoris sequence imply that the inclusion map above induces an isomorphism in H_0 and in π_1.

Now, we use the contractibility of $\tilde{X}$ to conclude that condition (HC_n) implies that $f^{-1}(t_0)$ is homologically n-connected, and condition (C_1) implies that $f^{-1}(t_0)$ is 1-connected. Note, that the Hurewicz theorem will tell us that conditions (HC_n) and (C_1) together imply that $f^{-1}(t_0)$ is n-connected. $\square$

It is a much harder problem in general to determine when the failure of one of the conditions (HC_n) or (C_n) implies that $f^{-1}(t_0)$ fails to be (homologically) n-connected. One situation where results are possible is when one can make *top dimensional* homology arguments. If $\tilde{X}$ is an n-dimensional complex and $\sigma \subset \tilde{X}$ is a horizontal simplex, then the $\partial\sigma * Lk_e(\sigma, \tilde{X})$ are $(n-1)$-complexes. The following theorem shows when the failure of the *top dimensional* connectivity properties of these complexes translates into the failure of top dimensional connectivity properties of $f^{-1}(t_0)$. It can be viewed as a negative local to global result. Note that we need to be in the (cocompact) equivariant setting for this result.

Theorem 2.8. *Let $f : \tilde{X} \to T$ be a ϕ-equivariant Morse function as in Definition 2.3. Let $\tilde{X}$ be a contractible n-complex, and let $t_0 \in T$ be a base vertex. If f satisfies conditions (HC_{n-2}) but not condition (HC_{n-1}), then $f^{-1}(t_0)$ is homologically $(n-2)$-connected, but $H_{n-1}(f^{-1}(t_0), \mathbb{Z})$ is not finitely generated as a $\mathbb{Z}Ker(\phi)$-module.*

Proof. The homological $(n-2)$-connectivity of $f^{-1}(t_0)$ was established in Theorem 2.7. In fact, it was established that $f^{-1}(B_r(t_0))$ is homologically $(n-2)$-connected for every integer $r \geq 0$, and that the inclusion maps

$$f^{-1}(B_r(t_0)) \to f^{-1}(B_s(t_0)) \qquad \text{(for integers } 0 \leq r \leq s\text{)}$$

induce epimorphisms in H_{n-1}.

We shall derive a contradiction from the finite generation of $H_{n-1}(f^{-1}(t_0), \mathbb{Z})$ over $\mathbb{Z}Ker(\phi)$ as follows. First we show that there exists a positive integer r such that

$$H_{n-1}(f^{-1}(B_r(t_0)), \mathbb{Z}) = 0\,.$$

Then we use the failure of condition (HC_{n-1}) to establish that there exists a nontrivial n-cycle in the contractible n-complex $\tilde{X}$ (in fact this cycle is found in $f^{-1}(B_{r+d}(t_0))$ where d is the diameter of the quotient graph Γ). Here are the details.

Suppose that $H_{n-1}(f^{-1}(t_0), \mathbb{Z})$ is finitely generated as a $\mathbb{Z}Ker(\phi)$-module. Thus there exists a finite number of $(n-1)$-cycles, $z_1, \ldots, z_l$, whose $Ker(\phi)$ translates generate $H_{n-1}(f^{-1}(t_0), \mathbb{Z})$. Since $\tilde{X}$ is contractible, each of the z_i bounds an n-chain c_i. Now there exists a positive integer r such that the chains $c_1, \ldots, c_l$ are all contained in $f^{-1}(B_r(t_0))$. Thus, all of the $Ker(\phi)$ translates of the c_i will be contained in $f^{-1}(B_r(t_0))$, and so the inclusion induced map

$$H_{n-1}(f^{-1}(t_0), \mathbb{Z}) \to H_{n-1}(f^{-1}(B_r(t_0)), \mathbb{Z})$$

is trivial. But it is also an epimorphism. Thus, $H_{n-1}(f^{-1}(B_r(t_0)), \mathbb{Z}) = 0$.

Now let d be the diameter of the quotient Γ of T. Since condition (HC_{n-1}) fails, there exists a horizontal simplex $\sigma \subset f^{-1}(B_{r+d}(t_0)) - f^{-1}(B_r(t_0))$ with $H_{n-|\sigma|-1}(Lk_e(\sigma, \tilde{X}), \mathbb{Z}) \neq 0$. Now $H_{n-1}(f^{-1}(B_{f(\sigma)-1}(t_0)), \mathbb{Z}) = 0$. Recall from Theorem 2.7 that we build up $f^{-1}(B_{f(\sigma)}(t_0))$ from $f^{-1}(B_{f(\sigma)-1}(t_0))$ by constructing intermediate complexes Y_k, labeled by the maximum dimension of the horizontal simplices that we have attached. The conditions (HC_{n-2}) ensure that $H_{n-1}(Y_k, \mathbb{Z}) = 0$. Now take a nontrivial $(n-1)$-cycle in the complex $\partial\sigma * Lk_e(\sigma, \tilde{X})$. This bounds an n-chain $c_1 \subset Y_k$ by the above homology computation. It also bounds the n-chain

$$c_2 = \hat{\sigma} * (\partial\sigma * Lk_e(\sigma, \tilde{X}))$$

where $\hat{\sigma}$ denotes the barycentre of the horizontal simplex σ. Note that this chain does not lie in Y_k, and so

$$z = c_1 - c_2$$

is a nontrivial n-cycle in the contractible n-complex $\tilde{X}$, giving our contradiction. $\square$

Corollary 2.9. *Let $l : X \to \Gamma$ be a linear map from a finite n-dimensional Δ-complex X to a finite graph Γ with nontrivial π_1. Suppose that the universal cover $\tilde{X}$ is a contractible simplicial complex, and that l satisfies the condition (HC_{n-2}) but not the condition (HC_{n-1}). Then the kernel of the map*

$$l_* : \pi_1(X) \to \pi_1(\Gamma)$$

is of type FP_{n-1} but not FP_n.

Proof. The hypotheses ensure that l lifts to a tree valued Morse function f on the universal covering $\tilde{X}$ which satisfies the condition (HC_{n-2}) but not the condition (HC_{n-1}). Theorems 2.7 and 2.8 imply that the point preimages $f^{-1}(t)$ are homologically $(n-2)$-connected and that $H_{n-1}(f^{-1}(t), \mathbb{Z})$ is not finitely generated as a $\mathbb{Z}Ker(l_*)$-module. Now $Ker(l_*)$ acts cocompactly by restriction of deck transformations on $f^{-1}(t)$. It follows (from the definition of FP_m) that $Ker(l_*)$ is of type FP_{n-1}, and it follows from Schanuel's Lemma (see [Bro] Lemma 4.4) that $Ker(l_*)$ is not of type FP_n. □

§3 The construction

In this section we will give proofs for Theorems A and B. As we mentioned in the introduction, we start with some model 2-complexes and then take branched coverings. In order to preserve connectivity of the e-links we have to be careful about how we take the branched coverings. This is the point of Theorem 3.4. We build up to this result with some background material on graph coverings. First, we have an algebra lemma which will be used to choose cyclic covers in Lemma 3.2.

Lemma 3.1. *Let A be a finitely generated free abelian group and let X be a finite subset of $A - \{0\}$. Then there is a homomorphism $\gamma : A \to \mathbb{Z}$ such that $\gamma(x) \neq 0$ for each $x \in X$.*

Proof. There are a number of ways of proving this. Here is a geometric proof. Endow $\mathbb{R}^n$ with the standard euclidean inner product, and identify A with the integer lattice $\mathbb{Z}^n \subset \mathbb{R}^n$. Note that the conclusion of the lemma follows once we establish the existence of a *rational* vector $v \in \mathbb{R}^n$ with the property that the orthogonal projection of $\mathbb{R}^n$ onto $Span(v)$ does not map any element of X to 0.

Now, in order that $x \in X$ does not map to 0 we have to choose v to not lie in the orthogonal complement of $Span(x)$. This complement is a closed codimension 1 subspace of $\mathbb{R}^n$. Finally, we have to verify that there is a rational vector v in the complement of a finite union of such spaces. But this is trivial since the union is a closed proper subspace of the topological space $\mathbb{R}^n$ and the rational points are dense in $\mathbb{R}^n$. □

Let D be a finite graph, and let $M(D)$ denote the length of a shortest (homotopically) nontrivial edgepath loop in D.

Lemma 3.2. *Let D be a finite connected graph which is not a tree, and let $\mathcal{L}$ be a collection of connected subgraphs of D which are not trees. Then there is a finite cyclic covering*

$g : \widetilde{D} \to D$ *with* $M(\widetilde{D}) \geq 2M(D) + 2$ *so that* $g^{-1}(L)$ *is connected for each* $L \in \mathcal{L}$.

Proof. Suppose that p is a loop in D which is homologically trivial but not homotopically trivial. The image of p is a connected subgraph $P \subset D$, and, since its fundamental group can't be abelian, P is neither homotopy equivalent to a point nor to a circle. Then there is a nontrivial simple loop in P and an edge outside that loop, so P has at least $M(D)+1$ distinct unoriented edges. Being homologically trivial, p must traverse each unoriented edge of P at least twice, once in each direction. Thus the length of p is no less than $2M(D) + 2$. This shows that any homologically trivial loop in D with length less than $2M(D) + 2$ must also be homotopically trivial.

Now there are finitely many loops in D which are homotopically nontrivial and have length less than $2M(D) + 2$. Let $\alpha_1, \ldots, \alpha_l$ be the elements of $H_1(D)$ which these loops represent. By the above, none of the α_i's is trivial in $H_1(D)$. Let $\mathcal{L} = \{L_1, \ldots, L_k\}$. For each $j = 1, \ldots, k$, we choose a nontrivial simple closed curve in L_j and denote by β_j the (nontrivial) element of $H_1(D)$ which it represents. By lemma 3.1 there is a homomorphism $\gamma : H_1(D) \to \mathbb{Z}$ for which $\gamma(\alpha_i) \neq 0$ when $i = 1, \ldots, l$ and $\gamma(\beta_j) \neq 0$ when $j = 1, \ldots, k$. Choose a positive integer N which is relatively prime to each of the $\gamma(\beta_j)$'s and does not divide any of the $\gamma(\alpha_i)$'s. Now let $g : \widetilde{D} \to D$ be the cyclic covering corresponding to the epimorphism $\widetilde{\gamma} : \pi_1(D) \to H_1(D) \to \mathbb{Z} \to \mathbb{Z}/N\mathbb{Z}$. By construction, the image under $\widetilde{\gamma}$ of each loop in D with length less than $2M(D) + 2$ is nontrivial. It follows that no loop in $\widetilde{D}$ has length less than $2M(D) + 2$. Similarly the image of each $\pi_1(L_j)$ under $\widetilde{\gamma}$ is $\mathbb{Z}/N\mathbb{Z}$ since the image of $\gamma(\beta_j)$ generates $\mathbb{Z}/N\mathbb{Z}$. □

Lemma 3.3. *Let D be a finite connected graph which is not a tree, and let $\mathcal{L}$ be a collection of connected subgraphs of D which are not trees. For any integer $M > 0$ there is a finite covering $g : \widetilde{D} \to D$ so that $M(\widetilde{D}) \geq M$ and $g^{-1}(L)$ is connected for each $L \in \mathcal{L}$.*

Proof. Choose n so that $2^n M(D) + 2^{n+1} - 2 \geq M$. Let $D_0 = D$. Applying lemma 3.2 inductively yields cyclic coverings $g_i : D_i \to D_{i-1}$ with $M(D_i) \geq 2M(D_{i-1}) + 2$ for $i = 1, \ldots, n$. Put $\widetilde{D} = D_n$ and let g be the covering $g_1 \circ \cdots \circ g_n : \widetilde{D} \to D$. Then $M(\widetilde{D}) \geq 2^n M(D) + 2^{n+1} - 2 \geq M$. □

Let X and Y be 2-dimensional Δ-complexes. A linear map $g : Y \to X$ is a *branched covering* if its restriction to $Y - Y^{(0)}$ is a covering map from $Y - Y^{(0)}$ onto $X - X^{(0)}$. Note that a branched covering g preserves the dimensions of simplices. In addition, g maps the link of each simplex σ in Y onto the link of $g(\sigma)$ in X. If σ is a vertex of Y then this map on links is a covering map, and if the dimension of σ is larger than 0 then it is an isomorphism.

Let X be a finite, connected 2-dimensional Δ-complex and assume that the links of the vertices of X are connected. The link of the 0-skeleton $X^{(0)}$ is denoted by L_X, and its components are the links of the individual vertices of X. Let X_1 be the complement in X of an open regular neighborhood of X_0. Then X_1 is a deformation retract of $X - X^{(0)}$ which contains L_X. We define D_X to be the graph which is the full subcomplex of the barycentric subdivision X' spanned by the barycenters of the 1- and 2-dimensional simplices of X. There is a canonical deformation retraction from X_1 onto D_X. So each covering map π over D_X uniquely determines a covering map π over X_1, and this covering map over X_1 in turns extends to a branched covering map over X by attaching a cone to each component of $\pi^{-1}(L_X)$ and extending π conically. This is the procedure which we will use in theorem 3.4 to construct branched coverings. This same construction was used in [G] and more details can be found in §2 of that paper.

There is a natural graph morphism η_X from the barycentric subdivision L'_X onto D_X. It is a graph immersion and so induces a monomorphism on fundamental groups. See figure 2.

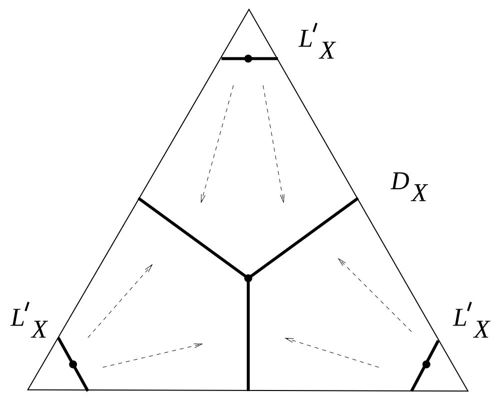

FIGURE 2. A view of the morphism η_X inside a 2-cell of X.

Theorem 3.4. *Let X be a finite, connected 2-dimensional Δ-complex for which the links of vertices are connected. Let $l : X \to \Gamma$ be a linear map onto a graph Γ. Suppose that l satisfies condition (HC_0) and that each e-link $Lk_e(v, X)$ of each vertex v of X is not a tree. There is a finite branched covering $g : Y \to X$ where Y is a negatively curved piecewise hyperbolic complex and $l \circ g$ satisfies condition (HC_0). We may choose the piecewise hyperbolic structure on Y so that each 2-simplex is a hyperbolic equilateral triangle.*

Proof. Recall that a linear map from a 2-dimensional Δ-complex to a graph as described in the statement of the theorem satisfies condition (HC_0) if and only if (i) each e-link of a horizontal 1-simplex is nonempty and (ii) each e-link of a vertex is 0-connected.

Let $g : Y \to X$ be any branched covering. A 1-simplex in Y is horizontal with respect to $l \circ g$ if and only if its image under g is a horizontal 1-simplex in X with respect to l. Let σ be either a vertex in Y or a horizontal 1-simplex in Y. In either case, g restricts to a covering map from $Lk(\sigma, Y)$ onto $Lk(g(\sigma), X)$, and so it also restricts to a covering map from $Lk_e(\sigma, Y)$ onto $Lk_e(g(\sigma), X)$ for any edge e of Γ which originates at the vertex

$l(g(\sigma))$. Since l satisfies condition (HC_0), $Lk_e(g(\sigma), X)$ is nonempty, and so $Lk_e(\sigma, Y)$ is nonempty also. From this it follows that $l \circ g$ will satisfy condition (HC_0) if and only if $Lk_e(\sigma, Y)$ is connected for each vertex σ in Y. This property will automatically hold if, for example, $g^{-1}(Lk_e(v, X))$ is connected for each vertex v in X and edge e of Γ originating at $l(v)$.

Consider the graph D_X, which is connected since X is connected, and not a tree since it contains the barycentric subdivisions of the e-links of vertices which are not trees by hypothesis. Let $L_1, \dots, L_k$ be the subgraphs of D_X which are the images under η_X of the barycentric subdivisions of $Lk_e(v, X)$ as v ranges through the vertices of X and e ranges through the edges in Γ originating at $l(v)$. None of the L_j's are trees by hypothesis. We now apply lemma 3.3 to obtain a covering map $\widetilde{g} : \widetilde{D} \to D_X$ where $M(\widetilde{D}) > 12$ and $\widetilde{g}^{-1}(L_j)$ is connected for each j. By the process described above, this covering map extends to a branched covering $g : Y \to X$ where Y is a 2-dimensional Δ-complex with $D_Y = \widetilde{D}$. If L_Y contains a nontrivial simple loop with length no larger than six then the image under η_Y is a homotopically nontrivial loop in D_Y with length no larger than 12, but this is impossible since $M(D_Y) > 12$. We conclude that $M(L_Y) > 6$ and that Y has a negatively curved piecewise hyperbolic structure where the 2-simplices are realized by *small* equilateral hyperbolic triangles (angles are close to $\pi/3$). If v is a vertex of X and e is an edge of Γ originating at $l(v)$ then $g^{-1}(Lk_e(v, X))$ is connected its barycentric subdivision is one of the $\widetilde{g}^{-1}(L_j)$'s. It follows that $l \circ g$ satisfies condition (HC_0), and this completes the proof. $\square$

The assumption that each e-link of each vertex of X is not a tree is essential in Theorem 3.4. For example, if some vertex v has one e-link which is a tree and another which contains a nontrivial loop with length less than seven then $l \circ g$ will fail the (HC_0) condition for any branched covering $g : Y \to X$ where $M(L_Y) \geq 7$.

Now we describe the procedure which we will employ for constructing linear maps l satisfying the hypotheses of Theorem 3.4.

Procedure 3.5. Given any finite connected graph Γ consider the graph of spaces where each vertex space is a circle, each edge space is the wedge of two circles, and the morphisms from an edge space to its two incident vertex spaces are the quotient maps collapsing down one circle of the wedge in one direction and the other circle in the other direction. The total space X_Γ of this graph of spaces, as described for example in [SW], can be given an obvious 2-dimensional Δ-complex structure. If l is the natural projection

$$l : X_\Gamma \to \Gamma$$

then it is easy to verify that

(1) there are no horizontal 2-simplices in X_Γ,
(2) the e-links of horizontal 1-simplices are single points, and
(3) the e-link of each vertex of X_Γ is a circle.

In particular, l satisfies condition (HC_0), and the even stronger assumption of Theorem 3.4. The specific examples analyzed and depicted in §4 below will be of this type.

Remark 3.6. It is clear that Procedure 3.5 above can be generalized considerably. For example, given the finite connected graph Γ one could consider the following graph of spaces. The edge and vertex spaces are finite connected graphs satisfying the following conditions. The maps from the edge graphs to each of the two (endpoint) vertex graphs are required to be quotient maps with the following additional property: each edge of the edge graph will be killed (mapped to a vertex) under exactly one of the two quotient maps. This ensures that the resulting graph of spaces inherits a natural Δ-complex structure. One has a lot of freedom to build in interesting link conditions, by choosing the edge graphs and the vertex (quotient) graphs.

Let X be a nonpositively curved piecewise euclidean or piecewise hyperbolic Δ-complex and let $\widetilde{X}$ be its universal cover. If a simplex is not embedded in $\widetilde{X}$ then there is a 1-simplex not embedded in $\widetilde{X}$ which can be subdivided into a geodesic triangle which contradicts the CAT(0) inequality. If two simplices intersect in more than one face then there is a geodesic bigon which also contradicts the CAT(0) inequality. This argument shows that the Δ-complex $\widetilde{X}$ is actually a simplicial complex, and this will allow us to apply the results of section 2.

We are now ready for the proofs of Theorems A and B.

Proof of Theorem A. Let Γ be a finite, connected graph which is not a tree. Choose a linear map $l : X \to \Gamma$ from a finite, connected 2-dimensional Δ-complex X to Γ such that condition (HC_0) is satisfied and each e-link $Lk_e(v, X)$ of each vertex v of X is not a tree. (For example, we might choose $l : X_\Gamma \to \Gamma$ as described above.) By Theorem 3.4 there is a finite branched covering $g : Y \to X$ so that Y is a nonpositively curved piecewise hyperbolic complex and $l \circ g$ satisfies condition (HC_0). Let w be a vertex of Y. The restriction of g to an e-link of w (with respect to $l \circ g$) is a covering map onto an e-link of $g(w)$ (with respect to l). Since g is finite-to-one and the e-link of $g(w)$ is not a tree, the e-link of w cannot be a tree. Thus $l \circ g$ fails condition (C_1) since condition (C_1) requires the e-links of vertices to be 1-connected.

From this, we obtain a short exact sequence

$$1 \to K \to \pi_1(Y) \to F_n \to 1$$

where $\pi_1(\Gamma) = F_n$, Y is a compact, nonpositively curved piecewise hyperbolic complex, and the kernel is finitely generated but not finitely presented by Corollary 2.9. $\square$

Proof of Theorem B. The proof will proceed exactly as the proof of Theorem A, but then we will restrict to a special subcomplex $\widehat{Y}$ of the branched cover complex Y. Here are the details.

Let Γ be a finite, connected graph which is not a tree. Put $X = X_\Gamma$ and let $l : X \to \Gamma$ be the natural projection as described in Procedure 3.5. Then X is a finite, connected 2-dimensional Δ-complex and l satisfies condition (HC_0). Moreover, there are no horizontal 2-simplices in X, the preimage under l of each vertex in Γ is a circle with one vertex (these are the horizontal 1-simplices of X), and each e-link $Lk_e(v, X)$ of each vertex v of X is also a circle with one vertex. By Theorem 3.4 there is a finite sheeted branched covering $g : Y \to X$ so that Y is a negatively curved piecewise hyperbolic complex and $l \circ g$ satisfies condition (HC_0).

We describe Y in more detail, before choosing the subcomplex $\widehat{Y}$. Note that the vertices of X are in one-to-one correspondence with the vertices of Γ by the construction of X. Now, since we ensured that the preimage of e-links remained connected when choosing our branched covering, we see that the vertices of Y are also in one-to-one correspondence with those of Γ. Thus the preimage of a vertex of Γ under the composite $l \circ g$ will have exactly one vertex and N edges where N is the degree of the branched covering map. That is, the preimages of vertices of Γ are wedges of N circles in Y.

Let $\widehat{Y}$ be the Δ-subcomplex of Y obtained by removing, for each vertex $w \in \Gamma$, a single edge $\sigma \subset (l \circ g)^{-1}(w)$ and all 2-cells which contain σ as a face. Let $\widehat{l}$ denote the restriction of $l \circ g$ to $\widehat{Y}$. Now $\widehat{Y}$ is still a nonpositively curved piecewise hyperbolic Δ-complex (its links are still large). Note that $\widehat{l} : \widehat{Y} \to \Gamma$ has the property that all its e-links are trees (intervals of length $(N-1)$ actually). To see this observe that a choice of vertex v in Y and e-link of v determines a 2-chain in Y whose boundary is the preimage under $l \circ g$ of a vertex in Γ. Thus, it follows from an application of Theorem 2.7 (2) to the universal covering of $\widehat{Y}$, that $Ker(\widehat{l}_*)$ acts on a 1-connected 1-dimensional level set (i.e. a tree) with compact quotient, and so is a finite rank free group. In fact we see that the rank of this kernel is $(N-1)$, since the $\widehat{l}$-preimage of a vertex of Γ is a wedge of $(N-1)$ circles. Thus $\pi_1(\widehat{Y})$ is of the form $F_{N-1} \rtimes F_n$. □

§4 Specific examples

The construction in theorem B produces an abundance of CAT(-1) free-by-free groups. In this section we present two specific examples which illustrate some ways that this construction can be made more generic. In the discussions we will describe some general variations to the techniques of section 3.

Example 4.1. As our first example, we take $l : X \to \Gamma$ to be the natural projection where Γ is the theta graph and $X = X_\Gamma$ as defined in procedure 3.5. Thus X has two 0-simplices, five 1-simplices and six 2-simplices. Two of the 1-simplices are horizontal, and these are labeled x and y in figure 3. Since X has two vertices, L_X has two components (see figure 4) and its rank is 10. Each vertex of X has three e-links, all of which are circles with one vertex. These six circles are labeled by $A_1, A_2, A_3, B_1, B_2, B_3$ as indicated in figure 4. Note that figure 5 shows the graph D_X, which is connected and has rank 8. The map $\eta_X : L'_X \to D_X$ sends x_+ and x_- to $\widehat{x}$, y_+ and y_- to $\widehat{y}$, and sends the barycentric subdivisions of the A_i and B_i loops to the loops a_i and b_i. If we denote the homology classes of the a_i's and b_i's again by a_i and b_i then a generating set for $H_1(D_X)$ consists of $a_1, a_2, a_3, b_1, b_2, b_3$ and two additional elements which project onto generators for $H_1(\Gamma)$.

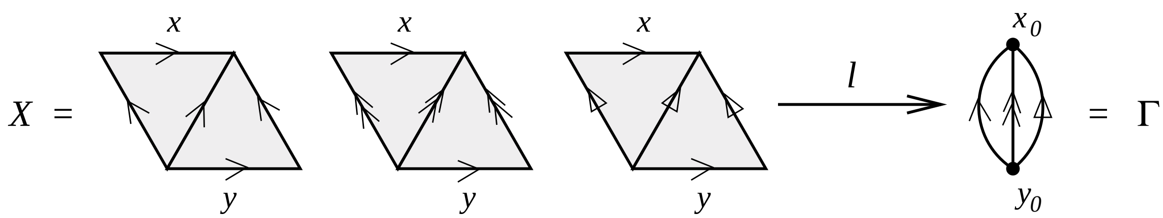

FIGURE 3. The map $l : X \to \Gamma$ of example 4.1.

We are going to produce a branched covering $g : Y \to X$ from a cyclic covering of D_X. To guarantee that Y has a nonpositively curved piecewise hyperbolic structure it is enough to just require that the image of all short loops in L'_X map under η_X to loops which open up in the cyclic covering of D_X. (In the procedure for the proof of theorem B we required that all short loops in D_X open up in the covering.)

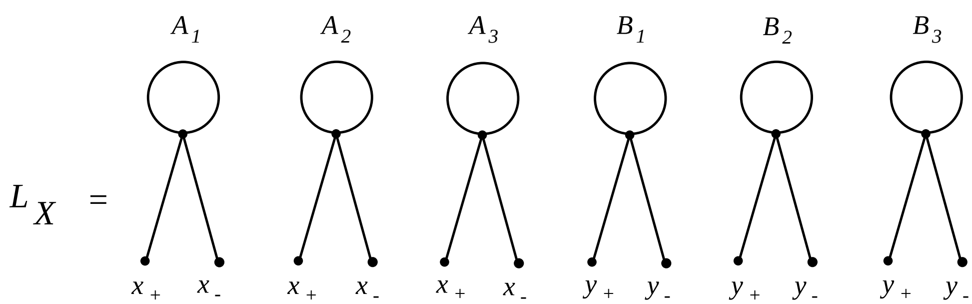

FIGURE 4. The link graph L_X of example 4.1.

In our example with the theta graph, the image of $H_1(L_X)$ under η_{X*} can be seen to be generated by $\{a_1, a_2, a_3, b_1, b_2, b_3\}$. We now investigate the images in $H_1(D_X)$ of the homotopically nontrivial loops in L_X with length less than seven. We find that (up to multiplication by -1) there are 108 such elements, and they are

(1) $ka_i, \quad i \in \{1,2,3\},\ k \in \{1,2,3,4,5,6\}$
(2) $kb_i, \quad i \in \{1,2,3\},\ k \in \{1,2,3,4,5,6\}$
(3) $a_i \pm a_j, \quad i,j \in \{1,2,3\},\ i < j$
(4) $b_i \pm b_j, \quad i,j \in \{1,2,3\},\ i < j$
(5) $a_i - a_j \pm kb_i, \quad i,j \in \{1,2,3\},\ i < j,\ k \in \{1,2\}$
(6) $a_i - a_j \pm kb_j, \quad i,j \in \{1,2,3\},\ i < j,\ k \in \{1,2\}$

(7) $b_i - b_j \pm k a_i, \quad i,j \in \{1,2,3\},\ i<j,\ k \in \{1,2\}$
(8) $b_i - b_j \pm k a_j, \quad i,j \in \{1,2,3\},\ i<j,\ k \in \{1,2\}$
(9) $a_i - a_j \pm (b_i + b_j), \quad i,j \in \{1,2,3\},\ i<j$
(10) $b_i - b_j \pm (a_i + a_j), \quad i,j \in \{1,2,3\},\ i<j$
(11) $a_i - a_j \pm (b_i - b_j), \quad i,j \in \{1,2,3\},\ i<j.$

Note that the list contains generators for the first homology group of the six vertex e-links of ℓ. A homomorphism $\gamma : H_1(D_X) \to \mathbb{Z}$ is determined by arbitrarily specifying integers $\gamma(a_i)$, $\gamma(b_i)$ for $i \in \{1,2,3\}$ and mapping the two additional generators in any manner. In accordance with lemma 3.2, we would like to specify these integers so that γ sends each of the 108 elements listed above to a nonzero integer. By lemma 3.1, there are infinitely many ways to achieve this. Having chosen one we then would like to find a positive integer N which is relatively prime to the $\gamma(a_i)$'s and $\gamma(b_i)$'s, and does not divide the image under γ of any of the 108 listed elements (again there will be infinitely many possibilities for N). As it turns out, the smallest possible value for N is $N = 15$, and this can be achieved by respectively taking $\gamma(a_1)$, $\gamma(a_2)$, $\gamma(a_3)$, $\gamma(b_1)$, $\gamma(b_2)$ and $\gamma(b_3)$ to be 1, 4, 7, 4, 13 and 7. Through the proof of theorem B, we obtain from this a CAT(-1) structure on the semidirect product $F_{14} \rtimes F_2$.

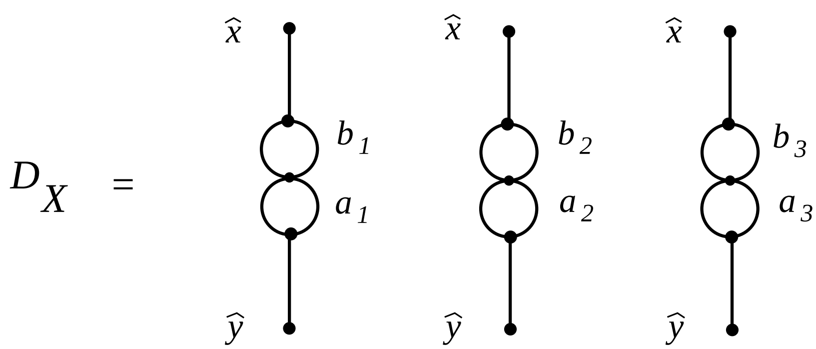

FIGURE 5. The graph D_X of example 4.1.

Example 4.2. For our second example, we take $l : X \to \Gamma$ to be the projection of Procedure 3.5 where Γ is the wedge of two circles and $X = X_\Gamma$ (see figure 6). In this case X has one 0-simplex, three 1-simplices (of which the one labeled x is horizontal) and four 2-simplices. The one vertex in X has four e-links, which are the circles labeled A_1, B_1, A_2, B_2 in figure 7. The immersion η_X takes x_+ and x_- to $\widehat{x}$ and takes the barycentric subdivisions of A_1, B_1, A_2, B_2 to the loops a_1, b_1, a_2, b_2 respectively. When a maximal tree is chosen, these loops in D_X determine elements of the free group $\pi_1(D_X)$ which we also denote by a_1, b_1, a_2, b_2. In fact, these four elements together with two additional elements, labeled s_1 and s_2, form a basis for $\pi_1(D_X)$. The elements s_1 and s_2 project to a basis for $\pi_1(\Gamma)$ and are respectively represented by loops based at $\widehat{x}$ which move down vertically in the

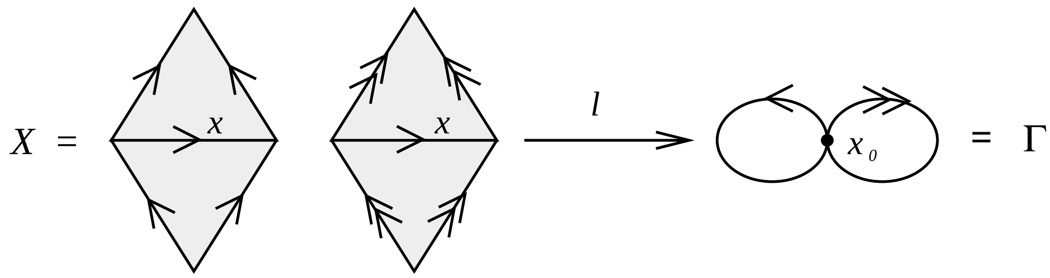

FIGURE 6. The map $l : X \to \Gamma$ of example 4.2.

left and right portions of D_X pictured in figure 7. The kernel of l_* is normally generated by $\{a_1, b_1, a_2, b_2\}$.

The loop of length 12 in L'_X which starts at x_+, traverses $\overline{A_1}$, passes through x_- then traverses B_1 and returns to x_+ is homotopically nontrivial in L'_X but its image under η_{X*} is $a_1^{-1}b_1^{-1}s_1^{-1}a_1b_1s_1$ which is homologically trivial. This means that a branched covering g as provided by theorem 3.4 cannot arise from a cyclic covering of D_X. Instead, the process of theorem 3.4 yields g as the composition of two cyclic coverings. (In general, assuming that $l : X_\Gamma \to \Gamma$ is defined as in Procedure 3.5, there will be a similar picture in L'_X and D_X if Γ contains an edge which is a loop. So a cyclic covering of D_X will suffice if and only if there are no such loops in Γ.)

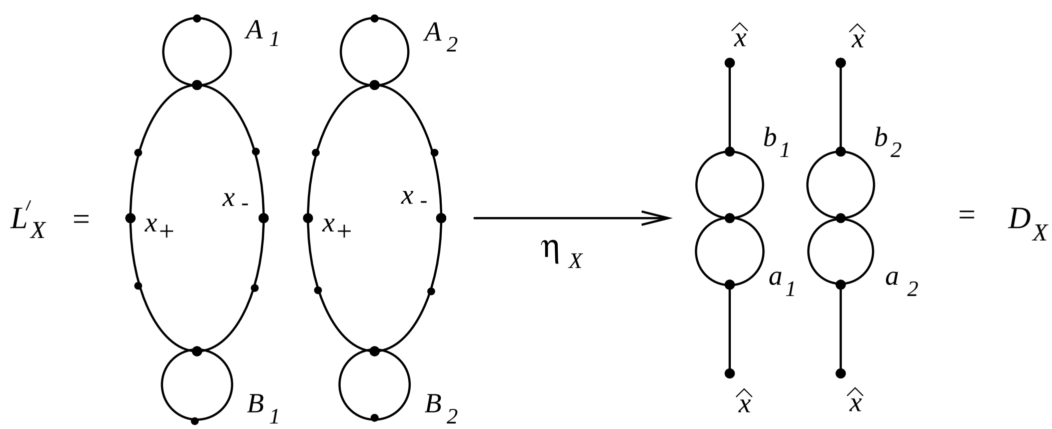

FIGURE 7. The map η_X of example 4.2.

Rather than taking a composition of two cyclic coverings, it is also possible to look directly for a non-abelian N-sheeted branched cover $g : Y \to X$. Such a branched covering corresponds to, and is determined by, a transitive permutation representation $\rho : \pi_1(D_X) \to S_N$. The proof of theorem 3.4 can be adapted to show that if ρ satisfies

(1) the restriction of $\rho \circ \eta_{X*}$ to the fundamental group of each e-link of a vertex of X

is transitive, and

(2) the image under $\rho \circ \eta_{X*}$ of each homotopically nontrivial loop with length less than 13 has no fixed element in its action on $\{1, \dots, N\}$

then $l \circ g$ satisfies (HC_0) and Y is a nonpositively curved piecewise hyperbolic Δ-complex with equilateral hyperbolic triangles. (We note that suppression of any basepoints in statements (1) and (2) is justified because the cycle (orbit) structure of permutations is invariant under conjugation.) When l is the projection $X_\Gamma \to \Gamma$ of Procedure 3.5, all of the e-links of vertices are circles and (1) is equivalent to requiring that ρ sends the elements of $\pi_1(D_X)$ represented by these circles to N-cycles in S_N.

The computational complexity of finding a permutation representation ρ satisfying (1) and (2) is significantly greater than that of finding a homomorphism into $\mathbb{Z}_N$ as in example 4.1. However there is a simplified approach which can be employed. We will describe this approach now, before returning to the specific example 4.2. Suppose l is a projection $X_\Gamma \to \Gamma$ as described in Procedure 3.5. In this situation, each vertex e-link of l is a circle. For each oriented edge e, denote the e-link at the vertex of X_Γ which projects to the initial vertex of e by A_e, and denote the $\bar{e}$-link at the vertex which projects to the terminal vertex of e by B_e. The image under η_X of the barycentric subdivisions of A_e and B_e are denoted by a_e and b_e respectively. When a maximal tree is chosen in D_X these loops determine elements of $\pi_1(X)$ which we also denote by a_e and b_e. Choose an orientation E_+ for Γ and let E'_+ consist of those elements of E_+ which are loops. Let n denote the rank of the subgraph obtained from Γ by removing all of the edge which are loops. Then $\pi_1(D_X)$ has a basis of the form

$$\{a_e,\, b_e \mid e \in E_+\} \cup \{s_e \mid e \in E'_+\} \cup \{t_1, \dots, t_n\}$$

where each s_e is a loop in D_X which projects onto the oriented edge-loop e and the set $\{s_e \mid e \in E'_+\} \cup \{t_1, \dots, t_n\}$ projects to a basis for $\pi_1(\Gamma)$. Let H be the subgroup of $\pi_1(D_X)$ (freely) generated by

$$\{a_e,\, b_e \mid e \in E_+\} \cup \{s_e^{-1} a_e s_e,\, s_e^{-1} b_e s_e \mid e \in E'_+\}\ .$$

It is not hard to see that η_{X*} sends the fundamental group of each component of L'_X in to H (up to conjugacy). In order to build a permutation representation $\rho : \pi_1(X) \to S_N$ we first construct a homomorphism $\gamma : H \to \mathbb{Z}$ by selecting integers $\alpha_e = \gamma(a_e)$ and $\beta_e = \gamma(b_e)$ for each $e \in E_+$, and integers $\alpha'_e = \gamma(s_e^{-1} a_e s_e)$ and $\beta'_e = \gamma(s_e^{-1} b_e s_e)$ for each $e \in E'_+$. We choose the values of these integers so that

$$\alpha'_e \beta_e = \alpha_e \beta'_e,$$

for each $e \in E'_+$, and so that the image under $\gamma \circ \eta_{X*}$ of any homotopically nontrivial short loop in L'_X is not equal to 0. We now find a positive integer N which is relatively prime to all α_e, β_e, α'_e and β'_e, and which does not divide the image under $\gamma \circ \eta_{X*}$ of any nontrivial short loop in L'_X. Let σ be an N-cycle in S_N. Define $\rho : \pi_1(D_X) \to S_N$ by taking $\rho(h) = \sigma^{\gamma(h)}$ for $h \in H$, by taking $\rho(t_k) = 1$ for $k = 1, \ldots, n$ and by choosing $\rho(s_e)$ to be a permutation which satisfies $\rho(s_e)^{-1}\sigma^{\alpha_e}\rho(s_e) = \sigma^{\alpha'_e}$. for each $e \in E'_+$. Observe that the compatibility equation $\alpha'_e\beta_e = \alpha_e\beta'_e$ guarantees that $\rho(s_e)^{-1}\sigma^{\beta_e}\rho(s_e) = \sigma^{\beta'_e}$, so that ρ is well-defined. The reader will note that the permutation representation $\rho : \pi_1(D_X) \to S_N$ that we have produced has its image in the normalizer of the cyclic subgroup generated by σ.

Using the procedure given above, we can now describe a specific representation ρ which works for the example where Γ is the wedge of two circles. Let N be an integer which is not divisible by 2 or 3, and let σ be an N-cycle. Recall that $\pi_1(D_X)$ is freely generated by $\{a_1, b_1, a_2, b_2, s_1, s_2\}$. Take

$$\alpha_1 = 6,\ \beta_1 = 16,\ \alpha'_1 = 3,\ \beta'_1 = 8,\ \alpha_2 = 4,\ \beta_2 = 9,\ \alpha'_2 = 8,\ \beta'_2 = 18.$$

By the process described above, this determines a representation $\rho : \pi_1(D_X) \to S_N$ with

$$\rho(a_1) = \sigma^6,\ \rho(b_1) = \sigma^{16},\ \rho(a_2) = \sigma^4,\ \rho(b_2) = \sigma^9,$$

where $\rho(s_1)$ and $\rho(s_2)$ are permutations satisfying

$$\rho(s_1)^{-1}\sigma^6\rho(s_1) = \sigma^3 \quad \text{and} \quad \rho(s_2)^{-1}\sigma^4\rho(s_2) = \sigma^8.$$

By the choice of N, the image under ρ of each a_i and b_i is an N-cycle, and this guarantees that condition (1) holds. One can now list the loops in L_X with short length and verify that if N is chosen appropriately large then condition (2) will hold. For example, $N = 43$ works, as will any larger integer which is not divisible by 2 or 3. The choice of $N = 43$ yields a CAT(-1) structure on a semidirect product $F_{42} \rtimes F_2$.

References

[BB] M. Bestvina and N. Brady, *Morse theory and finiteness properties of groups*, Invent. Math. **129** (1997), 445–470.

[BF] M. Bestvina and M. Feighn, *A combination theorem for negatively curved groups*, J. Diff. Geom. **35:1** (1992), 85–101.

[BFH] M. Bestvina, M. Feighn and M. Handel, *Laminations, trees, and irreducible automorphisms of free groups*, GAFA, geom. funct. anal. **7** (1997), 215-244.

[Bo] B. H. Bowditch, *A short proof that a subquadratic isoperimetric inequality implies a linear one*, Mich. J. Math **42** (1995), 103–107.

[BH] M. Bridson and A. Haefliger, *Metric spaces of nonpositive curvature*, manuscript of a book, to appear in 1999.

[Bra] N. Brady, *Branched coverings of cubical complexes and subgroups of hyperbolic groups*, To appear in Jour. LMS.

[Bro] K. Brown, *Cohomology of Groups*, GTM, vol. 87, Springer-Verlag, New York, 1982.

[Ger] S. M. Gersten, *Branched coverings of 2-complexes and diagramatic reducibility*, TAMS **303:2** (1987), 689–706.

[GdH] E. Ghys and P. de la Harpe (eds.), *Sur les groupes hyperboliques d'apres Mikhael Gromov*, Prog. in Math., vol. 83, Birkhauser, Boston, 1990.

[Gr1] M. Gromov, *Infinite groups as geometric objects*, Proceedings of the I.C.M. Warsaw 1, 1984, pp. 385–392.

[Gr2] M. Gromov, *Hyperbolic groups*, Essays in group theory (S.M. Gersten, ed.), MSRI series, vol. 8, Springer Verlag, 1987.

[Ha] A. Hatcher, *Algebraic Topology I*, Available online at http://math.cornell.edu/~hatcher, 1999.

[McW] J. McCammond and D. Wise, *Coherence, local quasiconvexity and the perimeter of 2-complexes*, preprint (1999).

[Mo] L. Mosher, *A hyperbolic-by-hyperbolic hyperbolic group*, Proc. AMS **125** (1997), 3447–3455.

[Ol] A. Yu Ol'shanski, *Hyperbolicity of groups with subquadratic isoperimetric inequalities*, Intl. J. alg. comp. **1** (1991), 281–290.

[P] P. Papasoglu, *On the sub-quadratic isoperimetric inequality*, Geometric group theory (R. Charney, M. Davis, M. Shapiro, eds.), de Gruyter, Berlin-New York, 1995.

[SW] P. Scott and C. T. C. Wall, *Topological methods in groups theory*, Homological Group Theory (ed. C. T. C. Wall), LMS Lecture Note Series, vol. 36, CUP, 1979, pp. 137–203.

[W] D. Wise, *Incoherent negatively curved groups*, Proc. AMS (to appear).

Department of Mathematics, University of Oklahoma, Norman, OK 73019, USA

E-mail address: `nbrady@math.ou.edu` and `amiller@math.ou.edu`