Complete square complexes

Daniel T. Wise

Abstract. We study groups which act cocompactly and properly discontinuously on the direct product of two trees. This class of groups turns out to be much richer than one might expect. An interplay is developed between the immersed flats in the complex and the subgroup separability properties of its fundamental group. This link between algebra and geometry leads to the solution of several problems concerning the residual properties of automatic groups and small-cancellation groups. In particular an explicit example is given of a compact non-positively curved square complex whose fundamental group is not residually finite. A more complicated such example is given with the property that its fundamental group has no finite quotients. The universal covers of these examples are isomorphic to the direct product of two trees. Other examples include a $C(4)$–$T(4)$ small-cancellation group which is not virtually torsion-free.

Contents

1. Introduction

The theory of small-cancellation groups was developed in the sixties, though its origins may be traced back much further. With the increased interest in geometric group theory, the various classes of small-cancellation groups have taken their places as examples of automatic and, in some cases, word-hyperbolic groups. In addition,

many simply-connected small-cancellation complexes admit metrics of non-positive curvature so their fundamental groups act properly-discontinuously and cocompactly on a CAT(0) space, and are 'CAT(0) groups'. Perhaps the simplest examples of 2-dimensional non-positively curved spaces that also have the small-cancellation property, are the *non-positively curved square complexes*. A non-positively curved square 2-complex is a 2-complex formed by gluing squares together, which has the property that the link of each vertex contains no cycles of length < 4. In addition to being CAT(0) groups, the fundamental groups of compact non-positively curved square complexes are finitely presented $C(4)$–$T(4)$ groups and thus biautomatic by [GS90], [GS91].

Finitely presented small-cancellation groups have numerous agreeable properties, and as usual, researchers sought to compare them with other classes of agreeable groups. One such class, is the class of *residually finite* groups which are groups with the property that every non-trivial element lies outside some finite index subgroup.

And so it was asked for some time whether all finitely presented small-cancellation groups are residually finite [Sch73], [Wal79], [Pri89]. More recently, the same question was asked about automatic groups [Ger92], and groups which act properly discontinuously and cocompactly on CAT(0) spaces.

The main purpose of this paper is to construct examples of non-residually finite groups belonging to each of these classes. These examples are the fundamental groups of certain natural spaces defined below. These results appeared as [Wis96b, Part II], and circulated earlier in [Wis].

Shortly thereafter, Burger and Mozes announced their now-famous work constructing infinite simple groups of the same type that I studied. We refer the reader to [BM97] for their announcement and to [BM00] for a complete account of their work. Later, Rattaggi resourcefully incorporated the powerful method of Burger–Mozes together with the small example X presented here, to produce other relatively small examples with infinite simple fundamental groups [Rat07]. Now that the dust has cleared, I hope that there is still some merit in my elementary approach which contrasts with the deeper and more powerful approach of Burger–Mozes.

The principal objects of study in this paper are *complete square complexes* (or CSCs for short). A CSC is a square complex with the property that the link of each vertex is isomorphic to a complete bipartite graph. Section 2 described the background and basic properties of $\mathcal{VH}$-complexes, which are a somewhat more general class of non-positively curved 2-complexes, and form a more natural context to describe some of the basic results and language that we will later need. Section 3 introduces CSCs and provides many of the basic theorems and definitions about CSCs that we shall need. For instance, Theorem 3.8 asserts that a square complex is a CSC if and only if its universal cover is isomorphic to the direct product of two trees. It is easy to see that CSCs satisfy the combinatorial non-positive curvature condition for square complexes, and therefore satisfy the more general $C(4)$–$T(4)$

small cancellation condition. It thus follows (by [GS90]) that the fundamental groups of compact CSCs are biautomatic. A simple proof of this result for the case of CSCs is sketched in Proposition 3.10.

The easiest examples of CSCs are the direct products $\Gamma_1 \times \Gamma_2$ of two graphs. Other examples include complexes with a finite cover that is isomorphic to $\Gamma_1 \times \Gamma_2$. In Section 4, I describe a more exotic example called X which is the main example in this paper. It is a CSC formed by gluing together only 6 squares. Most of the examples described later in the paper are constructed using X in some way, and many of their properties are inherited from X.

In Section 5, we describe an aperiodically tiled plane called an *anti-torus*. It is shown that this aperiodically tiled plane appears in the universal cover of X. This anti-torus lies behind most of our main results.

In Section 6, the anti-torus is used to show that the covering spaces of X are of a very limited nature. In algebraic terms, $\pi_1 X$ is not subgroup separable with respect to the subgroup of $\pi_1 X$ corresponding to the fundamental group of a certain subspace V of X contained in X^1. This subspace V, actually corresponds to one of the trees in the isomorphism between $\widetilde{X}$ and the direct product of two trees. Corollary 6.8 asserts that $\pi_1 X$ is not virtually $F_n \times F_m$. In geometric terms, we assert that X does not have a finite cover which is the product of two graphs. One might have expected that any group acting properly discontinuously and cocompactly on the direct product of two trees is virtually the direct product $F_n \times F_m$ of two free groups, but $\pi_1 X$ is an easy example to the contrary.

In Section 7, some algebra is used to show that when X is doubled along V to obtain a complex D, the resulting fundamental group $\pi_1 D$ is not residually finite. Since V is a totally geodesic subspace which is also contained in the 1-skeleton of X, the complex D is itself a CSC. This is quite surprising, since it means that $\pi_1 D$ and $F_4 \times F_3$ are algebraically very different but geometrically very similar since their (non-labeled) Cayley graphs are isomorphic.

In Section 8 we use D to build a compact CSC E such that $\pi_1 E$ has no non-trivial finite quotient (nor any non-trivial torsion quotient). The complex E is built up through a sequence of related complexes, which carry immersed copies of D, and which are glued together in such a way that the non-residual finiteness of $\pi_1 D$ leads to the related property in $\pi_1 E$. That this can be done is not surprising, since for instance, one would only have to embed $\pi_1 D$ so that it is not contained in any proper (finite index) normal subgroup.

In Section 9 various examples are given of biautomatic groups and CAT(0) groups which are not virtually torsion-free. Also, it is shown using similar ideas that the question of whether or not word-hyperbolic groups are virtually torsion-free is equivalent to the question of whether or not word-hyperbolic groups are residually finite. At the end of Section 9 there is a sketch of a possible method to attempt to produce examples of word-hyperbolic groups which are not residually finite.

In Section 10 I pose a collection of problems about compact CSCs and their fundamental groups. The problems tend to involve either algorithmic or algebraic issues. Because of the anti word-hyperbolic behavior of their fundamental groups, I feel strongly that CSCs are an excellent testing ground for some conjectures related to biautomatic and CAT(0) groups, and I hope that some of the problems will prove fruitful.

I am grateful to Martin Bridson who was my advisor during 1993–1995, and who encouraged me to study Gromov's problem of determining whether the universal cover of a non-positively curved square complex contains a flat if and only if it contains a periodic flat. My futile attempts at solving this problem ultimately led me to the example X and its anti-torus. It took a while to figure out how to utilize X to build non-residually finite fundamental groups, and it took a bit longer to realize that $\widetilde{X}$ was the product of two trees, but I eventually got there…. Finally, I am grateful to Tim Hsu who has been a close collaborator over the last eight years, and has persistently pressed me to finally publish this work. Indeed, Tim and I have applied these results in [HW98], [HW99] to form non-residually finite non-positively curved polygons of finite groups, and applied these results in [HW04] to construct CSCs whose fundamental groups have infinitely many non-isomorphic fixed subgroups.

2. $\mathcal{VH}$-complexes

2a. Definitions and basic properties of $\mathcal{VH}$-complexes

Notation 2.1. Throughout this paper, we let $\mathbb{R} = (-\infty, \infty)$ and $\mathbb{R}^+ = [0, \infty)$ with the usual structures as graphs, with 0-cells at each $n \in \mathbb{Z}$ and open 1-cells at each $(n, n+1)$. We let $I_n \subset \mathbb{R}$ denote the subgraph $[0, n]$, and we let $I = I_1 = [0, 1]$.

Definition 2.2 (Square complex $\mathcal{VH}$-complex). A *square complex X* is a combinatorial 2-complex whose 2-cells are attached by combinatorial paths of length 4. Thus, we think of each 2-cell as a square attached to X^1.

A square complex X is a $\mathcal{VH}$-*complex* if the 1-cells of X are partitioned into two classes V and H called *vertical* and *horizontal* edges respectively, and as in the first square of Figure 1, the attaching map of each 2-cell of X alternates between edges in V and H.

We let $V_X = V \cup X^0$ denote the *vertical* 1-*skeleton* and $H_X = H \cup X^0$ denote the *horizontal* 1-*skeleton*. For a 0-cell $x \in X^0$, we let V_x denote the component of V_X containing x. We define H_x similarly.

Remark 2.3 (Bipartite links). Let $x \in X^0$ where X is a 2-complex. We let $\mathrm{Link}(x)$ denote the *link* of x in X which is a graph whose vertices and edges correspond to the ends of 1-cells and corners of 2-cells incident with x. Note that $\mathrm{Link}(x)$ is topologized so that it looks like the ϵ-sphere about x in X.
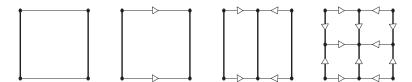
Figure 1. Squares: The squares in the figure above are meant to suggest (from left to right) a $\mathcal{VH}$-square, a horizontally directed $\mathcal{VH}$-square, a subdivided $\mathcal{VH}$-square (to obtain two horizontally directed subsquares), and the first barycentric subdivision of a $\mathcal{VH}$-square.

Recall that a graph Γ is *bipartite* if Γ^0 is partitioned into two disjoint classes such that each edge of Γ connects vertices from distinct classes.

Let X be a $\mathcal{VH}$-complex and let $x \in X^0$. The partition of the 1-cells of X into two classes V and H, induces a partition of the vertices of $\mathrm{Link}(x)$. Furthermore, since attaching maps of the squares of X alternate between 1-cells in V and 1-cells in H, we see that the edges of $\mathrm{Link}(x)$ connect vertices from different classes. Therefore, the $\mathcal{VH}$-structure on X induces a bipartite structure on $\mathrm{Link}(x)$ for each $x \in X^0$.

This motivates the following definition:

Definition 2.4 (Locally $\mathcal{VH}$). The square complex X is *locally $\mathcal{VH}$* if for each $x \in X^0$ there is a chosen bipartite structure on the graph $\mathrm{Link}(x)$. As in Remark 2.3, a $\mathcal{VH}$-complex has an induced local $\mathcal{VH}$-structure.

Note that X may admit several $\mathcal{VH}$-structures because a graph admitting a bipartite structure actually admits 2^c such structures, where c is the number of connected components of the graph.

Example 2.5 (Not global $\mathcal{VH}$). The simplest example of a square complex with a local $\mathcal{VH}$-structure which is not consistent with any (global) $\mathcal{VH}$-structure is a loop with one 0-cell and one 1-cell, where the bipartite structure on the link of the 0-cell has one vertex in each class. While the underlying complex in this example does have two $\mathcal{VH}$-structures, both are inconsistent with the local $\mathcal{VH}$-structure that we chose.

A square complex which has a local $\mathcal{VH}$-structure but no (global) $\mathcal{VH}$-structure, is obtained from a square by identifying two of its sides as in Figure 2.

Figure 2. Local but not global: The complex obtained by identifying two sides of a square as indicated above, has a local $\mathcal{VH}$-structure but no global $\mathcal{VH}$-structure.

While Example 2.5 shows that there are many examples of local $\mathcal{VH}$-complexes with no consistent (global) $\mathcal{VH}$-structure, we do have the following theorem which is analogous to the existence of orientable double covers of manifolds. See [Wis96b] or [Wis05] for the proof, which is straightforward.

Theorem 2.6 ($\mathcal{VH}$ double cover). *Let X be a square complex which is locally $\mathcal{VH}$. Then there is a double cover $\widehat{X} \to X$ such that the induced local $\mathcal{VH}$-structure on $\widehat{X}$ is consistent with a global $\mathcal{VH}$-structure.*

Definition 2.7 (Non-positive curvature). A square complex X is *non-positively curved* if all immersed cycles in the link of each 0-cell of X have length at least 4. A square complex satisfying this combinatorial link condition admits a locally CAT(0) metric (see [Gro87] or [BH99]).

Remark 2.8 (Locally-$\mathcal{VH}$ and curvature). If X is a locally-$\mathcal{VH}$ square complex, then X satisfies the combinatorial non-positive curvature condition if and only if there are no cycles of length 2 in the links of 0-cells of X. This is because all the cycles in a bipartite graph have even length, and so we may rule out the short cycles of length 1 or 3.

Definition 2.9 (Directed $\mathcal{VH}$-complex). We view the attaching map of each square of a square complex X, as a map from the boundary of the unit square $I \times I$ to X^1. We orient both horizontal edges of the unit square from left to right as illustrated in the second square of Figure 1. Let X be a $\mathcal{VH}$-complex and suppose that H_X is a directed graph. The $\mathcal{VH}$-structure on X is *horizontally directed* if the attaching map of each square of X is orientation preserving on its horizontal edges. We define *vertically directed* similarly. We shall use the term *directed* to mean *horizontally directed*.

Remark 2.10 (Subdividing). There is little loss of generality in considering only directed $\mathcal{VH}$-complexes. This is because, given a $\mathcal{VH}$-complex, we may subdivide H_X and subdivide each square of X by adding a vertical edge connecting the centers of its horizontal edges. If we orient all horizontal edges towards the new 0-cells then we obtain a directed $\mathcal{VH}$-complex. See the third square of Figure 1. Similarly, we can subdivide each square both vertically and horizontally so that X is both vertically and horizontally oriented. See the fourth square of Figure 1.

2b. The graph of spaces decomposition of a $\mathcal{VH}$-complex

Definition 2.11 (Vertical foliation and V_x). We now define a "singular vertical foliation" on the $\mathcal{VH}$-complex X. The unit square $I \times I$ is foliated by vertical line segments. Similarly, the image of a square in X is foliated by *vertical segments*

parallel to the pair of vertical edges on its boundary. For an arbitrary point $x \in X$ we define the *leaf* V_x to be the smallest subset of X having the property that $x \in V_x$ and that V_x contains any vertical segment which intersects it. This definition of V_x is consistent with the definition given earlier for $x \in X^0$.

Thinking of X as being foliated by these vertical subspaces, it is natural to take the quotient of X in which each leaf is identified to a point. This quotient is a graph denoted by Γ_X which we shall discuss in Definition 2.14.

Remark 2.12 (Singular leaves and directed $\mathcal{VH}$). When x is a point in the interior but not in the center of a horizontal edge, then for y very close to x, the leaf V_x is isomorphic to the leaf V_y by an isomorphism induced by sliding it along in X. However, when x is the center of a horizontal edge, the leaf V_x may be different from the surrounding leaves, in which case they correspond to double covers of it. This situation can occur only if X is not directed. It is convenient to add all such singular leaves V_x to the vertical 1-skeleton of X. This corresponds to subdividing certain squares of X. The resulting complex has a directed $\mathcal{VH}$-structure. Note that if H_X can be oriented so that V_X is directed then for any points x and y in the same horizontal edge, the leaves V_x and V_y are isomorphic by a translation isomorphism.

Example 2.13 (Möbius strip). A Möbius strip obtained by identifying the top and bottom horizontal edge of a square with a twisted identification map, has an obvious $\mathcal{VH}$-structure. The circle at the center of the Möbius strip is singular.

Definition 2.14 (The decomposition graph $X \to \Gamma_X$). Given a directed $\mathcal{VH}$-complex X, we define a graph Γ_X and a map $\varrho \colon X \to \Gamma_X$. The vertices of Γ_X correspond to the connected components of V_X, which are called *vertex spaces*. Note that each vertex space arises as V_x for some 0-cell $x \in X^0$. The edges of Γ_X correspond to the connected components of $X - V_X$, which are called *edge spaces*.

If x and y are in the same edge space, then V_x and V_y are isomorphic graphs and there is a natural isomorphism between them. Indeed, if x is a point in some edge space C, then $C \cong V_x \times (0, 1)$. It is natural to think of C as a subspace of $V_x \times [0, 1]$ which is a square complex which we denote by $\overline{C}$. We will also refer to $\overline{C}$ as an edge space.

For each edge space C, the inclusion $C \hookrightarrow X$ uniquely extends to a combinatorial map $\overline{C} \to X$. Since $\overline{C} \cong V_x \times I$, we obtain induced maps $V_x \times \{0\} \to V_X$ and $V_x \times \{1\} \to V_X$ which we call the attaching maps of the edge space C. Note that these are combinatorial maps and in case X is non-positively curved then they are immersions (i.e. local injections) and are therefore π_1-injective [Sta83]. The components of V_X that the ends of $\overline{C}$ are mapped to correspond to the vertices of Γ_X that the edge of Γ_X corresponding to C is attached to.

Finally, the map $X \to \Gamma_X$ is the quotient map induced by identifying each vertical leaf of Definition 2.11 to a point.

Construction 2.15 (Constructing X from data). In order to further understand the map $\varrho \colon X \to \Gamma_X$ we show how X can be built up from the information encoded in Γ_X and the associated data. This is a special case of the notion of a graph of spaces [SW79].

Consider a graph Γ_X, and suppose that for each vertex $v \in \Gamma_X^0$ we have an associated *vertex space* X_v which is a graph, and for each edge $e \in \mathrm{Edges}(\Gamma_X)$ we have an associated *edge space* $X_e \times I$ where X_e is a graph. Suppose that for each edge e which is attached to the vertices e_0 and e_1, there are corresponding maps $\phi_{e0} \colon X_e \times \{0\} \to X_{e_0}$ and $\phi_{e1} \colon X_e \times \{1\} \to X_{e_1}$.

Using this data we may construct a $\mathcal{VH}$-complex X as follows: Let V_X be the disjoint union of the set of vertex spaces, that is

$$V_X = \Big\{ \bigcup_{v \in \Gamma_X^0} X_v \Big\}.$$

We form X by attaching $X_e \times I$ along its ends to V_X for each edge e of Γ_X, so that

$$X = \Big\{ V_X \bigcup_{e \in \mathrm{Edges}(\Gamma_X)} X_e \times I \Big\}.$$

Theorem 2.16 (Graph of free groups). *Suppose that X is a non-positively curved $\mathcal{VH}$-complex. Then the map $\varrho \colon X \to \Gamma_X$ determines a splitting of $\pi_1 X$ as a graph of free groups. Specifically, for each vertex $v \in \Gamma_X$, $\pi_1 v = \pi_1 X_v$ and for each edge e of Γ_X, $\pi_1 e = \pi_1 X_e$. For each edge e, attached to the vertices e_0 and e_1, the inclusion $\pi_1 X_e \to \pi_1 X_{e_i}$ is induced by the maps ϕ_{ei} described above.*

Proof. The non-positive curvature hypothesis implies that the attaching maps ϕ_{ei} are immersions. They therefore induce π_1-injections [Sta83]. $\square$

We close this section with the following:

Definition 2.17 (Map of $\mathcal{VH}$-complexes). Let X and Y be $\mathcal{VH}$-complexes, and let $\phi \colon X \to Y$ be a cellular map. Then ϕ is a $\mathcal{VH}$-*map* provided that $\phi(V_X) \subset V_Y$ and $\phi(H_X) \subset H_Y$. We will be dealing with combinatorial maps, and so ϕ is a $\mathcal{VH}$-map provided it maps vertical edges to vertical edges and horizontal edges to horizontal edges. Note that for a $\mathcal{VH}$-map ϕ, there is an induced map $\phi_\star \colon \Gamma_X \to \Gamma_Y$.

3. Complete square complexes

Definition 3.1 (Complete square complex). A *complete bipartite graph* is a graph whose vertices are partitioned into two classes V and H such that there is exactly one edge joining v, h for each $(v, h) \in V \times H$, and there are no other edges.

A connected square complex X is a *complete square complex* (*CSC*) if Link(v) is a complete bipartite graph for each $v \in X^0$.

Remark 3.2 (CSCs are non-positively curved). If X is a CSC then X obviously satisfies the combinatorial non-positive curvature condition for square complexes because for each vertex $v \in X^0$, each cycle in Link(v) has even length, and there is no length 2 cycle.

Example 3.3 (Not global-$\mathcal{VH}$). While a CSC X is obviously a locally $\mathcal{VH}$-complex, there are examples of CSCs which do not have a global $\mathcal{VH}$-structure. For instance, the Klein bottle is the standard 2-complex of $\langle a, b \mid a^2 b^2 \rangle$, and this 2-complex is a CSC with no $\mathcal{VH}$-structure.

Lemma 3.4 (CSC has $\mathcal{VH}$ double cover). *Let X be a CSC, then there is a double cover $\widehat{X} \to X$ such that $\widehat{X}$ is a CSC which is also a $\mathcal{VH}$-complex.*

Proof. This follows from Theorem 2.6 because X is a locally $\mathcal{VH}$-complex, but there is also a more direct proof. Indeed, the automorphism group of $\widetilde{X}$ has a subgroup of index 2 which preserves the local $\mathcal{VH}$-structure. Its intersection with the covering transformation group gives the desired subgroup of index ≤ 2 of $\pi_1 X$. Note that some extra care must be taken in the degenerate case where X is a graph. $\square$

Lemma 3.5 (Trees). *Let $\widetilde{X}$ be a simply connected non-positively curved $\mathcal{VH}$-complex. Then for each point $x \in \widetilde{X}$, the spaces V_x and H_x are trees.*

Proof. The map $V_x \to \widetilde{X}$ is a local-isometry and therefore lifts to an isometry. (I first learned of this idea from [Mos95], and it has had a large influence on me.) Note that this reasoning shows that local-isometries lift to isometries between universal covers and hence induce π_1-injections, so $\pi_1 V_x$ is trivial, and hence the graph V_x is a tree. $\square$

Lemma 3.6 (Connecting with vertical and horizontal). *Any two points w and z in a connected simply connected CSC $\widetilde{X}$ can be uniquely connected by a path $\sigma_v \cdot \sigma_h$, which is the concatenation of a vertical and horizontal path. Similarly, any two points can be uniquely connected by a path $\sigma'_h \cdot \sigma'_v$ which is the concatenation of a horizontal and vertical path.*

Proof. By subdividing $\widetilde{X}$ we may assume that w and z are in $\widetilde{X}^0$. Since $\widetilde{X}$ is connected we may connect w to z by a path of minimal length σ in $\widetilde{X}^1$. We now show that σ is path homotopic to a path $\sigma_v \cdot \sigma_h$ which is the concatenation of a vertical path beginning at w, followed by a horizontal path ending at z. To see this, note that any length 2 subpath say $h_1^{\epsilon_1} v_2^{\epsilon_2}$ of a path λ, determines a unique square s of $\widetilde{X}$, whose attaching map can be thought of as $h_1^{\epsilon_1} v_2^{\epsilon_2} = v_3^{\epsilon_3} h_4^{\epsilon_4}$ for some choices of h_4, v_3, ϵ_3, and ϵ_4. But then we can homotope λ through the square s to a path of the same length by homotoping the subpath $h_1^{\epsilon_1} v_2^{\epsilon_2}$ to $v_3^{\epsilon_3} h_4^{\epsilon_4}$, and leaving everything outside the subpath $h_1^{\epsilon_1} v_2^{\epsilon_2}$ unchanged.

Define the *complexity* of a path λ in $\widetilde{X}^1$ as follows: For each horizontal edge a in λ, let K_a denote the number of vertical edges in λ coming after a. We let the complexity of λ equal $\sum_{a \in \text{Horizontal Edges}(\lambda)} K_a$.

Observe that each square homotopy as above reduces the complexity of the path. It follows that after finitely many such homotopies, we obtain a path of complexity 0 which is therefore of the form $\sigma_v \cdot \sigma_h$. □

Lemma 3.7 ($V_w \cap H_z$ is one point). *Let w, z be points in the connected simply connected CSC $\widetilde{X}$, then the intersection of V_w and H_z consists of a single point.*

Proof. To see that $V_w \cap H_z$ has at least one point, note that by Lemma 3.6, there is a path $\sigma_v \cdot \sigma_h$ which connects w to z. But σ_v ends at the point where σ_h begins, and so this point is in $V_w \cap H_z$, because σ_v is a path in V_w and σ_h is a path in H_z.

We will apply the Combinatorial Gauss–Bonnet Theorem (see, for instance, [Ger87]). Let σ_v and σ_h be paths in V_x and H_x with the same endpoints. Let $D \to X$ be a reduced disc diagram whose boundary path is $\sigma_v \sigma_h^{-1}$. Assigning angles of $\frac{\pi}{2}$ to each corner of each 2-cell, we see that each 2-cell and interior 0-cell has non-positive curvature, and each boundary 0-cell (except for the endpoints of σ_v and σ_h) has non-positive curvature. If σ_v and σ_h are non-trivial then the curvature at their endpoints is $\leq \frac{\pi}{2}$ which is impossible. If one is non-trivial and the other is trivial then the two endpoints are identified and that point has curvature $\leq \pi$ which is impossible. Therefore, both σ_v and σ_h are trivial. □

As motivation for the following theorem, observe that the condition that X is a CSC is equivalent to X being locally isomorphic to the direct product of two trees.

Theorem 3.8 (X is CSC $\Leftrightarrow$ $\widetilde{X} \cong$ Tree $\times$ Tree). *Let $\widetilde{X}$ be the universal cover of the square complex X. $\widetilde{X}$ is a CSC if and only if X is a CSC. Specifically, let $x \in \widetilde{X}$, and consider the two trees V_x and H_x. If X is a CSC then $\widetilde{X}$ is isomorphic to $V_x \times H_x$.*

Proof. Note that by Lemma 3.4, $\widetilde{X}$ is a $\mathcal{VH}$-complex, so the statement of the theorem makes sense. We define a map $\phi \colon \widetilde{X} \to V_x \times H_x$ defined by $y \mapsto (V_x \cap H_y, H_x \cap V_y)$. The map ϕ is well-defined by the previous lemma.

Note that ϕ has an inverse, $V_x \times H_x \to \widetilde{X}$ defined by $(a, b) \mapsto H_a \cap V_b$. It is easy to see that these two maps are inverse to each other because $H_a = H_y$ and $V_b = V_y$. It is easy to check that ϕ is cellular as well and so it is an isomorphism. $\quad\square$

Remark 3.9 (Factoring π_1). When x is the only 0-cell of X, then one can deduce from Lemma 3.6 or Theorem 3.8 that $\pi_1 X$ factors as $\pi_1 X = \pi_1(V_x)\pi_1(H_x) = \pi_1(H_x)\pi_1(V_x)$. Note that we always have $\pi_1(V_x) \cap \pi_1(H_x) = \{1\}$, but $\pi_1(V_x)\pi_1(H_x)$ is not always a subgroup when X has more than one 0-cell.

Proposition 3.10 (Biautomatic: the Tree $\times$ Tree language). *Let X be a compact CSC. Then $\pi_1 X$ is biautomatic.*

Using biautomatic structures for groupoids (see [ECH$^+$92]), it is easy to show that the language consisting of paths which are vertical followed by horizontal, is the language of a biautomatic structure for the fundamental group of a compact CSC. Note that it is obvious that this is a language of geodesics which satisfies the fellow traveller property. A similar language works for the higher dimensional cases (that is Tree $\times$ Tree $\times$ Tree, etc.). Biautomatic structures have been obtained for more general classes of complexes: For finite $C(4)$–$T(4)$-complexes this is proven in [GS90], and for compact non-positively curved cubulated complexes this is proven in [NR98].

Lemma 3.11 (CSC $\Leftrightarrow$ attaching maps are covering maps). *A directed $\mathcal{VH}$-complex X is a CSC if and only if each attaching map of each edge space in the decomposition of X is a covering map.*

Proof. If all the attaching maps are covering maps, then it is easy to see that the links are complete bipartite graphs. To see the converse, we note that for cellular maps between graphs, a covering map is just a local homeomorphism. But this local condition is satisfied by the attaching maps of the edge spaces in the case of CSCs because the links of vertices are complete bipartite graphs. So, in case X is a CSC, all the attaching maps of edge spaces are covering maps. $\quad\square$

Remark 3.12 (Weighted graph). Because of Lemma 3.11, we are able to add some structure to the decomposition graph Γ_X. When X is a directed CSC, we label the ends of edges in Γ_X with numbers indicating the covering degrees of the associated attaching maps.

4. The main example: X

Example 4.1 (X as six squares). We define X to be the complex consisting of the six squares indicated in Figure 3. The squares are glued together as indicated by

the labeled directed edges. It is easily verified that X is a $\mathcal{VH}$-complex and that it is a CSC. Note that X has a unique 0-cell. Let $H = H_X$, denote the subcomplex consisting of the 2 horizontal edges, and let $V = V_X$ denote the subcomplex consisting of 3 vertical edges.
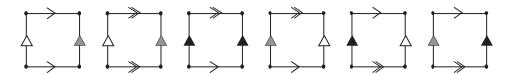

Figure 3. X as six squares: The figure above indicates the gluing pattern for the six squares of X. The three vertical edges colored white, grey, and black are denoted a, b, and c respectively. The two horizontal edges, single and double arrow, are denoted x and y respectively.

Explanation 4.2 (The vertical decomposition of X). Since X is both vertically and horizontally directed, X may be described in terms of both its vertical and horizontal decompositions. We now describe the vertical decomposition of X.

Let $\rho_L \colon L \to V$ and $\rho_R \colon R \to V$ be the covering spaces of V indicated in Figure 4. Denote by $\tau \colon L \to R$ the obvious (translation) graph isomorphism. Note that τ is *not* a covering space isomorphism.
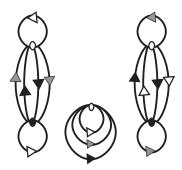

Figure 4. Vertical decomposition of X: The bouquet of 3 circles in the center is V. The graph on the left and right are denoted by L and R respectively. There are covering maps $\rho_L \colon L \to V$ and $\rho_R \colon R \to V$ induced by the colorings of the graphs.

Using the information contained in ρ_L, ρ_R and τ we may form X as follows: Γ_X will be a graph with one vertex v, and one edge e. Let Λ_e denote a graph isomorphic to R and thus isomorphic to L via τ. The vertex $v \in \Gamma_X$ corresponds to $\Lambda_v = V \subset X$. The edge space $\Lambda_e \times I$ corresponds to the edge e in Γ_X. Identify $\Lambda_e \times \{0\}$ and $\Lambda_e \times \{1\}$ with L and R respectively, in such a way that the isomorphism $\Lambda_e \times \{0\} \cong \Lambda_e \times 1$

is the same as τ. Now, glue each end of $\Lambda_e \times I$ to V using the maps ρ_L and ρ_R, respectively. So we have

$$X = \big((\Lambda_e \times I) \cup V\big)\big/\{(x, 0) \sim \rho_L(x),\ (x, 1) \sim \rho_R(x) \mid x \in \Lambda_e\}.$$

Notice that this describes $\pi_1 X$ as an HNN extension of a free group of rank 3, associating subgroups of index 2. The decomposition in the horizontal direction provides a decomposition of $\pi_1 X$ as an HNN extension of a free group of rank 2, associating subgroups of index 3.

5. The anti-torus

Definition 5.1 (Anti-torus). Let Y be a CSC with a $\mathcal{VH}$-structure, and let $\widetilde{Y}$ denote the universal cover of Y. Let $y \in Y^0$ and let $\tilde{y} \in \widetilde{Y}$ denote a preimage of y.

Let $s_h \to H_y \to Y$ and $s_v \to V_y \to Y$ denote immersed circles based at y. Let $\tilde{s}_h \hookrightarrow \widetilde{Y}$ and $\tilde{s}_v \hookrightarrow \widetilde{Y}$ denote the induced maps between universal covers based at $\tilde{y}$.

Let $\tilde{s}_h \times \tilde{s}_v$ denote the convex hull of $\tilde{s}_h$ and $\tilde{s}_v$ in $\widetilde{Y}$. Note that since $\widetilde{Y} \cong H_{\widetilde{Y}} \times V_{\widetilde{Y}}$, the convex hull of $\tilde{s}_h \cup \tilde{s}_v$ is the product subspace $\tilde{s}_h \times \tilde{s}_v \hookrightarrow H_{\widetilde{Y}} \times V_{\widetilde{Y}}$.

We say that the plane $\tilde{s}_h \times \tilde{s}_v \subset \widetilde{Y}$ is an *anti-torus* if it is not tiled periodically by the preimages of squares of Y. More precisely, it is an anti-torus if the map $\tilde{s}_h \times \tilde{s}_v \to Y$ does not factor through a torus $T^2 \to Y$.

We shall now describe an immersed anti-tori in a compact CSC which will play a fundamental role in this work.

Example 5.2 (The anti-torus Π immersed in X). Let $\tilde{c}$ denote the infinite periodic vertical line in $\widetilde{X}$ which is the component containing the basepoint $\tilde{p} \in \widetilde{X}^0$ of the preimage in $\widetilde{X}$ of the loop labeled by c in X. Define $\tilde{x}$ similarly. Let Π denote the plane in $\widetilde{X}$ that is the convex hull of $\tilde{c}$ and $\tilde{x}$. The plane Π, is tiled by the six orbits of squares in $\widetilde{X}$ as partially illustrated in Figure 5.

The remainder of this section is devoted to proving that:

Theorem 5.3. Π *is not doubly periodic and therefore* Π *is an anti-torus.*

Let the basepoints l_0, v_0, and r_0 of L, V, and R be represented by the white vertices. Let the second vertex of L be denoted by l_1.

Given a path $\sigma : I \to V$ with $\sigma(0) = x_0$, let $\tilde{\sigma}^o$ and $\tilde{\sigma}^1$ denote the lifts of σ to L based at l_0, and l_1 respectively.

Define $\phi_i(\sigma) : I \to V$ to be $\phi_i(\sigma) = \rho_r \circ \tau \circ \tilde{\sigma}^i$.

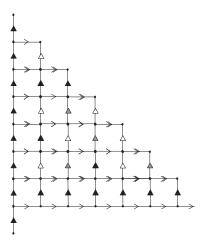Let $\iota : V \to V$ be the orientation-preserving automorphism that permutes the edges a and c and fixes b.

Figure 5. The anti-torus Π: The plane Π above is the convex hull of two periodically labeled lines in $\widetilde{X}$. A small region of the northeast quadrant has been tiled by the squares of X.

We are interested in the case were σ is a combinatorial path, and so is naturally represented as a word in $\{a^{\pm 1}, b^{\pm 1}, c^{\pm 1}\}$. As usual, the product of two paths $\sigma_1 \cdot \sigma_2$ is represented by the concatenation of words.

Define $\#_a \colon \pi_1 V \to \mathbb{Z}$ to be the homomorphism induced by $a \mapsto 1$, $b \mapsto 0$, and $c \mapsto 0$. Note that this merely counts the number of signed occurrences of a in the reduced word representing the element. Define $\#_b$ and $\#_c$ similarly, and let $\#_{b+c}$ denote $\#_b + \#_c$.

Lemma 5.4. (1) $\sigma \in \pi_1(L) \iff \#_{b+c}(\sigma) \equiv_2 0$.

(2) $\phi_1 = \iota \circ \phi_0$.

(3) $\sigma \notin \pi_1(L) \implies \phi_0(\sigma^2) = \phi_0(\sigma) \cdot \phi_1(\sigma)$.

Proof. The statements follow easily from the definitions. □

Lemma 5.5 (Doubling lemma). $\sigma \notin \pi_1(L) \implies \phi_0(\sigma^2) \notin \pi_1(L)$.

Proof.
$$
\begin{aligned}
\#_{b+c}(\phi_0(\sigma^2)) &\equiv_2 \#_a(\phi_0(\sigma^2)) \\
&= \#_a(\phi_0(\sigma) \cdot \phi_1(\sigma)) \\
&= \#_a(\phi_0(\sigma)) + \#_a(\iota \circ \phi_0(\sigma)) \\
&= \#_a(\phi_0(\sigma)) + \#_c(\phi_0(\sigma)) \\
&= \#_b(\sigma) + \#_c(\sigma) \\
&= \#_{b+c}(\sigma) \\
&\equiv_2 1.
\end{aligned}
$$

The first equality follows because $|\sigma^2| \equiv_2 0$. The second follows from Lemma 5.4.3. The third follows from Lemma 5.4.2 and the fact that $\#_a$ is a homomorphism. The fourth follows from the definition of ι. The fifth follows because the letters that may be transformed by ϕ_0 into either a or c are the letters b and c (not respectively). The sixth follows by definition. The seventh follows from Lemma 5.4.1 and the assumption that $\sigma \notin \pi_1(L)$. $\qquad\square$

Corollary 5.6. $\phi_0^n(c^{2^n}) \notin \pi_1(L)$.

Proof. We use Lemma 5.5 to prove this by induction. When $n = 0$ we have $\phi^0(c^{2^0}) \notin \pi_1(L)$ since this just means that $c \notin \pi_1(L)$. Suppose that the statement is true for $n = k$, we shall prove it for $n = k + 1$.

By Lemma 5.5, $\phi_0^k(c^{2^k}) \notin \pi_1(L)$ implies that $\phi_0\big((\phi_0^k(c^{2^k}))^2\big) \notin \pi_1(L)$. But we shall show below that $\big(\phi_0^k(c^{2^k})\big)^2 = \phi_0^k(c^{2^{k+1}})$ and so substituting proves the lemma.

We now prove by induction on k that $\phi_0^k(w^{2^{k+1}}) = [\phi_0^k(w^{2^k})]^2$ for any word w in $a^{\pm 1}, b^{\pm 1}, c^{\pm 1}$. The base case where $k = 0$ holds because $\phi_0^0(w^2) = w^2 = [w^1]^2 = [\phi_0^0(w^1)]^2$. Now the result follows from the following equalities:

$$\phi_0^k(w^{2^{k+1}}) = \phi_0^{k-1}\big(\phi_0(w^{2^{k+1}})\big) = \phi_0^{k-1}\big([\phi_0(w^2)]^{2^k}\big)$$
$$= \big[\phi_0^{k-1}\big(\phi_0(w^2)^{2^{k-1}}\big)\big]^2 = \big[\phi_0^{k-1}[\phi_0(w^{2^k})]\big]^2 = \big[\phi_0^k(w^{2^k})\big]^2.$$

The first and last equalities hold because $\phi_0^{k-1} \circ \phi^0 = \phi_0^k$. The second and fourth equalities hold because $[\phi_0(u^2)]^m = \phi_0(u^{2m})$ for any u, since $u^2 \in \pi_1(L, l_0)$. The third equality holds by induction. $\qquad\square$

Remark 5.7. Notice that $aba^{-1}b^{-1} \xrightarrow{\phi_0} bab^{-1}a^{-1} \xrightarrow{\phi_0} aba^{-1}b^{-1}$. So Lemma 5.5 cannot be strengthened to say that for any σ, there exists n such that, $\phi_0^n(\sigma) \notin \pi_1(L)$.

Remark 5.8 (Geometric interpretation). Consider the horizontal geodesic in $\widetilde{X}$ based at $\tilde{p}$ corresponding to x^n. Given a vertical geodesic based at $\tilde{p}$ represented by some word σ in $\{a^{\pm 1}, b^{\pm 1}, c^{\pm 1}\}$, the convex hull of the union of these two paths at $\tilde{p}$, is a rectangle. The word $\phi^n(\sigma)$ is simply the label (beginning at the endpoint of x^n) of the side opposite σ in this rectangle. In a sense, all of the 'calculations' in the lemmas above are really occurring inside of $\widetilde{X}$, and Corollary 5.6 is a statement about Π. Using this point of view, it is easy to show that Π is not periodic, merely as a consequence of Lemma 5.5. However, instead of going through that relatively simple argument, we will look more carefully at the tiling of Π, and obtain a somewhat more informative proof of aperiodicity.

Let $W_n(m)$ denote the horizontal positive path in Π of length n beginning at the endpoint of the vertical path c^m (see Figure 6). Thus, $W_n(m)$ determines a word

which is the label of the side opposing x^n in the rectangle which is the convex hull of x^n and c^m. We will also use $W_n(m)$ to denote this associated word.
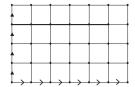

Figure 6. Computations in Π: The figure above represents the upper right quadrant of Π. The bold horizontal line, represents the path corresponding to $W_5(3)$.

Proposition 5.9 (Period doubling). *For each n, the words $\{W_n(m) \mid 0 \le m \le 2^n - 1\}$ are all distinct, and thus, every positive word of length n is $W_n(m)$ for some m.*

Proof. We begin by observing that a vertical path represents an element of $\pi_1(L)$ if and only if the convex hull of its base lift together with the horizontal edge labeled by y or x yields a rectangle whose top edge is labeled by y or x respectively.

Now, arguing by induction, we assume that the theorem is true for $n = k$, so that the $W_k(m)$ are distinct for $0 \le m \le 2^k - 1$. In particular, it follows that

$$W_k(m) = W_k(m + 2^k).$$

Consequently the first k letters of the words $W_{k+1}(m)$ and $W_{k+1}(m + 2^k)$ are equal, therefore it suffices to show that the last letters of these words are distinct. However, the vertical word connecting the endpoints of the paths associated with $W_k(m)$ and $W_k(m + 2^k)$ is just a subword of length 2^k of the word $\left(\phi_0^k(a^{2^k})\right)^2$, so it is a conjugate of $\phi_0^k(a^{2^k})$ which, by Corollary 5.6, is not an element of $\pi_1(L)$. It follows from our observation at the beginning of the proof, that the last letters of $W_{k+1}(m)$ and $W_{k+1}(m + 2^k)$ are distinct, and we are done. $\square$

We can now prove Theorem 5.3 which we restate as follows:

Proposition 5.10. Π *is an anti-torus.*

Proof. To see that Π is aperiodic, observe that Proposition 5.9 shows that the width n, infinite vertical strip in Π bounded on the left by the $\tilde{c}$ axis has period 2^n. Alternatively, the proof of Proposition 5.9 shows that every finite positive word in x and y appears in Π, and so Π cannot possibly be periodic. $\square$

Question 5.11 (Finding anti-tori). It would be remarkable if there were an algorithm to determine whether or not two perpendicular elements in the fundamental group of

a CSC do not have powers that commute, or in other words, to decide whether the convex hull of their axes is an anti-torus or a periodic plane. It seems very likely that CSCs typically have immersed anti-tori. An algorithm in the context of certain non-positively curved square complexes which are rather different from CSCs is described in [Wis05].

6. The covering spaces of X

Let Y denote a directed CSC and let $Y \to \Gamma_Y$ denote the associated vertical decomposition. A covering space $\widehat{Y}$ of Y has an induced directed $\mathcal{V}\mathcal{H}$-structure, and the covering map gives rise to a commutative diagram:

$$
\begin{array}{ccc}
\widehat{Y} & \longrightarrow & \Gamma_{\widehat{Y}} \\
\rho \downarrow & & \downarrow \rho_\star \\
Y & \longrightarrow & \Gamma_Y
\end{array}
$$

It is natural to label each end of each edge of Γ_Y with the degree of the associated attaching map in X. So, if the initial vertex of $\varepsilon \in \text{Edges}(\Gamma_Y)$ is v, then we label the end of ε at v with the degree of the associated attaching map $Y_\varepsilon \times \{0\} \to Y_v$ of the edge space corresponding to ε.

Keeping track of this additional information, the map $\rho_\star \colon \Gamma_{\widehat{Y}} \to \Gamma_Y$ is *locally weight conserving*. That is, the link of each vertex of $\Gamma_{\widehat{Y}}$ and Γ_Y has a weighting by natural numbers, and for each $v \in \Gamma_{\widehat{Y}}$, the map $\text{Link}(v) \to \text{Link}(\rho_\star(v))$ has the property that the sum of the weights of the preimage of a point is equal to the weight of that point. This is a rather simple restriction, but it will enable us to limit the types of covering spaces that may arise. We record this as:

Lemma 6.1. *The map $\rho_\star$ is weight conserving.*

Remark 6.2 (Restrictions on covering spaces of X). Since the attaching maps of the sole edge space of X are both of degree 2, (that is, the maps $\rho_L \colon L \to V$ and $\rho_R \colon R \to V$ are each of degree 2, as above), there are only four local models for the map $\Gamma_{\widehat{X}} \to \Gamma_X$ induced by a covering space $\widehat{X} \to X$. These are illustrated in Figure 7 below. Using the anti-torus Π, we will be able to show that under certain circumstances only case (iv) may occur.

Lemma 6.3 (Anti-torus obstruction). *Suppose that $\widehat{X} \to X$ is a regular covering, and suppose there exists n such that the closed path c^n lifts to a closed path in $\widehat{X}$. Then only case* (iv) *of Figure 7 may occur. That is, $\bigl(\pi_1 V \cap \pi_1(\widehat{X})\bigr) \not\subset \pi_1(L)$, and $\bigl(\pi_1 V \cap \pi_1(\widehat{X})\bigr) \not\subset \pi_1(R)$.*
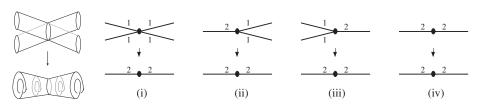
Figure 7.   The four cases: Diagrams (i), (ii), (iii), and (iv) represent the four different local possibilities for the map $\Gamma_{\widehat{X}} \to \Gamma_X$. The figure on the left is meant to suggest the type of local behavior of a map $\widehat{X} \to X$ that would produce case (i).

It is worth mentioning that in the case where $\widehat{X}$ is a finite cover of X, cases (ii) and (iii) are automatically excluded because components of $\rho^{-1}(R)$ and $\rho^{-1}(L)$ must be isomorphic as graphs and so it is impossible for one of them to be isomorphic to a proper cover of $\widehat{V}$ while the other one is isomorphic to the trivial cover. A different proof that there is no finite cover of type (i) can be obtained by combining Proposition 6.6 and Corollary 6.8 below.

Proof of Lemma 6.3. Suppose that case (iv) does not occur, then in the decomposition of $\widehat{X}$, the attaching maps of edge spaces corresponding to at least one of $L \to V$ or $R \to V$ have degree 1.

The idea of the proof is to assume that a cover $\widehat{X} \to X$ exists, and produce a contradiction, namely that either the right or left half of Π is (singly) periodic. We do this by showing that $\widehat{\Pi}$ factors through an immersion of a cylinder $C \to \widehat{X}$, where $C = S^1 \times [0, \infty)$ and where the map $S^1 \to X$ represents c^n and the map $[0, \infty) \to X$ represents $\tilde{y}$. This seems is very plausible since $C \to X$ is determined by its restriction to these subspaces, but I have included a more detailed explanation below.

Let us consider the case where the attaching maps of the edge spaces of $\widehat{X}$ corresponding to the map $L \to V$ are of degree 1.

From the point of view of a preimage $\widehat{V}$ of V, this means that outgoing horizontal arrows point towards a graph isomorphism between $\widehat{V}$ and a cross section of the respective edge space, that is, towards a component of the preimage L. This means that given a closed path in $\widehat{V}$ one is able to push it in the direction of the horizontal arrows, towards another copy of $\widehat{V}$. More precisely, given a map of a graph $\Gamma \to \widehat{V}$ one can extend it through this edge space to a combinatorial map of $\Gamma \times I \to \widehat{X}$. Since each y-labeled outgoing edge points towards such a graph isomorphism, we are able to extend the map of Γ so that an edge $\gamma_0 \times I \subset \Gamma \times I$ is sent to an outgoing y-labeled edge.

Denote by C the semi-infinite cylindrical square-complex $S^1 \times [0, \infty)$ where S^1 consists of n edges and $[0, \infty)$ is subdivided into edges in the usual way. Also, orient all the edges of $[0, \infty)$ so that they are pointing away from 0.
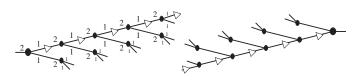
Figure 8. Trajectory of C in $\widehat{X}$: The figures on the left and right represent $\Gamma_{\widehat{X}}$ corresponding to cases (ii) and (iii) of Figure 7. The sequence of arrowed edges in the figure on the left, is meant to suggest the image of the map $\Gamma_C \to \Gamma_{\widehat{X}}$ that would be induced by the map $C \to \Gamma_X$ constructed in the proof. The sequence of arrowed edges in the figure on the right is meant to suggest the same for a similar proof in case (iii).

We identify S^1 with $S^1 \times \{0\} \hookrightarrow C$ and, letting s be a basepoint 0-cell of S^1, we identify $[0, \infty)$ with $s \times [0, \infty) \hookrightarrow C$. In addition, let $(s, 0)$ denote the basepoint of C.

We will show that there is a $\mathcal{V}\mathcal{H}$-immersion $\phi\colon C \to \widehat{X}$, so that $S^1 \hookrightarrow C$ gets sent to a closed path in $\widehat{X}$ with label c^n, and so that $[0, \infty) \hookrightarrow C$ gets sent to the infinite path $\tilde{y}$. This will result in a contradiction as explained below, so it follows that our assumption that there is a covering space $\widehat{X}$ falling into case (iii) above is incorrect.

We begin by using our assumption that c^n lifts to a closed path in order to define $\phi\colon S^1 \to \widehat{X}$ to represent such a closed path. Then, as above, we may extend the map ϕ restricted to $S^1 = S^1 \times \{0\}$ to $S^1 \times [0, 1)$ so that the edge $s \times [0, 1]$ is sent to a y-labeled outgoing edge. Similarly we extend ϕ to $S^1 \times [0, 2)$, and so on to define ϕ on all of C.

Now, since there are only finitely many labels, and because the "trajectory" of C in $\widehat{X}$ is periodic, the induced tiling of C is periodic. But since the anti-torus Π factors through C, or rather, factors through the induced map $C \to X$, it would follow that Π itself is periodic. This would be a contradiction.

Cases (i) and (ii) of Figure 7 are handled similarly, so that the only possible covers $\widehat{X}$ of X which contain a lift of c^n are of type (iv). $\qquad\square$

Note that statements similar to Lemma 6.3 can be proven for any compact CSC containing an anti-torus.

A subgroup K of a group G is *separable* if K is the intersection of finite index subgroups of G. In this case, G is referred to as K-*separable*. We say the subgroup K can be separated from an element g if there is a finite index subgroup $H \subset G$ such that $K \subset H$ but $g \notin H$.

Note that K is separable from g if and only if there is a finite quotient $G \to \overline{G} = G/N$ such that $\overline{g} \notin \overline{K}$. Indeed $gN \not\subset KN$ if $\overline{g} \notin \overline{K}$, so we could let $H = KN$. Conversely, the permutation representation of G on the left cosets of H gives a finite quotient such that $\overline{g} \notin \overline{H}$ if $g \notin H$.

The following result explains that $\pi_1 V$ is not separable from certain elements of $\pi_1 X$ in certain quotients, where we use the word "separable" analogously to its use in the above discussion. In particular, $\pi_1 V$ cannot be separated from certain elements by a finite index subgroup.

Corollary 6.4 ($\pi_1 X$ is not $\pi_1 V$-separable). *Consider the homomorphism $\psi : \pi_1 X \rightarrow \mathbb{Z}$ induced by $a \mapsto 0, b \mapsto 0, c \mapsto 0, x \mapsto 1, y \mapsto 1$. For each $n \geq 1$ we have the inclusion*

$$\ker(\psi) \subset \pi_1 V \langle\langle c^n \rangle\rangle.$$

Consequently, $\pi_1 V$ is not separable from any element of $\mathrm{Kernel}(\psi)$ in any quotient in which c has finite order. In particular, for each $n \geq 1$ we have

$$xy^{-1} \in \pi_1 V \langle\langle c^n \rangle\rangle.$$

We use $\langle\langle g \rangle\rangle$ to denote the normal closure of the element g in the group G.

Proof. Let $\widehat{X}$ denote the cover of X corresponding to the normal subgroup $\langle\langle c^n \rangle\rangle$. Let $\widehat{V}$ denote the component of the preimage of V in $\widehat{X}$ containing the basepoint. Consider the decomposition $\widehat{X} \rightarrow \Gamma_{\widehat{X}}$. By Lemma 6.3, $\Gamma_{\widehat{X}}$ is a graph homeomorphic to $\mathbb{R}$. Furthermore $\Gamma_{\widehat{X}}$ can be oriented so that $\widehat{X} \rightarrow \Gamma_{\widehat{X}}$ is orientation preserving on the horizontal edges, and of course, the vertical edges are sent to 0-cells. Now, $g \in \pi_1 V \langle\langle c^n \rangle\rangle$ exactly when the closed path corresponding to g in X lifts to a path ending at $\widehat{V}$ in $\widehat{X}$. Looking at the projection $\varrho : \widehat{X} \rightarrow \Gamma_{\widehat{X}}$, this is equivalent to requiring that the path corresponding to g projects to a closed path in $\Gamma_{\widehat{X}}$. But since ϱ is orientation preserving on $H_{\widehat{X}}$, the elements of $\mathrm{Kernel}(\psi)$ are projected to closed paths and we are done. $\square$

Definition 6.5 (Clean). A CSC is *clean* if all attaching maps of edge spaces are degree 1 covers. This is a special case of the notion of a clean $\mathcal{VH}$-complex which is studied in [Wis02], [Wis06].

Proposition 6.6 (Clean $\Rightarrow$ virtually $F_n \times F_m$). *Let K be a compact CSC that is clean. Then K has a finite cover $\widehat{K} \rightarrow K$ such that $\widehat{K}$ is isomorphic to the cartesian product of two graphs, and so $\pi_1 \widehat{K}$ is isomorphic to the cartesian product of free groups.*

Proof. Choose a basepoint $b \in \Gamma_K^0$. Any closed loop in Γ_K based at b determines an element of $\mathrm{Aut}(\Lambda_b)$. This is because each edge e of Γ_K attached to the vertices u and v, corresponds to an edge space C_α with attaching maps of degree 1, and therefore induces an isomorphism $\Lambda_u \rightarrow \Lambda_v$. Likewise, a closed path based at b in Γ induces an automorphism of Λ_b, which is the composition of isomorphisms. In this way we obtain a homomorphism

$$\pi_1(\Gamma_K) \rightarrow \mathrm{Aut}(\Lambda_b).$$

Composing this with the projection $K \to \Gamma_K$ we obtain (choosing $k \in \Lambda_b$) the homomorphism μ, which is the composition

$$\pi_1(K, k) \to \pi_1(\Gamma_K, b) \to \mathrm{Aut}(\Lambda_b).$$

Since K is compact, Λ_b is a finite graph, so $\mathrm{Aut}(\Lambda_b)$ is finite as well. It follows that the kernel of μ is of finite index in $\pi_1(\widehat{K}, k)$, and consequently, the cover $\widehat{K}$ that it corresponds to is of finite degree. It is easy to see that

$$\widehat{K} \cong \Lambda_b \times \Gamma_{\widehat{K}},$$

since we have just chosen it to be a trivial Λ_b bundle over $\Gamma_{\widehat{K}}$.   □

Proposition 6.7 ($F_n \times F_m \Rightarrow$ clean). *If K is a CSC, and $\pi_1(K) \cong F_n \times F_m$ then K is clean. That is, the attaching maps of the edge spaces of K are all degree* 1.

Sketch of proof. Consider first the case where both n and m are ≥ 2. The key point is that a diagonal element, that is, an element which is not represented by a totally vertical or horizontal closed geodesic path, has $Z \times Z$ as its centralizer. This is true because for a line ℓ that is neither vertical nor horizontal, the non-positive curvature condition determines a unique plane containing ℓ.

From this we conclude that the F_n and F_m factors act vertically and horizontally or vice-versa. And then it is not difficult to show that the complex is clean and in fact the direct product of the two graphs corresponding to F_n and F_m.

In the case where $n \geq 2$ and $m = 1$, using the same idea we can determine that the F_1 subgroup is either vertical or horizontal, and from there it is easy to show that it is clean and in fact a product though the second graph factor might correspond to another F_2 subgroup.

The case of $F_1 \times F_1$ is easily seen to be clean, but here neither factor is necessarily vertical or horizontal.

Various more general results giving product decompositions for non-positively curved spaces can be found in [BH99].   □

Corollary 6.8. *$\pi_1 X$ is isometric to $F_2 \times F_3$ but not commensurable with it.*

Proof. $\pi_1 X$ is isometric to $F_2 \times F_3$ because by Theorem 3.8, $\widetilde{X}$ is isometric to the direct product of a regular tree with valence 4 with a regular tree of valence 6.

To see that $\pi_1 X$ is not commensurable with the direct product of two free groups, first note that as in Proposition 6.7, if some finite cover $\widehat{X}$ of X had $\pi_1(\widehat{X}) \cong F_n \times F_m$, then F_n and F_m would be vertical and horizontal. This means that they would be subgroups of finite index of $\pi_1 V$ and $\pi_1(H)$. But then there would be finite integers q and r so that $c^q \in F_n$ and $y^r \in F_m$ and then $[c^q, y^r] = 1$. This cannot be.

Indeed, according to the "flat torus theorem" for cocompact actions on CAT(0) spaces, maximal free-abelian subgroups act on a flat subspace in the same way as lattices on E^n [BH99]. Therefore, if c and y had powers that commuted, then there would be a periodic plane with a pair of axis corresponding to $\tilde{c}$ and $\tilde{y}$. On the other hand, the non-positive curvature condition implies that the anti-torus Π is the unique plane containing the $\tilde{c}$ and $\tilde{y}$ axes. Consequently, there are no integers $m, n \geq 1$ such that $[c^m, y^n] = 1_{\pi_1 X}$. $\square$

The following conjecture is surely overly ambitious, but I do not know how to construct a counter-example.

Conjecture 6.9 (Virtually clean $\Leftrightarrow$ no anti-torus). If a compact CSC Y has no immersed anti-torus then there is a finite clean cover $\widehat{Y} \to Y$, and so Y has a finite cover that is a product.

Before closing this section, we mention that Corollary 6.8 resolves negatively the following question which was posed in [Tuc90].

Question 6.10. Suppose that K_1 and K_2 are finite simplicial complexes and there is a simplicial complex K which simplicially covers both K_1 and K_2. Must there be a finite simplicial complex which also simplicially covers both K_1 and K_2?

Corollary 6.8 asserts that both X and the direct product $B_2 \times B_3$ of bouquets of circles, have the same universal cover, but they do not have isomorphic finite covers. Note that the above question has an affirmative answer in case K_1 and K_2 are graphs (see [Tuc90].

7. Non-residually finite examples

In this section we form a CSC D by "doubling" the example X of Section 4. The main result of this section is that $\pi_1 D$ is not residually finite.

We will use the following lemma in the proof of Theorem 7.2.

Lemma 7.1 (Invariant subgroup). *If G is finitely generated then the subgroup*

$$N = \bigcap_{[G:H] \leq n} H$$

is of finite index and is also fully invariant.

Proof. To see that it is fully invariant let $\psi\colon G \to G$ be an endomorphism, then if $[G : H] \leq n$ then $[G : \psi^{-1}(H)] \leq n$ and so

$$N = \bigcap_{[G:H]\leq n} H \subseteq \bigcap_{[G:H]\leq n} \psi^{-1}(H) = \psi^{-1}\left(\bigcap_{[G:H]\leq n} H\right) = \psi^{-1}(N). \qquad \square$$

The following theorem and corollary below were observed independently by Long and Niblo in the case of an automorphism. Their motivation was proving the separability of boundary surface subgroups in a 3-manifold group [LN91]. Their proof is slightly different.

Theorem 7.2 (Fixed subgroup of an endomorphism). *Let ϕ be an endomorphism of a finitely generated residually finite group G, and let* Fix_ϕ *denote the subgroup of elements fixed by ϕ. Then G is* Fix_ϕ*-separable.*

Proof. Let $a \in G - \mathrm{Fix}_\phi$. Choose $N \not\ni a^{-1}\phi(a)$ to be normal and ϕ-invariant. The map $\phi\colon G \to G$ induces a map $\bar{\phi}\colon G/N \to G/N$. By choice of N, the fixed subgroup of $\bar{\phi}$ does not contain $\bar{a}$. It does, however, contain the image of Fix_ϕ, and so we are done. $\qquad \square$

Corollary 7.3 (Fixed subgroup of double). *Let $M = G *_{H=\bar{H}} \bar{G}$ be the double of G along H. If M is residually finite then G is H-separable.*

Proof. Let ι be the involution of M which switches G and $\bar{G}$. Observe that $\mathrm{Fix}_\iota = H$. Now apply the previous theorem. $\qquad \square$

Remark 7.4 (Finite order automorphism). Note that if ϕ is an automorphism of finite order, then we do not have to assume that G is finitely generated in Theorem 7.2. This is because given a finite index subgroup $H \subset G$, if ϕ has order r, then we can form a ϕ-invariant finite index subgroup N by taking the intersection $\bigcap_{0 \leq i < r} \phi^i(H)$. And the remainder of the proof of Theorem 7.2 goes through unchanged. So for instance, Corollary 7.3 does not require that G be finitely generated.

Main Theorem 7.5 (The double of X is not residually finite). *Let D denote the CSC which is the double of X along V. That is, letting $\bar{X}$ denote an isomorphic copy of X, let*

$$D = (X \cup \bar{X})/\{V = \bar{V}\}.$$

Then $\pi_1 D$ is not residually finite.

Proof. By Corollary 6.4, $\pi_1 X$ is not $\pi_1 V$-separable. Let $\iota\colon D \to D$ be the involution interchanging X and $\bar{X}$. Applying Corollary 7.3, to $\iota_\star\colon \pi_1 D \to \pi_1 D$ we see that $\pi_1 D$ is not residually finite. In particular, the element $(xy^{-1})^{-1}(\bar{x}\bar{y}^{-1})$ maps to the identity in any finite quotient. $\qquad \square$

We now obtain a slightly stronger result, but the essential argument is the same.

Theorem 7.6 (Torsion quotients). *For each $n \geq 1$, $(xy^{-1})^{-1}(\bar{x}\bar{y}^{-1}) \in \langle\!\langle c^n \rangle\!\rangle$. Thus $(xy^{-1})^{-1}(\bar{x}\bar{y}^{-1})$ maps to 1 in any quotient of $\pi_1 D$ in which the image of c has finite order.*

Proof. Consider the quotient map $f : \pi_1 D \to \pi_1 D/\langle\!\langle c^n \rangle\!\rangle$. It follows from Corollary 6.4 that $f(xy^{-1}) \in f(\pi_1 V)$. By symmetry, $f(\bar{x}\bar{y}^{-1}) = f(xy^{-1})$. Therefore $f\big((xy^{-1})^{-1}(\bar{x}\bar{y}^{-1})\big) = 1$.

We make the "symmetry" mentioned above more precise by using the fixed subgroup proof as in Corollary 7.3. Let $\iota_\star : \pi_1 D \to \pi_1 D$ be induced by the involution of the double D, and let $\ell : \pi_1 D/\langle\!\langle c^n \rangle\!\rangle \to \pi_1 D/\langle\!\langle c^n \rangle\!\rangle$ be the induced involution of $\pi_1 D/\langle\!\langle c^n \rangle\!\rangle$. Then $f(\pi_1 V \langle\!\langle c^n \rangle\!\rangle) \subset \mathrm{Fix}_\ell$ because both $f(\pi_1 \bar{V})$ and $f(\langle\!\langle \bar{c}^n \rangle\!\rangle)$ are contained in Fix_ℓ.

But if $(xy^{-1})^{-1}(\bar{x}\bar{y}^{-1}) \notin \langle\!\langle c^n \rangle\!\rangle$ then $(xy^{-1})\langle\!\langle c^n \rangle\!\rangle \neq (\bar{x}\bar{y}^{-1})\langle\!\langle c^n \rangle\!\rangle$ which means that $f(xy^{-1}) \notin \mathrm{Fix}_\ell$ which contradicts Corollary 6.4. $\qquad\square$

8. A CSC with no finite covers

At this stage, it is not difficult to obtain (by starting with X and using several HNN extensions and amalgamated free products) a non-positively curved square complex which has no finite quotients. In fact, we had implemented this in an early version of this work [Wis]. However, it is worthwhile to get a bit more involved and obtain a CSC with no finite covers, this being somewhat remarkable because it is geometrically very similar to $F_n \times F_m$, and yet algebraically very different.

Guide. The plan in this section is to take a subdivided copy X' of X (see Figure 11), and immerse it in a new CSC W (see Figure 12). We then immerse a subdivided copy W' of W into a CSC A_1 (see Figure 13). The CSC A_1 is a subcomplex of a the CSC A that is illustrated in Figure 9. We then form a CSC M by doubling A along a vertical subgraph. We then amalgamate several copies of M along horizontal subgraphs to obtain our goal E, which is a CSC with no non-trivial connected finite covering space. The reader may wish to look ahead at the complexes X' and W since A_1 and then A is constructed "around" them.

Since our complexes will get rather large, we begin with a gluing lemma to ensure that they are CSCs.

Lemma 8.1 (Gluing CSCs). *Let X and Y be CSCs with a $\mathcal{VH}$-structure, let V_x and V_y be connected components of V_X and V_Y and suppose that there is an isomorphism $\phi : V_x \to V_y$. Then the space $Z = (X \cup Y)/\phi$ obtained by identifying V_x with V_y using ϕ, is itself a CSC with a $\mathcal{VH}$-structure.*

Proof. To see that Z is a CSC, note that we only have to check this locally. The required fact is that if B and C are complete bipartite graphs, whose vertices are divided into the classes, B_0, B_1 and C_0, C_1 respectively, then if there is a bijection $B_0 = C_0$, then the amalgamated union $(B \cup C)/(B_0 \cong C_0)$ is also a complete bipartite graph.

The fact that the resulting complex is a $\mathcal{VH}$-complex follows because the gluing identifies vertical edges, and so the $\mathcal{VH}$-structure of Z is inherited from X and Y. Slightly more general results about constructing new CSCs from old ones by gluing them together along certain gluing maps can be obtained, the details are similar. $\square$

Example 8.2 (Example A). We first describe a CSC called A which admits a local isometry $X'' \to A$, where X'' is the first barycentric subdivision of X. Both the vertical and horizontal decompositions of A correspond to amalgamated free products. We describe A through its vertical decomposition.

Let $\rho_L \colon L \to B$ and $\rho_R \colon R \to B$ denote the covering maps from the left and right labeled graphs to copies of the smaller graph in the center of Figure 9. Let $\tau \colon L \to R$ denote the graph isomorphism between R and L induced by the obvious translation.
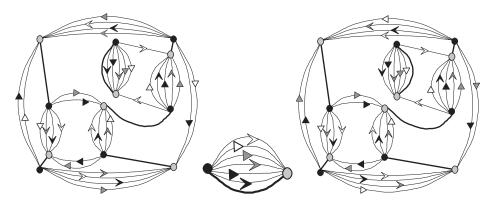

Figure 9. The complex A: Let $\rho_L \colon L \to B$ and $\rho_R \colon R \to B$ be the covering maps of the central graph B by the large graphs on the left and right. The complex A has underlying graph Γ_A consisting of one edge and two vertices. The vertex spaces V_0 and V_1 are each isomorphic copies of B. The single edge space is isomorphic to $L \times I$ and the attaching maps on either side are modelled on ρ_L and ρ_R.

We form a CSC called A as follows: The underlying graph Γ_A is isomorphic to I. Let the vertex graphs corresponding to the vertices of I, denoted by V_0 and V_1, be isomorphic copies of B. Let V_e denote a graph isomorphic to L and thus to R via τ. The attaching maps of the edge space $V_e \times I$ to each of V_0 and V_1 are just the compositions

$$V_e \times \{0\} \to L \to B \cong V_0 \quad \text{and} \quad V_e \times \{1\} \to R \to B \cong V_1.$$

The horizontal decomposition of A. Let H_0 and H_1 be the two components of H_A, where H_0 is the union of the horizontal edges in $V_e \times I$ corresponding to grey vertices of V_e, and H_1 is the union of the horizontal edges of $V_e \times I$ corresponding to black vertices of V_e. Both H_0 and H_1 are bipartite graphs with two vertices and 6 edges. Thus the horizontal decomposition of A is of the form $A \to \Gamma_A^h \cong I$ where H_0 and H_1 correspond to the two vertices of Γ_A^h.

The unique edge space $H_e \times I$ in the horizontal decomposition of A, has attaching maps represented by $\rho_n : H_e \to H_0$ and $\rho_s : H_e \to H_1$, and each of these maps is a degree 7 covering map.

Lemma 8.3 (Normal closure of H_e). *The normal closure of the image of $\pi_1 H_e$ in $\pi_1 H_1$ is all of $\pi_1 H_1$. Consequently, the normal closure of $\pi_1 H_0$ in $\pi_1 A$ is all of $\pi_1 A$.*

Proof. The covering map $H_e \to H_1$ has degree 7. Since the corresponding subgroup has index 7, it follows that the only subgroups containing it are itself and all of $\pi_1 H_1$. We will show that $\pi_1 H_e$ is not a normal subgroup of $\pi_1 H_1$ and thus since the normal closure of $\pi_1 H_e$ in $\pi_1 H_1$ properly contains $\pi_1 H_e$, we see that it must be all of $\pi_1 H_1$.

To see that the subgroup corresponding to $\pi_1 H_e$ is not normal, we will show that the associated covering space is not regular. Indeed there is a closed length 2 path in H_1 which has both a lift of degree 1 and a lift of degree 6 to the covering space $H_e \to H_1$. The two edges of this closed path correspond to the circled black vertices of Figure 10. The two edges of the degree 1 lift of the closed path correspond to the ends of edges marked by a bold **1** in Figure 10. The twelve edges in the degree 6 lift
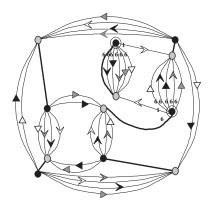

Figure 10. $H_e \to H_1$ is not regular: The graph above represents a cross section (near one end) of the vertical edge space $V_e \times I$. The horizontal graph H_1 corresponds to the black vertices in the graph above. The two circled black vertices determine a closed cycle of length 2 in H_1. It is easy to check that the two ends of edges labeled by a bold **1** determine a degree 1 lift of the closed cycle, and that the twelve ends of edges labeled by bold **6** determine a degree 6 lift of the closed cycle.

of the closed path correspond to the ends of edges marked by a bold **6** in Figure 10.

To see that $\pi_1(A)$ is the normal closure of $\pi_1(H_0)$, note that

$$\pi_1(A) = \pi_1(H_0) *_{\pi_1(H_e)} \pi_1(H_1). \qquad \square$$

Lemma 8.4 (Local isometry $X \to A$). *Let X'' denote the first barycentric subdivision of X. There is a $\mathcal{VH}$-immersion $X'' \to A$. Furthermore, the path xy^{-1} in X projects to a path in $H_0 \subset A$ which is the concatenation of four distinct edges.*

Let X' denote the vertical subdivision of X, and let X'' denote the horizontal subdivision of X', so X'' is the first barycentric subdivision of X.

In order to describe the $\mathcal{VH}$-immersion $X'' \to A$, we first describe a $\mathcal{VH}$-immersion $X' \to W$, where W is a certain CSC described below. Afterwards, we describe a $\mathcal{VH}$-immersion $W' \to A$ where W' denotes the horizontal subdivision of W. The desired $\mathcal{VH}$-immersion $X'' \to A$ is then obtained by taking the composition $X'' \to W' \to A$.

We first subdivide X vertically by subdividing V_X and adding a horizontal edge in the middle of each square of X. We call the resulting complex X'. The vertical decomposition of the complex X' is illustrated in Figure 11. Note that the edges of X' are oriented away from the new vertices of X'.
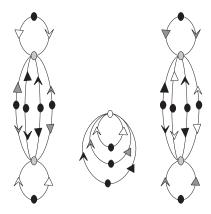

Figure 11. The vertical decomposition of the vertical subdivision X' of X.

The vertical decomposition of W is illustrated in Figure 12. There is an obvious $\mathcal{VH}$-immersion $X' \to W$ which preserves the labels of vertical edges of X' and W. One can think of this immersion as the quotient of X' obtained by identifying certain horizontal edges of X'.

We now form the horizontal subdivision W' of W, and describe an immersion $W' \to A$. To understand this immersion, consider the subcomplex A_1 of A obtained by removing the interior of each vertical bold black edge of A, and removing the
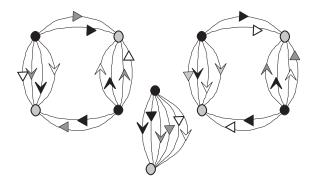
Figure 12. The vertical decomposition of W.

interior of each square having a bold black vertical edge on its boundary. The vertical decomposition of this subcomplex is depicted in Figure 13. Note that there are two edge spaces each of which is attached at both ends to the subgraphs of V_0 and V_1 not containing any bold black edge. There is an obvious $\mathcal{VH}$-immersion $W' \to A_1$ and thus an immersion $X'' \to W' \to A_1 \to A$.
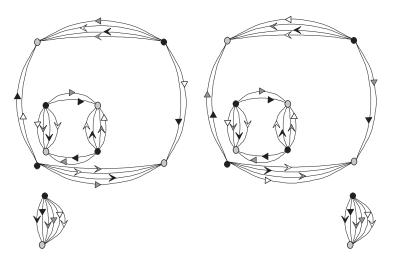

Figure 13. The vertical decomposition of the subcomplex A_1 of A.

Remark 8.5. In order to suppress mentioning the basepoints in the arguments concerning π_1 of the various spaces below (all of which are constructed from copies

of A), whenever π_1 is mentioned we will assume that an edge of $V_0(A)$ has been contracted and that the resulting vertex is the basepoint.

Example 8.6 (The complex M). Let M denote the CSC which is the double of A along V_0. So

$$M = (A \cup \bar{A}) / \{V_0 = \bar{V}_0\}.$$

Note that $\Gamma_M^h \cong I$ and that Γ_M^v is a bouquet of 2 edges. See Figure 14 for a depiction of the 1-skeleton of M. Let $H_0(M)$ denote the component of H_M which equals $H_0 \cup \bar{H}_0$.
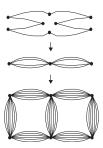

Figure 14. The 1-skeleton of M: The lower graph of the figure above represents the 1-skeleton of M. The subgraph $H_0(M)$ of M^1 is represented by the twelve horizontal edges at the top of M^1. The path ω is represented by an immersion into $H_0(M)$ of the cycle of length 8 illustrated by the top graph of the figure. The immersion $\omega \to H_0(M)$ factors through the subgraph of $H_0(M)$ indicated by the graph in the middle of the illustration, in the obvious way.

Observe that there is a $\mathcal{VH}$-immersion

$$X \bigcup_{V = \bar{V}} \bar{X} \ \to \ A \bigcup_{V_0 = \bar{V}_0} \bar{A} = M.$$

We let ω denote the path in $H_0(M)$ obtained by taking the projection of the path $(xy^{-1})^{-1}(xy^{-1})'$ in the double of X. By Theorem 7.6, the element of $\pi_1 M$ represented by ω is trivial whenever the element c of $\pi_1 M$ has finite order. For later use, note that as in Figure 14, ω is the concatenation of two paths in H_0 and $\bar{H}_0$ each consisting of 4 distinct edges.

Lemma 8.7. $\pi_1 M = \langle\!\langle \pi_1 (H_0(M)) \rangle\!\rangle.$

Proof. $\pi_1 M = \langle \pi_1 A \cup \pi_1 \bar{A} \rangle = \left\langle \langle\!\langle \pi_1 H_0 \rangle\!\rangle_{\pi_1 A} \cup \langle\!\langle \pi_1 \bar{H}_0 \rangle\!\rangle_{\pi_1 \bar{A}} \right\rangle \subset \langle\!\langle \pi_1 (H_0 \cup \bar{H}_0) \rangle\!\rangle_{\pi_1 M} = \langle\!\langle \pi_1 (H_0(M)) \rangle\!\rangle_{\pi_1 M}.$ □

Let $V_0(M)$ denote the component of V_M containing the basepoint 0 of A. In addition, let $0 \in A \subset M$ be chosen as the basepoint of M.

Lemma 8.8 (Generators for $\pi_1(H_0(M))$). *The group $\pi_1(H_0(M))$ is generated by elements represented by closed paths of the form: $ab^{-1}cd^{-1}wx^{-1}yz^{-1}$, where a, b, c, d denote distinct edges of H_0 and w, x, y, z denote distinct edges of $\bar{H}_0$.*

Proof. Observe that $\pi_1(H_0(M))$ is generated by elements represented by closed paths of length 2 beginning at the basepoint. It is therefore sufficient to show that every closed path of length 2 beginning at the basepoint is the product of two length 8 paths of the form indicated above. Note that we use the fact that H_0 (and hence $\bar{H}_0$) has more than 4 edges. By symmetry, it is sufficient to indicate how this is done for one such path of length 2.

$$ab^{-1} = (ae^{-1}cd^{-1}wx^{-1}yz^{-1})(zy^{-1}xw^{-1}dc^{-1}eb^{-1})$$

and

$$wx^{-1} = (wv^{-1}yz^{-1}ab^{-1}cd^{-1})(dc^{-1}ba^{-1}zy^{-1}vx^{-1}). \qquad \square$$

Example 8.9 (The CSC E with no finite covers). We will form a CSC by amalgamating several copies of M along their $H_0(M)$ subspaces.

First let K be a graph which is isomorphic to $H_0(M)$ with basepoint corresponding to the basepoint of $H_0(M)$. Let $\{\sigma_1, \ldots, \sigma_r\}$ denote the full set of length 8 generators of $\pi_1 K$ which are of the form indicated in Lemma 8.8. Choose r copies $\{M_1, \ldots, M_r\}$ of M. For each i let $H_0(M_i)$ denote the subspace of M_i corresponding to $H_0(M)$. For each i let ω_i be the path in M_i corresponding to the path ω in M.

For each i, choose a graph isomorphism $\psi_i: H_0(M_i) \to K$ which sends ω_i to σ_i.

Now let E be the union of the M_i identified along $H_0(M_i)$ using the maps ψ_i. That is,

$$E = K \cup_{\psi_i} M_i.$$

Note that by Lemma 8.8, $\pi_1 K$ is generated by the σ_i and therefore $\pi_1 K$ is generated by the ω_i, because $\sigma_i = \omega_i$ in E. Furthermore, E is a CSC by Lemma 8.1.

Theorem 8.10 (E has no finite covers). *The fundamental group of E has no non-trivial torsion quotient. In particular, $\pi_1 E$ has no non-trivial finite quotient.*

Proof. First observe that $\pi_1 E$ is the normal closure of $\pi_1 K$. This is because $\pi_1 E = \langle \bigcup \pi_1 M_i \rangle$ and by Lemma 8.7 for each i, $\pi_1 M_i = \langle\langle \pi_1 H_0(M_i) \rangle\rangle$.

Thus it is sufficient to show that $\pi_1 K$ is trivialized in every finite quotient. But for each i, the element ω_i is trivial in any finite quotient, and so since $\pi_1 K$ is generated by the ω_i, we see that $\pi_1 K$ is trivial in any finite quotient.

The proof that $\pi_1 E$ has no torsion quotients is almost the same, but we use Theorem 7.6 to show that ω_i is trivialized in any torsion quotient. $\qquad \square$

We close this section with an application of Theorem 7.5.

Theorem 8.11 (Non-residually finite manifold). *There exists a closed non-positively curved manifold whose fundamental group is not residually finite.*

Proof. In [Hu95], it is shown that any compact non-positively curved complex, is a retract of a closed non-positively curved manifold. Applying this to the example D of Theorem 7.5 yields the result. □

9. Non-virtually torsion-free examples

In this section we give some examples of non-virtually torsion-free groups which are either biautomatic or the fundamental group of a compact non-positively curved 2-complex (or both). We also show that two well-known questions concerning the residual finiteness of word-hyperbolic groups are equivalent. We close this section by outlining a possible approach towards finding a non-residually finite, word hyperbolic group.

Lemma 9.1. *Let $g \in G - \{1_G\}$ have the property that g belongs to the kernel of any finite quotient. Suppose that for some n, g is not in the kernel of the quotient $G \to G/\langle\langle g^n \rangle\rangle$. Then $G/\langle\langle g^n \rangle\rangle$ is not virtually torsion-free.*

Proof. We show by contradiction that $\bar{g}$ is contained in every finite index subgroup of $\bar{G} = G/\langle\langle g^n \rangle\rangle$. Suppose $\bar{g} \notin H$ where $[\bar{G} : H] < \infty$, then using the left coset representation, we obtain a finite quotient $\bar{G} \to F$ such that $\bar{g}$ is not sent to the identity in F. But now, consideration of the map $G \to \bar{G} \to F$ shows that the element $\bar{g} \in \bar{G}$ gets sent to the identity. □

The following proposition implies that two well-known questions popularized by Gromov, concerning residual properties of word-hyperbolic groups, are equivalent.

Proposition 9.2. *If all word-hyperbolic groups are virtually torsion-free, then all word-hyperbolic groups are residually finite.*

Note that the converse is true because a given word-hyperbolic group has only finitely many different torsion elements up to conjugacy. One can deduce this easily from the properties of the action on the Rips Complex of the group (see [Aea91]). Thus, if the group is residually finite, then one may separate each of these finitely many conjugacy classes from the identity in a finite quotient, and the kernel of this homomorphism is torsion-free and of finite index.

A related result was proven more recently in [KW00], where we showed that every word-hyperbolic group is residually finite if and only if every word-hyperbolic group has a proper finite index subgroup.

Before giving the proof, we quote a special case of a result of Delzant, Gromov, Ol'shanskĭ. See [Del96], [Gro87], [Ol'93].

Theorem 9.3. *Let G be word-hyperbolic. For each infinite order element $g \in G$, there exists $n > 1$ such that $G/\langle\langle g^n \rangle\rangle$ is itself word-hyperbolic and $\bar{g} \neq 1$.*

Proof of Proposition 9.2. What we actually do is show that if one had a non-residually finite word-hyperbolic group, then one could construct a non-virtually torsion-free word-hyperbolic group. So suppose G is word hyperbolic, and suppose that $g \in G$ does not survive in any finite quotient of G. Then if g is of finite order, then G itself is not virtually torsion-free. If g is of infinite order, then we apply Theorem 9.3 to form a quotient $G \to G/\langle\langle g^n \rangle\rangle = \bar{G}$ such that $\bar{G}$ is word hyperbolic, and such that $\bar{g}$ is non-trivial in $\bar{G}$. But now Lemma 9.1 shows that $\bar{G}$ is not virtually torsion-free.

Avoiding Theorem 9.3, we may argue as follows: Given a group G with an element $g \in G$ that maps to the identity in any finite quotient, we may construct a new group $K = \langle G, t \mid [g, t]^n \rangle$ which is easily seen to be word hyperbolic for $n \geq 2$. Now $[g, t]$ has order n, but $[g, t]$ is contained in every proper normal finite index subgroup since g does. Thus K is not virtually torsion-free, since if it had a torsion-free finite index subgroup it would have a torsion-free normal finite index subgroup. □

The method used above to construct a non-virtually torsion-free group from a non-residually finite group is a bit trickier to implement in the case of automatic groups or CAT(0) groups which are not word hyperbolic. One difficulty is that a version of Theorem 9.3 does not exist in this case.

Define the *turning angle* of a geodesic path γ at a point $p \in \gamma$ to be the distance between α and β in $\text{Link}(p)$, where α and β are the points in $\text{Link}(p)$ corresponding to the incoming and outgoing rays of γ at p. Note that $\text{Link}(p)$ has the path metric locally isometric to the "angle metric" but we do not cap the distance by π.

Example 9.4. Here we give an example of an automatic group G and an element $b \in G$ such that there are no automatic $G/\langle\langle b^n \rangle\rangle$ quotients. Consider the group G, presented by:
$$\langle a, b, t \mid [a, b], a^t = a^2 b \rangle.$$

It is isomorphic to
$$\langle a, \beta, t \mid [a, \beta], a^t = \beta \rangle,$$

which is automatic since the standard 2-complex of this presentation is a non-positively curved square complex (see [GS90]).

However, for all $n > 0$ the quotient $Q = G/\langle\langle b^n \rangle\rangle$ is not automatic. One can prove this by first noting that the subgroup generated by $\langle a^n, t \rangle$ is isomorphic to the Baumslag–Solitar group presented by $\langle \alpha, t \mid \alpha^t = \alpha^2 \rangle$. Then by using standard methods, one deduces that Q has an exponential isoperimetric function and is therefore not automatic (see [ECH$^+$92]).

Lemma 9.5 (Modding out by n-th powers for CAT(0) groups). *Suppose that X is a non-positively curved 2-complex, and that $\pi_1 X$ has an element g whose conjugacy class is represented by a closed geodesic path γ which has a point with a large turning angle $\theta > \pi$. Then for large n, the element g survives in the quotient $\pi_1 X \rightarrow \pi_1 X / \langle\!\langle g^n \rangle\!\rangle$, and $\pi_1 X / \langle\!\langle g^n \rangle\!\rangle$ is the fundamental group of a 2-dimensional non-positively curved orbihedron.*

Proof. We choose n so that for a regular Euclidean polygon P with n sides, and thus interior angle $\frac{(n-2)\pi}{n}$, we would have $\frac{(n-2)\pi}{n} + \theta \geq 2\pi$. In other words, we choose $n \geq \frac{2\pi}{\theta - \pi}$. We first subdivide X so that γ is contained in the 1-skeleton. Now we attach P along γ^n to X, so that each vertex of P corresponds to the point of γ with large turning angle. The group $\mathbb{Z}_n$ acts on P in the obvious way, and so since the attaching map of P is $\mathbb{Z}_n$-equivariant we can form the non-positively curved orbihedron $Y = X \cup P / \mathbb{Z}_n$. To see that Y is a non-positively curved orbihedron, note that the link condition is still satisfied at the vertex of X where γ has the large turning angle. Furthermore, the interior of P, and in particular the cone point, is modelled on $P / \mathbb{Z}_n$. The fundamental group of $X \cup P / \mathbb{Z}_n$, is isomorphic to $\pi_1 X / \langle\!\langle \gamma \rangle\!\rangle$. (See [Sta91], [Hae91], [Cor92], [BH99] for information about non-positively curved orbihedra) $\qquad\square$

In order to use Lemma 9.5 to understand the non-virtually torsion-free group of Example 9.7, we will need the following lemma.

A group G is *Hopfian* if every surjective endomorphism $G \rightarrow G$ is an automorphism. It is a theorem of Malćev [LS77], that a finitely generated residually finite group is Hopfian. We give a version of this below.

Lemma 9.6 (Non-Hopf element). *Let G be a finitely generated group, and let $\phi \colon G \rightarrow G$ be a surjective endomorphism. Suppose that $g \in \mathrm{Kernel}(\phi)$, then g is contained in every finite index subgroup of G.*

Proof. Suppose that g where contained outside a finite index subgroup $H \subset G$, then by Lemma 7.1 there would be a ϕ-invariant finite index normal subgroup N such that $g \notin N$. But then if we consider the diagram below,

$$
\begin{array}{ccc}
G & \xrightarrow{\ \phi\ } & G \\
\downarrow & & \downarrow \\
G/N & \xrightarrow[\ \bar{\phi}\ ]{} & G/N
\end{array}
$$

since $\bar{\phi}$ is a surjective endomorphism of finite groups, we see that $\bar{\phi}$ is an automorphism, and so $\bar{g}$ is not mapped to the identity by $\bar{\phi}$. But since the diagram is commutative, and we assumed that $\phi(g) = 1_G$, this is a contradiction. $\qquad\square$

Here is an example which is obtained using Lemma 9.5.

Example 9.7 (A CAT(0) example). A relatively simple example of a group which acts properly discontinuously and cocompactly on a CAT(0) space, but is not virtually torsion-free, may be constructed by adapting the following non-Hopfian example of [Wis96a]:

$$G = \langle a, b, s, t \mid [a, b] = 1, \ a^s = (ab)^2, \ b^t = (ab)^2 \rangle.$$

One can see that G is non-Hopfian by considering the endomorphism ψ of G induced by $a \mapsto a^2$, $b \mapsto b^2$, $s \mapsto s$, $t \mapsto t$. A non-trivial element of G which is in the kernel of ψ is $\gamma = [(ab)^{s^{-1}}, (ab)^{t^{-1}}]$.

Following [Wis96a], the presentation for G given above is equivalent to the following presentation by using a couple of Tietze transformations:

$$G = \langle a, b, c, s, t \mid c = ab, c = ba, \ a^s = c^2, \ b^t = c^2 \rangle.$$

Now to see that G is π_1 of a non-positively curved 2-complex C, we just use the standard 2-complex of the second presentation given above. We metrize the edges a, b, s, and t as length 1 intervals, and we metrize the edge c as length $\frac{1}{2}$. The cells of the complex C are metrized as Euclidean isosceles triangles and rectangles as in Figure 15 below.
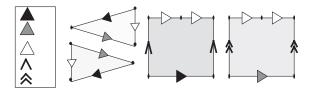

Figure 15. *C* has a non-positively curved metric.

Consider the commutator $\gamma = [(ab)^{s^{-1}}, (ab)^{t^{-1}}]$. Substituting c for ab, we can rewrite γ as $\gamma = [c^{s^{-1}}, c^{t^{-1}}]$ and thus think of γ as a path in C. Figure 16 illustrates that γ is homotopic to a closed geodesic, which we will call γ'. To see that γ' is a local geodesic, observe that the turning angle of γ' at each point that is sent to the 0-cell of C is either $\pi + 2\arcsin(1/4) - 2\arctan(1/8) > 3.39 > \pi$ or $\pi + 2\arcsin(1/4) \sim 3.646 > \pi$.

Incidentally, note that the lengths of the edges s and t can be varied independently and the metrics of the corresponding rectangles can be varied accordingly, so that the resulting metric on C is still non-positively curved. As we change the metric on C by increasing the lengths of s and t and the associated rectangles, each of the four turning angles of γ' approaches $\pi + 2\arcsin(1/4) \sim 3.646$.
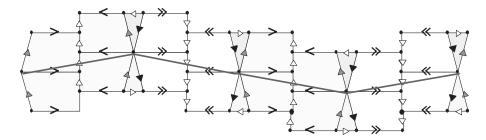
Figure 16. The geodesic γ' homotopic to γ.

By Lemma 9.6, we see that γ is contained in every finite index subgroup of G. By Lemma 9.5, we see that the quotient $G/\langle\!\langle \gamma^{13} \rangle\!\rangle$ is the fundamental group of a non-positively curved 2-complex of groups.

Example 9.8. Let $n \geq 2$ and let G be as in Example 9.7, and form the group $J = \langle G, t \mid [\gamma, t]^n \rangle$. Then J is the fundamental group of a non-positively curved orbihedron, and J is not virtually torsion-free.

Proof. The element $[\gamma, t]$ is contained in every finite index subgroup because γ is. Now we vary the construction of Lemma 9.5 slightly to obtain the result. First we add a loop t at the basepoint to C such that t has the same length as γ'. Next we let P be a regular polygon with $4n$ sides, and we attach P to $C \cup t$ along $[\gamma', t]^n$ in the obvious way. It is easy to see that $C \cup t \cup (P/\mathbb{Z}_{4n})$ is non-positively curved. The resulting orbihedron has $\pi_1 = J$. Since $[\gamma, t]$ survives in the quotient (it has order n) we see by Lemma 9.1 that J is not virtually torsion-free. $\square$

The difficulty with the preceding examples is that there are not many general results except for certain small cancellation conditions, providing automaticity for groups acting on CAT(0) 2-complexes. Therefore we construct the following example.

Theorem 9.9 (A non-virtually torsion-free biautomatic group). *There is a group G such that*

(1) *G is not virtually torsion-free;*

(2) *G is a $C(4)$–$T(4)$ group;*

(3) *G acts cocompactly and properly discontinuously on a CAT(0) square complex;*

(4) *G is biautomatic.*

Proof. Let D be the CSC of Theorem 7.5. Let γ denote the closed horizontal geodesic path $(xy^{-1})^{-1}(\bar{x}\bar{y}^{-1})$ in the 1-skeleton of D which represents an element of $\pi_1 D$ that is trivial in every finite quotient. We will use γ to refer to the word $(xy^{-1})^{-1}(\bar{x}\bar{y}^{-1})$.

We form the group $G = (\pi_1(D) * t)/\langle\langle[\gamma, t]^n\rangle\rangle$. Combining Lemma 9.5 with Theorem 7.5, we see that G is not virtually torsion-free.

To see that G is a $C(4)$–$T(4)$ group, note that G is π_1 of the complex $L = D \cup t \cup P$ where P is a polygon with $10n$ sides, that is attached to $D \cup t$ along the combinatorial path $[\gamma, t]^n$. It is easily verified that L satisfies the $C(4)$–$T(4)$ small cancellation condition.

To see that G acts cocompactly on a CAT(0) square complex we first represent G as the fundamental group of a non-positively curved orbihedron, by taking the quotient $L \to L/\mathbb{Z}_n$ where $\mathbb{Z}_n$ acts on P as in Lemma 9.5. We now subdivide $L/\mathbb{Z}_n$ so that it is a square complex. We subdivide the t edge of $L/\mathbb{Z}_n$ so that it is of length 2, and we subdivide $P/\mathbb{Z}_n$ so that it is isomorphic to the direct product $I_2 \times I_4$. The cone point of order n of $P/\mathbb{Z}_n$ is therefore at the vertex at its center. The resulting subdivision L' of L, is a non-positively curved square orbihedron, and therefore its universal cover is a CAT(0) square complex.

That G is biautomatic follows either by [GS90] because G is a $C(4)$–$T(4)$ group, or by the results of [NR98] because G acts properly discontinuously and cocompactly on a CAT(0) square complex.   □

Remark 9.10 (Word-hyperbolic suggestion). One possible way of using the method of this paper to work towards finding a non-residually finite word-hyperbolic group is as follows:

(1) Find a word hyperbolic group G and a malnormal subgroup M such that G is not M-separable.

(2) Prove that M is quasiconvex.

(3) Double G along M to obtain a non-residually finite word hyperbolic group as in Corollary 7.3. See [BF95].

Part (3) of the plan is already clear.

Part (2) should not be difficult. In fact, there is currently no known example of a finitely generated malnormal subgroup of a word-hyperbolic group that is not quasiconvex. In fact, if $M \subset G$ were a f.g. malnormal subgroup of a word-hyperbolic group that is not quasiconvex, then M would be distorted, and thus the double $G *_M G$ would be a nice example of a non-word hyperbolic group with no Baumslag–Solitar subgroups.

Part (1) is where the difficulty lies in this approach. To begin with, I know of no (finitely presented) residually finite group G which is not M-separable (for M a f.g. malnormal subgroup). I therefore pose this as a problem. Examples of word-hyperbolic groups which are not subgroup separable with respect to an arbitrary subgroup are easy to construct, yet in all examples that I am aware of, the subgroup is very far from being malnormal.

Problem 9.11. Is there a (finitely presented) residually finite group G and a f.g. malnormal subgroup M, such that G is not M-separable?

Perhaps Problem 9.11 is as difficult as the question of residual finiteness for word-hyperbolic groups. In both cases, the difficulties appear to stem from an avoidance of 'long cylinders'. I hope that some progress on Problem 9.11 will yield some insight. Of course the main goal would be to find a word hyperbolic solution to Problem 9.11.

10. A collection of problems about CSCs

We close with a set of problems about CSCs. The CSC territory has proven to be a natural testing ground for the possibility of generalizing various properties of word-hyperbolic groups to biautomatic or CAT(0) groups. Unless stated otherwise, all CSCs below are assumed to be compact. Most of the problems center around algorithmic questions. Two of the problems in [Wis96b] were solved before anybody knew I had asked them: Namely, it has been shown that CSCs are quasi-isometrically rigid, (proven in [Ahl02] using results from [KL97] and [MSW03]), and that there are CSCs with infinite simple fundamental groups [BM97]. In the second case, we can still ask for:

Problem 10.1. Give an elementary proof and construction of CSCs with infinite simple fundamental group.

Problem 10.2. Does there exist an algorithm which determines whether two CSCs have isomorphic finite covers? Does there exist an algorithm which determines whether two CSCs have commensurable fundamental groups?

Problem 10.3. Is the isomorphism problem solvable for fundamental groups of CSCs?

A notable special case of Problem 10.2 is:

Problem 10.4. Is there an algorithm which determines whether a CSC is virtually clean?

Problem 10.5. Is there an algorithm which determines whether a CSC has residually finite fundamental group? Linear fundamental group?

Problem 10.6. Is there an algorithm that determines whether a CSC contains an immersed anti-torus?

Problem 10.7. Given a CSC X, is there an algorithm which takes as input vertical and horizontal paths v and h, and determines whether or not $\tilde{v}$ and $\tilde{h}$ are the axes of an anti-torus in $\widetilde{X}$?

Problem 10.8. Given vertical and horizontal paths v and h, is the subgroup $\langle v^n, h^n \rangle$ isomorphic to either F_2 or $\mathbb{Z}^2$ for sufficiently large n? In particular, suppose that $\tilde{v}$ and $\tilde{h}$ are the axes of an anti-torus. Is it possible for $\langle v^n, h^n \rangle \cong F_2$ for some n?

Some very interesting partial results on the above three problems where obtained by Rattaggi in [Rat05].

Every CSC that is not a graph, contains a $\mathbb{Z}^2$ subgroup in its fundamental group. In fact, such subgroups correspond to period flats, and these periodic flats are dense in the set of flats [Wis05], [Wis96b]. Call a CSC *elementary* if it is virtually the product of two graphs one of which has euler characteristic ≥ 0. (So the fundamental group is either F_n or $F_n \times \mathbb{Z}$.)

Problem 10.9. If the CSC X is non-elementary, does $\pi_1 X$ contain an $F_2 \times F_2$ subgroup? Is there an algorithm to determine if $F_2 \times F_2$ embeds in $\pi_1 X$? For CSCs X and Y, is there an algorithm to determine if there is an immersion $Y \hookrightarrow X$? To determine if $\pi_1 Y$ embeds as a subgroup of $\pi_1 X$?

The first part of the above problem has been answered by Rattaggi–Robertson in [RR05]. They produced a marvellous CSC R such that the only non-trivial product embedded in $\pi_1 R$ is isomorphic to $\mathbb{Z}^2$.

A group is *coherent* if all its finitely generated subgroups are finitely presented. A well-known group that is not coherent is $F_2 \times F_2$.

Problem 10.10. Let X be a non-elementary CSC. Is $\pi_1 X$ incoherent?

A partial result towards Problem 10.10 is made in [BW99] where it is shown that if X contains no v and h such that $\langle v, h \rangle \cong F_2$, then every finitely presented subgroup of $\pi_1 X$ is either free or equals $\pi_1 Y$ where Y is a compact CSC and $Y \hookrightarrow X$ is an immersion.

Problem 10.11. Compute $\mathrm{Out}(\pi_1 X)$ for a CSC X.

Problem 10.12. Does every non-elementary CSC have a proper (infinite index?) finitely generated subgroup that is not isomorphic to either F_n or $\mathbb{Z}^2$?

On the other hand we can ask specifically:

Problem 10.13. Is there a CSC which contains a closed surface (quasiconvex?) subgroup of genus ≥ 2?

Note that surface groups do embed in non-compact CSCs. Indeed, they act freely on the product of two trees as does the fundamental group of any non-positively curved directed $\mathcal{V}\mathcal{H}$-complex [Wis06].

Another striking property of $F_2 \times F_2$ is that it has finitely generated subgroups with undecidable membership problem.

Problem 10.14. Is there a non-elementary CSC X such that $\pi_1 X$ has decidable membership problem for each finitely generated subgroup?

Problem 10.15. Is $\pi_1 X$ hopfian for every CSC X?

Problem 10.16. Describe CSCs with 'interesting' self-immersions $X \looparrowright X$. This implies that $\pi_1 X$ is not cohopfian and can also lead to interesting 'self-similar' constructions of anti-tori.

I have proven that every directed non-positively curved $\mathcal{V}\mathcal{H}$-complex Y is a subcomplex of a CSC X. This motivates:

Problem 10.17. Let Y be a compact non-positively curved $\mathcal{V}\mathcal{H}$-complex. Is there always a CSC X such that $\pi_1 Y$ embeds as a subgroup in $\pi_1 X$?

Problem 10.18. Does every CSC immerse π_1-injectively (quasiconvexly?) in a closed non-positively curved cubulated manifold. How many dimensions are needed?

I suspect the following problem about the main example in this paper has a negative solution, but have been unable to solve it:

Problem 10.19. Let X be the six square CSC presented in Example 4.1. Is $\pi_1 X$ residually finite?

References

[Ahl02]   A. Reiter Ahlin, The large scale geometry of products of trees. *Geom. Dedicata* **92** (2002), 179–184. Zbl 1009.20033 MR 1934017

[Aea91]   J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, H. Short, Notes on word hyperbolic groups. In *Group theory from a geometrical viewpoint* (Trieste, 1990), ed. by É. Ghys, A. Haefliger, and A. Verjovsky, World Sci. Publishing, River Edge, N.J., 1991, 3–63. Zbl 0849.20023 MR 1170363

[BF95]   Mladen Bestvina and Mark Feighn, Stable actions of groups on real trees. *Invent. Math.* **121** (2) (1995), 287–321. Zbl 0837.20047 MR 1346208

[BH99]   Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, Springer-Verlag, Berlin 1999. Zbl 0988.53001 MR 1744486

[BW99]   Martin R. Bridson and Daniel T. Wise, $\mathcal{VH}$ complexes, towers and subgroups of $F \times F$. *Math. Proc. Cambridge Philos. Soc.* **126** (3) (1999), 481–497. Zbl 0942.20009 MR 1684244

[BM97]   Marc Burger and Shahar Mozes, Finitely presented simple groups and products of trees. *C. R. Acad. Sci. Paris Sér. I Math.* **324** (7) (1997), 747–752. Zbl 0966.20013 MR 1446574

[BM00]   Marc Burger and Shahar Mozes, Lattices in product of trees. *Inst. Hautes Études Sci. Publ. Math.* **92** (2000), 151–194 Zbl 1007.22013 MR 1839489

[Cor92]   Jon Michael Corson, Complexes of groups. *Proc. London Math. Soc.* (3) **65** (1) (1992), 199–224. Zbl 0792.57004 MR 1162493

[Del96]   Thomas Delzant, Sous-groupes distingués et quotients des groupes hyperboliques. *Duke Math. J.* **83** (3) (1996), 661–682. Zbl 0852.20032 MR 1390660

[ECH$^+$92] David B. A. Epstein, James W. Cannon, Derek F. Holt, Silvio V. F. Levy, Michael S. Paterson, and William P. Thurston, *Word processing in groups*. Jones and Bartlett Publishers, Boston, MA, 1992. Zbl 0764.20017 MR 1161694

[Ger87]   S. M. Gersten, Reducible diagrams and equations over groups. In *Essays in group theory*, Springer-Verlag, New York, Berlin 1987, 15–73. Zbl 0644.20024 MR 0919828

[Ger92]   S. M. Gersten, Problems on automatic groups. In *Algorithms and classification in combinatorial group theory* (Berkeley, CA, 1989), Springer-Verlag, New York 1992, 225–232. Zbl 0781.20022 MR 1230636

[GS90]   S. M. Gersten and H. B. Short, Small cancellation theory and automatic groups. *Invent. Math.* **102** (2) (1990), 305–334. Zbl 0714.20016 MR 1074477

[GS91]   S. M. Gersten and H. Short, Small cancellation theory and automatic groups. II. *Invent. Math.* **105** (3) (1991), 641–662. Zbl 0734.20014 MR 1117155

[Gro87]   M. Gromov, Hyperbolic groups. In *Essays in group theory*, Math. Sci. Res. Inst. Publ. 8, Springer-Verlag, New York 1987, 75–263. Zbl 0634.20015 MR 0919829

[Hae91]   André Haefliger, Complexes of groups and orbihedra. In *Group theory from a geometrical viewpoint* (Trieste, 1990), ed. by É. Ghys, A. Haefliger, and A. Verjovsky, World Sci. Publishing, River Edge, N.J., 1991, 504–540. Zbl 0858.57013 MR 1170375

[HW98]   Tim Hsu and Daniel T. Wise, Embedding theorems for non-positively curved polygons of finite groups. *J. Pure Appl. Algebra* **123** (1–3) (1998), 201–221. Zbl 0892.20026 MR 1492901

[HW99]   Tim Hsu and Dani Wise, A non-residually finite square of finite groups. In *Groups St. Andrews 1997 in Bath, I*, Cambridge University Press, Cambridge 1999, 368–378. Zbl 0940.20028 MR 1676633

[HW04]   Tim Hsu and Daniel T. Wise, Groups with infinitely many types of fixed subgroups. *Israel J. Math.* **144** (2004), 93–107. Zbl 02147123 MR 2121535

[Hu95]   B. Hu, Retractions of closed manifolds with nonpositive curvature. In *Geometric group theory* (Columbus, OH, 1992), Ohio State Univ. Math. Res. Inst. Publ. 3, Walter de Gruyter, Berlin, 1995, 135–147. Zbl 0842.57003 MR 1355108

[KW00]   Ilya Kapovich and Daniel T. Wise, The equivalence of some residual properties of word-hyperbolic groups. *J. Algebra* **223** (2) (2000), 562–583. Zbl 0951.20029 MR 1735163

[KL97]   Bruce Kleiner and Bernhard Leeb, Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings. *Inst. Hautes Études Sci. Publ. Math.* **86** (1997), 115–197. Zbl 0910.53035 MR 1608566

[LN91]   D. D. Long and G. A. Niblo, Subgroup separability and 3-manifold groups. *Math. Z.* **207** (2) (1991), 209–215. Zbl 0711.57002 MR 1109662

[LS77]   Roger C. Lyndon and Paul E. Schupp, *Combinatorial group theory*. Ergeb. Math. Grenzgeb. 89, Springer-Verlag, Berlin 1977. Zbl 0368.20023 MR 0577064

[Mos95]   Lee Mosher, Geometry of cubulated 3-manifolds. *Topology* **34** (4) (1995), 789–814. Zbl 0869.57015 MR 1362788

[MSW03]   Lee Mosher, Michah Sageev, and Kevin Whyte, Quasi-actions on trees. I. Bounded valence. *Ann. of Math.* (2) **158** (1) (2003), 115–164. Zbl 1038.20016 MR 1998479

[NR98]   G. A. Niblo and L. D. Reeves, The geometry of cube complexes and the complexity of their fundamental groups. *Topology* **37** (3) (1998), 621–633. Zbl 0911.57002 MR 1604899

[Ol'93]   A. Yu. Ol'shanskiĭ, On residualing homomorphisms and *G*-subgroups of hyperbolic groups. *Internat. J. Algebra Comput.* **3** (4) (1993), 365–409. Zbl 0830.20053 MR 1250244

[Pri89]   Stephen J. Pride, Some problems in combinatorial group theory. In *Groups—Korea 1988* (Pusan, 1988), Lecture Notes in Math. 1398, Springer-Verlag, Berlin 1989, 146–155. Zbl 0685.20032 MR 1032822

[Rat05]   Diego Rattaggi, Anti-tori in square complex groups. *Geom. Dedicata* **114** (2005), 189–207. Zbl 02227563 MR 2174099

[Rat07]   Diego Rattaggi, Three amalgams with remarkable normal subgroup structures. *J. Pure Appl. Algebra* **210** (2) (2007), 537–541. Zbl 05167562

[RR05]   Diego Rattaggi and Guyan Robertson, Abelian subgroup structure of square complex groups and arithmetic of quaternions. *J. Algebra* **286** (1) (2005), 57–68. Zbl 1077.22017 MR 2124808

[Sch73]   Paul E. Schupp, A survey of small cancellation theory. In *Decision problems and the Burnside problem in group theory*, Stud Logic Found. Math. 71, North-Holland, Amsterdam 1973, 569–589. Zbl 0292.20034 MR 0412289

[SW79]   Peter Scott and Terry Wall, Topological methods in group theory. In *Homological group theory* (Proc. Sympos., Durham, 1977), London Math. Soc. Lecture Note Ser. 36, Cambridge University Press, Cambridge 1979, 137–203. Zbl 0423.20023 MR 0564422

[Sta83]   John R. Stallings, Topology of finite graphs. *Invent. Math.* **71** (3) (1983), 551–565. Zbl 0521.20013 MR 0695906

[Sta91]   John R. Stallings, Non-positively curved triangles of groups. In *Group theory from a geometrical viewpoint* (Trieste, 1990), ed. by É. Ghys, A. Haefliger, and A. Verjovsky, World Sci. Publishing, River Edge, N.J., 1991, 491–503. Zbl 0843.20033 MR 1170374

[Tuc90]   Thomas W. Tucker, Some topological graph theory for topologists: a sampler of covering space constructions. In *Topology and combinatorial group theory* (Hanover, NH, 1986/1987; Enfield, NH, 1988), Lecture Notes in Math. 1440, Springer-Verlag, Berlin 1990, 192–207. Zbl 0713.57001 MR 1082993

[Wal79]   C. T. C. Wall (ed.), List of problems. In *Homological group theory* (Proc. Sympos., Durham, 1977), Cambridge University Press, Cambridge 1979, 369–394. Zbl 0409.00004 MR 0564417

[Wis96a]   Daniel T. Wise, A non-Hopfian automatic group. *J. Algebra* **180** (3) (1996), 845–847. Zbl 0848.20028 MR 1379212

[Wis96b]   Daniel T. Wise, Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups. PhD thesis, Princeton University, 1996.

[Wis02]   Daniel T. Wise, The residual finiteness of negatively curved polygons of finite groups. *Invent. Math.* **149** (3) (2002), 579–617. Zbl 1040.20024 MR 1923477

[Wis05]   Daniel T. Wise, Approximating flats by periodic flats in CAT(0) square complexes. *Canad. J. Math.* **57** (2) (2005), 416–448. Zbl 02182628 MR 2124924

[Wis06]   Daniel T. Wise, Subgroup separability of the figure 8 knot group. *Topology* **45** (3) (2006), 421–463. Zbl 1097.20030 MR 2218750

[Wis]   Daniel T. Wise, A non-positively curved squared complex with no finite covers. Preprint, 1995.

Daniel T. Wise, Department of Mathematics & Statistics, McGill University, 805 Sherbrooke Street West, Montreal, Quebec, Canada H3A 2K6

E-mail: wise@math.mcgill.ca