Michigan Math. J. 53 (2005)

On the Classification of CAT(0) Structures for the 4-String Braid Group

JOHN CRISP & LUISA PAOLUZZI

1. Introduction

This paper is concerned with the class of so-called CAT(0) groups, namely, those groups that admit a *geometric* (i.e., properly discontinuous, co-compact, and isometric) action on some CAT(0) space. More precisely, we are interested in knowing to what extent it is feasible to classify the geometric CAT(0) actions of a given group (up to, say, equivariant homothety of the space). A notable example of such a classification is the flat torus theorem, which implies that the minimal geometric CAT(0) actions of the free abelian group $\mathbb{Z}^n$ ($n \geq 1$) are precisely the free actions by translations of Euclidean space $\mathbb{E}^n$.

Typically, however, a given group will have uncountably many nonequivalent actions, making any chance of a complete classification rather slim. It is therefore reasonable to consider, for example, only those actions of a group on a space of minimal possible dimension—namely, the geometric dimension of the group. Thus, the geometric actions of the free group F_n, $n \geq 2$, on 1-dimensional CAT(0) spaces ($\mathbb{R}$-trees) are classified by the compact metric graphs of Euler characteristic $1 - n$. Even so, the variety of, say, the 2-dimensional CAT(0) structures for closed surface groups would seem to be vast, with many of these structures being nonplanar (see Section 5.1). By contrast, we are able to cite examples of CAT(0) groups of geometric dimension 2 that have no 2-dimensional CAT(0) structure [7; 3]. Other results in a similar vein are to be found in [2], [12], and [11].

Consider the n-string braid group B_n defined by the following presentation:

$$B_n = \langle a_1, a_2, \ldots, a_{n-1} \mid a_i a_{i+1} a_i = a_{i+1} a_i a_{i+1} \text{ for } i = 1, \ldots, n-2;\ a_i a_j = a_j a_i \text{ if } |i-j| \geq 2 \rangle.$$

The braid group B_n has geometric dimension $n - 1$, and it is known to be CAT(0) (with a structure in dimension $n - 1$) if $n \leq 5$ [5; 4; 6]. The question of existence of a CAT(0) structure is open for $n \geq 6$. We note that each braid group has infinite cyclic center (see [9]).

The group B_3 acts geometrically on the product $T \times \mathbb{R}$, where T is a trivalent tree, and up to an independant homothety of each factor of the product this is the unique minimal CAT(0) structure. This fact was used by Hanham in his thesis [12]. We give a proof here (see Theorem 4) that simply reduces the question to

Received November 26, 2003. Revision received July 27, 2004.

studying 1-dimensional actions of the virtually free group $C_2 * C_3$, which is the quotient of B_3 by its center. Note that Theorem 4 and its proof generalize easily to Artin groups of type $I_2(p)$, p prime, but a similar result is certainly not true for Artin groups of type $I_2(2k)$ (a number of different CAT(0) structures for these latter groups are presented in [12]).

The main focus of this paper is the group

$$B_4 = \langle a, b, c \mid aba = bab,\ bcb = cbc,\ ac = ca \rangle, \tag{1}$$

or (more precisely) the quotient of this group by its center, which we shall write as

$$G = B_4/Z(B_4).$$

Note that $Z(B_4)$ is generated by the element $z = (bac)^4$. By analogy with the case of B_3, our first observation (see Proposition 5) is that the classification of 3-dimensional CAT(0) structures for the braid group B_4 is essentially equivalent to the classification of 2-dimensional CAT(0) structures for G.

In [4], T. Brady showed that the braid group B_4 is a CAT(0) group. His construction, which we describe in detail in Section 3, yields a geometric action of the group G on a CAT(0) 2-complex that we denote X_0. The complex X_0 carries a piecewise Euclidean metric and is triangulated by equilateral triangles. We shall refer to X_0 as *Brady's complex* and to the action of G on X_0 as *Brady's action,* or simply the *standard action*. The main result of this paper is the following theorem.

THEOREM 1. *Suppose that G acts geometrically on a* 2*-dimensional* CAT(0) *space X. Then there exists a uniquely determined map $F\colon X_0 \to X$ with the properties that*:

(i) *F is G-equivariant with respect to the standard action on X_0; and*
(ii) *F is locally injective and, up to a constant scaling of the metric on X (or X_0), may be supposed to be locally isometric on the complement of the* 0*-skeleton of X_0.*

REMARKS. (1) Note that we do not restrict X to be a simplicial or polyhedral complex. Here and throughout the paper, "n-dimensional" means "covering dimension n". We refer to [13] for the theory of covering dimensions.

(2) We note that our proof of Theorem 1 actually needs only the slightly weaker hypothesis that the action of G on X is properly discontinuous and semisimple (rather than geometric). This is because our main tool, the flat torus theorem, uses only these hypotheses. Fujiwara, Shioya, and Yamagata have recently given a refinement of the flat torus theorem [11, Sec. 5.1] that would, alternatively, allow us to consider arbitrary proper isometric actions of G on a 2-dimensional space X, as long as we require that the space X be proper (i.e., that closed balls in X be compact).

It is not too hard to construct a map F as in Theorem 1 that fails to be a local isometry (by using the fact that the vertex links in X_0 have diameter strictly greater

than π). One might wonder, however, whether the canonical map F is still injective in all cases. The following result shows that the statement of Theorem 1 is really the best possible.

THEOREM 2. *There exist geometric actions of G on 2-dimensional* CAT(0) *spaces such that the map F of Theorem 1 fails to be injective. More precisely, for any positive real R, one can find a geometric action of G on a* CAT(0) *piecewise Euclidean 2-complex X such that*:

(i) *the canonical map $F\colon X_0 \to X$ of Theorem 1 is a local isometry on the complement of the 0-skeleton of X_0;*
(ii) *the map F is not injective and identifies two distinct orbits of X_0 that are a distance at most $1/R$ apart; and*
(iii) *there exists an element of G whose translation length on X_0 is greater than R but whose translation length on X is at most $1/R$.*

Finally, as an application of Theorem 1, we are able to show the following.

THEOREM 3. *The group $G = B_4/Z(B_4)$ is co-Hopfian (i.e., every monomorphism $G \to G$ is an isomorphism).*

The proof of this statement reduces, by Theorem 1, to showing that it is not possible to isometrically embed an enlarged copy of Brady's complex X_0 back into X_0.

In a similar vein, one obtains from Theorem 1 that the group G admits a unique geometric action on X_0 (up to G-equivariant isometry). One obtains as an immediate corollary the fact that B_4 has outer automorphism group of order 2 (generated by the automorphism that simply inverts the standard generators). This fact is well known for *all* braid groups and was first proved by Dyer and Grossman [10].

Note added. It has since been proved by Bell and Margalit [1] that $B_n/Z(B_n)$ is co-Hopfian for all $n \geq 4$. Their method consists of viewing these groups as mapping class groups of punctured spheres and then building on the techniques developed by Ivanov and McCarthy in their treatment of mapping class groups of closed surfaces.

The structure of this paper is as follows. In Section 2 we discuss the 3-string braid group and prove the classification of CAT(0) structures in this case. In Section 3 we introduce the 4-string braid group, describe Brady's action, and explain the reduction of the classification of CAT(0) structures to the study of the central quotient G. In Section 4 we give the proof of Theorem 1. In Section 5 we describe a technique for obtaining new CAT(0) structures for a group by essentially distorting a given structure and hence obtain Theorem 2. Finally, in Section 6 we show that the group G is co-Hopfian (Theorem 3).

ACKNOWLEDGMENT. We would like to thank Luis Paris for raising the question of co-Hopficity of the group G, and Chris Hruska who gave us some valuable help with the proof of Theorem 3.

2. The 3-String Braid Group B_3

We begin by recalling that the group $B_3 = \langle a, b \mid aba = bab \rangle$ is isomorphic to

$$B(u, v, w) = \langle u, v, w \mid uv = vw = wu \rangle.$$

The isomorphism is given by $a \leftrightarrow u$, $b \leftrightarrow v$. The group presentation $B(u, v, w)$ may be represented schematically by a triangle in the plane with edges labeled u, v, w, as shown in Figure 1, the relations being read in the clockwise direction always. (Note that one obtains an isomorphism between $B(u, v, w)$ and the braid group on the vertices of the triangle of Figure 1 by mapping each generator to the positive braid twist on the corresponding edge of the triangle—that is, a braid twist performed in an arbitrarily small neighborhood of the edge that exchanges its endpoints.)

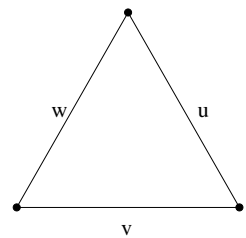

Figure 1 The triangular relations $uv = vw = wu$ for $B(u, v, w)$

The easiest way to exhibit a proper action of B_3 on a CAT(0) space is via the following construction of Brady and McCammond [5]. Let K denote the presentation complex for the presentation $\langle u, v, w, t \mid uv = t, vw = t, wu = t \rangle$. Let $0 < \theta < \pi/2$, and suppose that K is built of three Euclidean isosceles triangles with base angle θ, so that the edges u, v, and w have length 1 and the edge t has length $2 \cos \theta$. This is illustrated in Figure 2.

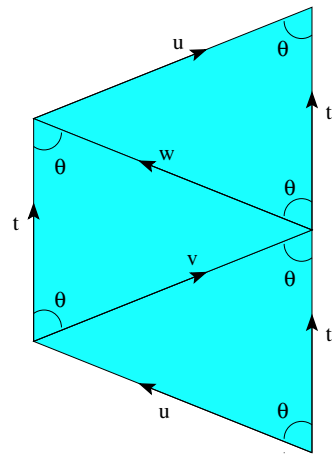

Figure 2 A locally CAT(0) Eilenberg–MacLane space K for $B(u, v, w)$

NOTATION. We write $X_\theta(u, v, w)$ (or simply $X(u, v, w)$) to denote the universal cover of K endowed with its piecewise Euclidean metric and equipped with the standard action of $B(u, v, w)$ by covering translations. This action is proper, co-compact, and isometric.

The metric space $X_\theta(u, v, w)$ is easily seen to be a CAT(0) space. In fact, it is the direct product $\mathbb{R} \times T_\theta$ of a real line and a regular trivalent tree T_θ of edge length $\sin\theta$. Projection onto the tree factor induces an action $\alpha\colon B(u, v, w) \to \mathrm{Isom}(T_\theta)$ in which u, v, w are represented by hyperbolic isometries of translation length $\sin\theta$, as shown in Figure 3.

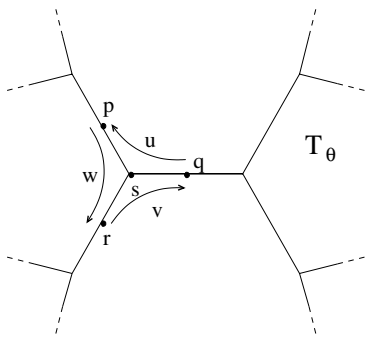

Figure 3 Action of $B(u, v, w)$ on T_θ

In the figure, s denotes a vertex of the tree and p, q, r the midpoints of the three edges adjacent to s. The action is such that $p = u(q)$, $q = v(r)$, and $r = w(p)$, and the element $uv = vw = wu$ leaves the vertex s fixed while mapping $p \mapsto q \mapsto r \mapsto p$. The kernel of α is the center of $B(u, v, w)$, the cyclic subgroup generated by $(uv)^3$, and the image of α is a free product $C_3 \star C_2$ (generated by $\alpha(uv)$ and $\alpha(uvu)$; the element uvu acts by an involution fixing the point q).

The following result appears in [12], where it is deduced largely from Proposition 2.2 of [3]. Note that, if a group G acts on a CAT(0) space X, then we may suppose (by a linear scaling of the metric on X) that any given hyperbolic element of G has translation length 1.

THEOREM 4 [12]. *Any minimal proper semisimple action of* $B_3 = B(u, v, w)$ *on a* 2*-dimensional* CAT(0) *space in which the element* u *has translation length* 1 *is, for some* $0 < \theta < \pi/2$, *isometric to the standard action of* $B(u, v, w)$ *on* $X_\theta(u, v, w)$.

Proof. Suppose one has such an action on a 2-dimensional CAT(0) space X. Recall that the center of $B(u, v, w)$ is generated by the element $\zeta = (uv)^3$. Since the action on X is minimal, we have $X = \mathrm{Min}(\zeta)$. It follows (cf. [8; 3, Prop. 1.2]) that $X = \mathbb{R} \times T$, where T is a 1-dimensional CAT(0) space (an $\mathbb{R}$-tree), the central element acts by translation along the $\mathbb{R}$-fibres, and projection onto T induces a proper

semisimple and minimal action of $B(u,v,w)/\langle\zeta\rangle$ on T. Let $x_3 = uv$ and $x_2 = uvu$ denote the elements of order 3 and 2, respectively, that generate $B(u,v,w)/\langle\zeta\rangle$ as a free product $C_3 \star C_2$. Let F_3 denote the fixed point set of x_3 in T, with F_2 that of x_2. These are closed convex sets, and they are disjoint from one another because the action is proper and can have no global fixed point. Hence there is a nontrivial closed interval $J = [s,q]$ of length L in T such that $J \cap F_3 = \{s\}$ and $J \cap F_2 = \{q\}$. The union of all translates of J now forms a subtree of T (necessarily convex in dimension 1) that is equivariantly isometric to T_θ, where $\sin(\theta) = 2L$. By minimality of the action we must have that $T = T_\theta$. Finally, the fact that u, v, w are conjugate elements implies that they have equal translation length when projected to either factor of $X = \mathbb{R} \times T$, in particular when projected to $\mathbb{R}$. Since uv is a root of the central element, it translates nontrivially in this central direction; we conclude that u, v, and w must all translate in the same direction along $\mathbb{R}$. In fact, positive axes for these elements span an angle of θ with the positive central axes. It now follows that the action of $B(u,v,w)$ is isometric to the standard action on $X_\theta(u,v,w)$. □

3. The 4-String Braid Group B_4 and Brady's Action

We begin with a brief discussion of the structure of the group B_4. It is easily seen that the presentation (1) for the 4-string braid group is equivalent to the following presentation, where a, b, c represent the same elements as before:

$$\begin{aligned} B_4 = \langle a,b,c,d,e,f \mid\ & ba = ae = eb,\ de = ec = cd, \\ & bc = cf = fb,\ df = fa = ad, \\ & ca = ac,\ ef = fe\rangle. \end{aligned}$$

Note that, in fact, we have simply introduced generators $e = a^{-1}ba$, $f = c^{-1}bc$, and $d = (ac)^{-1}b(ac)$. This presentation corresponds to the complete labeled graph Γ shown in Figure 4(i), where (as with the triangular presentation of B_3) one may view B_4 as the braid group on the four vertices of this complete graph by associating each generator a, b, c, d, e, f with the positive braid twist on the corresponding edge.

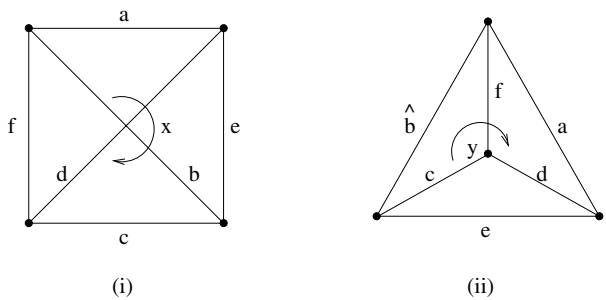

Figure 4 Presentations of B_4: (i) x-symmetric; (ii) y-symmetric

The foregoing presentation has the advantage of being symmetric with respect to the conjugation action of $x = bac$. This element acts on Γ by a quarter clockwise turn, so

$$x\colon a \mapsto e \mapsto c \mapsto f \mapsto a \quad \text{and} \quad x\colon b \leftrightarrow d.$$

We can modify the presentation of B_4 yet again to one that reveals a 3-fold symmetry. This new presentation is described schematically by the labeled graph Γ' of Figure 4(ii), where $\hat{b} = c^{-2}bc^2 = f^2bf^{-2}$. The conjugation action of the element $y = xc = fx$ induces a one-third clockwise rotation of Γ', so

$$y\colon c \mapsto f \mapsto d \mapsto c \quad \text{and} \quad y\colon a \mapsto e \mapsto \hat{b} \mapsto a.$$

Recall that $x^4 = y^3 = z$ generates the center Z of B_4 and that x, y and their powers represent all conjugacy classes of finite-order elements in $G = B_4/Z$. This last fact may be observed by considering the action of G on the space X_0 described below.

Observe that the subgroup of B_4 generated by b, c, and f is an isomorphic image of the group $B(b, c, f)$ by the obvious map. The subgroup in question is actually the standard parabolic subgroup generated by b, c. Moreover, this subgroup intersects the center of B_4 trivially and so maps injectively to G under the quotient by Z. (To see this, note that an arbitrary central element of the subgroup is written $(bc)^{3k}$ for some integer k and is conjugate in B_4 to $(ba)^{3k}$ (by the element x^2) but not equal to $(ba)^{3k}$.) By considering the conjugations by x and y, it is clear that similar statements may be made for each of the triangles appearing in either graph Γ or Γ'.

NOTATION. Henceforth, we shall write $B(u, v, w)$ for the subgroup of $G = B_4/Z$ generated by the images of elements u, v, w that form a triangle in either Γ or Γ'. Moreover, we make no distinction in our notation between an element in B_4 and its image in G. This will create no confusion, as from now on we shall be concerned only with the latter group.

We conclude this section with a description of Brady's action of B_4 on Y_0 (and of the quotient action of G on X_0). Note that there are exactly sixteen ways of writing x as a product of three of the generators $a, \ldots, f$ (see [4]). These expressions are exactly the length of three subwords of the following words of length 12 (viewed cyclically):

$$W_1 = bcadefbacdfe; \qquad W_2 = faecfaecfaec.$$

(Note that these words are expressions for the central element $z = x^4 = y^3$, where $faec = xc = y$).

In [4], Brady builds a locally CAT(0) $K(G, 1)$ for B_4 as follows. For each of the sixteen expressions $x = a_1a_2a_3$ ($a_i \in \{a, b, c, d, e, f\}$) just described, one constructs a Euclidean tetrahedron with edge labels as illustrated in Figure 5 for the case $x = bca$. Every face of the tetrahedron is a right-angle triangle with $\ell(a) = \ell(b) = \ell(c) = 1$, $\ell(bc) = \ell(ca) = \sqrt{2}$, and $\ell(x) = \sqrt{3}$. Six such tetrahedra can

be assembled to form a unit cube with the edge labeled x running a long diagonal. The space $\overline{Y_0}$ is defined to be the union of these sixteen labeled tetrahedra with all identically labeled pairs of faces identified and all vertices identified to a single vertex; it is equipped with the induced piecewise Euclidean metric. One can check that $\pi_1(\overline{Y_0}) = B_4$. The complex $\overline{Y_0}$ is locally CAT(0), and B_4 acts by covering isometries on its universal cover Y_0, which is a 3-dimensional CAT(0) space. Choosing some vertex as a basepoint (labeled 1), the 1-skeleton of Y_0 may be identified with the Cayley graph of B_4 with respect to an enlarged generating set.

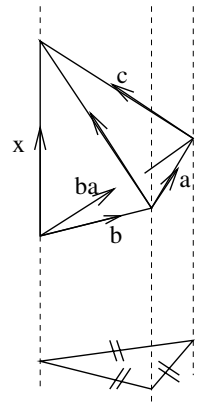

Figure 5 A labeled tetrahedra in Y_0 and its projection in X_0

The space Y_0 decomposes as a metric product $X_0 \times \mathbb{R}$. We now describe the factor X_0, which should be thought of as the result of projecting along the z-axes. Note that the z-axis through the base vertex 1 contains vertices corresponding to the central cosets $\langle z\rangle$, $x.\langle z\rangle$, $x^2.\langle z\rangle$, and $x^3.\langle z\rangle$. It follows that the vertices of X_0 correspond naturally to the left cosets of the cyclic subgroup $\langle x\rangle$ of order 4 in G. Each tetrahedron in Y_0 projects to an equilateral triangle in X_0 as shown in Figure 5, and these triangles decompose X_0 as a metric simplicial complex.

Let Γ denote the Cayley graph of G with respect to the generating set $\{a, b, c, d, e, f\}$. Then the projection of the 1-skeleton of Y_0 to that of X_0 factors through a 4-to-1 simplicial covering $\Gamma \to X_0^{(1)}$ (where the preimage of each vertex is a left coset of $\langle x\rangle$). An edge path in $X_0^{(1)}$ may be labeled by lifting the path to Γ. This, of course, depends on a choice of representative for the coset of $\langle x\rangle$ corresponding to the initial vertex of the path. There are two distinct types of triangles in X_0. Reading an edge path that traverses four times the perimeter of any triangle of the first type yields a cyclic permutation of the word W_1 (or its inverse). Reading four times around any triangle of the second type yields a cyclic permutation of the word W_2 (or its inverse). The action of G on X_0 is determined by its action by left multiplication on the left cosets of $\langle x\rangle$. In particular, we observe that the center of any triangle of the second type is the fixed set of some conjugate of y (of order 3 in G). Also, each vertex in X_0 is the fixed set of some conjugate of x and has stabilizer of order 4.

A fundamental region for the action of G on X_0 is shown in Figure 6 together with the link of an arbitrary vertex of X_0. The labels of vertices in the link correspond to edges entering and leaving the vertex. These vertices are necessarily labeled with respect to a choice of representative of the coset of $\langle x \rangle$. Changing the choice of representative simply relabels the link according to the action of an element of the vertex stabilizer. By the theory of complexes of groups (see [8]), the information given in Figure 6 (including the fact that y acts by a rotation of order 3 on the triangle of the second type) is enough to determine the action of G on X_0 up to isometry. This is because the given information leads to a description of G as the fundamental group of a nonpositively curved complex of groups—we leave the details to the reader.

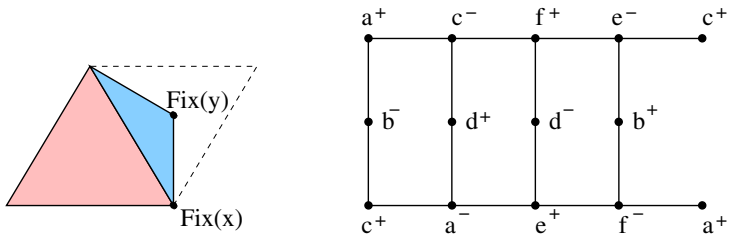

Figure 6 A fundamental region for the action of G on X_0 and the link of a vertex in X_0

We conclude this section with the following proposition, which shows that (as stated in the Introduction) the classification of 3-dimensional CAT(0) structures for the braid group B_4 is equivalent to the classification of minimal 2-dimensional CAT(0) structures for G.

PROPOSITION 5. *If the braid group B_4 acts geometrically and minimally on a* CAT(0) *space Y then $Y = \mathbb{R} \times X$, where X is* CAT(0), $\dim(X) = \dim(Y) - 1$, *and the action of B_4 induces a geometric action of G on X. On the other hand, any geometric action of G on a* CAT(0) *space X lifts to a proper action of B_4 on $\mathbb{R} \times X$, and this lifted action is unique up to a linear scaling of the $\mathbb{R}$-factor.*

Proof. The first statement follows immediately from standard results concerning the minset of the central element (cf. [8, Chap. II.6]), with the exception of the dimension calculation, which is a consequence of [14].

To prove the second statement it suffices to classify the nontrivial isometric actions of B_4 on $\mathbb{R}$. Take such an action. Since a, b, c are mutually conjugate (see Section 3), it follows that $\ell_{\mathbb{R}}(a) = \ell_{\mathbb{R}}(b) = \ell_{\mathbb{R}}(c) = \ell_0$, say, where $\ell_0 \neq 0$ by nontriviality of the action. Also $x^4 = y^3 = z$ implies that $4\ell_{\mathbb{R}}(x) = 3\ell_{\mathbb{R}}(y)$, while $y = xc$ implies $\ell_{\mathbb{R}}(y) = \ell_{\mathbb{R}}(x) \pm \ell_0$. Consequently, $\ell_{\mathbb{R}}(x = bac) = 3\ell_0$, in which case each of a, b, and c translates in the same direction on $\mathbb{R}$. The action of B_4 therefore depends, up to isometry, only on the choice of ℓ_0, giving the desired result. □

4. Proof of Theorem 1

For the purposes of this section we suppose that we are given a geometric (or, more generally, a proper semisimple) action of $G = B_4/Z(B_4)$ on some 2-dimensional CAT(0) space X. Moreover, we may suppose that the metric on X is scaled by a constant factor so that any one of the generators $a, \ldots, f$ has translation length 1. It follows, since they are all conjugate, that each of the generators has translation length 1. Our objective is to show that there exists a locally injective G-equivariant map $F\colon X_0 \to X$, which is locally isometric away from the 0-skeleton of X_0.

This is achieved by a combination of techniques developed in [3] and [12] together with observations about the structure of B_4. The strategy is to use the geometric and group-theoretic information contained in the previous sections in order to identify sufficient features of the action of G on X to be able to recognize the map F. More precisely, we shall define a particular region R in X that is the (possibly empty) intersection of minsets of a certain collection of elements of G. The proof then falls into a number of cases according to the structure of the set R. We obtain a contradiction in every case but one: where R is an equilateral triangle whose center is fixed by y and one of whose corners is a fixed point of x. In this case, we are able to reconstruct the G-equivariant image (under F) of the complex X_0 inside X.

4.1. Definition of the Region R

NOTATION. For any pair of edges u, v that are adjacent in Γ (or in Γ'), we write $z_{u,v}$ for the element $(uv)^3$. We note that $z_{u,v} = z_{v,u}$ and generates the center of the subgroup of G generated by u and v.

We define the (possibly empty) region R in X by

$$R = \mathrm{Min}(f) \cap \mathrm{Min}(z_{f,c}) \cap \mathrm{Min}(c) \cap \mathrm{Min}(z_{c,d}) \cap \mathrm{Min}(d) \cap \mathrm{Min}(z_{d,f}),$$

observing that R is invariant under the action of y.

Let $g, h \in G$ be commuting infinite-order elements; then, by the flat torus theorem [8] and the fact that X is 2-dimensional, there exists a unique $\langle g, h\rangle$-invariant isometrically embedded Euclidean plane in X, which we shall denote $\Pi(g, h)$. (Since the action is proper and semisimple, the quotient of $\Pi(g, h)$ by $\langle g, h\rangle$ is a flat torus.) We have, in fact, $\Pi(g, h) = \mathrm{Min}(g) \cap \mathrm{Min}(h)$. An alternative description of R is therefore as an intersection of flat planes, as follows. Define the "strips"

$$S(f, c) = \Pi(f, z_{f,c}) \cap \Pi(z_{f,c}, c),$$
$$S(c, d) = \Pi(c, z_{c,d}) \cap \Pi(z_{c,d}, d),$$
$$S(d, f) = \Pi(d, z_{d,f}) \cap \Pi(z_{d,f}, f).$$

Then

$$R = S(f, c) \cap S(c, d) \cap S(d, f).$$

Note that $y\colon S(f, c) \to S(c, d) \to S(d, f) \to S(f, c)$. By an application of Theorem 4 (to $B(c, f, b)$), we have that $S(f, c)$ and its y-translates are each

isometric to $\mathbb{R} \times I$, where I is a real interval of length $\sin\theta$ and where θ denotes the angle between the positive directions of any f-axis and any $z_{f,c}$-axis lying in the plane $\Pi(f, z_{f,c})$.

We shall now discuss some properties of the region R.

LEMMA 6. *The region R is a closed, convex, and bounded subset of X.*

Proof. Since it is the intersection of closed convex sets, R is clearly a closed and convex subset of X. We now show that R is bounded. If this is not the case then, as an unbounded subset of the strip $S(c, f)$, R must contain a *half-axis* of $z_{c,f}$: that is, an infinite geodesic ray ρ such that either $z_{f,c}(\rho) \subset \rho$ or $z_{f,c}^{-1}(\rho) \subset \rho$. Note that the only geodesic rays in $S(c, f)$ are half-axes of this type. Similarly, R contains a half-axis of $z_{f,d}$ and of $z_{d,c}$, and any geodesic ray is necessarily a half-axis of each of $z_{c,f}$, $z_{f,d}$, and $z_{d,c}$. In a proper action, however, two hyperbolic elements can share a half-axis only if they share a common power. A common power of $z_{c,f}$, $z_{f,d}$, and $z_{d,c}$ is central in G and hence trivial, while $z_{c,f}$, $z_{f,d}$, and $z_{d,c}$ each have infinite order in G. Therefore, R is bounded. □

NOTATION. We shall write $(u_1, u_2, u_3) = (c, f, d)$, $S_i = S(u_i, u_{i+1})$ for $i = 1, 2, 3$, and $z_i = z_{u_i, u_{i+1}}$ for $i = 1, 2, 3$. Indices are taken mod 3.

LEMMA 7. *Let $p \in R$ and suppose that p is in the interior of S_i for each $i = 1, 2, 3$. Then R is 2-dimensional and p is in the interior of R.*

Proof. Let $i = 1, 2, 3$. Since $p \in \text{int}(S_i)$, it follows that $\text{Lk}(p, S_i)$ is a circle of length 2π (as shown in Figure 7) and, since S_i is a convex subspace of X, $\text{Lk}(p, S_i)$ embeds as a subgraph of $\text{Lk}(p, X)$. (Note: in fact, this embedding is isometric because $\text{Lk}(p, S_i)$ has diameter π and $\text{Lk}(p, X)$ is CAT(1).) Therefore, $\text{Lk}(p, X)$ contains a subgraph (the union of the $\text{Lk}(p, S_i)$) that is a quotient of the graph

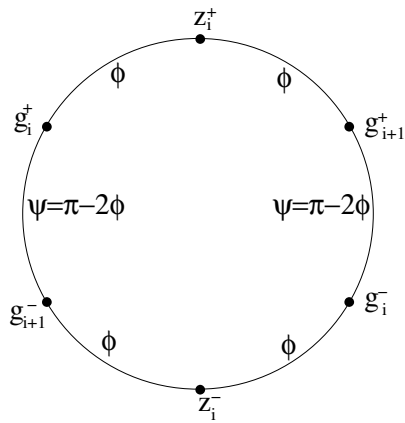

Figure 7 $\text{Lk}(p, S_i)$ is a circle: $0 < \psi, \phi < \pi$

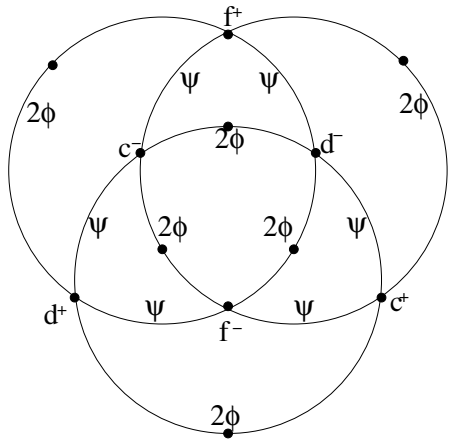

Figure 8 A graph that maps onto the union of the $\mathrm{Lk}(p, S_i)$ in $\mathrm{Lk}(p, X)$

pictured in Figure 8, under which all edges are mapped isometrically and, in fact, each of the three round circles is mapped isometrically. In both figures we have $0 < \psi, \phi < \pi$ and $\psi + 2\phi = \pi$.

Consider the triangle with vertices c^-, f^+, d^- and perimeter $\psi + \psi + 2\phi = 2\pi - 2\phi < 2\pi$. This maps to a geodesic triangle T (of the same dimensions) in $\mathrm{Lk}(p, X)$. Since $\mathrm{Lk}(p, X)$ is CAT(1) and T has perimeter strictly less than 2π, it follows that T is in fact a tripod with feet of length ϕ, ϕ, and $\psi - \phi \geq 0$. In particular, in the neighborhood of d^- there will be identification (in $\mathrm{Lk}(p, X)$) between the segments $[f^+, d^-]$ and $[c^-, d^-]$.

Considering in a similar fashion the triangle with vertices f^+, d^-, c^+, we observe also some local identification (near d^-) between $[f^+, d^-]$ and $[c^+, d^-]$ unless $\phi = \psi$. But such identifications would contradict the fact that $[c^-, d^-] \cup [d^-, c^+]$ maps to a geodesic of length π in $\mathrm{Lk}(p, X)$. Thus $\phi = \psi = \pi/3$ and the three circles $\mathrm{Lk}(p, S_i)$ are identified in $\mathrm{Lk}(p, X)$ with a single circle C. It follows that any point of C defines a direction in each of S_1, S_2, and S_3 and consequently a direction in R. Therefore, p is in the interior of the planar region R. □

The following is a consequence of Lemma 7 together with Lemma 6.

PROPOSITION 8. *Either*

(a) *R is empty*; *or*
(b) *R is nonempty but* $\dim(R) < 2$, *in which case R is either a single point or a closed bounded segment that is fixed pointwise by y (since y has order* 3); *or*
(c) *R is nonempty and* 2*-dimensional*.

In case (b), *R lies entirely in the boundary of each strip S_i for $i = 1, 2, 3$.*

In case (c), *R is either an equilateral triangle or a hexagon with a* 3*-fold rotational symmetry* (*and parallel opposite sides*), *and each side* (*or pair of parallel*

sides) of R lies in the boundary of one of the three strips S_1, S_2, or S_3. The element y acts on R by a nontrivial rotation of order 3 *in this case.*

4.2. A Useful Lemma

We now prove a geometrical lemma that will be useful in the subsequent development. We shall need to recall some standard terminology and refer the reader to [8, esp. Chap. II.9] for further details.

Recall that the visual boundary ∂X of a CAT(0) space (X, d) may be defined as the set of equivalence classes of infinite geodesic rays $[0, \infty) \to X$, where two rays c_1, c_2 are equivalent if there exist positive real numbers t_1, t_2 and a D such that $d(c_1(t_1+t), c_2(t_2+t)) \leq D$ for all $t > 0$. We note that, for each $p \in X$, there is a unique geodesic ray from p representing each point in the boundary. For $\xi \in \partial X$ and $p \in X$, we write $[p, \xi]$ for the infinite ray at p representing ξ. The angular metric $d_\angle$ on ∂X is defined as follows: $d_\angle(\xi, \zeta) = \sup_{p \in X}\{\angle_p([p, \xi], [p, \zeta])\}$. Observe that an isometric embedding of CAT(0) spaces always induces an isometric embedding of their boundaries with respect to this metric. (This is a consequence of [8, Prop. II.9.8(1)].)

LEMMA 9. *Suppose that g and h are commuting hyperbolic isometries of a* CAT(0) *space X and that p is a point lying on a g-axis γ. Denote by $\gamma^\pm$ the positive and negative parts, respectively, of γ (these are infinite rays based at p defining points $g^\pm$ respectively in ∂X). If ρ is an (or, rather, "the") infinite ray based at p and defining the point h^+ at infinity, then*

$$\angle_p(\gamma^+, \rho) = d_\angle(g^+, h^+), \qquad \angle_p(\gamma^-, \rho) = d_\angle(g^-, h^+),$$

and these two angles add to precisely π.

NOTATION. In general, we denote by g^+ (resp. g^-) the point in the visual boundary of a CAT(0) space X defined by the positive (resp. negative) half-axes of a hyperbolic element g acting on X. Whenever a g-axis passes through a point $p \in X$, we also use $g^\pm$ to denote the corresponding points in $\mathrm{Lk}(p, X)$ defined by the positive and negative half-axes based at p. Thus the positive half-axis $[p, g^+]$ for g leaves $\mathrm{Lk}(p, X)$ in the direction g^+ and is such that $g([p, g^+]) \subset [p, g^+]$.

Proof of Lemma 9. Since g and h commute we may measure the angles $d_\angle(g^\pm, h^+)$ in the flat plane $\Pi(g, h)$ to be exactly the angles formed between a g-axis and an h-axis lying in this plane. In particular, the two angles add to precisely π. On the other hand, the angles $\angle_p(\gamma^\pm, \rho)$ measured at p are (by definition of $d_\angle$) lower bounds respectively for the angular distances in ∂X, and they add to at least π because γ is a geodesic through p. The equalities now follow. □

4.3. The Case $R = \emptyset$

We consider the subgroup $B(c, \hat{b}, f) = B(c, f, b) < G$ acting on X (cf. Section 3). As a consequence of Theorem 4 there exists a $B(c, \hat{b}, f)$-equivariant isometric embedding

$$F_{c,f}\colon X_\theta(c,\hat{b},f)\to X.$$

The image of this map contains the planes $\Pi(c,z_{c,f})$ and $\Pi(z_{c,f},f)$ and is, in fact, the union of translates of these planes by elements of $B(c,\hat{b},f)$. It follows that $F_{c,f}$ is uniquely determined. We need only consider the following convex subset of $X_\theta(c,\hat{b},f)$:

$$\Sigma_1 = F_{c,f}^{-1}(\Pi(c,z_{c,f})\cup\Pi(z_{c,f},f)).$$

Conjugation by y yields the similar isometric embeddings

$$F_{f,d}\colon X_\theta(f,a,d)\to X,$$
$$F_{d,c}\colon X_\theta(d,e,c)\to X$$

as well as similarly defined subsets $\Sigma_2\subset X_\theta(f,a,d)$ and $\Sigma_3\subset X_\theta(d,e,c)$. Note that (by symmetry) the value θ, though yet to be determined, is the same for each of the three embeddings.

Let H_c denote the convex hull of the planes $\Pi(z_{d,c},c)$ and $\Pi(c,z_{c,f})$ in X, and write $F_c\colon H_c\to X$ for the inclusion map. Note that H_c embeds isometrically in $\mathrm{Min}(c)$ and has the form (tree) $\times\,\mathbb{R}$, where the $\mathbb{R}$-fibres correspond to c-axes in X and where "(tree)" is the convex hull of two geodesic lines in an $\mathbb{R}$-tree.

Again, conjugation by y yields the similar inclusions

$$F_f\colon H_f\to X,$$
$$F_d\colon H_d\to X.$$

The following construction follows closely that of [3, Sec. 3]. We define an identification space

$$Y=\coprod\{\Sigma_1,H_f,\Sigma_2,H_d,\Sigma_3,H_c\}/\approx,$$

where the equivalence relation is as follows. Let $u,v\in\{c,f,d\}$, and note that inclusion of the plane $\Pi(u,z_{u,v})$ into X factors through each of the maps F_u and $F_{u,v}$. We set $F_u^{-1}(p)\approx F_{u,v}^{-1}(p)$ for each point $p\in\Pi(u,z_{u,v})$, and this for each of the six choices of u,v. Identification over a flat plane occurs in this way between each pair of spaces that are adjacent in the cyclic order indicated in the preceding description of Y. We note also that Y is actually the union of its subspaces H_c, H_f, and H_d. Our description, however, is convenient for the proof of Lemma 10.

There is a naturally defined map $F\colon Y\to X$ induced by the family of isometric embeddings F_f, F_d, and F_c. We write $\overline{Y}$ for the image of this map.

LEMMA 10. *Either $\pi_1(Y)=\mathbb{Z}$, or R is a nonempty set.*

Proof. Let Y_0 denote the identification space defined exactly as for Y but neglecting the identification of points over the plane $\Pi(c,z_{c,f})$ (i.e., no identification between Σ_1 and H_c). We note that Y_0 is a CAT(0) space (hence contractible) since it is built up by a sequence of isometric gluings of pairs of CAT(0) spaces along convex subspaces (cf. [8]). Let $\Pi=F_{c,f}^{-1}(\Pi(c,z_{c,f}))$ and $\Pi'=F_c^{-1}(\Pi(c,z_{c,f}))$, both viewed as (convex) subspaces of Y_0. The space Y is obtained from Y_0 by identification of Π with Π'. Clearly, if Π and Π' are disjoint subsets of Y_0, then $\pi_1(Y)=\mathbb{Z}$.

On the other hand, suppose that Y_0 contains a point $p \in \Pi \cap \Pi'$. Such a point lies in both subspaces $\Sigma_1 \subset Y_0$ and $H_c \subset Y_0$. But, by the construction of Y_0, this is possible only if p also lies in each of the five remaining flat planes along which identifications were made. But then the point $\bar{p} = F(p)$ lies in each of the six planes $\Pi(u, z_{u,v})$ for $u, v \in \{c, f, d\}$ and therefore lies in R. □

The cases where R is a nonempty region will be treated in subsequent sections. We suppose for now that $R = \emptyset$ and, by Lemma 10, that $\pi_1(Y) = \mathbb{Z}$.

The generator of $\pi_1(Y)$ is represented by a locally geodesic circle γ in Y, which may be decomposed into a concatenation of three geodesic segments $\gamma_f \in H_f$, $\gamma_d \in H_d$, and $\gamma_c \in H_c$. Observe that, corresponding to the action of y on X, the space Y admits a three-fold symmetry: an isometry $\varphi\colon Y \to Y$ given by $\varphi = F^{-1} \circ y \circ F$. We may assume that the circle γ is φ-invariant and, moreover, that $\varphi\colon \gamma_c \to \gamma_f \to \gamma_d \to \gamma_c$. (If t denotes the generator of $\pi_1(Y)$, then γ is chosen from among the circular fibres of the locally convex subspace $\mathrm{Min}_{\tilde{Y}}(t)/\langle t\rangle$ of Y. This subspace has the structure $S^1 \times T$, where T is a tree, and the isometry φ induces an isometry on T that leaves at least one point fixed. We therefore choose γ to be the circle lying over such a fixed point.)

Since each of H_f, H_d, H_c embeds isometrically, the image of γ under the map F is a geodesic triangle $\Delta = \Delta(\bar{\gamma}_f, \bar{\gamma}_d, \bar{\gamma}_c)$ with vertices $\bar{p}_1, \bar{p}_2, \bar{p}_3$, where $\bar{p}_i = F(p_i)$ and $p_i \in \Sigma_i$ for $i = 1, 2, 3$. Clearly, Δ is also equilateral and of nonzero side length. In fact, it is invariant by the action of y, where $y\colon p_1 \mapsto p_2 \mapsto p_3 \mapsto p_1$.

By the arguments of [3, Lemmas 3.4–3.6], which apply precisely in this case, we deduce for $i = 1, 2, 3$ that the angle in the triangle at the point $\bar{p}_i$ is at least $\pi/3$. Because, in a CAT(0) space, angles in a triangle sum to at most π, we find that each of these three angles is exactly $\pi/3$.

Let us look a little more closely at the arguments from [3] to which we just appealed. It will be convenient for the latter part of our argument if we consider in particular the vertex p_3 of Δ. The situation is completely analogous for the other two vertices (particularly in view of the y-symmetry). The space of directions $\mathrm{Lk}(p_3, Y)$ is a metric graph. Define $L_3 = \mathrm{Lk}(p_3, H_c \cup \Sigma_3 \cup H_d)$, a subgraph of this graph. The space $\mathrm{Lk}(p_3, \Sigma_3)$ is either a circle or a theta-graph, as illustrated in Figure 9, depending on where p_3 lies in Σ_3. The graph L_3 is obtained from $\mathrm{Lk}(p_3, \Sigma_3)$ by adding 0, 1, or 2 π-arcs (closed intervals of length π attached by their endpoints) between c^+ and c^- as well as 0, 1, or 2 π-arcs between d^+ and

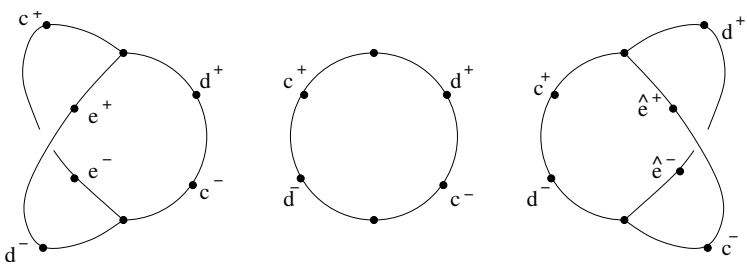

Figure 9 Possibilities for $\mathrm{Lk}(p_3, \Sigma_3)$

d^-. The number of π-arcs attached depends a priori on the structure of H_c and H_d (respectively) and on where p_3 lies in these sets. Let γ^+ and γ^- denote the points in L_3 where the circle γ leaves and enters the point p_3. Let ϕ_3 denote the angle at $\bar{p}_3$ in Δ. In [3, Lemma 3.6] it is shown that $\operatorname{diam}(L_3) \leq 5\pi/3$, and in [3, Lemma 3.5] it is shown that $d_{L_3}(\gamma^+, \gamma^-) + \phi_3 \geq 2\pi$, whence the lower bound on ϕ_3.

Since in our case $\phi_3 = \pi/3$ exactly, we must have $d_{L_3}(\gamma^+, \gamma^-) = 5\pi/3$. This is possible only if:

(i) $\operatorname{Lk}(p_3, \Sigma_3)$ is a theta-graph;
(ii) γ^+ and γ^- lie on π-arcs of different type; and
(iii) $\theta = \pi/3$ and $\gamma^\pm$ are each a distance $\pi/3$ from one end of its π-arc, as indicated in Figure 10.

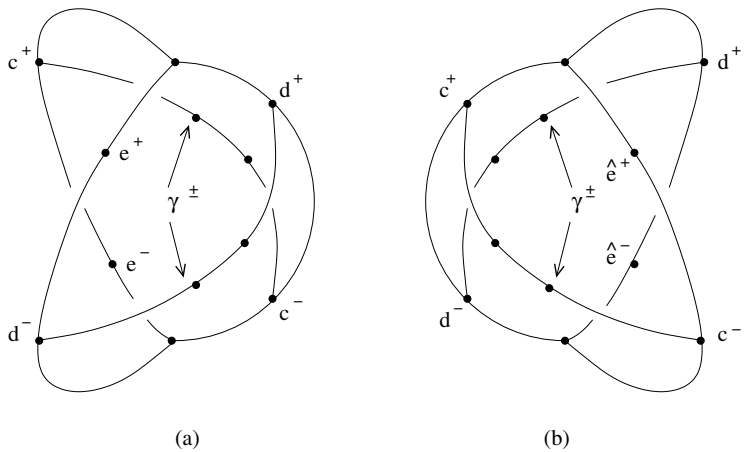

Figure 10 Possible forms of the graph L_3 (note: $\hat{e}^+ = d^2ed^{-2} = c^{-2}ec^2$)

We note that the argument of [3, Lemma 3.5] used the fact that F induces a local isometry $L_3 \to \operatorname{Lk}(\bar{p}_3, X)$, together with the link condition in a CAT(0) space. In fact, since $\operatorname{diam}(L) < 2\pi$, the map $L \to \operatorname{Lk}(\bar{p}_1, X)$ is necessarily an embedding—but not (of course) isometric, since in $\operatorname{Lk}(\bar{p}_3, X)$ the points γ^+ and γ^- are joined by an arc of length $\pi/3$.

Let ρ_0 denote an infinite geodesic ray from $\bar{p}_3$ to a point in ∂X as follows: in the case shown in Figure 10(a), ρ_0 follows the e-axis passing through $\bar{p}_3$ in the direction e^+; in the case shown in Figure 10(b), ρ_0 first crosses the strip $S(d, c)$, leaving $\bar{p}_3$ in the direction d^-, and then follows an e axis in the direction e^-.

Observe that, in each case, ρ_0 forms an angle π with the geodesic segment γ_c (which enters $L_3 \subset \operatorname{Lk}(\bar{p}_3)$ at the point γ^+ or γ^- that lies on a π-arc joining c^+ and c^-). Thus ρ_0 extends along γ_c to an infinite geodesic ray ρ from the point $\bar{p}_1$ to either e^+ or e^- in ∂X, depending on the case.

Consider now the sublink $L_1 \subset \operatorname{Lk}(\bar{p}_1, X)$, which appears exactly as in Figure 10 but with the labels $c^\pm$ and $d^\pm$ changed to $f^\pm$ and $c^\pm$. There is an f-axis

passing through $\bar{p}_1$ (defining the directions $f^{\pm}$ in L_1), and the ray ρ leaves $\bar{p}_1$ in whichever direction $\gamma^{\pm}$ that lies on a π-arc joining c^+ and c^-. It follows that ρ forms an angle π with one direction of the f-axis and an angle $4\pi/3$ with the other. However, since f and e are commuting hyperbolic elements, we obtain a contradiction in light of Lemma 9.

We proceed now to the cases where R is nonempty.

4.4. The Case $R \neq \emptyset$ and $\dim(R) \leq 1$

We consider in this section the case (b) of Proposition 8 in Section 4.1: R is either a single point or a closed geodesic segment, and each point of R is fixed by y and lies in the boundary of one of (and hence, by symmetry, each of) the S_i for $i = 1, 2, 3$.

Let p be a point in R. It follows that $\mathrm{Lk}(p, S_i) \subset \mathrm{Lk}(p, X)$ is one of the two graphs $\hat{L}_i$ and L_i pictured in Figure 11. Note that $\mathrm{Lk}(p, X)$ contains a subgraph K which is the union of the subgraphs $\mathrm{Lk}(p, S_i)$ for $i = 1, 2, 3$ and which is clearly invariant by the action of y. This graph K must be either one of the two graphs $\hat{L}$ and L composed from the links $\hat{L}_i$ and L_i, respectively (for $i = 1, 2, 3$), as shown in Figure 11 (where the vertices $c^{\pm}$ on the right are identified with those on the left), or at least a quotient of one of these two graphs in which the contributions $\hat{L}_i$ (resp. L_i) are mapped injectively. In fact, since each $\hat{L}_i$ (resp. L_i) has diameter π, each one is mapped isometrically into $\mathrm{Lk}(p, X)$.

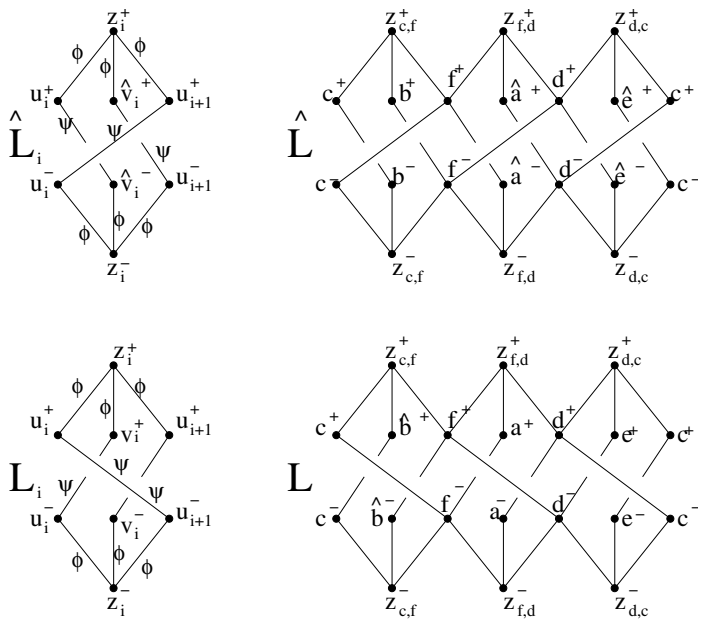

Figure 11 Building the graphs $\hat{L}$ and L from the contributions $\hat{L}_i$ and L_i (respectively) for $i = 1, 2, 3$

NOTATION. In the figures and in what follows, we write $(v_1, v_2, v_3) = (\hat{b}, a, e)$ and $(\hat{v}_1, \hat{v}_2, \hat{v}_3) = (b, \hat{a}, \hat{e})$, where $\hat{a} = d^{-2}ad^2$, $\hat{e} = c^{-2}ec^2$, and $\hat{b} = c^{-2}bc^2$. Indices are taken mod 3.

In the first situation ($K = \hat{L}$ or a quotient of this graph), no e-axis passes through p. However, it is easy to see that the segment $[p, c(p)] \subset S_3$ can be extended to a geodesic ray, say ρ, containing a positive half-axis of e. Denote by γ^+ and γ^- (respectively) the positive and negative half-axis of f passing through p. We have $\angle_p(\gamma^+, \rho) = 2\phi \neq 0$ and $\angle_p(\gamma^-, \rho) = \pi$ (this can be measured in the convex subgraph L_3). However, e and f are commuting hyperbolic elements and so we have a contradiction with Lemma 9, since these two angles do not add to exactly π.

In the second situation ($K = L$ or a quotient of this graph), there is an e-axis passing through p and we can measure $d_\angle(f^+, e^+)$ and $d_\angle(f^-, e^+)$ in the link of p. Suppose that $K = L$. We then have $d_K(f^+, e^+) = \min(4\phi, 2\psi) \leq \pi$ and $d_K(f^+, e^-) = \pi$. Note that the inclusion $K \subset \mathrm{Lk}(p, X)$ is a local isometry into a CAT(1) space, where local geodesics of length at most π are necessarily geodesic (see e.g. [8, Prop. II.1.4(2)]). Thus we have $0 < \angle_p(f^+, e^+) = d_K(f^+, e^+) \leq \pi$ and $\angle_p(f^-, e^+) = \pi$. Since the sum of these two angles is strictly greater than π, we again have a contradiction to Lemma 9. Thus K must be a nontrivial quotient of L.

We now consider how L may possibly collapse in the quotient map $L \to K \subset \mathrm{Lk}(p, X)$.

LEMMA 11. *Let q be a point in L_i and q' a point in L_{i-1}. Assume that q and q' are identified in the quotient. Then there are two geodesic arcs $\alpha \subset L_i$ and $\beta \subset L_{i-1}$, the first containing q and the second q', that are identified in the quotient and that both contain u_i^- or u_i^+.*

Proof. It is easy to see that at least one between u_i^- and u_i^+, say u_i^-, is at distance strictly less than π from q. Using the fact that the diameter of L_i is π, we have that u_i^- is at distance at most π from q'. For the quotient link to be CAT(1), it is then necessary for the geodesic arcs $[q, u_i^-]$ and $[q', u_i^-]$ to be identified. Now observe that if $d(q, u_i^-) \neq d(q', u_i^-)$ then a collapsing must occur inside L_i or L_{i-1}, which is contrary to the hypotheses. □

Lemma 11 means in particular that identifications must occur between pairs of segments emanating from one of the vertices $c^\pm$, $f^\pm$, or $d^\pm$ of the graph and belonging to two different subgraphs L_i and L_{i+1}. By symmetry and by invariance under the y-action, we can assume that the identifications take place at c^+, f^+, and d^+ and are of the same type at each vertex. There are four cases to be considered.

Case 1: Identify (part of) the edge $u_i^+ z_{i-1}^+$ with (part of) the edge $u_i^+ z_i^+$. For the link to remain CAT(1) after this identification, the arc $u_i^+ z_{i-1}^+ u_{i-1}^+ u_i^-$ is constrained to collapse onto the arc $u_i^+ z_i^+ v_i u_i^-$. It is now easy to see that the vertices f^+ and e^+ are identified, giving a contradiction. Indeed, by using Lemma 9 we

may measure $d_\angle(f^+, e^+) = 0$, but the elements f and e generate a rank-2 abelian group and so their axes in $\Pi(e, f)$ form nonzero angles.

Case 2: Identify (*part of*) *the edge* $u_i^+ v_{i-1}^-$ *with* (*part of*) *the edge* $u_i^+ u_{i+1}^-$. For the link to remain CAT(1) after this identification, the arc $u_i^+ v_{i-1}^- z_{i-1}^- u_i^-$ is constrained to collapse onto the arc $u_i^+ u_{i+1}^- z_i^- u_i^-$. It is now easy to see that the vertices f^- and e^- are identified, and the contradiction follows as in the previous case.

REMARK. We can thus now assume that, in $\mathrm{Lk}(p, X)$, the arc $u_i^+ z_{i-1}^+ u_{i-1}^+$ intersects the edge $u_i^+ z_i^+ u_{i+1}^+$ only in u_i^+. We therefore have a locally geodesic loop $u_1^+ z_1^+ u_2^+ z_2^+ u_3^+ z_3^+ u_1^+$ whose length must be at least 2π. An easy computation then yields the inequality $0 < \psi \le \pi/3 \le \phi < \pi/2$.

Case 3: Identify (*part of*) *the edge* $u_i^+ v_{i-1}^-$ *with* (*part of*) *the edge* $u_i^+ z_i^+$. For the link to remain CAT(1) after this identification, the arc $u_i^+ v_{i-1}^- z_{i-1}^- u_i^-$ is constrained to collapse onto the arc $u_i^+ z_i^+ v_i^+ u_i^-$ (the identifications are shown in Figure 12). Observe that the point $z_{c,f}^+$ is identified (in K) with a point on the geodesic segment $[z_{d,c}^-, c^+]$, a distance ψ from $z_{d,c}^-$. Similarly, $z_{d,c}^-$ is identified with a point on the geodesic segment $[z_{c,f}^+, c^-]$. We therefore have a geodesic segment $I = [z_{c,f}^+, z_{d,c}^-]$ of length ψ in $\mathrm{Lk}(p, X)$ that is common to both $\mathrm{Lk}(p, S_1) = L_1$ and $\mathrm{Lk}(p, S_3) = L_3$. We now construct a path $[d^-, z_{d,c}^-] \cup I \cup [z_{f,c}^+, f^+]$ in K that is a local geodesic (since it is composed of geodesic segments in L_1 and L_3, which overlap in I) and that has length exactly π. On the other hand, f^+ and d^- are joined by a geodesic of length $\psi < \pi$ in L_2. Each of these two local geodesics yields a true geodesic in $\mathrm{Lk}(p, X)$ (since the inclusion $K \subset \mathrm{Lk}(p, X)$ is a local isometry and local geodesics of length at most π in $\mathrm{Lk}(p, X)$ are necessarily geodesic). But this is a contradiction, because $\psi \neq \pi$.

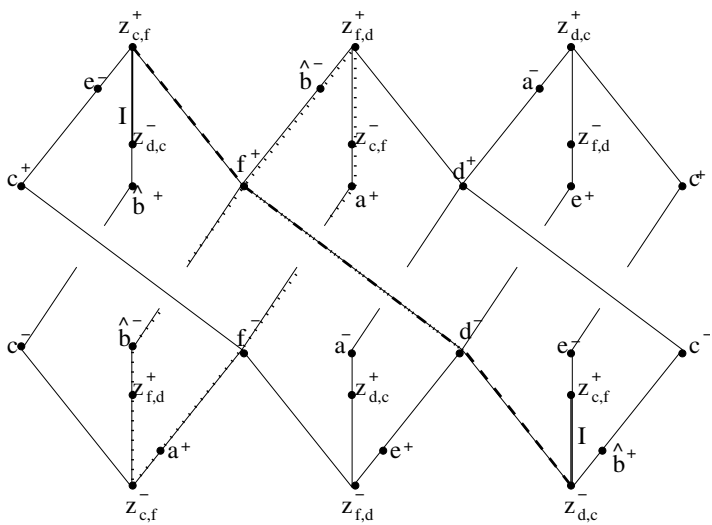

Figure 12 Collapsing of L in Case 3

Case 4: Identify (*part of*) *the edge* $u_i^+ z_{i-1}^+$ *with* (*part of*) *the edge* $u_i^+ u_{i+1}^-$. For the link to remain CAT(1) after this identification, the arc $u_i^+ z_{i-1}^+ u_{i-1}^+ u_i^-$ is constrained to collapse onto the arc $u_i^+ u_{i+1}^- z_i^- u_i^-$ (the identifications are shown in Figure 13). This time we observe that c^- is identified with a point in L_2 a distance $2\phi - \psi < \pi$ from f^+ (as measured in L_2). On the other hand, there is a geodesic of length π in L_1 joining c^- and f^+. Thus we have a segment in L_1 and a segment in L_2 that both yield geodesics in $\mathrm{Lk}(p, X)$ yet have different lengths—a contradiction.

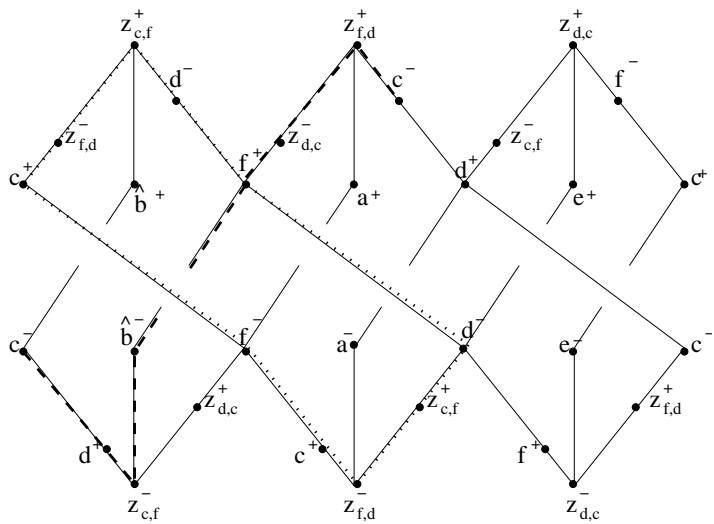

Figure 13 Collapsing of L in Case 4

4.5. The Case $R \neq \emptyset$ *and* $\dim(R) = 2$

Recall from Section 4.1 that if the region R is 2-dimensional then R is either an equilateral triangle or a hexagon with a 3-fold rotational symmetry and parallel opposite sides. The various possibilities are shown in Figure 14. These 2-dimensional situations can degenerate to the case when R is reduced to a single point (see cases (i) and (ix) in the figure); however, it was shown in Section 4.4 that these degenerate cases cannot occur. In fact, we shall now prove that six of the seven 2-dimensional cases cannot occur either.

LEMMA 12. *Assume that R is* 2*-dimensional. Then the following equality holds*: $\ell(ef) = \ell(e) = \ell(f) = 1$.

Proof. Take p in the interior of R; in particular, p is in the interior of S_i for $i = 1, 2, 3$. Thus there is a segment of the positive half-axis of f starting at p that coincides with a segment of the negative half-axis of $z_{d,c}$ starting at p. Let $[p, s]$ be a segment contained in the positive half-axis of c such that s lies on the boundary of

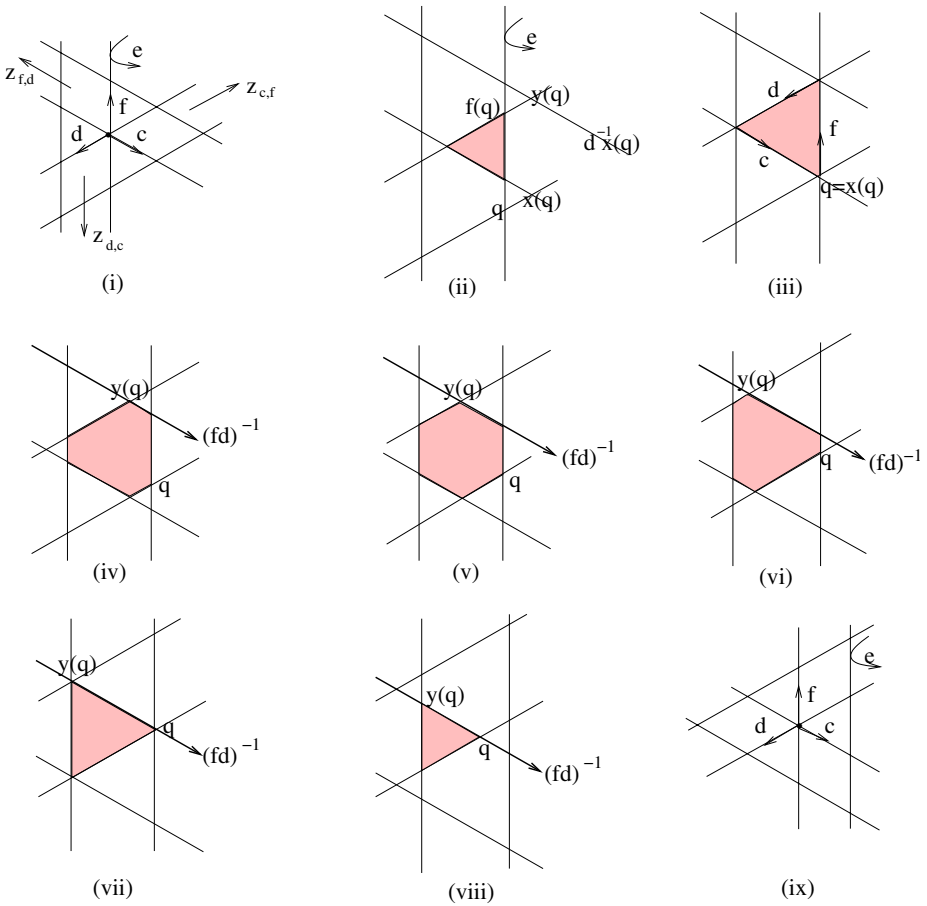

Figure 14 The various cases of R with $\dim(R) = 2$ (and degenerations of R to a point)

$S(d, c)$. Then (a) there is an e-axis passing through s and (b) $[p, s]$ extends along this axis to an infinite geodesic ray to the point e^+ in the visual boundary ∂X of X. This ray forms an angle of $2\pi/3$ with the positive direction of the f-axis at p and of $\pi/3$ with its negative direction. It follows from Lemma 9 that these are exactly the angles that one should observe between e- and f-axes in $\Pi(e, f)$. Namely, if $p \in \Pi(e, f)$ then the points $f^{-1}(p)$, p, and $e(p)$ are the vertices of an equilateral triangle in $\Pi(e, f)$, whence $\ell(ef) = \ell(e) = \ell(f) = 1$. □

PROPOSITION 13. *The cases* (ii) *and* (iv)–(viii) *of Figure 14 cannot occur.*

Proof. For each of the cases (iv)–(viii), we shall exhibit a point q (see the figure) whose distance from $d^{-1}x(p)$ is strictly less than 1. This contradicts the fact that $d^{-1}x = d^{-1}def = ef$ has translation length 1 (by Lemma 12). Note also that $d^{-1}x = (fd)^{-1}y$ and that the upper edge of $S(f, d)$ is an axis for fd. In cases

(iv)–(viii), q is a vertex of R and it is straightforward to see that $d((fd)^{-1}y(p), p)$ is strictly less than 1.

For case (ii), choose the point q to be the point of intersection between the lower edge of $S(c, f)$ and the f-axis containing one side of the triangle. The point $y(q)$ belongs to the upper edge of the strip $S(f, d) = y(S(c, f))$, and $x(q) = f^{-1}(y(q))$ belongs to the other side of the same strip. The point $d^{-1}x(q)$ also belongs to the strip $S(f, d)$; this point must coincide with $ef(q)$. Consider the point $f(q)$, which is the upper right vertex of the triangle. Notice that there is an e-axis passing through $f(q)$ so that $d(f(q), ef(q)) = 1$. On the other hand, it is easy to see that $d(f(q), d^{-1}x(q)) > 1$ and so $ef(q) \neq d^{-1}x(q)$; thus we reach a contradiction. □

We shall now show that in case (iii) we can construct a locally injective G-equivariant map $F\colon X_0 \to X$ that is locally isometric away from the 0-skeleton of X_0.

Take a fundamental region in X_0 for the action of G. As shown in Figure 6(i), such a region consists of two triangles. We may choose the vertices of the these triangles to be $\mathrm{Fix}(x)$, $\mathrm{Fix}(y)$, $c^{-1}(\mathrm{Fix}(x)) = \mathrm{Fix}(c^{-1}xc)$ and $\mathrm{Fix}(x)$, $\mathrm{Fix}(c^{-1}xc)$, $d(\mathrm{Fix}(x)) = \mathrm{Fix}(dxd^{-1})$. Now map these triangles to the corresponding triangles of X contained in $S(c, d)$ (see Figure 14(iii)). Observe that $S(c, d)$ as well as $\Pi(c, z_{c,d})$, $\Pi(z_{c,d}, d)$ are convex subsets of X. Thus, to see that the map just defined on the fundamental region of X_0 extends G-equivariantly to a local injection of X_0 into X, it suffices to show that $\mathrm{Lk}(\mathrm{Fix}(x), X_0)$ embeds into $\mathrm{Lk}(\mathrm{Fix}(x), X) = \mathrm{Lk}(q, X)$. Note that this embedding need not be π-distance preserving (i.e., we cannot assume that the map F is locally isometric at the vertex $\mathrm{Fix}(x)$).

As in Section 4.4, we see that $\mathrm{Lk}(q, X)$ must contain isometric copies of the three graphs $\mathrm{Lk}(q, S(f, c))$, $\mathrm{Lk}(q, S(c, d))$, and $\mathrm{Lk}(q, S(d, f))$ pictured in Figure 15.

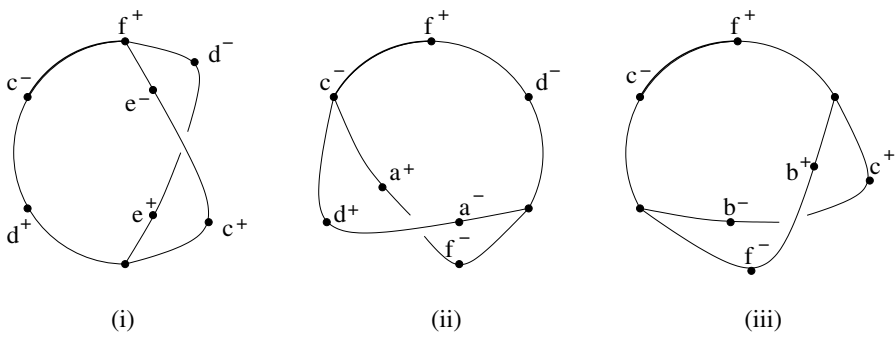

Figure 15 The pieces that form $\mathrm{Lk}(\mathrm{Fix}(x), X)$

Note now that the geodesic arcs of length $2\pi/3$ in $\mathrm{Lk}(q, S(f, c))$ and in $\mathrm{Lk}(q, S(d, f))$ from f^+ to c^+ must coincide, by the uniqueness of geodesics of length $< \pi$ in a CAT(1) space. For the same reason, the geodesic arcs of length $2\pi/3$ in $\mathrm{Lk}(q, S(c, d))$ and in $\mathrm{Lk}(q, S(d, f))$ from c^- to f^- must coincide as well.

Observe that c^- is a distance $\pi/3$ from each of d^+, f^+, and a^+. Since $\mathrm{Lk}(q, X)$ is left invariant by the action of x, we deduce that f^- is a distance $\pi/3$ from each of b^+, a^+, and e^+. In particular, there is a geodesic arc in $\mathrm{Lk}(q, X)$ between f^- and e^+ of length $\pi/3$. Such an arc—together with the geodesic segments from e^+ to d^- and from d^- to f^- contained in $\mathrm{Lk}(q, S(f, c))$ and $\mathrm{Lk}(q, S(c, d))$, respectively—gives a circuit in $\mathrm{Lk}(q, X)$ of total length $4\pi/3 < 2\pi$. Since $\mathrm{Lk}(q, X)$ is CAT(1), we deduce that e^+ is the midpoint of the geodesic segment from d^- to f^- contained in $\mathrm{Lk}(q, S(c, d))$. By a similar argument, we see furthermore that a^- is the midpoint of the geodesic segment between d^+ and c^+ lying in $\mathrm{Lk}(q, S(f, c))$, and we conclude that $\mathrm{Lk}(q, S(f, c))$ and $\mathrm{Lk}(q, S(c, d))$ have a circle $(c^- f^+ d^- e^+ a^- d^+)$ in common.

By realizing all these identifications among $\mathrm{Lk}(q, S(f, c))$, $\mathrm{Lk}(q, S(c, d))$, and $\mathrm{Lk}(q, S(d, f))$, we end up precisely with the graph of Figure 6(ii), that is, a graph isomorphic to $\mathrm{Lk}(\mathrm{Fix}(x), X_0)$. Observe also that no further identifications among $\mathrm{Lk}(q, S(f, c))$, $\mathrm{Lk}(q, S(c, d))$, and $\mathrm{Lk}(q, S(d, f))$ can occur in $\mathrm{Lk}(q, X)$, since these three subgraphs embed isometrically in $\mathrm{Lk}(q, X)$. It follows that the map F induces an injective map $\mathrm{Lk}(\mathrm{Fix}(x), X_0) \to \mathrm{Lk}(q, X)$.

Finally, we note that the map F is clearly the unique map $X_0 \to X$ satisfying the statement of Theorem 1.

5. A Method for Deforming CAT(0) Structures

In this section we show how to construct new CAT(0) structures for the group G that are topologically different from the one described by Brady. The construction is a bit involved, so we start by illustrating the main ideas on a simpler motivating example.

5.1. CAT(0) *Structures for Surface Groups*

Let $\Pi = \pi_1(S)$ be the fundamental group of the polygonal complex S obtained by identifying (in pairs) the edges of a Euclidean regular hexagon, as shown in Figure 16(i). Note that S is the nonorientable closed surface of Euler characteristic

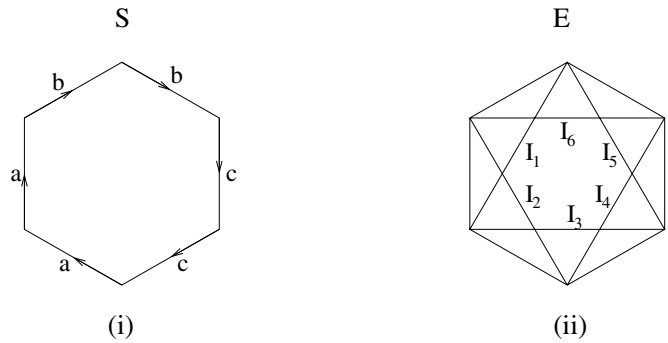

Figure 16 Polygonal description of the surface S, and intervals I_i

-1. Let $P_0 = \tilde{S}$ be the universal cover of S endowed with the piecewise Euclidean metric induced by S; then Π acts freely and properly discontinuously by isometries on P_0. Note that P_0 is locally flat outside the 0-skeleton while the link of every vertex is a circle of length 4π. We shall now construct a new CAT(0) structure P for Π with the properties that:

(1) P is a piecewise Euclidean complex;
(2) there exists a locally injective Π-equivariant map $f\colon P_0 \to P$ that is locally isometric outside the 0-skeleton of P_0; and
(3) f is not injective and identifies orbits of the Π-action that are arbitrarily close.

More precisely: for any $N, \varepsilon > 0$, there exist g in Π with $\ell_{P_0}(g) > N$ and $x, y \in P_0$ such that $d(x, y) < \varepsilon$ and $f(x) = f(g(y))$; consequently, $\ell_P(g) < \varepsilon$.

Identify the hexagon E of Figure 16(ii) with a choice of lift in P_0, and let $I_1, \ldots, I_6$ be the geodesic segments inside E as shown in the figure. Consider the infinite geodesic line in P_0 that contains I_1 and that is constrained to enter and leave any vertex of P_0 through diametrically opposed points in the link (i.e., points of path distance 2π from one another in a circle of length 4π). More concretely, the line λ is a concatenation of translates of the intervals I_i:

$$\ldots, (ab)^{-1}(I_6),\ I_1,\ a^2(I_2),\ a^2ca(I_3),\ a^2cac^2(I_4),$$
$$a^2cac^2bc(I_5),\ a^2cac^2bcb^2(I_6),\ a^2cac^2bcb^2ab(I_1),\ \ldots;$$

in particular, λ is an axis for the element $\gamma = a^2cac^2bcb^2ab$.

Let p denote the midpoint of $I_1 = [u, v]$, and choose x and y in E arbitrarily close to p such that the segment $[x, y]$ is bisected orthogonally at p by I_1. Let α, β denote the segments $\alpha = [x, v]$ and $\beta = [u, y]$, and let δ denote the angle at v between α and I_1 (respectively, the angle at u between β and I_1); see Figure 17(i). The geodesic in P_0 from x to $\gamma(y)$ consists of α and $\gamma(\beta)$ together with the segments on λ from v to $\gamma(u)$ that are the translates of $I_2, \ldots, I_6$ mentioned previously.

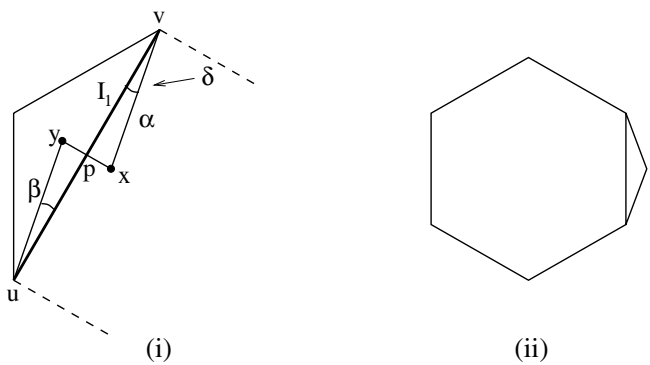

Figure 17 Construction of P from P_0: (i) segments α and β; (ii) polygon D

Let D denote the polygonal disk shown in Figure 17(ii), which consists of a regular hexagon of side length equal to the length of I_1 and with an isosceles triangle Δ of base angle δ attached along its base to one side of the hexagon.

Let P denote the polyhedral complex obtained from P_0 by:

(i) identifying the orbits of x and y equivariantly, so that $x \sim \gamma(y)$; and
(ii) attaching equivariantly an orbit $\Pi \cdot D$ of isometric copies of D, so that the boundary of $D = 1 \cdot D$ is identified by a piecewise isometry to the image of the geodesic path in P_0 from x to $\gamma(y)$.

PROPOSITION 14. *P is a 2-dimensional* CAT(0) *complex admitting a natural proper isometric Π-action such that the natural map $f: P_0 \to P$ satisfies conditions* (2) *and* (3).

Proof. The complex P has a polygonal description with eight orbits of vertices:

(i) the orbits of the six points of intersection in the interior of E between the segments $\alpha, \beta, I_2, \ldots, I_6$;
(ii) the orbit of the point $x = \gamma(y)$;
(iii) the orbit of the vertex v.

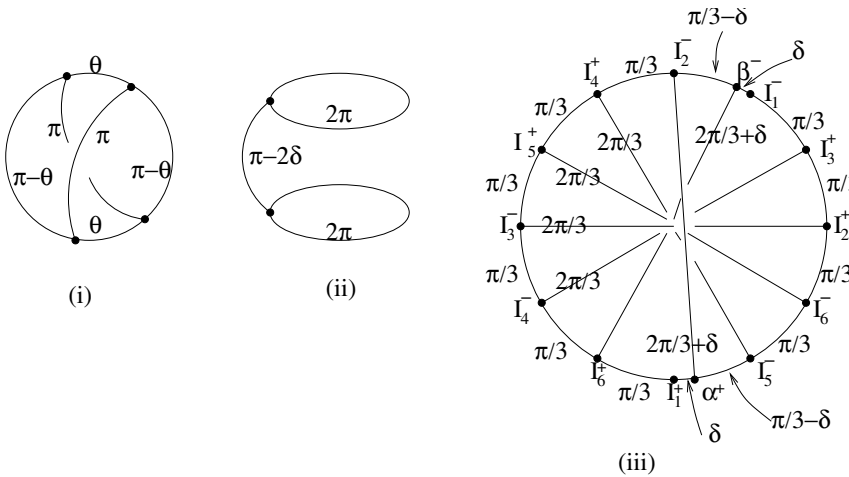

Figure 18 The three types of vertex links in P

The links of these three types of vertices are shown in Figure 18 and can be easily checked to be CAT(1) (i.e., to have no simple circuits of length $< 2\pi$). Thus P is CAT(0) by the link condition. The remaining claims can be easily checked. □

5.2. Proof of Theorem 2

We now present a construction analogous to that of the previous section, starting from the standard action of $G = B_4/Z$ on the Brady complex X_0 and resulting in the following statement.

THEOREM 15. *There exists a* 2*-dimensional* CAT(0) *complex* X*, together with an action of* G *on* X *and a* G*-equivariant map* $f\colon X_0 \to X$*, satisfying conditions* (2) *and* (3) *of Section 5.1.*

The key features of the construction of the previous section were:

(i) the choice of the axis λ (associated to the group element γ);
(ii) the modification of the space P_0 by addition of an orbit of polygons $\Pi \cdot D$; and, as a consequence,
(iii) the modification of the link of the typical vertex in P_0.

Recall from Figure 6 the structure of the link L_0 of a typical vertex v in X_0. We may choose v such that it lies on an axis for each element a, b, c, d, e, f, so that vertices of L_0 are labeled according to these axes. In particular, the circuit $(a^+, c^-, d^+, a^-, c^+, b^-)$ corresponds to $\mathrm{Lk}(v, \Pi(a, c))$ and circuit $(f^+, e^-, b^+, f^-, e^+, d^-)$ to $\mathrm{Lk}(v, \Pi(e, f))$.

The following lemma provides a source of geodesic segments in the flat planes $\Pi(a, c)$ and $\Pi(e, f)$ that will serve as analogues of the intervals $I_1, \ldots, I_6$ of the previous construction. Recall that $\Pi(a, c)$ is invariant by the action of the rank-2 free abelian group generated by a and c.

LEMMA 16. (i) *Given* $\varepsilon > 0$ *and* $\theta \in \mathrm{Lk}(v, \Pi(a, c))$*, there exists a* $g_{\varepsilon,\theta} \in \langle a, c\rangle$ *such that the segment* $I_{\varepsilon,\theta} = [v, g_{\varepsilon,\theta}(v)]$ *in the plane* $\Pi(a, c)$ *defines a point in* L_0 *of distance at most* ε *from* θ*.*

(ii) *By taking sufficient powers of* g_{ε,θ_i} *we may simultaneously approximate in this way any finite number of directions* $\{\theta_1, \ldots, \theta_k\}$ *by segments* $I_i = I_{\varepsilon,\theta_i}$ *whose lengths are arbitrarily close to some* (*large*) *fixed number* C*, in the sense that* $(\ell(I_i) - C)/C < \delta$ *for a given* $\delta > 0$*.*

Note that the stabilizer in G of the vertex v is the cyclic group of order 4 generated by x and that it acts on $L_0 = \mathrm{Lk}(v, X_0)$ via

$$x\colon a^+ \mapsto e^+ \mapsto c^+ \mapsto f^+ \mapsto a^+, \qquad b^+ \leftrightarrow d^+.$$

The G-equivariant modification we shall perform on X_0 will result in adding to L_0 a collection of new segments (or "chords") that will necessarily be $\langle x\rangle$-invariant. We wish to approximate the following system of chords.

Let n be a positive even integer and let $\phi = \pi/6n$. For $k = 1, \ldots, n$, let θ_k denote the point on the segment $[b^-, a^+]$ a distance $k\phi = k\pi/6n$ from b^- and let θ_{-k} be the point the same distance from b^- on the segment $[b^-, c^+]$. Observe that $x(\theta_k) \in [d^-, e^+]$ and $x(\theta_{-k}) \in [d^-, f^+]$ for each k. Also, $x^2(\theta_k) = \theta_{-k}$.

Let Γ denote the abstract bipartite graph with vertex set $\Theta = \{\theta_k, x(\theta_k) : k = \pm 1, \ldots, \pm n\}$ and edge set $\mathcal{C} = \{\{\theta_{\pm k}, x(\theta_{\pm l})\} : k + l \text{ odd}\}$. Then $\Gamma \cong K_{n,n} \sqcup K_{n,n}$, where $K_{n,n}$ denotes the complete bipartite graph on two sets of n vertices.

Note that Θ is an $\langle x\rangle$-invariant set and that the action of x induces a free graph automorphism of Γ. Let $\bar{\Gamma} = \Gamma/\langle x\rangle$. Then $\bar{\Gamma}$ is isomorphic to the graph obtained from $K_{n/2,n/2}$ by doubling each edge. In particular, $\bar{\Gamma}$ has $n^2/2$ edges and has valence n.

Since n is chosen to be even, it follows that we may find an Eulerian circuit in $\bar{\Gamma}$ (i.e., a circuit that visits each edge exactly once). This circuit determines an orientation and a cyclic ordering of the edges of $\bar{\Gamma}$ that we shall henceforth denote $E_1, E_2, \ldots, E_{n^2/2}$, where the terminal vertex of E_i is the initial vertex of E_{i+1}.

We construct in a similar fashion a graph $\Gamma'(\Theta', \mathcal{C}')$ where the vertices Θ' lie on the segments $[c^-, a^-]$ and $[f^-, e^-] = x([c^-, a^-])$. The antipodal maps on $\mathrm{Lk}(v, \Pi(a, c))$ and $\mathrm{Lk}(v, \Pi(e, f))$ determine a bijection $\Theta :\to \Theta'$ that induces an $\langle x \rangle$-equivariant graph isomorphism $\Gamma \to \Gamma'$. Let $\overline{\Gamma'} = \Gamma'/\langle x \rangle$ and write $\varphi\colon \bar{\Gamma} \to \overline{\Gamma'}$ for the induced isomorphism.

REMARK. If we modify L_0 by adding a segment joining points p, q for each edge $\{p, q\} \in \Gamma \sqcup \Gamma'$, then the result will be a CAT(1) metric graph provided the length of each added segment is at least $\pi - \phi$ (where $\phi = \pi/6n$). Moreover, if $0 < \varepsilon < \phi/2$, then we may allow ourselves to perturb the attaching points of the added segments within an ε-neighborhood provided the length of each segment is at least $\pi - \phi + 2\varepsilon$.

Now observe that the quotient map $X_0 \to X_0/G$ determines a map $\mathrm{Lk}(u, X_0) \to L_0/\langle x \rangle$, written $p \mapsto \bar{p}$, for any vertex $u \in X_0$. We may choose a family of segments I_θ in $\Pi(a, c)$ for each vertex θ of $\bar{\Gamma}$ such that I_θ approximates the angle θ as in Lemma 16. More precisely, if $I_\theta = [v, g(v)]$ for $g \in \langle a, c \rangle$ then it determines directions $p \in \mathrm{Lk}(v, X_0)$ and $q \in \mathrm{Lk}(g(v), X_0)$ such that $\bar{p}$ approximates $\theta \in \bar{\Gamma}$ and $\bar{q}$ approximates $\varphi(\theta) \in \overline{\Gamma'}$ in $L_0/\langle x \rangle$. In this way, the family of segments $\{I_\theta\}$ approximates all vertices of $\bar{\Gamma} \sqcup \overline{\Gamma'}$.

We now build an axis λ in X_0 (for some element $\gamma \in G$) by concatenating a sequence of interval segments

$$\ldots,\ I_0,\ g_1(I_1),\ g_2(I_2),\ g_3(I_3),\ \ldots,\ g_{n^2}(I_{n^2}) = \gamma(I_0),\ \gamma g_1(I_1),\ \ldots,$$

where the I_i are chosen from among the segments I_θ, $\theta \in \mathrm{Vert}(\bar{\Gamma})$. This axis λ is uniquely determined by the sequence of "transitions" from each interval to the next—namely, by the sequence of pairs $(\overline{p_i}, \overline{q_i})$, where p_i and q_i are the directions determined by segments $g_{i-1}(I_{i-1})$ and $g_i(I_i)$ (respectively) at the vertex $g_i(v)$. We choose the I_i from among the I_θ and orient them appropriately so that the corresponding sequence of transitions approximates the following sequence of oriented edges in $\bar{\Gamma} \sqcup \overline{\Gamma'}$:

$$E_1,\ \varphi(E_2),\ E_3,\ \varphi(E_4),\ \ldots,\ \varphi(E_{n^2/2}),$$
$$E_{n^2/2}^{-1},\ \varphi(E_{n^2/2-1}^{-1}),\ \ldots,\ E_4^{-1},\ \varphi(E_3^{-1}),\ E_2^{-1},\ \varphi(E_1^{-1}).$$

Note that each edge of $\bar{\Gamma} \sqcup \overline{\Gamma'}$ is visited exactly once by this sequence.

We now obtain an action of the group G on a new space X from the original action on X_0 by identifying a pair of orbits and attaching an orbit $G \cdot D$ of Euclidean polygonal disks, following exactly the procedure laid out in the previous section. The disk D is built from an n^2-sided figure, whose ith side has length equal to that of the segment I_i, by attaching a thin isosceles triangle Δ to one side.

By choosing the intervals as in Lemma 16(ii) and choosing the base angle of Δ sufficiently small, we may suppose that this figure approximates, as closely as we like, a regular n^2-gon and so has angles approximately $\pi - 2\pi/n^2$.

The foregoing process modifies the link of a vertex in X_0 by attaching to L_0 a family of segments corresponding (approximately) to the edges of $\Gamma \sqcup \Gamma'$ and each of length approximately $\pi - 2\pi/n^2$. As noted in the previous Remark, this results in a link that is still CAT(1) so long as $2\pi/n^2 \leq \phi + 2\varepsilon$, where $\phi = \pi/6n$. It suffices to choose $n = 14$ (and ε sufficiently small) to ensure that the resulting space X is CAT(0). (Note that, in addition to the modification of L_0 just described, there may also be the addition of segments of length exactly π that arise whenever one of the segments I_i contains a vertex of X_0 in its interior; but these added segments always join points that are already a distance exactly π in L_0 and so cannot introduce any short circuits into the link. Observe also that the axis λ and the polygonal disk D have been carefully chosen so that no double edges are introduced into the vertex link.) We leave the reader to verify that the remaining vertex links in the new space X are CAT(1) and hence that X is CAT(0).

It is clear that G acts geometrically on X and that the obvious map $X_0 \to X$ satisfies the statement of Theorem 15 (equivalently, Theorem 2).

6. Co-Hopficity of the Group $G = B_4/Z(B_4)$

In this section we use Theorem 1 to give a proof of the following statement.

THEOREM 17. *The group $G = B_4/Z(B_4)$ is co-Hopfian (i.e., every monomorphism $G \to G$ is an isomorphism).*

Proof. Given a monomorphism $\varphi\colon G \hookrightarrow G$, one obtains a new action on the space X_0 (by composing the standard action with φ). Though we may not suppose a priori that this new action is co-compact, we do know that it is semisimple because it arises as the action of a subgroup of a group acting co-compactly. However, the refinement of Theorem 1 to the semisimple case (see Remark (2) following the theorem's statement in Section 1) now gives a map $F\colon X_0 \to X_0$, which is a "homothetic" embedding induced by the monomorphism φ (in the sense that $\varphi(g) \circ F = F \circ g$ for all $g \in G$, where G is viewed as a subgroup of $\mathrm{Isom}(X_0)$ via the standard action). By "homothetic embedding" we mean that $F = F' \circ \lambda$, where $\lambda\colon X_0 \to \lambda X_0$ is a constant scaling of the metric on X_0 and $F'\colon \lambda X_0 \to X_0$ is locally isometric except possibly on the 0-skeleton. However, since the links of vertices in X_0 and λX_0 are identical, F' must be locally isometric on the 0-skeleton as well and hence must be an isometric embedding.

The theorem will now follow from our final proposition.

PROPOSITION 18. *The homothetic embedding $F\colon X_0 \hookrightarrow X_0$, induced by the monomorphism φ, has scaling factor $\lambda = 1$ and is therefore an isometry.*

Proof. The complex X_0 is built from Euclidean equilateral triangles of type I and II, as shown in Figure 19. The edges of X_0 are oriented and labeled in two

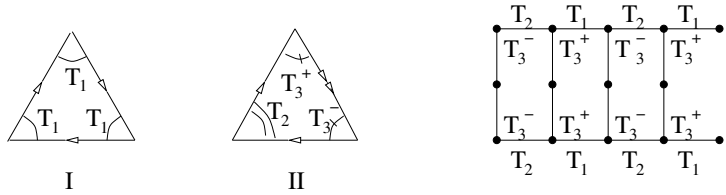

Figure 19 Description of the complex X_0

different ways: edges labeled with a single arrow have valence 3 and those with a double arrow have valence 2. There are four distinct types of corners, labeled T_1, T_2, T_3^+, T_3^-, which are arranged in the complex X_0 in such a way that the link of every vertex is as shown in the figure.

Suppose that $F\colon X_0 \hookrightarrow X_0$ is a homothetic embedding of scaling factor λ. Note that λ is constrained to be a positive integer and that each triangle in X_0 maps onto a union of λ^2 triangles. Let Δ_{II} denote the image under F of a triangle of type II.

CLAIM. *The triangle Δ_{II} is built out of smaller type-II triangles as in Figure 20(i), where the case $\lambda = 4$ is illustrated.*

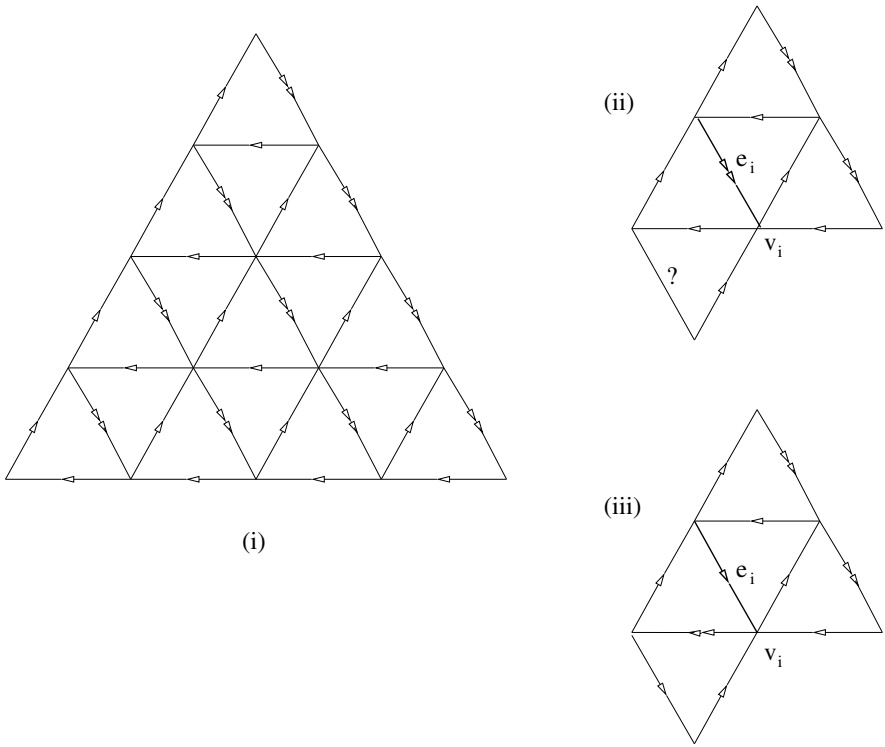

Figure 20 Structure of Δ_{II}; and the edge e_i of type (ii) or type (iii)

Proof. Observe that, under the standard action of G on X_0, the fixed point set of any conjugate of x^2 is a tree in the 1-skeleton consisting entirely of valence-2 edges. Since there is only one conjugacy class of order-2 elements in G (namely, that of x^2), it follows that the monomorphism φ must respect this conjugacy class. Thus the image of the valence-2 edge of any type-II triangle consists only of valence-2 edges. The image of the remaining two sides of the type-II triangle, because they are valence 3, must consist only of valence-3 edges and so the boundary of Δ_{II} is labeled as in Figure 20(i).

Now consider the interior edges of Δ_{II}. Label the first row of interior edges parallel to the valence-2 side of Δ_{II} by

$$e_1, e_2, \ldots, e_{\lambda-1}$$

going from top to bottom of the figure. Label the vertices in this row $v_0, v_1, \ldots, v_{\lambda-1}$ so that $e_i = [v_{i-1}, v_i]$ for $i = 1, \ldots, \lambda - 1$. Observe that, by considering the possibilities inside the link of vertex v_1, the neighborhood of the edge e_1 is constrained to appear either as in Figure 20(ii) or as in Figure 20(iii), depending on whether e_1 is valence 2 or 3 (respectively). We say that the edge e_1 is of "type" (ii) or (iii), respectively. By considering in the same way the possibilities in the link of vertex v_i, we see that if e_i is of type (ii) then e_{i+1} is of type (ii) or (iii), whereas if e_i is of type (iii) then so is e_{i+1}. However, since all edges running along the bottom of Δ_{II} are valence 3, it follows that $e_{\lambda-1}$ cannot be of type (iii). Therefore, every edge $e_1, e_2, \ldots, e_{\lambda-1}$ is of type (ii) and thus of valence 2. It now follows by induction on λ that every interior edge of Δ_{II} parallel to the valence-2 side is of valence 2 in X_0 and hence that Δ_{II} appears as in Figure 20(i). □

Resuming our proof of the proposition, suppose that $\lambda > 1$. We shall derive a contradiction. Consider an interior vertex v of one of the valence-3 sides of Δ_{II}. The subcomplex Δ_{II} contributes a path of type $T_2T_3^+T_3^+$ or of type $T_3^-T_3^-T_2$ in the link of v. Let Δ_{I} denote the image under F of a triangle of type I that is adjacent to Δ_{II} and contains the vertex v on its boundary. Since $\mathrm{Lk}(v, F(X_0))$ is a theta-graph embedded in $\mathrm{Lk}(v, X_0)$, the subcomplex Δ_{I} must contribute a path of type $T_1T_2T_1$ in the link of v. This implies that every edge in Δ_{I} which is adjacent to the boundary (or which lies on the boundary) is a valence-3 edge of X_0. But then any corner of Δ_{I} contains an equilateral triangle of side length 2 consisting of four type-I triangles, which contradicts the fact that no two triangles of type I can be adjacent along an edge of X_0. This completes the proof of Proposition 18 and hence of Theorem 17. □

References

[1] R. W. Bell and D. Margalit, *Braid groups are almost co-Hopfian,* preprint, arXiv:math.GT/0403145.

[2] N. Brady and J. Crisp, *On a torsion free hyperbolic group which is* CAT(0) *in dimension 2 but* CAT(−1) *only in dimension 3,* preprint, 2001.

[3] ———, *Two-dimensional Artin groups with* CAT(0) *dimension 3,* Geom. Dedicata 94 (2002), 185–214.

[4] T. Brady, *Artin groups of finite type with three generators,* Michigan Math. J. 47 (2000), 313–324.
[5] T. Brady and J. P. McCammond, *Three-generator Artin groups of large type are biautomatic,* J. Pure Appl. Algebra 151 (2000), 1–9.
[6] ———, *Four generator Artin groups of finite type* (in preparation).
[7] M. R. Bridson, *Length functions, curvature and the dimension of discrete groups,* Math. Res. Lett. 8 (2001), 557–567.
[8] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature,* Grundlehren Math. Wiss., 319, Springer-Verlag, Berlin, 1999.
[9] E. Brieskorn and K. Saito, *Artin–Gruppen und Coxeter–Gruppen,* Invent. Math. 17 (1972), 245–271.
[10] J. L. Dyer and E. K. Grossman, *The automorphism groups of the braid groups,* Amer. J. Math. 103 (1981), 1151–1169.
[11] K. Fujiwara, T. Shioya, and S. Yamagata, *Parabolic isometries of* CAT(0) *spaces and* CAT(0)*-dimensions,* preprint, 2003.
[12] P. Hanham, *The* CAT(0) *dimension of 3-generator Artin groups,* Ph.D. thesis, University of Southampton, 2002.
[13] W. Hurewicz and H. Wallmann, *Dimension theory,* Princeton Univ. Press, Princeton, NJ, 1941.
[14] K. Morita, *On the dimension of product spaces,* Amer. J. Math. 75 (1953), 205–223.

J. Crisp
I.M.B. (UMR 5584 du CNRS)
Université de Bourgogne
9 av. Alain Savary, BP 47 870
21078 Dijon
France

jcrisp@u-bourgogne.fr

L. Paoluzzi
I.M.B. (UMR 5584 du CNRS)
Université de Bourgogne
9 av. Alain Savary, BP 47 870
21078 Dijon
France

paoluzzi@u-bourgogne.fr