SEPARATING QUASICONVEX SUBGROUPS OF RIGHT-ANGLED ARTIN GROUPS

TIM HSU AND DANIEL T. WISE

Abstract. A *graph group*, or *right-angled Artin group*, is a group given by a presentation where the only relators are commutators of the generators. A graph group presentation corresponds in a natural way to a simplicial graph, with each generator corresponding to a vertex, and each commutator relator corresponding to an edge. Suppose that G is a graph group whose corresponding graph is a tree and H is a subgroup of G. We show that if H is *quasiconvex* with respect to either the word metric on G or the CAT(0) metric on the universal cover of the standard complex for G, then H is *separable*, that is, H is the intersection of finite index subgroups of G. We also discuss some consequences relating to certain 3-manifold groups.

1. Introduction

A *graph group*, or *right-angled Artin group*, is a group given by a presentation where the only relators are commutators of the generators. Such a presentation corresponds in a natural way to a simplicial graph, with each generator corresponding to a vertex, and each commutator relator corresponding to an edge. For example, the graph group corresponding to a triangle is presented by $\langle a, b, c \mid [a, b], [b, c], [a, c]\rangle$, and the graph group corresponding to a letter 'Y' graph is presented by $\langle a, b, c, d \mid [a, b], [a, c], [a, d]\rangle$. If the underlying graph of a graph group G is a tree, we call G a *tree group*.

Graph groups were first studied by Baudisch [Bau81], who was interested in their 2-generator subgroups. Since then, many of their properties have been investigated. For example, graph groups are biautomatic [VW94], [HM95], [NR98] and they have finitely generated automorphism groups [Ser89], [Lau95]. Two graph groups are isomorphic if and only if their corresponding graphs are isomorphic [Dro87b]. The property of being a graph group is often inherited by subgroups [Dro87c]. Most recently, in work of Bestvina and Brady [BB97], certain subgroups of graph groups have emerged as important examples because of their exotic homological properties.

Recall that a group G is said to be *residually finite* if the trivial subgroup of G is the intersection of finite index subgroups of G; more generally, a subgroup H of a group G is said to be *separable* if H is the intersection of finite index subgroups of G. If every finitely generated subgroup of a group G is separable, then G is said to be *subgroup separable*. Green [Gre90] showed that graph groups are residually finite (see also [Hum94] and [HW99]), and more generally, that every full subgroup (Definition 2.4) of a graph group is separable (see also [HW99]). In this article, we provide a more elaborate treatment of the residual properties of certain graph groups; in particular, we generalize the separability results mentioned above.

If we regard a group as a metric space, its *quasiconvex* subgroups are those that correspond in a coarse sense to convex subspaces. Previous work on groups of nonpositive curvature has suggested

Date: July 1, 2006.
1991 *Mathematics Subject Classification.* 20F32.
Key words and phrases. Graph groups, Artin groups, subgroup separability.
Wise was supported as an NSF postdoctoral fellow under grant no. DMS-9627506.

a connection between quasiconvexity and separability. For example, a separable subgroup of a finitely presented group has decidable membership problem [BN74] and a quasiconvex subgroup of a CAT(0) group has decidable membership problem (see [BH99, III.Γ.4.12, 5.15, 6.18]). More recently, it was shown in [Wis98] that the fundamental groups of certain nonpositively curved squared 2-complexes have the property that their quasiconvex subgroups are separable. In particular, the geometrically finite subgroups of the fundamental group of the figure 8 knot are separable.

The following theorem establishes a further connection between quasiconvexity and separability.

Main Theorem. *Every quasiconvex subgroup of a tree group is separable.*

More precisely, we will define the class of *quasifull* subgroups of a graph group, a class that includes both subgroups that are quasiconvex with respect to the word metric on G and subgroups that are quasiconvex with respect to the CAT(0) metric on the universal cover of the standard complex for G, and we will show that every quasifull subgroup of a tree group is separable.

Note that in general, in contrast with our main result, graph groups and similar groups of non-positive curvature can have quite pathological finitely generated subgroups. For example, the group $F_2 \times F_2$ is an innocuous looking graph group corresponding to a square. Nevertheless, Mikhailova showed (see [LS77]) that $F_2 \times F_2$ has finitely generated subgroups that have undecidable membership problem, and are therefore not separable. In fact, in an arbitrary group of non-positive curvature, even the quasiconvex subgroups need not be separable. For example, [Wis96] gives an example of a non-positively curved squared 2-complexes whose fundamental group is not residually finite. For non-separable examples more directly connected with our Main Theorem, see the discussion surrounding Theorems 1.1 and 1.2, below.

Recall that because of a possible impact on Thurston's geometrization conjecture, there is substantial interest in the problem of whether every compact irreducible 3-manifold with infinite fundamental group has a finite cover with an embedded π_1-injective surface of genus ≥ 1. Subgroup separability is of particular interest to this problem because it allows certain maps to lift to embeddings in a finite cover. More precisely, suppose that $\varphi : \Sigma \looparrowright M$ is a π_1-injective immersion of a compact surface into a 3-manifold such that φ has minimal self-intersection. Then if $\varphi_*(\pi_1\Sigma)$ is a separable subgroup of $\pi_1 M$, there exists a finite cover $\hat{M} \to M$ such that $\varphi : \Sigma \looparrowright M$ lifts to an embedding $\hat{\varphi} : \Sigma \hookrightarrow \hat{M}$.

Returning to graph groups, recall that every tree group is the fundamental group of a link complement [Dro87a]. (In fact, a graph group with a connected graph is a 3-manifold group if and only if its graph is either a tree or a triangle [Dro87a].) The projection of this link complement is easy to describe: First embed the tree in the plane, then draw a circle around each vertex, and finally, if two vertices are connected by an edge, then draw their corresponding circles so that they are linked exactly once. For example, the graph group

$$J = \overset{a}{\bullet}\!\!-\!\!\!-\!\!\overset{b}{\bullet}\!\!-\!\!\!-\!\!\overset{c}{\bullet}\!\!-\!\!\!-\!\!\overset{d}{\bullet} = \langle a, b, c, d \mid [a, b], [b, c], [c, d] \rangle \tag{1}$$

is the fundamental group of the complement of the link shown in Figure 1.

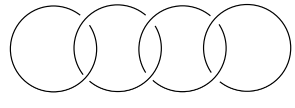

FIGURE 1. A link made from a chain of four circles

As shown in [NW], J is not subgroup separable and in fact, J is a subgroup of every known finitely generated 3-manifold group that is not subgroup separable. However, the following special case of the Main Theorem limits the failure of the subgroup separability of J.

Theorem 1.1. *Every quasiconvex subgroup of* $J = \langle a, b, c, d \mid [a, b], [b, c], [c, d]\rangle$ *is separable.*

For another consequence of the Main Theorem, recall that Burns, Karrass, and Solitar [BKS87] gave the first example of a finitely generated 3-manifold group that is not subgroup separable, namely,

$$B = \langle a, b, t \mid [a, b], a^t = b\rangle. \tag{2}$$

Later, [RW98] gave examples of π_1-injective immersed surfaces that do not lift to an embedding in any finite cover. In fact, it was shown in [NW98] that there are 3-manifold groups with proper finitely generated subgroups whose separability properties are as pathological as can be. For example, B contains a finitely generated proper subgroup that is not contained in any finite index proper subgroup.

Note that since the standard 2-complex of B is a nonpositively curved squared 2-complex K, we may consider the subgroups that are quasiconvex relative to the natural CAT(0) metric on the universal cover of K. As we will show, since B is very closely related to J, Theorem 1.1 implies:

Theorem 1.2. *Every quasiconvex subgroup of* $B = \langle a, b, t \mid [a, b], a^t = b\rangle$ *is separable.*

What is remarkable about Theorems 1.1 and 1.2 is that they demonstrate that the geometrically agreeable subgroups of B and J are separable, even though neither B nor J is subgroup separable. This lends credence to a theme that has been taking shape in recent years: namely, that the quasiconvex/geometrically finite subgroups of 3-manifold groups are separable.

We now briefly summarize the rest of this paper. After reviewing some background material (Section 2), we begin our proof of the Main Theorem by examining the quasiconvexity conditions we use in the proof (Section 3). Next, we examine the geometry of covering spaces of the standard complex of a graph group, concentrating on certain important subspaces of covering spaces (Section 4) and their intersections (Section 5). We then describe the notion of a *local isometry* between cubed complexes (Section 6), which is a sort of partial covering space possessing a certain convexity property.

As for the proof of the Main Theorem itself, our strategy for demonstrating the separability of a finitely generated subgroup H of a group $G = \pi_1 X$ has been used in several situations, most notably in M. Hall [Hal49] to show that free groups are subgroup separable, in Scott [Sco78] to show that surface groups are subgroup separable, and in Wise [Wis98] to show that the fundamental groups of many hyperbolic alternating link complements have separable geometrically finite subgroups. The principal new element here is that previously, the groups in question were either word-hyperbolic or similar to finite volume hyperbolic 3-manifold groups with cusps. The groups treated in this paper are distinguished by having a substantially richer collection of $\mathbb{Z} \times \mathbb{Z}$ subgroups; for example, the group J mentioned in Theorem 1.1 has subgroups isomorphic to $F_2 \times \mathbb{Z}$. Because of this richer collection of $\mathbb{Z} \times \mathbb{Z}$ subgroups, the analogous theorem in [Wis98] fails to apply.

Following Scott, the basic idea is to show that, for any compact subset C of the based covering space $\hat{X}$ of X corresponding to H, there is a finite covering space $\overline{X}$ of X in which C embeds. We do not construct $\overline{X}$ directly; rather, we first construct several other spaces, beginning with a certain *core*, each of which contains all of the important topology of $\hat{X}$ and C.

More precisely, given $C \subset \hat{X}$, we construct a finite cover of X in which C embeds by the following steps. First, in Section 7, we select a subspace S (the *locally convex core*) of $\hat{X}$ containing C that is not necessarily compact, but nevertheless embeds by a local isometry, and enjoys certain finiteness properties. Next, in Section 8, we close off the ends of S to obtain a compact quotient S_0 in which C still embeds, and which maps to X by a local isometry. In Section 9, we then extend the map $S_0 \to X$ to a finite cover of X. Finally, in Section 10, we assemble these results to obtain the proof of the Main Theorem. (A reader who is unfamiliar with this strategy may prefer to begin with Section 10 for a "top-down" version of the proof.)

We conclude with some applications and examples (Section 11) and a discussion of some open problems (Section 12).

Acknowledgement: We gratefully acknowledge Dick Canary for his help in paying for travel so we could meet at the University of Michigan to do this work. We also thank the members of the Claremont Colleges algebra/combinatorics and topology seminars and the UC Berkeley geometric group theory seminar for their helpful comments.

2. Background

In this section, we review some basic definitions and results.

Definition 2.1 (Graph groups)**.** Let Γ be a finite simplicial graph. The *graph group* (or *right-angled Artin group*) $G(\Gamma)$ is given by the presentation with a generator g_i for each vertex v_i of Γ, and a defining relation $[g_i, g_j] = 1$ for each edge between vertices v_i and v_j in Γ. The g_i are called the *standard generators* of $G(\Gamma)$.

In particular, if Γ is a tree, then $G(\Gamma)$ is called a *tree group*.

Note that each of the following "moves" changes a given word W in the standard generators $\{g_v\}$ to a word W' that is equal in $G(\Gamma)$ and has length less than or equal to the length of W.

1. (Cancel) Remove a subword of the form $g_i g_i^{-1}$ or $g_i^{-1} g_i$.
2. (Swap) If $g_i^{\epsilon_1}$ and $g_j^{\epsilon_2}$ ($\epsilon_1, \epsilon_2 = \pm 1$) commute, replace $g_i^{\epsilon_1} g_j^{\epsilon_2}$ with $g_j^{\epsilon_2} g_i^{\epsilon_1}$.

If g is represented by a word W that cannot be changed to a shorter word using any sequence of the above moves, then W is said to be a *normal form* for g.

We recall the following theorem of Green [Gre90]:

Theorem 2.2 (Normal forms)**.** *Two normal forms represent the same element of $G(\Gamma)$ if and only if they are equivalent via a sequence of swaps. Furthermore, the shortest possible words representing an element g of $G(\Gamma)$ are precisely the normal forms for g.* □

For a diagrammatic proof of Theorem 2.2 in the case $g = 1$, see [HW99]. We leave it as an exercise for the reader to deduce Theorem 2.2 from this special case.

Definition 2.3 (Standard complex)**.** Let $v_1, \dots, v_n$ be an ordering of the vertices of Γ, and let $g_1, \dots, g_n$ be the corresponding generators of $G(\Gamma)$. Since $\langle g_1, \dots, g_{k+1} \rangle$ is an HNN extension of $\langle g_1, \dots, g_k \rangle$ by the stable letter g_{k+1} centralizing the subgroup generated by those elements of $g_1, \dots, g_k$ that commute with g_{k+1}, it follows that $G(\Gamma)$ is itself an iterated HNN extension of the trivial group.

The geometric version of this iterated HNN construction of $G(\Gamma)$ yields an n-dimensional cell complex X that we call the *standard complex* of $G(\Gamma)$. Note that the 2-skeleton of X is the 2-complex for the presentation of $G(\Gamma)$, which means that $\pi_1 X = G(\Gamma)$. Note also that X may be described as the cubed n-complex such that:

1. X has precisely one 0-cell;
2. X has one 1-cell corresponding to each generator of $G(\Gamma)$; and
3. For every complete subgraph of Γ with vertices $v_1, \dots, v_k$, there is a k-cube glued to the oriented 1-cells corresponding to $v_1, \dots, v_k$, in the manner of a k-dimensional torus.

X is therefore independent of the ordering of the generators of $G(\Gamma)$, which justifies our calling it *the* standard complex of $G(\Gamma)$.

In fact, if we give the cubes of the standard complex X of a graph group $G(\Gamma)$ the geometric structure of Euclidean cubes, then the universal cover $\tilde{X}$ of X is a CAT(0) space on which $G(\Gamma)$ acts properly and cocompactly by isometries [MV95].

Definition 2.4 (Full subgroup/subcomplex)**.** Let $G(\Gamma)$ be a graph group, and let F be a subgroup of $G(\Gamma)$ generated by some subset of the standard generators. F is then said to be a *full subgroup* of $G(\Gamma)$. Furthermore, let X_F denote the subcomplex of the standard complex X consisting of the 0-cell of X and precisely those cells of X that contain only 1-cells corresponding to the standard generators in F. X_F is called the *full subcomplex* of X corresponding to F.

Definition 2.5 (Dimension)**.** Note that if the largest possible complete subgraph of a graph Γ has n vertices, then n is the largest dimension of a cell of the standard complex of $G(\Gamma)$, and $G(\Gamma)$ contains no $\mathbb{Z}^{n+1}$ full subgroups. We therefore say that such a graph group is *n-dimensional*.

Finally, we will also need the following notion of Scott and Wall [SW79]. Let I denote the closed interval $[0, 1]$.

Definition 2.6 (Graph of spaces)**.** A *graph of spaces* is:

(1) A graph ξ (not necessarily finite or simplicial);
(2) For each vertex v of ξ, a topological space X_v;
(3) For each edge e of ξ, a topological space X_e, and continuous maps (called the *attaching maps* of e) from X_e to the spaces corresponding to the initial and terminal vertices of e; and
(4) A topological space X formed by taking the disjoint union of the X_v and then, for each edge e, gluing a copy of $X_e \times I$ to the spaces corresponding to the initial and terminal vertices of e, as indicated by the attaching maps of e.

For brevity, we often refer to a graph of spaces by the space X that it determines. The X_v are called the *vertex spaces* of X; the $X_e \times I$ are called the *edge spaces* of X; and if $X_e \times I$ is an edge space of X, then the subspace $X_e \times (0, 1)$ of $X_e \times I$ is called the *interior* of $X_e \times I$.

Note that if all of the attaching maps of the edge spaces of a graph of spaces X are π_1-injections, then an application of van Kampen's theorem gives a splitting of $\pi_1 X$ as a graph of groups. Conversely, any graph of groups can be represented geometrically in this way, by choosing based spaces corresponding to the vertex and edge groups, and choosing attaching maps representing inclusions of edge groups. For example, an HNN extension is the fundamental group of a graph of spaces made from one vertex space and one edge space with π_1-injective attaching maps; and a free product with amalgamation is the fundamental group of a graph of spaces made from two vertex spaces connected by an edge space with π_1-injective attaching maps.

We also have the following general construction.

Definition 2.7 (Induced decomposition of a cover)**.** Let $\hat{X}$ be a covering space of a graph of spaces X. We define the *induced decomposition* of $\hat{X}$ as a graph of spaces by taking the vertex spaces of $\hat{X}$ to be the components of the inverse images of the vertex spaces of X, and the edge spaces of $\hat{X}$ to correspond in an obvious way to the components of the preimages of the interiors of edge spaces of X.

Note that there is always an induced cellular map from the underlying graph of $\hat{X}$ to the underlying graph of X. However, this map is, in general, not a cover of the underlying graph of X.

3. Quasiconvexity conditions

In this section, we define the classes of subgroups in which we are most interested, and examine some of their properties. We begin by defining the following important subspaces of a covering space of the standard complex of a graph group.

Definition 3.1 (Components)**.** Let $G = G(\Gamma)$ be a graph group, let F be a full subgroup of G (Definition 2.4), and let $\hat{X}$ be a cover of the standard complex X of G. An *F-component* of $\hat{X}$ is defined to be a component of the inverse image of the full subcomplex X_F (Definition 2.4).

Definition 3.2 (Quasifull subgroups and covers)**.** Let $G = G(\Gamma)$ be a graph group with standard complex X, and let H be a subgroup of G with corresponding covering space $\hat{X}$. We say that H is *quasifull*, and that $\hat{X}$ is a *quasifull cover*, if either of the following (equivalent) conditions hold:

(1) For every full subgroup F of G, every F-component of $\hat{X}$ has finitely generated fundamental group.
(2) For every conjugate xFx^{-1} of a full subgroup of G, $H \cap xFx^{-1}$ is finitely generated.

We immediately observe that:

Lemma 3.3. *Let X be the standard complex of a graph group $G = G(\Gamma)$. Then:*

(1) *The quasifull condition for subgroups of G is conjugacy invariant.*
(2) *If $\hat{X}$ is a quasifull cover of X, and F is a full subgroup of G, then every F-component of $\hat{X}$ is a quasifull cover of X_F.* □

We recall the following standard notions of quasiconvexity. (See [Sho91] and [BH99, III.Γ.4.11] for more details.)

Definition 3.4. A subgroup H of a group G with chosen generators g_i is said to be *word-quasiconvex* if there exists a constant K such that, in the Cayley graph of G with respect to the generators g_i, any geodesic path representing an element in H is contained in the K-neighborhood of H in G.

Definition 3.5. Let G be a group acting properly and cocompactly by isometries on a CAT(0) space $\tilde{X}$ (for example, a graph group acting on the universal cover of its standard complex). We say that a subgroup $H \leq G$ is *CAT(0)-quasiconvex* if there exists $x_0 \in \tilde{X}$ and a constant K such that for all $h \in H$, the geodesic in X between x_0 and hx_0 is contained in the K-neighborhood of Hx_0.

We have used the name quasifull in Definition 3.2 because, as we will show (Theorem 3.9), every full subgroup is quasiconvex (in both senses), and every quasiconvex subgroup (in either sense) is quasifull. We begin with Lemmas 3.6–3.8, throughout which we assume that G is a graph group with standard complex X.

Lemma 3.6. *Every conjugate of a full subgroup of G is word-quasiconvex.*

Proof. We claim that if F is a full subgroup of G and $x = x_1 \dots x_k$ is a word of length k in the standard generators, then xFx^{-1} is quasiconvex in G with constant $2k$. To prove our claim, let $w = x_1 \dots x_k g_1 \dots g_m x_k^{-1} \dots x_1^{-1}$ be an arbitrary element of xFx^{-1}, where the g_i are standard generators of the full subgroup F. Note that any normal form for w may be obtained by the following algorithm (compare Theorem 2.2):

(1) (Cancel with ghost) For a generator g, replace a subword of the form gg^{-1} with the "ghost" (marked symbol representing the identity) (g, g^{-1}).
(2) (Swap) Swap two commuting generators, a generator and a ghost, or two ghosts.

The resulting "reduced normal form with ghosts" r will just be a standard normal form with ghosts inserted at various places.

So now, let r be a "reduced normal form with ghosts" for w, and let r_0 be an initial portion of r. It remains to show that there exists a word x' of length $\leq 2k$ in the standard generators such that $r_0 x' \in xFx^{-1}$. Let r_1 be the word formed from r by the rule:

> If g is a non-ghost letter of r occurring after r_0, and g is not one of the $x_i^{\pm 1}$'s, then replace g with the ghost (g).

Since the only non-ghost letters in r_1 that are not in r_0 are the $x_i^{\pm 1}$'s occurring after r_0, we see that $r_1 = r_0 x'$, where x' is a word of length $\leq 2k$ in the x_i's. Furthermore, since changing a non-ghost letter to a ghost does not affect either swapping or anti-cancelling (g, g^{-1}) ghosts, we may take the reverse of the process used to reduce w to r, and apply it to the word r_1, thereby obtaining a word w_1 such that $w_1 = r_1 = r_0 x'$ in G. However, since w_1 can also be obtained from w by deleting some of the g_i's, $w_1 = r_0 x'$ is also an element of xFx^{-1}. The theorem follows. □

Lemma 3.7. *Let H be a CAT(0)-quasiconvex subgroup of G. Then the quasiconvexity of H does not depend on the chosen basepoint in $\tilde{X}$, and every conjugate of H is also CAT(0)-quasiconvex.*

Proof. The first statement was proven in [AB95, §10]. As for the second statement, an easy argument shows that for any $g \in G$, if H is CAT(0)-quasiconvex relative to the basepoint x_0, then gHg^{-1} is CAT(0)-quasiconvex relative to the basepoint $x_1 = gx_0$. □

Lemma 3.8. *Every conjugate of a full subgroup of G is CAT(0)-quasiconvex.*

Proof. Since every full subcomplex of X is totally geodesic, every full subgroup of G is CAT(0)-quasiconvex. Therefore, Lemma 3.7 implies that every conjugate of a full subgroup of G is CAT(0)-quasiconvex. □

Theorem 3.9. *A word-quasiconvex (resp. CAT(0)-quasiconvex) subgroup H of a graph group G is quasifull.*

Proof. As shown in Lemmas 3.6 and 3.8, using either sense of quasiconvexity, any conjugate of a full subgroup is quasiconvex. Therefore, since the intersection of quasiconvex subgroups is quasiconvex, and a quasiconvex subgroup of a finitely generated group is finitely generated (see [Sho91] and [BH99, III.Γ.4.12,4.13]), the intersection of any quasiconvex subgroup and any conjugate of a full subgroup is finitely generated. □

Example 3.10. Note that the converse of Theorem 3.9 is false for word-quasiconvexity. One well-known example (see [Sho91]) is the subgroup $H = \langle ab \rangle$ of the graph group 

. H is quasifull, since it is cyclic, but H is not word-quasiconvex, since the geodesic path $a^n b^n = (ab)^n \in H$ travels a distance n from H. On the other hand, H is finitely generated abelian, and therefore CAT(0)-quasiconvex [AB95, §9]; in fact, it may well be that CAT(0)-quasiconvexity is equivalent to quasifullness. See Section 12 for further discussion.

4. Flats and pencils

In this section, we examine some of the more important full components of covers of the standard complex of a graph group. For the rest of this section, let X be the standard complex of a graph group $G(\Gamma)$, and let $\hat{X}$ be a covering space of X.

Definition 4.1 (Flats)**.** If F is a full $\mathbb{Z}^n$ subgroup of G, then an F-component of $\hat{X}$ is called an *n-flat*, or more specifcally, an *F-flat*, of $\hat{X}$.

Since X_F is an n-torus when $F \cong \mathbb{Z}^n$, flats are always torus covers. For instance, a 1-flat is either a line or a circle, and a 2-flat is either a plane, a cylinder, or a torus. In particular:

Definition 4.2 (Cylinders)**.** If E is an $\langle a, b\rangle$-flat in $\hat{X}$ such that $\pi_1 E$ is generated by a conjugate of $a^m b^n \neq 1$, then E is said to be an *$\langle a, b\rangle$-cylinder.* If furthermore, both m and $n \neq 0$, E is said to be *twisted*; otherwise, E is said to be *untwisted.* (In the latter case, E is the product of a line and a circle, with the product cell structure.)

Figure 2 shows portions of an untwisted cylinder and a twisted cylinder.

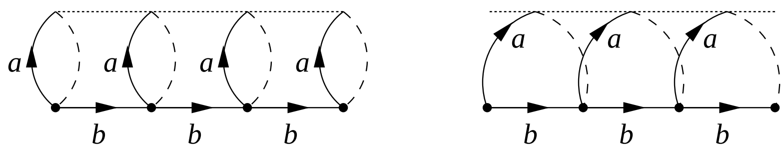

FIGURE 2. An untwisted cylinder and a twisted cylinder

Definition 4.3 (Pencil)**.** Let t be a standard generator of $G(\Gamma)$, and let $C(t)$ be the full subgroup generated by all standard generators that commute with t. (In fact, it follows from the normal form theorem that $C(t)$ is precisely the centralizer of t in $G(\Gamma)$.) A $C(t)$-component of $\hat{X}$ is called a *t-pencil.*

Note that since $X_{C(t)}$ is precisely the union of all cells of X that contain the 1-cell corresponding to t, a t-pencil in $\hat{X}$ is a connected union of flats containing $\langle t\rangle$-flats, and intersecting along those $\langle t\rangle$-flats. We have therefore chosen the name "pencil" for this kind of subspace because in classical Euclidean geometry, a pencil is a parallelism class of lines.

Now, for the rest of this section, let $G(\Gamma)$ be a 2-dimensional graph group, and let t be a standard generator of $G(\Gamma)$. In that case, since algebraically, $C(t) \cong F_n \times \langle t\rangle$, geometrically, $X_{C(t)}$ is the product of a circle and a bouquet of n circles. This product decomposition induces a bundle decomposition of each t-pencil in two different ways: namely, as a graph bundle over a circle or line, and as a circle or line bundle over a graph. In terms of the latter decomposition, a t-pencil whose t-fiber is a circle is called a *closed* pencil, and a t-pencil whose t-fiber is a line is called an *open* pencil.

We may now state the following lemma.

Lemma 4.4 (Open pencils of quasifull covers)**.** *Let $G = G(\Gamma)$ be a 2-dimensional graph group, let $\hat{X}$ be a quasifull cover of the standard complex X of G, let t be a standard generator of G, let P be an open t-pencil in $\hat{X}$, and let $\overline{P}$ be the base graph of the t-line bundle decomposition of P. Then either:*

1. *P is the direct product of the t-fiber and $\overline{P}$; or*
2. *$\pi_1\overline{P} \cong Z$, and all twisting in the t-line bundle decomposition occurs within the inverse image of the unique simple closed loop in $\overline{P}$.*

In particular, in case 2, if the unique reduced closed path in $\overline{P}$ is contained in a single 1-flat, then all of the twisting in P is contained within a single twisted cylinder.

Proof. Let $C(t) = F_n \times \langle t\rangle$. After conjugation, we may assume that the basepoint of $\hat{X}$ is in P, in which case $\pi_1 P$ is naturally isomorphic to $\pi_1\hat{X} \cap (F_n \times \langle t\rangle)$. Now, since the t-fiber of P is a line, $\pi_1 P \cap \langle t\rangle = 1$, and the projection p_1 of $\pi_1 P$ onto the first coordinate of $F_n \times \langle t\rangle$ is an isomorphism. Furthermore, P is a product if and only if the projection p_2 of $\pi_1 P$ onto the second coordinate of

$F_n \times \langle t \rangle$ is trivial. However, if p_2 is nontrivial, then $\ker p_2 = \pi_1 \hat{X} \cap F_n$ is a finitely generated (since $\hat{X}$ is quasifull) normal subgroup of infinite index in the free group $\pi_1 P$. Therefore, if P is not a product, then $\pi_1 P \cong Z$, and the lemma follows. □

Remark 4.5. Note that in Lemma 4.4, if $\pi_1 \hat{X}$ is actually word-quasiconvex in G, then only case 1 can occur. To see this, suppose that case 2 of Lemma 4.4 holds, and let $a \in F_n$ represent the unique simple closed loop in $\overline{P}$. Then $\pi_1 P = \pi_1 \hat{X} \cap (F_n \times \langle t \rangle) = \langle at^k \rangle$ for some k, and $\pi_1 P$ is not word-quasiconvex, since the geodesic word $a^n t^{kn} = (at^k)^n$ travels at least distance kn away from $\pi_1 P$. (Compare Example 3.10.)

5. Intersections of 2-flats

One key detail in the proof the Main Theorem is the following classification of possible intersections of 2-flats in a quasifull cover.

Lemma 5.1 (Intersection of 2-flats)**.** *Let X be the standard complex of a tree group $G(\Gamma)$, let $\hat{X} \to X$ be a quasifull cover of X, and let E_1 and E_2 be distinct intersecting 2-flats in $\hat{X}$. Then either:*

(1) *E_1 and E_2 intersect in finitely many 0-cells;*
(2) *E_1 and E_2 are in the same t-pencil and intersect in finitely many t-fibers of the pencil; or*
(3) *E_1 and E_2 intersect in finitely many 0-cells and an infinite periodic sequence of 0-cells inside a single twisted cylinder E. ("Periodic" means that there exists a translation of E that preserves $E_1 \cap E_2 \cap E$.)*

Proof. Let E_1 be an $\langle a, b \rangle$-flat, let E_2 be a $\langle c, d \rangle$-flat, and let F be the full subgroup $\langle a, b, c, d \rangle$. Without loss of generality we may assume that the basepoint of $\hat{X}$ is an intersection point of E_1 and E_2. Since Γ is a tree, after possibly renaming the generators, we have three possibilities for F:

(1) F is the free product of 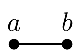 and 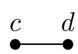.
(2) $b = c = t$, and F is the group 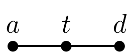.
(3) F is the group 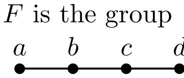.

Let $\hat{X}_F$ be the F-component of $\hat{X}$ containing both E_1 and E_2, and note that $\pi_1 \hat{X}_F$ is finitely generated, since $\hat{X}$ is quasifull.

Case 1. First, since E_1 and E_2 have no generators in common, they can only meet in 0-cells. Consider X_F as the edge of spaces corresponding to the obvious free product decomposition of F, and let Λ be the underlying graph of the induced graph of spaces decomposition of $\hat{X}_F$. E_1 and E_2 project down to vertices of Λ, and each intersection of E_1 and E_2 gives rise to a distinct edge in Λ between these two vertices, as shown in Figure 3 (in which E_1 and E_2 are represented by solid and dashed lines, respectively). Therefore, if E_1 and E_2 intersect infinitely many times, then $\pi_1 \Lambda$ has an infinitely generated free factor, which contradicts the finite generation of $\pi_1 \hat{X}_F$. We therefore have conclusion 1 of the lemma.

Case 2. First, since E_1 and E_2 intersect and share the generator t, they are part of the same t-pencil, and their intersection is precisely a disjoint union of t-fibers. Consider X_F as the edge of spaces corresponding to the decomposition of F as the free product of $\langle a, t \rangle$ and $\langle t, d \rangle$ amalgamating $\langle t \rangle$, and let Λ be the graph underlying the induced decomposition of $\hat{X}_F$. Again, E_1 and E_2 project to vertices of Λ, and each t-fiber in the intersection gives rise to a distinct edge in Λ between these two vertices, so the argument used to finish case 1 yields conclusion 2 of the lemma.

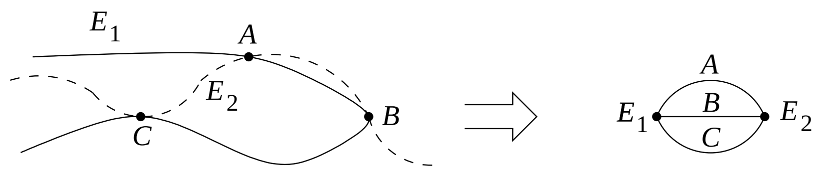

FIGURE 3. Projection to Λ of the intersection of two flats

Case 3. First, as in case 1, E_1 and E_2 intersect in 0-cells. Consider X_F as the edge of spaces corresponding to the decomposition of F as the free product of $\langle a, b, c\rangle$ and $\langle b, c, d\rangle$ amalgamating $\langle b, c\rangle$. By the same argument as before, we see that the intersection of E_1 and E_2 is contained in finitely many $\langle b, c\rangle$-flats, since each such flat corresponds to an edge in the underlying graph of spaces decomposition.

So now, let E be a $\langle b, c\rangle$-flat inside which E_1 and E_2 intersect. From case 2, we know that E_1 and E intersect in finitely many $\langle b\rangle$-flats, and E_2 and E intersect in finitely many $\langle c\rangle$-flats. Therefore, if E is a torus, a plane, or an untwisted cylinder, E_1 and E_2 meet in finitely many 0-cells inside E.

Finally, suppose E is a twisted cylinder (which can only happen once, from Lemma 4.4). By choosing the basepoint to be an intersection point of E_1 and E_2 inside E, we may assume that $\pi_1 E = \langle b^m c^{-n}\rangle$ for some integers m and n. In that case, E contains precisely $|n|$ $\langle b\rangle$-lines and $|m|$ $\langle c\rangle$-lines, and translation along E by b^m, or equivalently, by c^n, sends each $\langle b\rangle$-line of E to itself and each $\langle c\rangle$-line of E to itself. Since the intersections of E_1 and E_2 in E occur precisely where the $\langle b\rangle$-lines of $E_1 \cap E$ meet the $\langle c\rangle$-lines of $E_2 \cap E$, this translation preserves $E_1 \cap E_2 \cap E$. Conclusion 3 of the lemma follows. □

6. LOCAL ISOMETRIES

As our final preliminary before beginning the proof of the Main Theorem, in this section we give a combinatorial definition of local isometry, and we extend the definitions of F-components, flats, and pencils to this situation.

Definition 6.1 (Convex subcomplex)**.** Let L be a cell complex whose cells are simplices. A subcomplex K of L is said to be *convex* if every cell of L whose vertices are all contained in K is itself contained in K.

Definition 6.2 (Local isometry)**.** Let S and X be cubed n-complexes. Note that the link of a vertex of a cubed complex is a complex whose cells are simplices. We may therefore define a *local isometry* to be a combinatorial map $\varphi : S \to X$ such that, for every 0-cell $v \in S$, φ induces an injection of $\mathrm{Link}(v)$ onto a convex subcomplex of $\mathrm{Link}(\varphi(v))$. If $\varphi : S \to X$ is also an embedding, then we say that S is *locally convex* in X.

In the case where S and X are 2-complexes, $\varphi : S \to X$ is a local isometry precisely if:

(1) φ is an *immersion* (that is, the link of every 0-cell of S injects); and
(2) For all edges e_1, e_2 that meet at a vertex of S, if $\varphi(e_1)$ and $\varphi(e_2)$ bound a corner of a square f' in X, then e_1 and e_2 bound a square f in S such that φ maps the e_1-e_2 corner of f onto the $\varphi(e_1)$-$\varphi(e_2)$ corner of f'.

Note that any covering space $\hat{X} \to X$ is a local isometry. Conversely, if $\varphi : S \to X$ is a local isometry such that the link of every 0-cell $v \in S$ also *surjects* onto the link of $\varphi(v)$, then φ is a covering map. Furthermore, if $\hat{X}$ is a covering space of X, and S is a locally convex subspace of

$\hat{X}$, then since the composition of local isometries is a local isometry, the map $S \to X$ is also a local isometry.

Following Section 4, we now consider the restrictions of local isometries to several important types of subspaces. For the rest of this section, let X be the standard complex of a 2-dimensional graph group $G(\Gamma)$.

Definition 6.3 (Components)**.** Let F be a full subgroup of $G(\Gamma)$, and let $S \to X$ be a local isometry. An F-*component* of S is defined to be a component of the inverse image in S of the full subcomplex X_F (Definition 2.4). If $F \cong \mathbb{Z}^n$, then F-components are called F-flats, or n-flats; and if $F = C(t)$ for some standard generator t, then F-components are called t-pencils.

In particular, a 1-flat in S is either a line, a circle, or a *line segment.* Because of the possible "incompleteness" of the domain of a local isometry, line segments may actually have length 0; in fact, this will turn out to be the case in which we are most interested.

Note also that, as before, the product decomposition of $X_{C(t)}$ corresponding to the product $C(t) = F_n \times \langle t \rangle$ gives projections from every t-pencil in two directions; namely, a projection onto a graph with each fiber a 1-flat, and a projection onto a 1-flat with each fiber a graph. Now, *a priori*, it may not be clear that these projections actually give bundles, or in other words, that different fibers of the same pencil are related. However:

Lemma 6.4 (Segment pencils are products)**.** *Let P be a t-pencil of a local isometry $S \to X$, and let $\overline{P}$ be the base graph of the graph of spaces decomposition of P induced by splitting of $F_n \times \langle t \rangle$ corresponding to projection onto its first factor. Then P is locally a product (that is, P is actually a bundle). In particular, if the t-fiber of P is a line segment, then P is a product.*

Note that since t-pencils in local isometries are bundles, it makes sense to talk about "the" fiber of a t-pencil. We therefore say that a t-pencil is an *open pencil* if its t-fiber is a line, a *closed pencil* if its t-fiber is a circle, and a *segment pencil* if its t-fiber is a line segment.

Proof. Since $\overline{P}$ is connected, to show that P is locally a product, it is enough to show that the inverse image of an edge of $\overline{P}$ is a product. However, the local isometry condition implies that we may "parallel transport" each of the fibers onto each other isomorphically, as shown in Figure 4. (The heavy lines in Figure 4 represent fibers, and the lighter lines project down to an edge of $\overline{P}$.) It follows that the inverse image of an edge is a product. As for the final claim, if the fiber of P is a line segment, the local product must be a global product, since the orientation-preserving automorphism group of a line segment is trivial. □

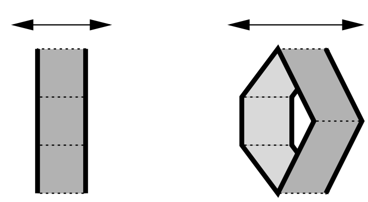

FIGURE 4. Parallel transport over an edge of $\overline{P}$

Remark 6.5. Lemma 6.4 can be generalized to pencils in n-dimensional graph groups; the main difference is that the base space $\overline{P}$ is an $n-1$-dimensional complex, not a graph. Otherwise, the lemma and its proof are essentially the same.

7. Locally convex cores

In this section, we begin to construct the cores we need to prove the Main Theorem.

Theorem 7.1 (Locally convex core)**.** *Let $\hat{X}$ be a quasifull cover of the standard complex X of a graph group $G(\Gamma)$ such that Γ is a tree (but not a point). For any finite set of 1-cells $C \subset \hat{X}$, there exists a subspace S of $\hat{X}$ such that:*

(1) *S is connected;*
(2) *S contains C;*
(3) *S is the union of finitely many 2-flats; and*
(4) *S is a locally convex subspace of $\hat{X}$.*

Now, by enlarging each of the 1-cells in C to the 1-flat containing it, we see that to obtain Theorem 7.1, it suffices to prove the following theorem.

Theorem 7.2. *For $n \geq 2$, let $\hat{X}_n$ be a quasifull cover of the standard complex X_n of a graph group G whose graph Γ_n is a tree with n vertices. For any finite collection L of 1-flats in $\hat{X}_n$, there exists a subspace S of $\hat{X}_n$ such that:*

(1) *S is connected;*
(2) *S contains L;*
(3) *S is the union of finitely many 2-flats; and*
(4) *S is a locally convex subspace of $\hat{X}_n$.*

Proof. Proceeding by induction, in the base case of $n = 2$, we may take $S = \hat{X}_n$ (a single 2-flat). Therefore, only the induction step remains.

First, to establish notation, since Γ_{n+1} is a tree, it can be built from some tree Γ_n with n vertices by adding a single vertex and a single edge from the new vertex to one of the old vertices. If X_n is the standard complex of the graph group $G_n = G(\Gamma_n)$, it follows that $G_{n+1} = G(\Gamma_{n+1})$ is an HNN extension of G_n by a stable letter t centralizing exactly one previous generator a. In other words, X_{n+1} is a graph of spaces with one vertex space, X_n, and one edge space, an $\langle a, t\rangle$-cylinder whose ends are identified with the a-circle in X_n. Note also that this induces a graph of spaces decomposition of $\hat{X}_{n+1}$. Let Λ be the underlying graph of this graph of spaces decomposition.

Now, by connecting the 1-flats of L with finitely many paths, and then enlarging each 1-cell of those paths to a 1-flat, we may assume that L is a connected collection of 1-flats. Next, observe that any 1-flat in $\hat{X}_{n+1}$ is either entirely contained in a unique vertex space, if it is a 1-flat corresponding to one of the generators of G_n, or entirely contained in a unique $\langle a, t\rangle$-flat. Let $\{V_i\}$ be the (finite) set of all vertex spaces of $\hat{X}_{n+1}$ that contain at least one of the 1-flats in L, and let $\{E_j\}$ be the (finite) set of $\langle a, t\rangle$-flats that entirely contain a $\langle t\rangle$-flat of L.

Now, for each V_i and each E_j, E_j intersects V_i in only finitely many $\langle a\rangle$-flats, for if E_j intersected V_i in infinitely many $\langle a\rangle$-flats, Λ would contain infinitely many independent loops (see Figure 5), $\pi_1\Lambda$ would contain an infinitely generated free factor, and $\pi_1\hat{X}_{n+1}$, which surjects onto $\pi_1\Lambda$, would not be finitely generated. Therefore, if L_i is the union of the set of all 1-flats in L that are entirely contained in V_i, and the set of all 1-flats that are in the intersection of V_i and some E_j, then L_i is a finite set of 1-flats in V_i. Furthermore, since V_i is a quasifull cover of X_n (Lemma 3.3), by induction, we may choose a subspace S_i of V_i that is connected, contains L_i, is the union of finitely many 2-flats, and is locally convex.

So now, let S be the union of the S_i and the E_j, as shown in Figure 6. We claim that S satisfies the required conditions. That is:

(1) S is connected, since each of the S_i and the E_j are connected to L, and L is connected.

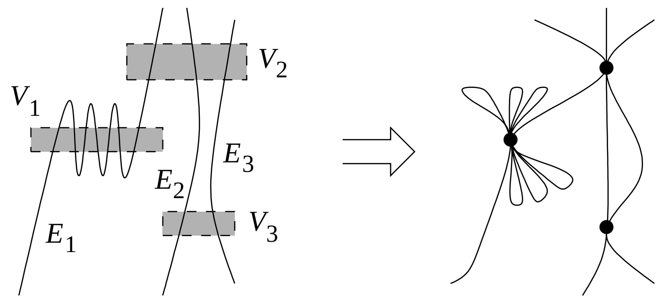

FIGURE 5. Projection to Λ of the E_j

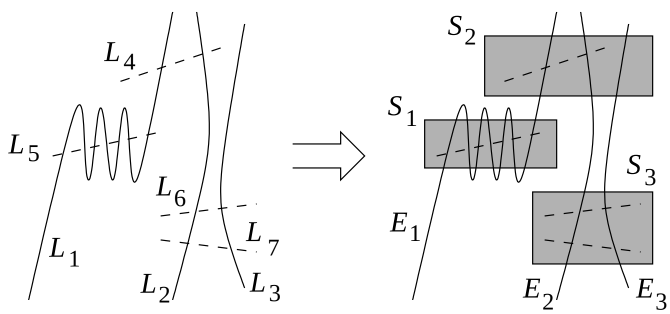

FIGURE 6. Completing L to a locally convex core

(2) S contains L, since every 1-flat in L is contained in either an S_i or an E_j.
(3) S is the union of finitely many 2-flats, since there are finitely many S_i, each of which is the union of finitely many 2-flats, and there are finitely many E_j.
(4) Finally, to see that S is locally convex, suppose that x and y are 1-cells of S that meet at a 0-cell and form the corner of some 2-cell of $\hat{X}_{n+1}$. We have two cases:
    (a) If either x nor y is a t 1-cell, then both x and y must be contained in some S_i. Therefore, the 2-cell they span must also be in $S_i \subset S$, since S_i is locally convex.
    (b) If x is a t 1-cell, then y is an a 1-cell, and both x and y must be contained in some E_j. However, since E_j is an entire 2-flat, it certainly contains the 2-cell spanned by x and y.

The theorem follows. □

Remark 7.3. Interestingly enough, Theorem 7.1 fails for $G(\Gamma) \cong F_2 \times F_2$. This is explained below in Example 11.4.

8. Closing off locally convex cores

Having constructed a locally convex core in Theorem 7.1, we now close it up to form a compact space that maps to X by a local isometry.

Theorem 8.1 (Closing up flats)**.** *Let X be the standard complex of a tree group G; let $\hat{X}$ be a quasifull cover of X; let $S \subseteq \hat{X}$ be connected, locally convex, and the union of finitely many 2-flats; and let C be the closure of a finite union of 1-cells in S. There exists a space S_0 and a combinatorial surjection $\varphi : S \to S_0$ such that:*

(1) *S_0 is a compact union of 2-flats;*
(2) *The local isometry $S \to X$ induced by the inclusion of S in $\hat{X}$ factors through a local isometry $\pi : S_0 \to X$; and*
(3) *The map $S \to S_0$ restricts to an embedding of C in S_0.*

Since S is the union of finitely many pencils, our strategy is to close S one open pencil at a time, or two, if necessary. Note that since every 2-flat of S is a 2-flat in $\hat{X}$, and every open pencil of S is a subspace of an open pencil in $\hat{X}$, our hypothesis and Lemmas 4.4 and 5.1 imply that the following conditions hold in S:

FF: S is the union of finitely many 2-flats.
TC: Every open t-pencil P of S contains at most one twisted cylinder.
FI: If two 2-flats E_1, E_2 intersect in S then either:
(a) E_1 and E_2 intersect in finitely many 0-cells;
(b) E_1 and E_2 are in the same t-pencil and intersect in finitely many t-fibers of the pencil; or
(c) E_1 and E_2 intersect in finitely many 0-cells and an infinite periodic sequence of 0-cells inside a single twisted cylinder.

Note that condition TC is slightly weaker than the conclusions of Lemma 4.4, in that P is also allowed to contain other (nontwisted) cylinders. We use this weaker condition because it is only the weaker version that will be preserved throughout our closing-up process.

In any case, it now follows by induction on the number of open pencils in S that to obtain Theorem 8.1, it suffices to prove:

Theorem 8.2. *Let $\pi' : S' \to X$ be a local isometry satisfying conditions FF, TC, and FI, and let C be the closure of a finite union of 1-cells of S'. Then if S' has open pencils, we may close up one or two open pencils of S' to obtain a quotient space S'' such that π' factors through a local isometry $\pi'' : S'' \to X$, conditions FF, TC, and FI hold in S'', and the inclusion of C in S' induces an inclusion of C in S''.*

Before we prove Theorem 8.2, we need the following lemma. We say that a flat inside a t-pencil is *transverse* if it contains no t 1-cells.

Lemma 8.3. *Let P be an open t-pencil of a local isometry $S' \to X$ satisfying condition FI, and for $i = 1, 2$, let $E_i \subset P$ be either a 0-cell or a transverse 1-flat. Furthermore, if E_i is a 1-flat ($i = 1, 2$), assume that E_i is not contained in a twisted cylinder. Then there can only be finitely many different lengths of t-paths between E_1 and E_2.*

Proof. Certainly this is clear if E_1 and E_2 are both 0-cells. Furthermore, if E_1 is a 0-cell and E_2 is a transverse 1-flat, then since E_2 is not contained in a twisted cylinder, there are no t-paths from E_2 to itself, and so there is at most one t-path from E_1 to E_2.

In the remaining case, if E_1 and E_2 are contained in the same 2-flat E, then every t-path between them has the same length, since E is not a twisted cylinder. On the other hand, if E_1 is an $\langle a \rangle$-flat and E_2 is a $\langle b \rangle$-flat, then the $\langle a, t \rangle$-flat containing E_1 and $\langle b, t \rangle$-flat containing E_2 intersect in only finitely many t-fibers (condition FI), and there can be only finitely many t-paths between E_1 and E_2. (Note that it is quite possible to have more than one kind of t-path between the same pair of 1-flats; for example, in Figure 7, compare any transverse 1-flat in the shaded plane with any transverse 1-flat in the white plane.) □

Proof of Theorem 8.2. Let P be an open t-pencil of S'. Since S' satisfies condition TC, we have two cases: either P contains no twisted cylinders, or P contains exactly one twisted cylinder.

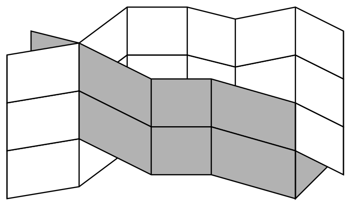

FIGURE 7. Different types of t-paths are possible

Case 1. Suppose P contains no twisted cylinders. Then from condition FI, the intersection of P and all 2-flats of S' not contained in P is a set of finitely many 0-cells and finitely many transverse 1-flats. Let D (the *deviant set* of P) be the union of these 0-cells and 1-flats and the 0-cells of $C \cap P$. The essential feature of the deviant set is that when we take a quotient of S' that closes off P, as long as we do not identify two deviant points, we will not be forced to make any identifications other than those inside P. From Lemma 8.3, we see that there is some integer M larger than the length of any t-path between any two 0-cells of D.

Let $Z = \langle z \rangle$ be an infinite cyclic group, and define an action of Z on S' by the rule that for $x \in S'$:

$$zx = \begin{cases} t^M x & \text{if } x \in P, \\ x & \text{if } x \notin P. \end{cases} \tag{3}$$

Note that the action may not be continuous. It may be helpful to think of taking the quotient S'/Z in two stages: namely, truncating all of the t-fibers of P to length M, and then gluing together the resulting loose ends. This process is shown in Figure 8, in which the deviant set D is represented by the two thickened line segments and the thickened 0-cell. (Note that the truncation in the t direction need not be uniform.)

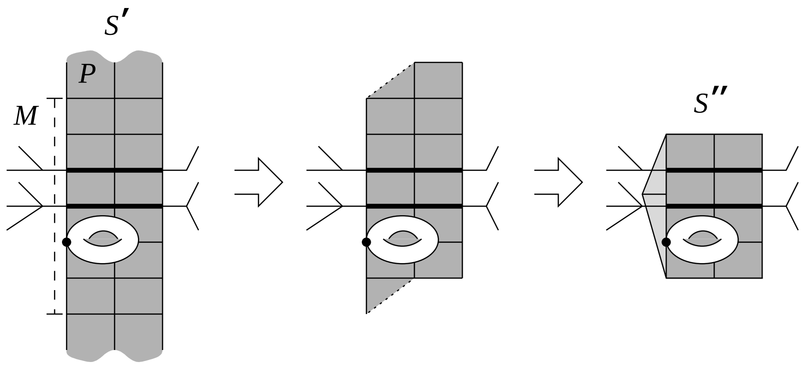

FIGURE 8. The quotient map $S' \to S''$

In any case, let $S'' = S'/Z$. We claim that S'' satisfies the desired conditions. First, since the t-distance between any two 0-cells in D is less than M, 0-cells in D are only identified with 0-cells outside of D. Therefore, every 0-cell of S'' is locally modelled on a point of S', and the induced map $\pi'' : S'' \to X$ is a local isometry. Second, since S' is a finite union of flats, and S'' is obtained from S' by replacing planes and cylinders with cylinders and tori, S'' satisfies condition FF. Third, since the quotient procedure only produces untwisted cylinders, condition TC still holds in S''.

Fourth, since the quotient procedure produces no new intersections between 2-flats, condition FI still holds in S''. Finally, since the quotient procedure does not identify any two 0-cells of D, the set C embeds in S''.

Case 2. Suppose P contains a single twisted cylinder E. Let $t_1 = t$, let $P_1 = P$, let E be the twisted $\langle t_1, t_2\rangle$-cylinder in P_1, and let P_2 be the t_2-pencil of S' that contains E. Now, for $i = 1, 2$, outside of the twisted cylinder E, the intersection of P_i and all 2-flats of S' not contained in P_i is a set of finitely many 0-cells and finitely many transverse 1-flats, as before. Let D_i (the deviant set of P_i) be the union of these 0-cells, these transverse 1-flats *not* contained in E, and the 0-cells of $C \cap P_i$. (Transverse 1-flats coming from intersections in E are not included in D_i, because they will not force any identifications other than the ones we are already planning to make.)

Now, recall from Lemma 5.1 that for $i = 1, 2$, E is periodic with some period m_i in the t_i direction. Applying Lemma 8.3, we choose some constant c such that if $M_i = cm_i$ $(i = 1, 2)$, then M_i is greater than the t_i-distance between any two 0-cells in D_i (the deviant set of P_i). We then define an action of the group $Z = \langle z_1, z_2\rangle \cong \mathbb{Z}^2$ on S' by the rule that for $x \in S'$:

$$z_i x = \begin{cases} t_i^{M_i} x & \text{if } x \in P_i, \\ x & \text{otherwise.} \end{cases} \tag{4}$$

Note that the action may not be continuous. It remains to show that (4) gives a well-defined action of Z on S' and that the resulting quotient satisfies the desired conditions.

To show that (4) defines an action of Z on S', it suffices to show that $z_1z_2x = z_2z_1x$ for all $x \in S'$. Now, since t_1 and t_2 commute, the local isometry condition on S' implies that $P_1 \cap P_2$ is a union of $\langle t_1, t_2\rangle$-flats. (More precisely, $P_1 \cap P_2$ is the union of the twisted cylinder E and finitely many $\langle t_1, t_2\rangle$-planes.) Therefore, since the only points of P_2 moved by z_1 are points inside $P_1 \cap P_2$, and since the action of z_1 sends each $\langle t_1, t_2\rangle$-flat bijectively to itself, we see that z_1 sends P_2 bijectively to itself. It follows that for all $x \in S'$, z_2 moves z_1x if and only if it moves x, and by the same motion. By symmetry, we also see that for all $x \in S'$, z_1 moves z_2x if and only if it moves x, and by the same motion. We conclude that $z_1z_2x = z_2z_1x$ for all $x \in S'$, and that the action of Z is well-defined.

Let $S'' = S'/Z$ be the resulting quotient. As in the untwisted case, we may obtain this quotient by truncating all of the t_i-fibers of P_i to length M_i, and then gluing together the resulting loose ends. Note that since we have chosen $M_i = cm_i$, and since the twisted cylinder E is periodic in the t_i-direction with period m_i, the truncation on E can be done in a manner consistent with the truncation on both P_1 and P_2. We then see that $S'' \to X$ is a local isometry satisfying conditions FF, TC, and FI, and embedding C, by essentially the same reasoning as in the untwisted case. The theorem follows. □

9. Completing local isometries to finite covers

Having obtained a compact space S_0 that embeds C, we now complete S_0 to a covering space of X.

Theorem 9.1 (Extending to a finite cover)**.** *Let X be the standard complex of a 2-dimensional graph group, and let $\varphi : S_0 \to X$ be a local isometry where S_0 is compact. Then S_0 may be completed to a finite cover of X.*

Proof. If S_0 is not already a covering space, then it must contain at least one segment pencil P. From Lemma 6.4, we see that P is the product of a graph Λ and a line segment σ of length n. Therefore, by adding 1-cells in the t direction and the appropriate 2-cells, we may complete P to the product of Λ and a circle of length $n + 1$, as shown in Figure 9. We thereby obtain a space S_0' containing S_0, and a natural extension $\varphi' : S_0' \to X$ of the map $\varphi : S_0 \to X$.

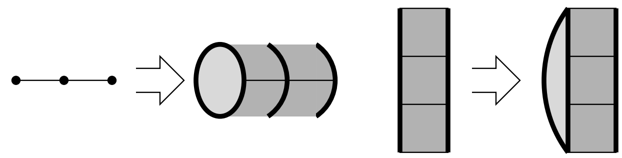

FIGURE 9. Completing pencils whose fibers are segments of length 0 and 3

Now, clearly φ' is an immersion, so to show that φ' is a local isometry, we need only show that if e_1 and e_2 are 1-cells that meet at a vertex of S_0', and $\varphi'(e_1)$ and $\varphi'(e_2)$ bound a corner of a square f' in X, then e_1 and e_2 bound a square f in S_0' that maps to f'. However, since this condition is true if both 1-cells are contained in S_0, we need only check this condition in the case where one of the 1-cells is one of the new 1-cells in the t direction. In that case, since e_1 and e_2 are both contained in the completion of P, the condition is satisfied, as the completion of P is the product of two graphs, and is therefore a local isometry.

We therefore conclude that there exists a space S_0' that contains S_0 and has one fewer segment pencil than S_0, since we have created no new pencils, and have "closed up" one of the old ones. The theorem follows by induction on the number of segement pencils in S_0. □

Remark 9.2. Note that since Lemma 6.4 is all we really use in the proof of Theorem 9.1, and Lemma 6.4 works in n dimensions, and not just 2 dimensions (see Remark 6.5), Theorem 9.1 may also be generalized to n dimensions. Note also that in the proof of the Main Theorem, the only segment pencils arising from the output of Theorems 7.1 and 8.1 are pencils whose fibers are segments of length 0, so we really only use a very special case of Theorem 9.1.

10. Proof of the Main Theorem

We now collect the above results to obtain the Main Theorem. In particular, because of Theorem 3.9, the following implies the Main Theorem.

Theorem 10.1 (Separable)**.** *If H is a quasifull subgroup of a tree group G, then H is separable.*

Proof. Let X be the standard complex of G; let $\hat{X}$ be the based cover of X corresponding to H; let $\{h_i\}$ be a finite set of based paths in $X^{(1)}$ representing generators for H; let g be a closed path in $X^{(1)}$ that represents an element of $G - H$; and let C be the union of the based lifts $\hat{h}_i$ and $\hat{g}$ of these paths to $\hat{X}$. Note that each of the $\hat{h}_i$ is closed, but $\hat{g}$ is not.

Since $\mathbb{Z}$ is subgroup separable, we may assume that $G \not\cong \mathbb{Z}$. In that case, Theorem 7.1 gives a connected locally convex finite union of 2-flats $S \subset \hat{X}$ such that C is contained in S. Applying Theorem 8.1 to S and C, we obtain a local isometry $\pi : S_0 \to X$ such that S_0 is a compact union of 2-flats and C embeds in S_0 in a manner compatible with its embedding in $\hat{X}$. Finally, applying Theorem 9.1 to S_0, we obtain a finite cover $\overline{X}$ of X in which C embeds.

Let $G_0 = \pi_1\overline{X}$. Since $\overline{X}$ is a finite cover of X, G_0 has finite index in G. However, since C embeds in $\overline{X}$, we see that the paths h_i lift to closed paths in $\overline{X}$, but the path g does not lift to a closed path in $\overline{X}$. Consequently, $H \leq G_0$ and $g \notin G_0$. The theorem follows. □

11. Examples

In this section, we discuss some examples of interest, both to illustrate applications of the Main Theorem (Examples 11.1 and 11.2) and also to describe some possible obstructions to its generalization (Example 11.4).

Example 11.1. As mentioned in the introduction, if $J =$ 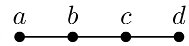

, then the Main Theorem implies that every quasiconvex subgroup of J is separable. In contrast, recall that J is not subgroup separable, even though J is a 3-manifold group; in fact, J is a subgroup of every currently known non-subgroup separable 3-manifold group [NW].

Example 11.2. Let

$$B = \langle a, b, t \mid [a, b], a^t = b \rangle, \tag{5}$$

as mentioned in the introduction. Note also that K, the standard 2-complex of B, is a nonpositively curved squared 2-complex, which means that we may consider its CAT(0)-quasiconvex subgroups. We now prove Theorem 1.2 from the introduction, in the following form.

Theorem 11.3. *Every CAT(0)-quasiconvex subgroup of B is separable.*

Proof. Let X be the standard 2-complex of the presentation of the graph group J from Example 11.1. It was shown in [NW] that there exists a double cover $\hat{K} \to K$ and a local isometry $\varphi : \hat{K} \to X$, which implies that there exists an isometric embedding of universal covers $\tilde{K} \hookrightarrow \tilde{X}$. It follows that φ_* embeds every CAT(0)-quasiconvex subgroup of $\pi_1 \hat{K}$ as a CAT(0)-quasiconvex subgroup of $\pi_1 X$; in particular, we may consider $\pi_1 \hat{K}$ to be a CAT(0)-quasiconvex subgroup of $\pi_1 X$.

Let Q be a CAT(0)-quasiconvex subgroup of $\pi_1 K$. Since CAT(0)-quasiconvexity is invariant under passing to subgroups of finite index [AB95, §10], $Q \cap \pi_1 \hat{K}$ is a CAT(0)-quasiconvex subgroup of $\pi_1 \hat{K}$ and hence of $\pi_1 K$. As shown above, we see that $Q \cap \pi_1 \hat{K}$ is also CAT(0)-quasiconvex in $\pi_1 X$. The Main Theorem then implies that $Q \cap \pi_1 \hat{K}$ is the intersection of finite index subgroups of $\pi_1 X$, and therefore, the intersection of finite index subgroups of $\pi_1 \hat{K}$ (take intersections with $\pi_1 \hat{K}$). In other words, $Q \cap \pi_1 \hat{K}$ is separable in $\pi_1 \hat{K}$. Since $\pi_1 \hat{K}$ has finite index in $\pi_1 K$, it follows easily that Q is separable in $\pi_1 K$ and we are done. □

Example 11.4. Finally, we consider an example showing that Theorem 7.1, as stated, does not apply to the 2-dimensional graph group

$$G = \langle a, b, c, d \mid [a, b], [b, c], [c, d], [d, a] \rangle \cong F_2 \times F_2. \tag{6}$$

As usual, let X be the standard complex of G.

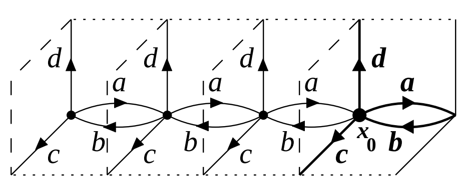

FIGURE 10. Failure of Theorem 7.1 for $F_2 \times F_2$

Let $H = \langle ab \rangle$, let $\hat{X} \to X$ be the cover corresponding to H, and let C be the union of the a, b^{-1}, c, and d 1-cells of $\hat{X}$ that touch the basepoint x_0, as indicated by the heavy solid 1-cells in Figure 10. Let S be a locally convex subset of $\hat{X}$ that contains C. Since H is cyclic, H is CAT(0)-quasiconvex, and therefore quasifull, but as we shall see, S is not contained in any finite union of 2-flats.

First, applying local convexity, we see that S must contain the shaded 2-cells of Figure 10. Next, let A (resp. B) be the a-line (resp. b-line) through x_0, and note that $A \cup B$ is a chain of

ab-loops. For each such ab-loop in S, local convexity implies that the two $\langle a, b\rangle$ 2-cells attached to the ab-loop are contained in S, which in turn implies that the adjacent ab-loops are also contained in S. It follows that all of $A \cup B$ is contained in S. So now, applying local convexity up and down $A \cup B$, we see that S must contain all of the $\langle a, d\rangle$ 2-cells touching A and all of the $\langle b, c\rangle$ 2-cells touching B, as indicated by the dotted lines in Figure 10. Finally, for each 0-cell v in $A \cup B$, local convexity implies that the $\langle c, d\rangle$ 2-cell at v must also be contained in S, as indicated by the dashed lines in Figure 10. Therefore, since no two of these $\langle c, d\rangle$ 2-cells are contained in the same 2-flat, S is not contained in a finite union of 2-flats.

12. Open questions

We conclude with some open questions.

(1) What is the precise relationship among the quasiconvexity conditions defined in Section 3? In a graph group, CAT(0)-quasiconvexity does not imply word-quasiconvexity; does word-quasiconvexity imply CAT(0)-quasiconvexity? Every CAT(0)-quasiconvex subgroup of a graph group is quasifull; is the converse true?

(2) Does the Main Theorem generalize to graph groups with *chordal* graphs? Recall that a graph Γ is said to be *chordal* if Γ contains no cycles as full subgraphs, or equivalently, if any cycle in Γ of length ≥ 4 has a "shortcut". We believe that:

Conjecture 12.1. *Theorem 10.1 is true for graph groups whose underlying graphs are chordal. That is, if Γ is a chordal graph, and H is a quasifull subgroup of Γ, then H is separable.*

Note that this class of graph groups does have some special properties. Most notably, a graph group is *coherent* (every finitely generated subgroup finitely presented) if and only if its graph is chordal [Dro87a].

It is conceivable that Conjecture 12.1 can be approached using some elaboration of the method of this paper. More specifically, given a suitable generalization of the material in Sections 4 and 5, it seems that the basic strategy used in Sections 7–10 should go through.

(3) Does the Main Theorem generalize to all graph groups? In fact, we believe that:

Conjecture 12.2. *Theorem 10.1 is true for any graph group. That is, if H is a quasifull subgroup of a graph group, then H is separable.*

Besides the modifications to Sections 4 and 5 mentioned above, Conjecture 12.2 would also require a suitable replacement for the material in Section 7, due to Example 11.4. Now, it is worth noting that in Example 11.4, even though S cannot be chosen to be a finite union of 2-flats, it can be chosen to be periodic; perhaps, in the general case, the analogue of the core obtained in Section 7 can be chosen to be the union of disjoint periodic subspaces and finitely many flats. On the other hand, the geometry of Example 11.4 seems to indicate that our inductive approach in Section 7 does not generalize, so a new, non-inductive approach would be required as well.

(4) Does the Main Theorem generalize in other directions? For example, does it hold for Artin groups whose graphs are trees?

References

[AB95] Juan M. Alonso and Martin R. Bridson. Semihyperbolic groups. *Proc. London Math. Soc. (3)*, 70:56–114, 1995.

[Bau81] A. Baudisch. Subgroups of semifree groups. *Acta Math. Acad. Sci. Hungar.*, 38(1-4):19–28, 1981.

[BB97] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.*, 129(3):445–470, 1997.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.

[BKS87] R. G. Burns, A. Karrass, and D. Solitar. A note on groups with separable finitely generated subgroups. *Bull. Austral. Math. Soc.*, 36(1):153–160, 1987.

[BN74] Andreas Blass and Peter M. Neumann. An application of universal algebra in group theory. *Michigan Math. J.*, 21:167–169, 1974.

[Dro87a] Carl Droms. Graph groups, coherence, and three-manifolds. *J. Algebra*, 106(2):484–489, 1987.

[Dro87b] Carl Droms. Isomorphisms of graph groups. *Proc. Amer. Math. Soc.*, 100(3):407–408, 1987.

[Dro87c] Carl Droms. Subgroups of graph groups. *J. Algebra*, 110(2):519–522, 1987.

[Gre90] Elisabeth R. Green. *Graph Products of Groups*. PhD thesis, University of Leeds, 1990.

[Hal49] Marshall Hall, Jr. Coset representations in free groups. *Trans. Amer. Math. Soc.*, 67:421–432, 1949.

[HM95] Susan Hermiller and John Meier. Algorithms and geometry for graph products of groups. *J. Algebra*, 171(1):230–257, 1995.

[Hum94] S. P. Humphries. On representations of Artin groups and the Tits conjecture. *J. Algebra*, 169:847–862, 1994.

[HW99] Tim Hsu and Daniel T. Wise. On linear and residual properties of graph products. *Michigan Math. J.*, 46(2):251–259, 1999.

[Lau95] Michael R. Laurence. A generating set for the automorphism group of a graph group. *J. London Math. Soc. (2)*, 52(2):318–334, 1995.

[LS77] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Springer-Verlag, Berlin, 1977. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.

[MV95] John Meier and Leonard VanWyk. The Bieri-Neumann-Strebel invariants for graph groups. *Proc. London Math. Soc. (3)*, 71(2):263–280, 1995.

[NR98] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.

[NW] Graham A. Niblo and Daniel T. Wise. Subgroup separability, knot groups and graph manifolds. *Proc. Amer. Math. Soc.*

[NW98] Graham A. Niblo and Daniel T. Wise. The engulfing property for 3-manifolds. In *The Epstein birthday schrift*, pages 413–418 (electronic). Geom. Topol., Coventry, 1998.

[RW98] J. Hyam Rubinstein and Shicheng Wang. π_1-injective surfaces in graph manifolds. *Comment. Math. Helv.*, 73(4):499–515, 1998.

[Sco78] Peter Scott. Subgroups of surface groups are almost geometric. *J. London Math. Soc. (2)*, 17(3):555–565, 1978.

[Ser89] Herman Servatius. Automorphisms of graph groups. *J. Algebra*, 126(1):34–60, 1989.

[Sho91] Hamish Short. Quasiconvexity and a theorem of Howson's. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 168–176. World Sci. Publishing, River Edge, NJ, 1991.

[SW79] Peter Scott and Terry Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge, 1979.

[VW94] Leonard Van Wyk. Graph groups are biautomatic. *J. Pure Appl. Algebra*, 94(3):341–352, 1994.

[Wis96] Daniel T. Wise. *Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups*. PhD thesis, Princeton University, 1996.

[Wis98] Daniel T. Wise. Subgroup separability of the figure 8 knot group. Preprint, 1998.

Department of Mathematics, Pomona College, Claremont, CA 91711
E-mail address: timhsu@pccs.cs.pomona.edu

Department of Mathematics, Cornell University, Ithaca, NY 14853
E-mail address: daniwise@math.cornell.edu