ISRAEL JOURNAL OF MATHEMATICS **144** (2004), 93–107

GROUPS WITH INFINITELY MANY TYPES OF FIXED SUBGROUPS

BY

Tim Hsu

Department of Mathematics, San Jose State University
San Jose, CA 95192-0103, USA
e-mail: hsu@math.sjsu.edu

AND

Daniel T. Wise*

Department of Mathematics & Statistics, McGill University
Montreal, Quebec, H3A 2K6, Canada
e-mail: wise@math.mcgill.ca

ABSTRACT

It is a theorem of Shor that if G is a word-hyperbolic group, then up to isomorphism, only finitely many groups appear as fixed subgroups of automorphisms of G. We give an example of a group G acting freely and cocompactly on a CAT(0) square complex such that infinitely many non-isomorphic groups appear as fixed subgroups of automorphisms of G. Consequently, Shor's finiteness result does not hold if the negative curvature condition is relaxed to either biautomaticity or nonpositive curvature.

1. Introduction

Let $\phi: G \to G$ be an automorphism of a group. The **fixed subgroup of** ϕ, denoted by $\mathrm{Fix}(\phi)$, is defined as follows:

$$\mathrm{Fix}(\phi) = \{g \in G: \phi(g) = g\}. \tag{1}$$

For example, if there exists $x \in G$ such that $\phi(g) = xgx^{-1}$ (i.e., if ϕ is inner), then $\mathrm{Fix}(\phi) = C_G(x)$, the centralizer of x in G.

* D. T. Wise was supported by grants from FCAR and NSERC.
Received August 13, 2003

Finitely generated free groups have the remarkable property that their fixed subgroups are finitely generated. Specifically:

THEOREM 1.1: *If ϕ: $F_r \to F_r$ is an automorphism of a free group of rank r, then* $\text{rank}(\text{Fix}(\phi)) \leq r$.

Theorem 1.1 began as a conjecture attributed to P. Scott, and was initially proven in various geometric cases. Gersten [Ger87] was the first to prove that $\text{Fix}(\phi)$ is always finitely generated, and Bestvina and Handel [BH92] were the first to establish that $\text{rank}(\text{Fix}(\phi)) \leq r$. See [Ven02] for a survey of this area, which has maintained continual activity.

Note that the finite rank hypothesis is crucial in Theorem 1.1. Indeed, for each n, the infinite rank free group F_∞ has an automorphism ϕ_n such that $\text{Fix}(\phi_n) \cong F_n$.

Another example worth noting is the group

$$F_2 \times \mathbb{Z} \cong \langle a, b, t \mid [a, t], [b, t] \rangle, \tag{2}$$

as the automorphism ϕ induced by $a \mapsto at, b \mapsto bt, t \mapsto t$ has $\text{Fix}(\phi) \cong F_\infty \times \mathbb{Z}$. Indeed, it is not hard to see that $\text{Fix}(\phi)$ is the subgroup of $F_2 \times \mathbb{Z}$ consisting of those elements with zero exponent sum in a and b. Nevertheless, $F_n \times \mathbb{Z}$ has only finitely many isomorphism types of fixed subgroups; more precisely, we leave it as an exercise to show that if ϕ is an automorphism of $F_n \times \mathbb{Z}$, then either $\phi(t) = t$, in which case $\text{Fix}(\phi)$ is isomorphic to $F_r \times \mathbb{Z}$ with $r \leq n$ or $r = \infty$, or $\phi(t) = t^{-1}$, in which case $\text{Fix}(\phi)$ is free of rank $r \leq 2n - 1$.

Recently, Shore [Sho99] gave the following interesting generalization of Theorem 1.1:

THEOREM 1.2: *Let G be a word-hyperbolic group. Then up to isomorphism, only finitely many groups appear as fixed subgroups of automorphisms of G.*

In this paper, we give examples of finitely generated groups that have infinitely many non-isomorphic fixed subgroups. More precisely, we show that:

THEOREM 1.3: *There exists a group G that acts freely and cocompactly on the cartesian product of two trees such that infinitely many non-isomorphic groups appear as fixed subgroups of automorphisms of G. In particular, there exist centralizers in G whose abelianizations are free abelian of arbitrarily high finite rank.*

We note that G is the fundamental group of a compact nonpositively curved square complex, which means that G is both a CAT(0) group and a $C(4)$-

$T(4)$ group, and thus biautomatic [GS91]. Therefore, Theorem 1.3 demonstrates a contrast between the behavior of word-hyperbolic groups and their semi-hyperbolic generalizations.

We begin the construction of G by reviewing **complete square complexes** in Section 2. In Sections 3 and 4, we describe the structure of centralizers in the fundamental group of a complete square complex, and in Section 5, we specialize that discussion to produce the desired example.

Notation 1.4: If H is a subgroup of G and $x \in G$, then the centralizer of x in H is denoted by $C_H(x)$. The length of a reduced word w is denoted by $|w|$. The abelianization G/G' of a group G is denoted by G_{ab}.

2. Complete square complexes

Recall that a **square complex** is a combinatorial 2-complex whose 2-cells are squares. We are mainly interested in the following type of square complex, introduced in [Wis96].

Definition 2.1: A **complete square complex**, or **CSC**, is a square complex X such that:

1. The 1-cells of X are partitioned into two classes, X_V and X_H, that induce the structure of a complete bipartite graph on the link of any vertex of X. We think of X_V as the **vertical** 1-cells and X_H as the **horizontal** 1-cells.
2. The squares of X are oriented as shown in Figure 1, with vertical cells oriented "up" and horizontal cells oriented "right", and the 1-cells of X are oriented in a compatible manner.

By a slight abuse of terminology, we will use X_V and X_H to refer to the graphs formed by taking the union of the appropriate 1-cells and the 0-cells of X. We also call the elements of $H = \pi_1(X_H)$ (resp. $V = \pi_1(X_V)$) corresponding to the oriented 1-cells of X_H (resp. X_V) the **standard generators** of H (resp. V).

For expository convenience, our notation and terminology differs somewhat from [Wis96], but is essentially equivalent. In particular, what we are calling a CSC here is actually a **directed $\mathcal{VH}$ CSC**, in the terminology of [Wis96].

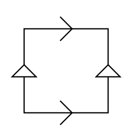

Figure 1. Orientation of the squares in a CSC

The fundamental groups of CSC's have many interesting properties. For example, the universal cover of a CSC is the product of two trees [Wis96, Thm 1.10], so the fundamental group of a compact CSC acts freely and cocompactly on a CAT(0) space. We also have the following decomposition.

Definition 2.2: It is not hard to see that a CSC X has the following structure as a graph of spaces [SW79]:

1. The vertex spaces of X are the connected components of the graph X_H.
2. The edge spaces of X correspond to the connected components of $X - X_H$, and each edge space of X has the form $\Lambda \times [0,1]$, where Λ is an oriented graph.
3. The ends $\Lambda \times 0$ and $\Lambda \times 1$ of an edge space $\Lambda \times [0,1]$ are attached to their corresponding vertex spaces by covering maps.

We call this graph of spaces the **horizontal decomposition** of X. Note that if X_H is connected, and there is exactly one edge space in the horizontal decomposition of X, then we can give $\pi_1(X)$ the structure of an HNN extension of $H = \pi_1(X_H)$ by taking the stable letter to be any of the standard generators of $V = \pi_1(X_V)$.

Let X be a CSC, let $H = \pi_1(X_H)$, and let $V = \pi_1(X_V)$. We will need the following observations about $\pi_1(X)$, all of which can be deduced easily from the fact that $\widetilde{X}$ is isomorphic to the product $\widetilde{X}_H \times \widetilde{X}_V$ (see [Wis96, Thm 1.10]).

Lemma 2.3 (Normal forms): *The inclusion of X_H and X_V in X induces embeddings of H and V in $\pi_1(X)$ such that $H \cap V = 1$. Furthermore, if X has one 0-cell, then for any $\sigma \in \pi_1(X)$, there exist unique $h \in H$ and $v \in V$ such that $\sigma = vh$.*

Lemma 2.4 ($HV = VH$): *Suppose X has only one 0-cell, and let $h_0 \in H$ and $v_0 \in V$ be reduced words. Then for the unique reduced words $h_1 \in H$ and $v_1 \in V$ such that $h_0v_0 = v_1h_1$, we have $|h_1| = |h_0|$ and $|v_1| = |v_0|$. Furthermore, if h_0 (resp. v_0) is a positive standard generator of H (resp. V), then so is h_1 (resp. v_1).*

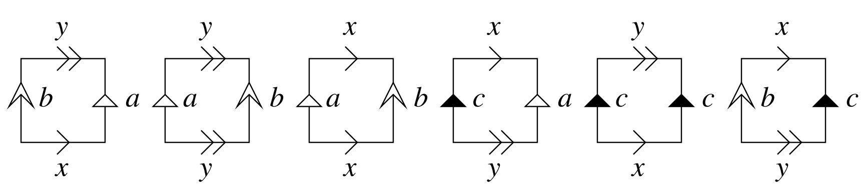

Figure 2. Example 2.5

Example 2.5: Let X be a square complex with one 0-cell, 1-cells $\{a,b,c,x,y\}$, and 2-cells defined by the six squares in Figure 2. Then X is a CSC with $X_V = \{a,b,c\}$ and $X_H = \{x,y\}$. The horizontal decomposition of X is illustrated in Figure 3. The bouquet of two circles on the right represents X_H, which is also the lone vertex space of the decomposition. The two graphs on the left are both combinatorially equivalent to a graph Λ, and the covering maps used to attach the edge space $\Lambda \times [0,1]$ to X_H are given by the labels on the two graphs. Note that the vertices in the left-hand graphs are precisely the endpoints of a, b, and c, as indicated, and the vertical edges $\{a,b,c\}$ are oriented from the lower graph to the upper one.

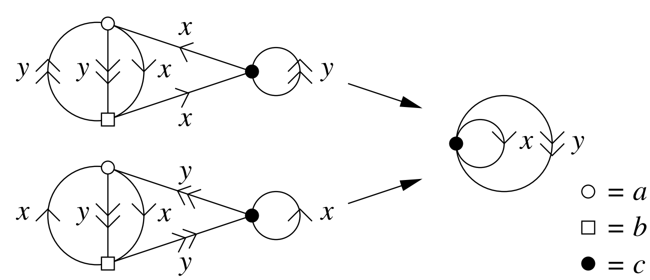

Figure 3. Horizontal decomposition of Example 2.5

3. Horizontal centralizers in CSC groups

Throughout this section, let X be a CSC with one vertex, let $G = \pi_1(X)$, let $H = \pi_1(X_H)$, let $V = \pi_1(X_V)$, and let c be a standard generator of V. We begin our study of centralizers in CSC groups by looking at $C_H(c^n)$, the “horizontal” part of the centralizer of c^n.

Let

$$H(m) = H \cap c^m H c^{-m}, \tag{3}$$

and let $\widetilde{X}_H(m)$ be the based cover of X_H corresponding to the subgroup $H(m) \leq H$. For example, when $m = 1$, $H(1)$ is one of the associated subgroups in the construction of G as an HNN extension of H, and $\widetilde{X}_H(1)$ is precisely the graph

Λ that appears in the horizontal decomposition of X (Definition 2.2).

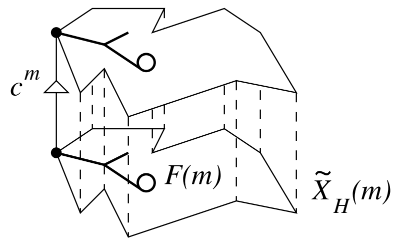

Figure 4. P arallel transporting $\widetilde{X}_H(m)$ along c^m

We consider $H(m)$ because on the one hand, clearly $C_H(c^m) \leq H(m)$, and on the other hand, it turns out that $\widetilde{X}_H(m)$ is the smallest cover of X_H that can be consistently "parallel transported" along the based path c^m, as sketched in Figure 4. (Figure 4 also gives an ov erview of the constructions in this section, including a subgraph $F(m) \subseteq \widetilde{X}_H(m)$ that will be explained later.)

In the rest of this section, we explain what is meant by "parallel transporting" $\widetilde{X}_H(m)$ along c^m. We begin by defining tw o useful functions.

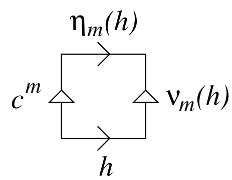

Figure 5. Definition of $\eta_m(h)$ and $\nu_m(h)$

Definition 3.1: By Lemma 2.4, for any $h \in H$, there exist unique $\eta_m(h) \in H$ and $\nu_m(h) \in V$ such that $h^{-1}c^m = \nu_m(h)\eta_m(h)^{-1}$. Therefore,

$$c^m \eta_m(h) = h\nu_m(h) \tag{4}$$

defines functions $\eta_m \colon H \to H$ and $\nu_m \colon H \to V$.

As illustrated in Figure 5, the main idea of Definition 3.1 is that $\eta_m(h)$ is supposed to be the "parallel transport" of the based path h along c^m.

The first goal of this section is to define the **cylinder of** $H(m)$ in Definition 3.5. We begin by showing in Lemmas 3.3 and 3.4 that the labels that we will use on the cylinder of $H(m)$ are well-defined.

Notation 3.2: In the sequel, we will freely identify any $h \in H$ with the unique based path it defines in $\widetilde{X}_H(m)$.

LEMMA 3.3: *For $h \in H(m)$, $\nu_m(h) = c^m$. More generally, for any based path $h \in H$, the value of $\nu_m(h)$ depends only on the endpoint of h in $\widetilde{X}_H(m)$.*

Proof: Suppose $h \in H(m)$. Since $h = c^m h' c^{-m}$ for some $h' \in H$, by Equation (4), we have

$$c^m \eta_m(h) = h\nu_m(h) = c^m h' c^{-m} \nu_m(h), \tag{5}$$

which means that $\eta_m(h) = h'(c^{-m}\nu_m(h))$. Therefore, since $\eta_m(h)$ and h' are in H and $c^{-m}\nu_m(h)$ is in V, the uniqueness of normal forms (Lemma 2.3) implies that $c^{-m}\nu_m(h) = 1$. The first assertion follows.

As for the second assertion, suppose h_1 and h_2 have the same endpoint in $\widetilde{X}_H(m)$. Then by Equation (4),

$$\begin{aligned} \nu_m(h_2)^{-1}\nu_m(h_1) &= (h_2^{-1}c^m\eta_m(h_2))^{-1}(h_1^{-1}c^m\eta_m(h_1)) \\ &= \eta_m(h_2)^{-1}c^{-m}h_2h_1^{-1}c^m\eta_m(h_1). \end{aligned} \tag{6}$$

However, since $h_2h_1^{-1} \in H(m)$, the first assertion of the lemma implies that $\nu_m(h_2h_1^{-1}) = c^m$, which means that $c^{-m}(h_2h_1^{-1})c^m = \eta_m(h_2h_1^{-1})$. Therefore,

$$\nu_m(h_2)^{-1}\nu_m(h_1) = \eta_m(h_2)^{-1}\eta_m(h_2h_1^{-1})\eta_m(h_1) \in H \cap V = 1, \tag{7}$$

and the second assertion follows. ∎

LEMMA 3.4: *For any $h \in H$ and any standard generator e of H, $\eta_m(h)^{-1}\eta_m(he)$ is also a standard generator of H, depending only on e and the endpoint of h.*

Proof: By Equation (4), we have

$$\nu_m(h) = h^{-1}c^m\eta_m(h) \quad \text{and} \quad \nu_m(he) = (he)^{-1}c^m\eta_m(he),$$

so

$$\nu_m(h)^{-1}e\nu_m(he) = \eta_m(h)^{-1}\eta_m(he). \tag{8}$$

Then, on the one hand, the left-hand side of Equation (8) depends only on e and the endpoint of h (Lemma 3.3); and on the other hand, since

$$e\nu_m(he) = \nu_m(h)(\eta_m(h)^{-1}\eta_m(he)), \tag{9}$$

by Lemma 2.4, $\eta_m(h)^{-1}\eta_m(he)$ is a standard generator of H. ∎

We may now define our main construction, which gives a precise meaning to the "parallel transport" sketch in Figure 4.

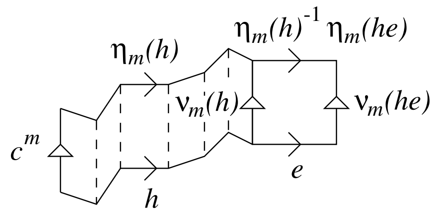

Figure 6. Labelling the cylinder of $H(m)$

Definition 3.5: The **cylinder of** $H(m)$ is the (labelled) square complex $C(m)$ constructed as follows.

1. Topologically let $C(m)$ be $\widetilde{X}_H(m) \times [0,1]$ (see Figure 4).
2. Label $\widetilde{X}_H(m) \times 0$ with the corresponding labels of $\widetilde{X}_H(m)$. (In fact, since $\widetilde{X}_H(m) \times 0$ is labelled like $\widetilde{X}_H(m)$, we will often treat them as identical in the sequel.)
3. For any based path $h \in H$, label the "vertical" edge starting at the endpoint of h in $\widetilde{X}_H(m) \times 0$ with the element $\nu_m(h)$, as shown in Figure 6. Note that by Lemma 3.3, this depends only on the endpoint of h. Furthermore, since $\widetilde{X}_H(m)$ is connected, this rule labels all vertical edges in $C(m)$.
4. For any based path $h \in H$ in $\widetilde{X}_H(m) \times 0$, and any standard generator e of H, label the edge of $\widetilde{X}_H(m) \times 1$ that lies above the last edge of the path he with the standard generator $\eta_m(h)^{-1}\eta_m(he)$ of H (see Lemma 3.4), as shown on the right-hand side of Figure 6.

For example, $C(1)$ is precisely the edge space of the horizontal decomposition of X, and the 1-cells of $C(1)$ are labelled as the corresponding 1-cells of that edge space.

Definition 3.6: In the notation of Definition 3.5, we define the **parallel transport** of any edge (resp. path) in $\widetilde{X}_H(m) \times 0$ to be the corresponding edge (resp. path) in $\widetilde{X}_H(m) \times 1$.

We now come to the key property of $C(m)$.

Lemma 3.7: *If $h \in H$ is a based path in $\widetilde{X}_H(m) \times 0$, then the parallel transport of h is the based path $\eta_m(h)$ in $\widetilde{X}_H(m) \times 1$.*

Proof: This follows from an easy induction on the length of h, as indicated by Figure 6. ∎

Definition 3.8: Let $F_0(m)$ be the union of the basepoint of $\widetilde{X}_H(m) \times 0$ and all edges in $\widetilde{X}_H(m) \times 0$ whose parallel transports in $\widetilde{X}_H(m) \times 1$ are labelled with the same generator. We define the **fixed graph** $F(m)$ to be the connected component of $F_0(m)$ that contains the basepoint of $\widetilde{X}_H(m) \times 0$.

The idea of $F(m)$ is illustrated in Figure 4. The following lemma gives two alternate descriptions of $F(m)$.

Lemma 3.9: *The fixed graph $F(m)$ is precisely the union of all based paths h in $\widetilde{X}_H(m) \times 0$ such that $\eta_m(h) = h$. In other words, a based path h is contained in $F(m)$ if and only if $h^{-1}c^m h \in V$.*

Proof: The first statement follows immediately from Lemma 3.7. The second statement follows from the first statement and the definition of η_m (Definition 3.1). ∎

We now come to our main tool for analyzing the centralizer of c^m.

Theorem 3.10: *The inclusion of the fixed graph $F(m)$ in $\widetilde{X}_H(m)$ induces an isomorphism from $\pi_1(F(m))$ onto $C_H(c^m) \leq H(m) = \pi_1(\widetilde{X}_H(m))$.*

Proof: On the one hand, let h be a based loop in $F(m)$. Since $h \in H(m)$, we have $\nu_m(h) = c^m$ (Lemma 3.3), and since h is in $F(m)$, we have $\eta_m(h) = h$ (Lemma 3.9). Therefore, by Equation (4), $c^m h = hc^m$, which means that $h \in C_H(c^m)$.

Conversely, let h be a reduced element of $C_H(c^m)$. Since h is also in $H(m)$, h defines a based loop in $\widetilde{X}_H(m)$ and $\eta_m(h) = c^m h c^{-m}$ (Lemma 3.3 and Equation (4)). However, if the path h ever leaves $F(m)$, then h has an edge that changes in its parallel transport, which means that $c^m h c^{-m} = \eta_m(h) \neq h$ (since $\eta_m(h)$ is also a reduced path). Therefore, h must be contained in $F(m)$. The theorem follows. ∎

Our last theorem of this section (Theorem 3.12) describes how the different subgraphs $F(m)$ fit together. We first need the following lemma.

LEMMA 3.11: *For $n > m$, we have $H(n) \leq H(m)$, which means that $\widetilde{X}_H(n)$ is a based cover of $\widetilde{X}_H(m)$.*

Proof: Thinking of G as an HNN extension of H with stable letter c, we note that for $h \in H$, $h \in H(n)$ if and only if $c^{-n}hc^n$ reduces to some $h' \in H$. However, by the normal form theorem for HNN extensions, for $0 < m < n$, if $c^{-m}hc^m$ cannot be reduced to an element of H, then $c^{-n}hc^n$ also cannot be reduced to an element of H. In other words, any element of $H(n)$ must also be an element of $H(m)$. ∎

THEOREM 3.12: *For $m \geq 1$, $F(m)$ is a subgraph of $F(2m)$. In particular, $F(1) \subseteq F(2) \subseteq \cdots \subseteq F(2^k) \subseteq \cdots$ is an ascending chain of graphs.*

Proof: Consider the cover $\widetilde{X}_H(2m) \to \widetilde{X}_H(m)$ from Lemma 3.11. The inclusion of $F(m)$ in $\widetilde{X}_H(m)$ lifts to a map from $F(m)$ into $\widetilde{X}_H(2m)$ if and only if every element of $\pi_1(F(m))$ lifts to a loop in $\pi_1(\widetilde{X}_H(2m))$. Therefore, since Theorem 3.10 implies

$$\pi_1(F(m)) = C_H(c^m) \leq C_H(c^{2m}) \leq H(2m) = \pi_1(\widetilde{X}_H(2m)), \tag{10}$$

we may consider $F(m)$ as a based subgraph of $\widetilde{X}_H(2m)$.

By Lemma 3.9, it now suffices to show that every based path h in $F(m)$ is also contained in $F(2m)$. So suppose h is contained in $F(m)$. In that case, $h^{-1}c^m h \in V$ (Lemma 3.9), which means that $(h^{-1}c^m h)^2 = h^{-1}c^{2m}h \in V$. Therefore, h is contained in $F(2m)$ (Lemma 3.9 again). The theorem follows. ∎

4. Full centralizers in CSC groups

In this section, retaining the notation and conventions of Section 3, we extend our analysis of $C_H(c^m)$ to the full centralizer $C_G(c^m)$. We begin by looking at how c acts on $C_H(c^m)$.

LEMMA 4.1: *For $m \geq 1$, there exists a graph automorphism γ_m of $F(m)$ such that:*

1. *The automorphism of $C_H(c^m) \cong \pi_1(F(m))$ induced by γ_m is precisely conjugation by c. (In particular, c normalizes $C_H(c^m)$.)*
2. *Considering $F(m)$ as a subgraph of $F(2m)$ (see Theorem 3.12), the restriction of γ_{2m} to $F(m)$ is precisely γ_m.*

Proof: First, following Definition 3.1, by Lemma 2.4, we may define functions $\phi\colon H \to H$ and $\psi\colon H \to V$ by the equation

$$c\phi(h) = h\psi(h). \tag{11}$$

Then, since $H(m) \leq H(1)$, the proof of Lemma 3.3 shows that $\phi(h)$ depends only on the endpoint of a based path h in $\widetilde{X}_H(m)$, and the proof of Lemma 3.4 shows that ϕ defines a relabelling of $\widetilde{X}_H(m)$ analogous to the one in Definition 3.5 and Lemma 3.7.

Next, consider any based path h contained in the fixed graph $F(m)$. By Lemma 3.9, we have that $\eta_m(h) = h$, which means that Equation (4) implies $h^{-1}c^m h = \nu_m(h)$. Applying Equation (11), we then see that

$$\phi(h)^{-1}c^m\phi(h) = \psi(h)^{-1}h^{-1}c^m h\psi(h) = \psi(h)^{-1}\nu_m(h)\psi(h) \in V, \tag{12}$$

which means that $\phi(h)$ is contained in $F(m)$ (Lemma 3.9 again). Therefore, the relabelling defined by ϕ preserves $F(m)$ setwise. Consequently, since $F(m)$ is a finite graph, ϕ must actually induce an automorphism γ_m of $F(m)$. Assertion 2 of the lemma follows immediately, since the automorphisms γ_m are induced by the relabelling ϕ, which is independent of m.

As for the rest of the lemma, suppose $h \in \pi_1(F(m))$. First, Lemma 2.4 and Equation (11) imply that $\psi(h)$ is a standard generator of V. Furthermore, since $h \in \pi_1(F(m)) \leq H(m)$, we see that $\nu_m(h) = c^m$, and Equation (12) implies

$$c^m\phi(h) = \phi(h)(\psi(h)^{-1}c^m\psi(h)). \tag{13}$$

Applying Lemma 2.4 again, we see that $\left|\psi(h)^{-1}c^m\psi(h)\right| = |c^m|$. Therefore, since $C_V(c^m) = \langle c\rangle$, we must have $\psi(h) = c$ and $\phi(h) = c^{-1}hc$. The lemma follows. ∎

Remark 4.2: Note that assertion 2 of Lemma 4.1 is equivalent to saying that we may combine the automorphisms γ_{2^k} to define an automorphism γ of the direct limit $F(1) \subseteq F(2) \subseteq \cdots \subseteq F(2^k) \subseteq \cdots$. We will refer to γ freely in the sequel.

LEMMA 4.3: *For $m \geq 1$, $C_G(c^m) = C_H(c^m) \rtimes \langle c\rangle$, where the semidirect product is defined by the action in Lemma 4.1.*

Proof: By Lemma 4.1, it suffices to show that $C_G(c^m) = \langle c\rangle\, C_H(c^m)$. So consider $g \in C_G(c^m)$. By the Normal Form Lemma, we have $g = g_v g_h$, where

g_v (resp. g_h) is a reduced word in V (resp. H). It will therefore be enough to show that if the first letter of g_v is not $c^{\pm 1}$, then $g_v = 1$.

So suppose that the first letter of g_v is not $c^{\pm 1}$. By Lemma 2.4, $g_h c^m = k_v k_h$ for some reduced word $k_h \in H$ and some reduced word $k_v \in V$ such that $|k_v| = |c^m| = m$. But then, since $g_v g_h \in C_G(c^m)$, it follows that

$$c^m g_v g_h = g_v g_h c^m = g_v k_v k_h. \tag{14}$$

Therefore, by the uniqueness of normal forms (Lemma 2.3), we have $c^m g_v = g_v k_v$, or in other words, $k_v = g_v^{-1} c^m g_v$. However, since $g_v^{-1} c^m g_v$ is reduced, and $|k_v| = m$, we must have $g_v = 1$. ∎

THEOREM 4.4: *Let $\Gamma = \langle \gamma \rangle$. We have that $(C_G(c^m))_{\mathrm{ab}} \cong \langle c \rangle \times \pi_1(\Gamma \backslash F(m))_{\mathrm{ab}}$. Consequently, if the rank of $\pi_1(F(2^k))$ is a strictly increasing function of k, then the groups $(C_G(c^m))_{\mathrm{ab}}$ are free abelian of arbitrarily high rank.*

Proof: On the one hand, since γ acts by conjugation on $\pi_1(F(m))$, the image of $\pi_1(F(m))$ in $(C_G(c^m))_{\mathrm{ab}}$ must certainly factor through $\pi_1(\Gamma \backslash F(m))_{\mathrm{ab}}$. On the other hand, γ has trivial action on $\Gamma \backslash F(m)$, so c commutes with elements of $\pi_1(\Gamma \backslash F(m))$. The first claim of the theorem follows.

As for the other claim, if the rank of $\pi_1(F(2^k))$ is a strictly increasing function of k, since the action of γ_m commutes with the direct limit $F(1) \subset F(2) \subset \cdots \subset F(2^k) \subset \cdots$, there will be new topology at every stage of the direct limit $\Gamma \backslash F(1) \subset \Gamma \backslash F(2) \subset \cdots \subset \Gamma \backslash F(2^k) \subset \cdots$. Therefore, the rank of $\pi_1(\Gamma \backslash F(2^k))$ is a strictly increasing function of k, and the theorem follows. ∎

5. The examples X and X^+

We now specialize the results of Sections 3 and 4 to two particular examples. We again retain the notation and conventions of Sections 3 and 4 in this section.

Consider the CSC X from Example 2.5. Since X has one 0-cell, the results of Sections 3 and 4 apply to X and the element $c \in \pi_1(X)$. Furthermore, we deduce from [Wis96] that X has the following remarkable property.

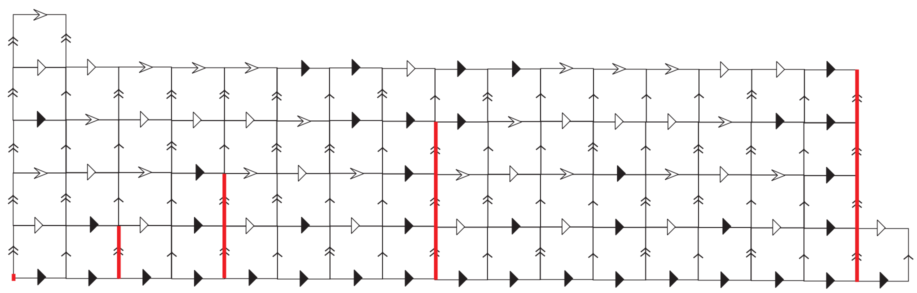

Figure 7. Positive quadrant of the antitorus A

THEOREM 5.1: *The fixed graphs $F(2^k)$ form a strictly increasing sequence $F(1) \subset F(2) \subset \cdots \subset F(2^k) \subset \cdots$, with new edges and vertices at each stage.*

Proof: By Theorem 3.12, we need only show that there are always new edges and vertices in the fixed graph each time we go from $F(2^k)$ to $F(2^{k+1})$, and this follows from the construction of the **antitorus** A in [Wis96, Section II.3]. Briefly, A is a flat plane in $\widetilde{X}$ whose axes are labelled by infinite paths of the form c^∞ and y^∞. The positive quadrant of A is illustrated in Figure 7 (reflected along the diagonal to conserve space), using the notation of Example 2.5 and Figure 2. For our purposes, the key feature of A is the fact that it is aperiodic; specifically, the period of each successive infinite horizontal strip is double the period of the previous strip. In particular, the theorem follows from the fact that for each $k \geq 0$, the path y^{k+1} lies in $F(2^{k+1})$ but not in $F(2^k)$, as indicated by the heavy lines in Figure 7. ∎

By Theorem 4.4, it then remains to show that the sequence $F(1) \subset F(2) \subset \cdots \subset F(2^k) \subset \cdots$ has new topology at each stage. This probably holds for the example X, but for the sake of brevity, we instead turn to a modified version of

X.

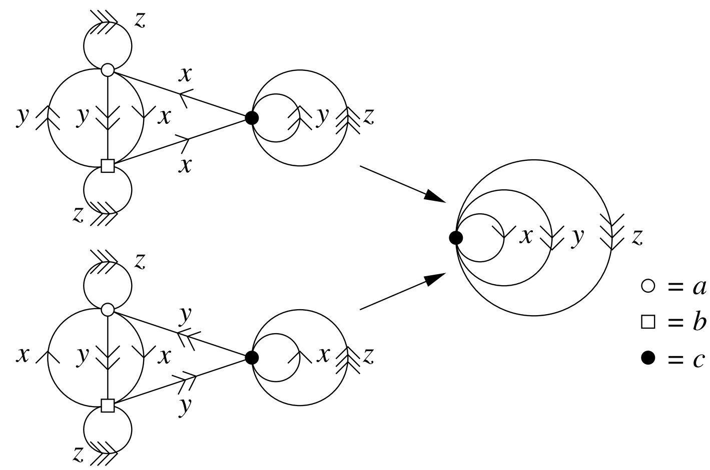

Figure 8. Horizontal decomposition of X^+

Proof of Theorem 1.3: Let X^+ be the CSC with one 0-cell, 1-cells $X_V^+ = \{a, b, c\}$ and $X_H^+ = \{x, y, z\}$, and the horizontal decomposition shown in Figure 8. Note that X^+ is just the example X plus an extra horizontal generator z that centralizes V. Let G^+, H^+, $H^+(m)$, $\widetilde{X}_H^+(m)$, η_m^+, ν_m^+, and $F^+(m)$ be the constructions for X^+ analogous to G, H, $H(m)$, $\widetilde{X}_H(m)$, η_m, ν_m, and $F(m)$ for X.

We first claim that each graph $\widetilde{X}_H^+(m)$ is precisely the graph $\widetilde{X}_H(m)$ with a z self-loop added at every vertex. This is equivalent to saying that for every $h \in H^+$, $hzh^{-1} \in H^+(m)$, which follows because Equation (4) and the fact that z centralizes V imply

$$c^{-m}(hzh^{-1})c^m = \eta_m^+(h)\nu_m^+(h)^{-1}z\nu_m^+(h)\eta_m^+(h)^{-1} = \eta_m^+(h)z\eta_m^+(h)^{-1} \in H^+, \tag{15}$$

as illustrated in Figure 9.

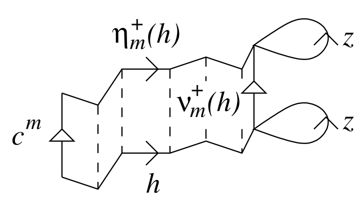

Figure 9. All z self-loops are fixed in parallel transport

We now observe that if h is contained in $F^+(m)$, then $\eta_m^+(h) = h$, and Equation (15) becomes $hzh^{-1}c^m = c^m hzh^{-1}$. It follows that every z self-loop at a vertex of $F^+(m)$ is included in $F^+(m)$, which means that $F^+(m)$ is precisely the graph $F(m)$ with a z self-loop added at every vertex. Therefore, from Theorem 5.1, we see that the strictly increasing sequence $F^+(1) \subset F^+(2) \subset \cdots \subset F^+(2^k) \subset \cdots$ has new topology at each stage. Theorem 4.4 then implies that the groups $(C_{G^+}(c^m))_{\mathrm{ab}}$ are free abelian of arbitrarily high finite rank, and Theorem 1.3 follows. ∎

References

[BH92] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, Annals of Mathematics (2) **135** (1992), 1–51.

[Ger87] S. M. Gersten, *Fixed points of automorphisms of free groups*, Advances in Mathematics **64** (1987), 51–85.

[GS91] S. M. Gersten and H. Short, *Small cancellation theory and automatic groups. II*, Inventiones Mathematicae **105** (1991), 641–662.

[Sho99] J. A. Shor, *On fixed subgroups of automorphisms in hyperbolic groups*, PhD thesis, Columbia University, 1999.

[SW79] P. Scott and T. Wall, *Topological methods in group theory* in *Homological Group Theory (Proc. Sympos., Durham, 1977)*, Volume 36 of London Mathematical Society Lecture Note Series, Cambridge University Press, Cambridge, 1979, pp. 137–203.

[Ven02] E. Ventura, *Fixed subgroups in free groups: a survey*, in *Combinatorial and Geometric Group Theory (New York, 2000/Hoboken, NJ, 2001)*, Volume 296 of Contemporary Mathematics, American Mathematical Society, Providence, RI, 2002, pp. 231–255.

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