EMBEDDING FINITELY GENERATED FREE-BY-CYCLIC GROUPS IN {FINITELY GENERATED FREE}-BY-CYCLIC GROUPS

MARCO LINTON

ABSTRACT. We refine Feighn–Handel's results on subgroups of mapping tori of free groups to the special case of free-by-cyclic groups. We use these refinements to show that any finitely generated free-by-cyclic group embeds in a {finitely generated free}-by-cyclic group. When the free-by-cyclic group is hyperbolic, it embeds in a hyperbolic {finitely generated free}-by-cyclic group as a quasi-convex subgroup. Combined with a result of Hagen–Wise, this implies that all hyperbolic free-by-cyclic groups are cocompactly cubulated.

1. Introduction

Feighn–Handel developed a powerful tool in [FH99] to study finitely generated subgroups of mapping tori $M(\psi)$ of free group monomorphisms $\psi\colon \mathbb{F} \to \mathbb{F}$. The purpose of this article is to refine Feighn–Handel's techniques for the special case in which ψ is an automorphism —that is, when $M(\psi)$ is a free-by-cyclic group $\mathbb{F} \rtimes_\psi \mathbb{Z}$. Here and throughout the article, we do not assume that $\mathbb{F}$ is finitely generated.

Feighn–Handel's main technical result stated that a finitely generated mapping torus of a free group $M(\psi)$ splits as a HNN-extension $F*_\phi$ with F a finitely generated free group and with ϕ an isomorphism identifying a free factor of F with another subgroup of F. Our first main result, Theorem 4.1, shows that a finitely generated free-by-cyclic group $\mathbb{F} \rtimes_\psi \mathbb{Z}$ splits as a HNN-extension $F*_\phi$ with F a finitely generated free factor of $\mathbb{F}$ and with ϕ an isomorphism identifying two free factors of F.

Using this, we characterise when a free-by-cyclic group $\mathbb{F} \rtimes_\psi \mathbb{Z}$ is finitely generated in terms of the automorphism ψ. Note that every non-trivial subgroup of a free-by-$\mathbb{Z}$ group is itself free-by-$\mathbb{Z}$.

Theorem 1.1. *Let $\mathbb{F}$ be a free group, let $\psi \in \mathrm{Aut}(\mathbb{F})$ be an automorphism and let $G = \mathbb{F} \rtimes_\psi \mathbb{Z}$. Then G is finitely generated if and only if there is a free product decomposition*

$$\mathbb{F} = A * (\ast_{i\in\mathbb{Z}} C_i)$$

where A and C_0 are finitely generated and where $C_i = \psi^i(C_0)$ for all $i \in \mathbb{Z}$.

In [Lin25] the author proved that finitely generated mapping tori of free groups have a canonical collection of maximal submapping tori with respect to which they are relatively hyperbolic and relatively locally quasi-convex. Using Theorem 1.1, we state a strengthening of this result, Theorem 4.3.

Chong–Wise proved in [CW24] that a finitely generated mapping torus of a free group embeds in a mapping torus of a finitely generated free group. They conjectured [CW24, Conjecture 1.2] that an analogous statement should hold for free-by-cyclic groups: that is, a finitely generated free-by-cyclic group embeds as a subgroup of a {finitely generated free}-by-cyclic group. Our second main result uses Theorem 1.1 to confirm their conjecture. A more general question, which we also answer, was also posed independently by Wu–Ye in [WY25, Question 7.4].

Theorem 1.2. *If $G = \mathbb{F} \rtimes_\psi \mathbb{Z}$ is a finitely generated free-by-cyclic group, then G embeds in a {finitely generated free}-by-$\mathbb{Z}$ group. More explicitly, there is some finitely generated free group F, an automorphism $\theta \in \mathrm{Aut}(F)$ and an embedding $\iota\colon \mathbb{F} \hookrightarrow F$ making the following diagram commute:*

$$\begin{tikzcd} \mathbb{F} \arrow[r, "\psi"] \arrow[d, hook, "\iota"] & \mathbb{F} \arrow[d, hook, "\iota"] \\ F \arrow[r, "\theta"] & F \end{tikzcd}$$

Moreover, if G is hyperbolic, then $F \rtimes_\theta \mathbb{Z}$ is also hyperbolic and G embeds in $F \rtimes_\theta \mathbb{Z}$ as a quasi-convex subgroup.

Hagen–Wise showed in [HW15] that a hyperbolic {fg free}-by-$\mathbb{Z}$ group acts properly and cocompactly on a CAT(0) cube complex. By a result of Agol [Ago13], this implies that hyperbolic {fg free}-by-$\mathbb{Z}$ groups are virtually compact special in the sense of Haglund–Wise [HW08]. Since compact special groups are precisely the quasi-convex subgroups of right angled Artin groups (see [HW08]), it follows that quasi-convex subgroups of hyperbolic virtually compact special groups are themselves virtually compact special. Combining these facts with Theorem 1.2 we obtain the following corollary.

Corollary 1.3. *A hyperbolic free-by-cyclic group is virtually compact special.*

Since every finitely generated subgroup of a hyperbolic free-by-cyclic group is hyperbolic by a result of Gersten [Ger96] (combined with coherence [FH99]), Corollary 1.3 implies that every finitely generated subgroup of a hyperbolic free-by-cyclic group is also virtually compact special.

An automorphism $\psi \in \mathrm{Aut}(\mathbb{F})$ is *fully irreducible* if there is no finitely generated non-trivial proper free factor $H \leqslant \mathbb{F}$ and integer $m \geqslant 1$ so that $\psi^m(H)$ is conjugate to H. The automorphism θ from Theorem 1.2 is never fully irreducible by its construction. It is likely that Theorem 1.2 can be improved when ψ is assumed fully irreducible.

Question 1.4. If $\psi \in \mathrm{Aut}(\mathbb{F})$ is a fully irreducible automorphism so that $\mathbb{F} \rtimes_\psi \mathbb{Z}$ is finitely generated, is there a finitely generated free group F, an embedding $\iota\colon \mathbb{F} \to F$ and an automorphism $\theta \in \mathrm{Aut}(F)$ so that $\theta \circ \iota = \iota \circ \psi$?

The analogue of Question 1.4 for monomorphisms $\mathbb{F} \to \mathbb{F}$ was established by Chong–Wise [CW24].

Acknowledgements. This work has received support from the grant 202450E223 (Impulso de líneas científicas estratégicas de ICMAT).

2. Preliminaries

2.1. **Graphs, folds and subgroups of free groups.** A *graph* is a 1-dimensional CW-complex Γ. The 0-cells of Γ are the *vertices* and the 1-cells are the *edges*. Maps $\Gamma \to \Lambda$ between graphs will always be assumed to be combinatorial —that is, vertices map to vertices and edges map homeomorphically to edges. A map of graphs is an *immersion* if it is locally injective. A pointed graph is a pair (Γ, v) with Γ a graph and $v \in \Gamma$ a vertex. A pointed map $(\Gamma, v) \to (\Lambda, u)$ is a map of graphs $\Gamma \to \Lambda$ sending the basepoint v to the basepoint u.

If $f\colon (\Gamma, v) \to (\Lambda, u)$ is a graph map, write $f_*\colon \pi_1(\Gamma, v) \to \pi_1(\Lambda, u)$ for the induced map on fundamental groups. When f is understood, we use $\Gamma^{\#}$ to denote the image $f_*(\pi_1(\Gamma, v))$.

A map of graphs $f\colon \Gamma \to \Lambda$ is a *fold* if it identifies a single pair of oriented edges $e_1, e_2 \subset \Gamma$ with a common origin. It is a classical observation of Stallings' [Sta83] that any graph map $f\colon \Gamma \to \Lambda$, with Γ a finite graph, factors as a finite sequence of folds $f_i\colon \Gamma_{i-1} \to \Gamma_i$ followed by an immersion $h\colon \Gamma_k \to \Lambda$:

$$\begin{tikzcd} \Gamma = \Gamma_0 \arrow[r, "f_1"'] \arrow[rrrr, bend left, "f"] & \Gamma_1 \arrow[r, "f_2"'] & \dots \arrow[r, "f_k"'] & \Gamma_k \arrow[r, "h"'] & \Lambda \end{tikzcd}$$

Although the sequence of folds f_i may not be unique, the graph Γ_k and the map h is unique.

A graph Γ is *core* if it is the union of the images of immersed cycles $S^1 \to \Gamma$. A pointed graph (Γ, v) is *pointed core* if it is the union of the images of immersed loops $I \to \Gamma$ at v. If Γ is a graph, the *core of* Γ is the subgraph $\mathrm{Core}(\Gamma) \subset \Gamma$ consisting of the union of all images of immersed cycles in Γ (it may be empty). If (Γ, v) is a pointed graph, the *pointed core of* (Γ, v) is the pointed subgraph $\mathrm{Core}(\Gamma, v) \subset \Gamma$ consisting of the union of all images of immersed loops at v.

If $H \leqslant \pi_1(\Lambda, u)$ is any subgroup, then there is a unique immersion $f\colon (\Gamma, v) \to (\Lambda, u)$ so that (Γ, v) is pointed core and $f_*(\pi_1(\Gamma, v)) = H$. This (pointed) graph, which we shall denote by $\Gamma(H)$ and call the *pointed subgroup graph for* H, is the pointed core subgraph of the pointed cover of Λ associated to H.

If $[H]$ is the conjugacy class of a subgroup $H \leqslant \pi_1(\Lambda, u)$, then there is a unique immersion $f\colon \Gamma \to \Lambda$ so that Γ is core and $[f_*(\pi_1(\Gamma, v))] = [H]$, where $v \in f^{-1}(u)$. This graph, which we shall denote by $\Gamma[H]$ and call the *subgroup graph for* $[H]$, is the core subgraph of the cover of Λ associated to H.

If $\delta \in \pi_1(\Lambda, u)$, we shall often write $\Gamma(\delta)$ or $\Gamma[\delta]$ instead of $\Gamma(\langle\delta\rangle)$ or $\Gamma[\langle\delta\rangle]$.

2.2. **Graph pairs and graph triples.** Let $f_Z\colon (Z, v) \to (R, u)$ be a pointed map of connected graphs. If $X \subset Z$ is a connected subgraph containing the basepoint, we call (Z, X), together with the map f_Z, which we often suppress, a *graph pair*. If $Y, X \subset Z$ are connected subgraphs such that $X \cap Y$ is connected and contains the basepoint, we call (Z, Y, X), together with the map f_Z, which we often suppress, a *graph triple.* A graph pair (Z, X) or a graph triple (Z, Y, X) is *tight* if the map f_Z is an immersion.

The *relative rank* of a graph pair (Z, X) is defined as

$$\mathrm{rr}(Z, X) = \mathrm{rk}(\pi_1(Z, v)/\langle\langle \pi_1(X, v)\rangle\rangle),$$

where $\mathrm{rk}(G)$ denotes the *rank* of the group G. Since $X \subset Z$, the subgroup $\pi_1(X, v)$ is a free factor of $\pi_1(Z, v)$ and so the quotient $\pi_1(Z, v)/\langle\langle \pi_1(X, v)\rangle\rangle$ is a free group. When $Z - X$ consists of finitely many 0-cells and 1-cells, we have

$$\mathrm{rr}(Z, X) = \mathrm{rk}(\pi_1(Z, v)) - \mathrm{rk}(\pi_1(X, v)) = -\chi(Z, X).$$

Let (Z_1, X_1) and (Z_2, X_2) be graph pairs with $f_1 \colon Z_1 \to R$, $f_2 \colon Z_2 \to R$ the associated maps. A *map of graph pairs* is a pair

$$q = (q_Z, q_X) \colon (Z_1, X_1) \to (Z_2, X_2)$$

of graph maps $q_Z \colon Z_1 \to Z_2$, $q_X \colon X_1 \to X_2$ making the following diagram commute:

$$\begin{tikzcd}
X_1 \arrow[r, hook] \arrow[d, "q_X"] & Z_1 \arrow[d, "q_Z"] \arrow[dr, "f_1"] & \\
X_2 \arrow[r, hook] & Z_2 \arrow[r, "f_2"] & R
\end{tikzcd}$$

Similarly, if (Z_1, Y_1, X_2) and (Z_2, Y_2, X_2) are graph triples, a *map of graph triples* is a triple

$$q = (q_Z, q_Y, q_X) \colon (Z_1, Y_1, X_1) \to (Z_2, Y_2, X_2)$$

of graph maps $q_Z \colon Z_1 \to Z_2$, $q_Y \colon Y_1 \to Y_2$, $q_X \colon X_1 \to X_2$ making the appropriate diagram commute.

We shall usually assume that our graph pairs and triples and maps of graph pairs and triples are pointed. For graph pairs (Z, X), the basepoint is required to be in X. For graph triples (Z, Y, X), the basepoint is required to be in $X \cap Y$.

Remark 2.1. Note that if (Z, Y, X) is a graph triple, then $(Y, X \cap Y)$ and $(X, X \cap Y)$, together with the induced maps, are both graph pairs.

2.3. **Invariant graph pairs.** The main objects of study in Feighn–Handel's article [FH99] were (ψ-)invariant graph pairs, which we now define.

Let $\mathbb{F}$ be a free group, let $\psi \colon \mathbb{F} \to \mathbb{F}$ be a monomorphism and assume we have identified $\mathbb{F}$ with $\pi_1(R, u)$ for some pointed graph (R, u). Recall that the *mapping torus of* ψ is the group:

$$M(\psi) = \langle \mathbb{F}, t \mid t^{-1} f t = \psi(f),\ f \in \mathbb{F}\rangle.$$

Let (Z, X) be a graph pair with associated graph map $f_Z \colon Z \to R$. Say the pair (Z, X) is *ψ-invariant* (or just invariant) if

$$Z^{\#} = \langle X^{\#}, \psi(X^{\#})\rangle.$$

If $H \leqslant M(\psi)$ is a subgroup with $t \in H$, say (Z, X) is an *ψ-invariant graph pair for H* if (Z, X) is a ψ-invariant graph pair and if

$$H = \langle X^{\#}, t\rangle.$$

If $H \leqslant M(\psi)$ is a finitely generated subgroup with $t \in H$, say a finite graph pair (Z, X) is *minimal for H* if (Z, X) is an invariant graph pair for H and if

$$\mathrm{rr}(Z, X) \leqslant \mathrm{rr}(Z', X')$$

for all finite invariant graph pairs (Z', X') for H.

The main technical result in [FH99] states that a finitely generated subgroup H of a mapping torus (that contains t) has a particularly nice HNN-extension decomposition which can be read off from a minimal invariant graph pair for H. See [Lin25, Theorem 4.6] for the version stated below.

Theorem 2.2. *Let (Z, X) be a finite ψ-invariant graph pair for H with $(f_Z)_*$ injective and let $C \leqslant Z^\#$ so that $Z^\# = X^\# * C$. The following are equivalent:*

1. *The pair (Z, X) is minimal.*
2. *The map*
$$\theta_n \colon \pi_1\left(X \vee \bigvee_{i=0}^{n} \Gamma(\psi^i(C)), v_Z\right) \to \pi_1(R, v_R)$$
is injective for all $n \geqslant 0$.
3. *We have*
$$H \cong \langle Z^\#, t \mid t^{-1}xt = \psi(x), \forall x \in X^\# \rangle.$$

In particular, if any of the above hold, then $\chi(H) = -\mathrm{rr}(Z, X)$.

2.4. **Bi-invariant graph triples.** We now define the main objects of study of this article, bi-invariant graph triples.

Let $\mathbb{F}$ be a free group, let $\psi \colon \mathbb{F} \to \mathbb{F}$ be an isomorphism and assume we have identified $\mathbb{F}$ with $\pi_1(R, u)$ for some pointed graph (R, u). Then the mapping torus of ψ is the free-by-cyclic group:

$$M(\psi) \cong \mathbb{F} \rtimes_\psi \mathbb{Z}.$$

Let (Z, Y, X) be a graph triple with associated map $f_Z \colon Z \to R$. Say the triple (Z, Y, X) is *ψ-bi-invariant* (or just bi-invariant) if $(Y, X \cap Y)$ is ψ-invariant and $(X, X \cap Y)$ is ψ^{-1}-invariant. Note that $(Y, X \cap Y)$ being ψ-invariant implies that (Z, X) is ψ-invariant. Similarly, $(X, X \cap Y)$ being ψ^{-1}-invariant implies that (Z, Y) is ψ^{-1}-invariant.

Lemma 2.3. *If (Z, Y, X) is a ψ-bi-invariant graph triple, then the induced homomorphism*

$$\psi \mid_{X^\#} \colon X^\# \to Y^\#$$

is an isomorphism.

Proof. We have that $(Y, X \cap Y)$ is a ψ-invariant graph pair and $(X, X \cap Y)$ is a ψ^{-1}-invariant graph pair. In particular,

$$\begin{aligned} Y^\# &= \langle (X \cap Y)^\#, \psi((X \cap Y)^\#) \rangle \\ X^\# &= \langle (X \cap Y)^\#, \psi^{-1}((X \cap Y)^\#) \rangle. \end{aligned}$$

Thus, $\psi(X^\#) = Y^\#$ and $\psi^{-1}(Y^\#) = X^\#$. Since ψ is an automorphism of $\mathbb{F}$, it follows that $\psi \mid_{X^\#}$ is an isomorphism as claimed. ☐

If $H \leqslant M(\psi)$ is a subgroup with $t \in H$, say (Z, Y, X) is a *ψ-bi-invariant graph triple for H* if (Z, Y, X) is ψ-bi-invariant and if

$$H = \langle (X \cap Y)^{\#}, t \rangle.$$

This last condition is equivalent to:

$$\left\langle \bigcup_{i \in \mathbb{Z}} \psi^i(Z^{\#}) \right\rangle = H \cap \mathbb{F}.$$

If $H \leqslant M(\psi)$ is a finitely generated subgroup with $t \in H$, say a finite graph triple (Z, Y, X) is *minimal for H* if (Z, Y, X) is a bi-invariant graph triple for H and if

$$\mathrm{rr}(Z, X) \leqslant \mathrm{rr}(Z', X') \quad \text{and} \quad \mathrm{rr}(Z, Y) \leqslant \mathrm{rr}(Z', Y')$$

for all finite bi-invariant graph triples (Z', Y', X') for H.

Remark 2.4. If $H \leqslant M(\psi)$ is a subgroup with $t \in H$, there is a subgroup $L \leqslant \mathbb{F}$ so that $H = \langle L, t \rangle$. In particular, the graph triple

$$(\Gamma(L) \vee \Gamma(\psi(L)) \vee \Gamma(\psi^{-1}(L)), \Gamma(L) \vee \Gamma(\psi(L)), \Gamma(L) \vee \Gamma(\psi^{-1}(L)))$$

is a bi-invariant graph triple for H. If H is finitely generated, then L can be taken to be finitely generated so that the graph triple above will be finite.

In Section 4.1 we shall prove the analogue of Theorem 2.2 for graph triples.

3. Tightening graph triples

We fix the same notation as in Section 2.4. Namely, $\mathbb{F}$ is a free group, $\psi\colon \mathbb{F} \to \mathbb{F}$ is an isomorphism, (R, u) is a pointed graph and there is an identification $\mathbb{F} \cong \pi_1(R, u)$. The aim of this section is to define a version of Feighn–Handel's tightening procedure for graph triples and prove Proposition 3.17.

We begin by identifying the various types of folds for a graph triple. There will be more cases than there were for the graph pair case.

3.1. Types of folds.

Let (Z, Y, X) be a ψ-bi-invariant graph triple for a subgroup H with basepoint $v_Z \in X \cap Y$. Let $e_1, e_2 \subset Z$ be two oriented edges with common origin which map to the same edge in R. Denote by Z_1 the graph obtained from Z by folding e_1 and e_2 and denote by $q\colon Z \to Z_1$ the folding map. Denote by v, w_1 and v, w_2 the two endpoints of e_1, e_2 respectively.

Denote also by $\check{q}_X\colon X \to q(X) = X_1$, $\check{q}_Y\colon Y \to q(Y) = Y_1$ and by $\check{q}_{X\cap Y}\colon X \cap Y \to q(X \cap Y)$ the induced maps on X, Y and $X \cap Y$. Just like in [FH99] we use the different notation for these maps to emphasise that $\check{q}_X, \check{q}_Y$ and $\check{q}_{X\cap Y}$ may not be folds, even though q is. Moreover, we may also have that $q(X \cap Y) \neq X_1 \cap Y_1$.

We will maintain this notation for the remainder of this section.

There are several cases to consider for q:

(1) q is an *ordinary fold* if $\check{q}_X, \check{q}_Y$ are folds or homeomorphisms. If q is ordinary, then say:
   (a) q is an *ordinary X-fold* if $\check{q}_X$ is a fold.
   (b) q is an *ordinary Y-fold* if $\check{q}_Y$ is a fold.

(2) q is a *single exceptional fold* if only one of $\check{q}_X, \check{q}_Y$ is a fold or a homeomorphism. If q is a single exceptional fold, then say:
   (a) q is an X*-exceptional fold* if $\check{q}_X$ is a fold or a homeomorphism.
   (b) q is a Y*-exceptional fold* if $\check{q}_Y$ is a fold or a homeomorphism.

(3) q is a *double exceptional fold* if neither $\check{q}_X$ nor $\check{q}_Y$ is a fold or a homeomorphism.

Note that an ordinary fold does not have to be an ordinary X-fold or an ordinary Y-fold as it can be a homeomorphism on both X and Y.

We immediately note the following.

Lemma 3.1. *We have:*

$$\mathrm{rr}(Z_1, X_1) \leqslant \mathrm{rr}(Z, X)$$
$$\mathrm{rr}(Z_1, Y_1) \leqslant \mathrm{rr}(Z, Y).$$

Proof. When Z is finite, this is [FH99, Lemma 4.5 (4)] applied to the ψ-invariant graph pair (Z, X) and the ψ^{-1}-invariant graph pair (Z, Y). The general case follows by the same reasoning. $\square$

Before analysing the three types of folds, we shall need a general criterion for when q is not a fold or a homeomorphism on a chosen subgraph. The following is immediate from the definitions.

Lemma 3.2. *If $A \subset Z$ is a subgraph, then $q\,|_A \colon A \to q(A)$ is not a fold or a homeomorphism if and only if the following three conditions hold:*

- $w_1, w_2 \in A$.
- $w_1 \neq w_2$.
- $e_1 \not\subset A$ *or* $e_2 \not\subset A$.

In particular, if $q\,|_A$ is not a fold or a homeomorphism, then it identifies a pair of vertices.

We now analyse the three cases separately.

3.2. **Ordinary folds.** The case in which q is an ordinary fold is the simplest to handle.

Lemma 3.3. *If q is an ordinary fold, then one of the following holds:*

(1) *q is an ordinary X-fold or an ordinary Y-fold. In this case $\check{q}_{X\cap Y}$ is a fold or a homeomorphism and $X_1 \cap Y_1 = q(X \cap Y)$.*

(2) *$w_1, e_1 \subset X - X \cap Y$ and $w_2, e_2 \subset Y - X \cap Y$ (or vice versa). In this case, $\check{q}_{X\cap Y}$ is a homeomorphism and $X_1 \cap Y_1 = q(X \cap Y) \sqcup q(e_1) \sqcup q(w_1)$.*

(3) *$e_1 \subset X - X \cap Y$, $e_2 \subset Y - X \cap Y$ (or vice versa) and $w_1 = w_2 \in X \cap Y$. In this case, $\check{q}_{X\cap Y}$ is a homeomorphism and $X_1 \cap Y_1 = q(X \cap Y) \sqcup q(e_1)$.*

Proof. If q is an X-fold and a Y-fold, then $q \mid X \cap Y$ is also a fold and so $e_1, e_2 \subset X \cap Y$ and $q(X \cap Y) = X_1 \cap Y_1)$. If q is an X-fold, but not a Y-fold, then $e_1, e_2 \subset X$ and either $w_1 = w_2$ or $w_1 \neq w_2$ and w_1, w_2 are not both in Y by Lemma 3.2. In particular, $q(X \cap Y) = X_1 \cap Y_1$. The case in which q is a Y-fold, but not an X-fold is handled similarly. This covers the first case.

Now suppose that q is neither an X-fold or a Y-fold. Then certainly $e_1 \subset X - X \cap Y$ and $e_2 \subset Y - X \cap Y$ (or vice versa). Since q is ordinary, $\check{q}_X, \check{q}_Y$

must be homeomorphism. Lemma 3.2 then implies that either $w_1, w_2 \in X \cap Y$. Thus, the two remaining cases correspond to whether $w_1 \neq w_2$ or $w_1 = w_2$. In the first case we have that $X_1 \cap Y_1$ consists of $q(X \cap Y)$ and $q(e_1) \sqcup q(w_1)$ while in the second case we have that $X_1 \cap Y_1$ consists of $q(X \cap Y)$ and $q(e_1)$. □

Lemma 3.4. *If q is an ordinary fold, then (Z_1, Y_1, X_1) is ψ-bi-invariant for H and*

$$\begin{aligned}\mathrm{rr}(Z_1, Y_1) &\leqslant \mathrm{rr}(Z, X),\\ \mathrm{rr}(Z_1, X_1) &\leqslant \mathrm{rr}(Z, Y).\end{aligned}$$

Moreover if $\mathrm{rr}(Z, X), \mathrm{rr}(Z, Y) < \infty$, then

$$\mathrm{rr}(Z_1, Y_1) + \mathrm{rr}(Z_1, X_1) < \mathrm{rr}(Z, Y) + \mathrm{rr}(Z, X)$$

if and only if $e_1 \cup e_2$ is a bigon not in $X \cap Y$.

Proof. By definition of ordinary folds, we have $Z^{\#} = Z_1^{\#}$ and $X^{\#} = X_1^{\#}, Y^{\#} = Y_1^{\#}$. In the first two cases of Lemma 3.3, we also have $(X \cap Y)^{\#} = (X_1 \cap Y_1)^{\#}$ and so in these cases (Z_1, Y_1, X_1) is ψ-bi-invariant since (Z, Y, X) was. Now suppose that we are in the third case of Lemma 3.3; that is, $e_1 \subset X - X \cap Y$, $e_2 \subset Y - X \cap Y$ (or vice versa) and $w_1 = w_2$. Let $\alpha, \beta \colon I \to X \cap Y$ be paths connecting the basepoint with $v, w_1 = w_2$ respectively. Then denote by $\delta_X = f_Z \circ (\alpha * e_1 * \overline{\beta})$ and $\delta_Y = f_Z \circ (\alpha * e_2 * \overline{\beta})$. We have that $[\delta_X] = [\delta_Y]$ and $(X_1 \cap Y_1)^{\#} = \langle (X \cap Y)^{\#}, [\delta_X] \rangle$. Since $[\delta_X] \in X^{\#}$, $[\delta_Y] \in Y^{\#}$, we see that $\psi([\delta_X]) \in Y^{\#}$ and $\psi^{-1}([\delta_X]) \in X^{\#}$ and so (Z_1, Y_1, X_1) is also ψ-bi-invariant in this case.

Lemma 3.1 gives us the desired bounds on reduced ranks.

Now suppose that $\mathrm{rr}(Z, X), \mathrm{rr}(Z, Y) < \infty$. Using the equality $\mathrm{rr}(Z, X) = -\chi(Z, X)$, the only way that we can have $\mathrm{rr}(Z_1, X_1) < \mathrm{rr}(Z, X)$ is if $Z_1 - X_1$ is obtained from $Z - X$ by removing an edge. This only can happen if $e_1, e_2 \subset Z - X$ and $w_1 = w_2$ or if $e_1 \subset X - X \cap Y$, $e_2 \subset Y$ (or vice versa). By Lemma 3.3, the latter case can only happen if $w_1 = w_2$ and so q folds a bigon not in $X \cap Y$. The case in which $\mathrm{rr}(Z_1, Y_1) < \mathrm{rr}(Z, Y)$ is handled similarly. □

3.3. **Single exceptional folds.** When q is a single exceptional fold, the new graph triple (Z_1, Y_1, X_1) may no longer be ψ-bi-invariant. This will be fixed in the subsequent lemmas. First, we identify precisely when a fold is a single exceptional fold.

Lemma 3.5. *If q is a single exceptional fold, then, up to exchanging e_1 and e_2, one of the following holds:*

(1) *q is a single X-exceptional fold. In this case, either*
   (a) *$\check{q}_X$ is a fold and so $\check{q}_{X \cap Y}$ identifies a pair of vertices and $X_1 \cap Y_1 = q(X \cap Y)$, or*
   (b) *$\check{q}_X$ is a homeomorphism and so $e_1 \subset X - X \cap Y$, $e_2 \subset Y - X \cap Y$, $w_1 \in X \cap Y$, $w_2 \in Y - X \cap Y$ and $X_1 \cap Y_1 = q(X \cap Y) \sqcup q(e_1)$.*

(2) *q is a single Y-exceptional fold. In this case, either*
   (a) *$\check{q}_Y$ is a fold and so $\check{q}_{X \cap Y}$ identifies a pair of vertices and $X_1 \cap Y_1 = q(X \cap Y)$, or*

(b) $\check{q}_Y$ *is a homeomorphism and so* $e_1 \subset X - X \cap Y$, $e_2 \subset Y - X \cap Y$, $w_1 \in X - X \cap Y$, $w_2 \in X \cap Y$ *and* $X_1 \cap Y_1 = q(X \cap Y) \sqcup q(e_1)$.

Proof. The four subcases are mutually exclusive and cover all possibilities by definition of single exceptional folds.

By Lemma 3.2, if $\check{q}_Y$ is not a fold or a homeomorphism, then $w_1 \neq w_2$, $w_1, w_2 \in Y$ and $e_1 \subset X - X \cap Y$ or $e_2 \subset X - X \cap Y$.

If additionally, $\check{q}_X$ is a fold, then we must have $e_1, e_2 \subset X$. In this case, $\check{q}_{X \cap Y}$ identifies a pair of vertices and $X_1 \cap Y_1 = q(X \cap Y)$. This is Case (1a).

If instead $\check{q}_X$ is a homeomorphism, then $e_1 \subset X - X \cap Y$ and $e_2 \subset Y - X \cap Y$ (or vice versa) and $X_1 \cap Y_1$ is obtained from $q(X \cap Y)$ by adjoining the edge $q(e_1)$. This is Case (1b).

Cases (2a) and (2b) are handled similarly. $\square$

In the next lemma we show how to fix ψ-bi-invariance of the triple (Z_1, Y_1, X_1) when q is a (single) X-exceptional fold. The case of a Y-exceptional fold is handled in the exact same way by exchanging the roles of X and Y and replacing ψ with ψ^{-1}.

Lemma 3.6. *Suppose that q is a single X-exceptional fold.*

- *If q is as in Case* (1a) *from Lemma 3.5, let $\alpha, \beta \colon I \to X \cap Y$ be an immersed path connecting v_Z with w_1, w_2 respectively and denote by $\delta = [(f_Z \circ \alpha) * (f_Z \circ \overline{\beta})] \in \mathbb{F}$.*
- *If q is as in Case (1b) from Lemma 3.5, let $\alpha, \beta \colon I \to X \cap Y$ be immersed paths connecting v_Z with w_1 and v respectively and denote by $\delta = [(f_Z \circ \alpha) * (f_Z \circ \overline{e_1}) * (f_Z \circ \overline{\beta})] \in \mathbb{F}$.*

If we put

$$X_2 = \begin{cases} X_1 & \textit{if } \psi^{-1}(\delta) \in X^{\#} \\ X_1 \vee \Gamma(\psi^{-1}(\delta)) & \textit{otherwise} \end{cases}$$
$$Y_2 = Y_1$$
$$Z_2 = X_2 \cup Y_2$$

then (Z_2, Y_2, X_2) is a ψ-bi-invariant graph triple for H and

$$\operatorname{rr}(Z_2, X_2) \leqslant \operatorname{rr}(Z, X)$$
$$\operatorname{rr}(Z_2, Y_2) \leqslant \operatorname{rr}(Z, Y).$$

If $\psi^{-1}(\delta) \in X^{\#}$ and $\operatorname{rr}(Z, Y) < \infty$, then $\operatorname{rr}(Z_2, Y_2) < \operatorname{rr}(Z, Y)$.

Proof. By Lemma 3.5 we have that $(X_1 \cap Y_1)^{\#} = \langle (X \cap Y)^{\#}, \delta \rangle$. Hence, since $X_2 \cap Y_2 = X_1 \cap Y_1$, we have $\psi^{-1}((X_2 \cap Y_2)^{\#}) \leqslant \psi^{-1}(Y_2^{\#}) \leqslant X_2^{\#}$. Thus, $(X_2, X_2 \cap Y_2)$ and (Z_2, Y_2) are ψ^{-1}-invariant. Moreover, since $\check{q}_X$ is a fold or a homeomorphism and (Z, X) was ψ-invariant, it also follows that $\psi((X_2 \cap Y_2)^{\#}) \leqslant \psi(X_2^{\#}) \leqslant Y_2^{\#}$ and so $(Y_2, X_2 \cap Y_2)$ and (Z_2, X_2) are ψ-invariant. Hence, (Z_2, Y_2, X_2) is ψ-bi-invariant.

By Lemma 3.1, we have that $\operatorname{rr}(Z_2, X_2) = \operatorname{rr}(Z_1, X_1) \leqslant \operatorname{rr}(Z, X)$. Since q is an X-exceptional fold, we have that Y_1 is obtained from Y by identifying two vertices and so $Z_1 - Y_1$ is obtained from $Z - Y$ by removing an edge. Thus, $\operatorname{rr}(Z_2, Y_2) \leqslant \operatorname{rr}(Z_1, Y_1) + 1 = \operatorname{rr}(Z, X)$. By construction, the inequality is strict precisely when $\psi^{-1}(\delta) \in X^{\#}$. $\square$

3.4. **Double exceptional folds.** Finally we analyse the double exceptional folds case. Just as in the case of a single exceptional fold, the graph triple (Z_1, Y_1, X_1) may no longer be ψ-bi-invariant.

Lemma 3.7. *If q is a double exceptional fold, then $w_1 \neq w_2$, $e_1 \subset X - Y \cap X$, $e_2 \subset Y - X \cap Y$ and $w_1, w_2 \in X \cap Y$ (or vice versa). Moreover, $\check{q}_{X\cap Y}$ is not a fold or a homeomorphism and $X_1 \cap Y_1 - q(X \cap Y) = q(e_1)$.*

Proof. This follows from applying Lemma 3.2 with $A = X, Y$. ☐

Lemma 3.8. *Suppose that q is a double exceptional fold. Let $\alpha, \beta\colon I \to X \cap Y$ be immersed paths connecting v_Z with w_1 and w_2 respectively and let $\delta = [(f_Z \circ \alpha) * (f_Z \circ \beta)] \in Z^\# \leqslant \mathbb{F}$. If*

$$X_2 = \begin{cases} X_1 & \text{if } \psi^{-1}(\delta) \in X^\# \\ X_1 \vee \Gamma(\psi^{-1}(\delta)) & \text{otherwise} \end{cases}$$

$$Y_2 = \begin{cases} Y_1 & \text{if } \psi(\delta) \in Y^\# \\ Y_1 \vee \Gamma(\psi(\delta)) & \text{otherwise} \end{cases}$$

$$Z_2 = X_2 \cup Y_2$$

then (Z_2, X_2, Y_2) is ψ-bi-invariant for H and

$$\mathrm{rr}(Z_2, Y_2) \leqslant \mathrm{rr}(Z, Y),$$
$$\mathrm{rr}(Z_2, X_2) \leqslant \mathrm{rr}(Z, X).$$

If $\psi^{-1}(\delta) \in X^\#$ and $\mathrm{rr}(Z, Y) < \infty$, then $\mathrm{rr}(Z_2, Y_2) < \mathrm{rr}(Z, Y)$.
If $\psi(\delta) \in Y^\#$ and $\mathrm{rr}(Z, X) < \infty$, then $\mathrm{rr}(Z_2, X_2) < \mathrm{rr}(Z, X)$.

Proof. The proof is very similar to the proof of Lemma 3.6 and so we leave the details to the reader. ☐

3.5. **Folding and adding loops if necessary.** We may now define folding and adding loops if necessary for graph triples.

Definition 3.9. If (Z, Y, X) is a ψ-bi-invariant graph triple, $q\colon Z \to Z_1$ is a fold and (Z_2, Y_2, X_2) is

- the ψ-bi-invariant graph triple (Z_1, Y_1, X_1) if q is an ordinary fold (see Lemma 3.4),
- the ψ-bi-invariant graph triple obtained in Lemma 3.6 if q is a single exceptional fold (after possibly exchanging the roles of X and Y),
- the ψ-bi-invariant graph triple obtained in Lemma 3.8 if q is a double exceptional fold,

then we say that (Z_2, Y_2, X_2) is obtained from (Z, Y, X) by *folding and adding loops if necessary.*

If $\check{q}_X$ is a fold, then we say that (Z_2, Y_2, X_2) is obtained from (Z, Y, X) by *folding X and adding a loop if necessary* (two loops cannot be added). If $\check{q}_Y$ is a fold, then we say that (Z_2, Y_2, X_2) is obtained from (Z, Y, X) by *folding Y and adding a loop if necessary.*

Remark 3.10. If the graph triple (Z_2, Y_2, X_2) is obtained from (Z, Y, X) by folding Y and adding a loop if necessary, then the graph pair $(Y_2, X_2 \cap Y_2)$ is obtained from $(Y, X \cap Y)$ by folding and adding a loop if necessary in

Feighn–Handel's sense [FH99]. Similarly, if the graph triple (Z_2, Y_2, X_2) is obtained from (Z, Y, X) by folding X and adding a loop if necessary, then the graph pair $(X_2, X_2 \cap Y_2)$ is obtained from $(X, X \cap Y)$ by folding and adding a loop if necessary in Feighn–Handel's sense (in the ψ^{-1} direction).

We summarise the most important properties of graph triples obtained by folding and adding loops if necessary that we have established below.

Proposition 3.11. *If $q\colon Z \to Z_1$ is a fold and if (Z_2, Y_2, X_2) is obtained from (Z_1, Y_1, X_1) by adding loops if necessary, then (Z_2, Y_2, X_2) is ψ-bi-invariant graph triple for H and the following holds:*

(1) *We have*

$$\mathrm{rr}(Z_2, Y_2) \leqslant \mathrm{rr}(Z, Y)$$
$$\mathrm{rr}(Z_2, X_2) \leqslant \mathrm{rr}(Z, X)$$

In particular, if (Z, Y, X) is minimal, then so is (Z_2, Y_2, X_2).

(2) *If one of the following holds:*
- *q is a subgraph fold with $e_1 \cup e_2$ a bigon not contained in $X \cap Y$.*
- *q is a single exceptional fold and no loop is added.*
- *q is a double exceptional fold and at most one loop is added.*

Then

$$\mathrm{rr}(Z_2, Y_2) + \mathrm{rr}(Z_2, X_2) < \mathrm{rr}(Z, Y) + \mathrm{rr}(Z, X)$$

if $\mathrm{rr}(Z, Y), \mathrm{rr}(Z, X) < \infty$.

(3) *If Z is finite and q is not as in (2) from Lemma 3.3, then either the number of vertices in $X_2 \cap Y_2$ is strictly less than the number of vertices in $X \cap Y$, or the number of edges in Z_2 is strictly less than the number of edges in Z.*

3.6. **Tightening X, Y and Z.** We are now ready to describe our variation of the Feighn–Handel tightening procedure for graph triples.

Definition 3.12. [*Tightening X, tightening Y*] Let (Z, Y, X) be a ψ-bi-invariant graph triple. If X is tight then do nothing. If $X \cap Y$ is not tight, then fold a pair of edges in $X \cap Y$. If $X \cap Y$ is tight, but X is not tight, then fold a pair of edges in X and add a loop if necessary. Repeat until we obtain a new ψ-bi-invariant graph triple $(\check{Z}, \check{Y}, \check{X})$ with $\check{X}$ tight. We say $(\check{Z}, \check{Y}, \check{X})$ is obtained from (Z, Y, X) by *tightening X*. We may similarly define the *tightening Y* procedure.

Remark 3.13. If (Z, Y, X) is a ψ-bi-invariant graph triple and $(\check{Z}, \check{Y}, \check{X})$ is obtained from (Z, Y, X) by the tightening Y procedure, then the graph pair $(\check{Y}, \check{X} \cap \check{Y})$ is obtained from the graph pair $(Y, X \cap Y)$ by tightening in the sense of Feighn–Handel [FH99], see Remark 3.10. The same may be said about the graph pair $(\check{X}, \check{X} \cap \check{Y})$, except that we replace the automorphism ψ with the inverse ψ^{-1}.

One could derive the following from Remark 3.10 and the remarks in [FH99, Definition 4.8]. We prove it directly using Proposition 3.11.

Lemma 3.14. *If (Z, Y, X) is a finite ψ-bi-invariant graph triple for H, then the tightening of X procedure terminates in a ψ-bi-invariant graph* *triple* $(\check{Z}, \check{Y}, \check{X})$ *for* H *with* $\check{X}$ *tight. Similarly, the tightening of* Y *procedure terminates in a* ψ*-bi-invariant graph triple* $(\check{Z}, \check{Y}, \check{X})$ *for* H *with* $\check{Y}$ *tight. In both cases, we have* $\mathrm{rr}(\check{Z}, \check{X}) \leqslant \mathrm{rr}(Z, X)$*,* $\mathrm{rr}(\check{Z}, \check{Y}) \leqslant \mathrm{rr}(Z, Y)$*.*

Proof. The fact that after each folding and adding loops if necessary we obtain a new finite ψ-bi-invariant graph triple for H is by Proposition 3.11. Letting

$$(Z_0, Y_0, X_0) \xrightarrow{q_0} (Z_1, Y_1, X_1) \xrightarrow{q_1} \dots \xrightarrow{q_{k-1}} (Z_k, Y_k, X_k) \xrightarrow{q_k} \dots$$

be the sequence of folds and adding loops if necessary from the procedure, where $(Z, Y, X) = (Z_0, Y_0, X_0)$. By Proposition 3.11, we have that $q_i(X_i \cap Y_i) = X_{i+1} \cap Y_{i+1}$ for all i. If q_i is an exceptional fold, then this implies that the number of vertices in $X_{i+1} \cap Y_{i+1}$ is strictly less than the number of vertices in $X_i \cap Y_i$. Hence, only finitely many exceptional folds are performed. If q_i is not an exceptional fold, then no loop is added and so the number of edges in Z_{i+1} is strictly less than the number of edges in Z_i. Thus, the procedure must terminate after finitely many steps. □

Definition 3.15. [*Tightening* X *and* Y] Let (Z, Y, X) be a ψ-bi-invariant graph triple. If X and Y are tight, then do nothing. If X is not tight, then tighten X or if Y is not tight, then tighten Y. Repeat until we obtain a new ψ-bi-invariant triple $(\check{Z}, \check{Y}, \check{X})$ with both $\check{Y}$ and $\check{X}$ tight.

Lemma 3.16. *If* (Z, Y, X) *is a finite* ψ*-bi-invariant graph triple for* H*, then the tightening of* X *and* Y *procedure terminates in a* ψ*-bi-invariant graph triple* $(\check{Z}, \check{Y}, \check{X})$ *for* H *with* $\check{Y}$ *and* $\check{X}$ *tight and with* $\mathrm{rr}(\check{Z}, \check{X}) \leqslant \mathrm{rr}(Z, X)$*,* $\mathrm{rr}(\check{Z}, \check{Y}) \leqslant \mathrm{rr}(Z, Y)$*.*

Proof. The proof is almost identical to that of Lemma 3.14. We simply need to note that throughout the procedure, only finitely many exceptional folds are performed (using Proposition 3.11) and so after all such fold is performed, each subsequent fold is ordinary and strictly decreases the number of edges. □

Combining Remark 3.13, [FH99, Proposition 5.4] and Lemma 3.16 we obtain the following.

Proposition 3.17. *Let* (Z, Y, X) *be a finite* ψ*-bi-invariant graph triple for* H *with* $\pi_1(X \cap Y, v_Z) \to \pi_1(R, v_R)$ *injective, but* $\pi_1(X, v_Z) \to \pi_1(R, v_R)$ *not injective. If* $(\check{Z}, \check{Y}, \check{X})$ *is obtained from* (Z, Y, X) *by the tightening* X *procedure, then* $\mathrm{rr}(\check{Z}, \check{Y}) < \mathrm{rr}(Z, Y)$ *and* $\mathrm{rr}(\check{Z}, \check{X}) \leqslant \mathrm{rr}(Z, X)$*.*

There is one final natural tightening procedure (which we shall not use) that one can define.

Definition 3.18. [*Tightening* Z] Let (Z, Y, X) be a ψ-bi-invariant graph triple. If Z is tight then do nothing. If X is not tight, then tighten X or if Y is not tight, then tighten Y. If X and Y are tight, but Z is not tight, then fold and add loops if necessary. Repeat until we obtain a tight ψ-bi-invariant triple $(\check{Z}, \check{Y}, \check{X})$.

We do not know whether there is an analogue of Lemma 3.16 for the tightening Z procedure.

Question 3.19. Does the tightening Z procedure always terminate after finitely many steps when given a finite ψ-bi-invariant graph triple?

In [Lin25] direct limits of maps of graph pairs are considered. Using the same ideas, one could take an appropriate direct limit of graph triples to obtain a tight ψ-bi-invariant graph triple for a given finitely generated subgroup. It will still have finite relative rank, but may not be finite itself.

4. Main theorems

4.1. **Presentations of free-by-cyclic groups.** Using Proposition 3.17 we may obtain a refinement of Theorem 2.2 for graph triples.

Theorem 4.1. *Let (Z, Y, X) be a finite ψ-bi-invariant graph triple for H with $(f_Z)_*$ injective and let $C, D \leqslant Z^\#$ so that $Z^\# = X^\# * C = D * Y^\# = D * (X \cap Y)^\# * C$. The following are equivalent:*

(1) *The triple (Z, Y, X) is minimal.*
(2) *The ψ-invariant pair (Z, X) is minimal.*
(3) *The ψ^{-1}-invariant pair (Z, Y) is minimal.*
(4) *The map*
$$\theta_n \colon \pi_1\left(\bigvee_{i=0}^{n} \Gamma(\psi^{-i}(D)) \vee (X \cap Y) \vee \bigvee_{i=0}^{n} \Gamma(\psi^{i}(C)), v_Z\right) \to \pi_1(R, v_R)$$
is injective for all $n \geqslant 0$.
(5) *We have*
$$H \cong \langle Z^\#, t \mid t^{-1}xt = \phi(x), \forall x \in X^\# \rangle.$$
where $\phi\colon X^\# \to Y^\#$ is the isomorphism given by $\phi = \psi\,|_{X^\#}$.

In particular, if any of the above hold, then $-\chi(H) = \mathrm{rr}(Z, X) = \mathrm{rr}(Z, Y)$.

Proof. The fact that Item 2 and Item 3 are both equivalent to Item 5 is Theorem 2.2 combined with Lemma 2.3. Using the fact that:
$$D * (X \cap Y)^\# = X^\# = \langle (X \cap Y)^\#, \psi^{-1}((X \cap Y)^\#)$$
$$(X \cap Y)^\# * C = Y^\# = \langle (X \cap Y)^\#, \psi((X \cap Y)^\#) \rangle$$
and that $\psi(D * (X \cap Y)^\#) = (X \cap Y)^\# * C$ by Lemma 2.3, by induction on n we see that
$$\pi_1\left(\bigvee_{i=0}^{n} \Gamma(\psi^{-i}(D)) \vee (X \cap Y) \vee \bigvee_{i=0}^{n} \Gamma(\psi^{i}(C)), v_X\right)$$
$$= \psi^{-n}\left(\pi_1\left((X \cap Y) \vee \bigvee_{i=0}^{2n} \Gamma(\psi^{i}(C)), v_X\right)\right),$$
and so Theorem 2.2 also implies that Item 5 is equivalent to Item 4. Finally, repeated application of Proposition 3.17 shows that Item 1 implies Item 4. Finally, since $-\chi(H) = \mathrm{rr}(Z, X) = \mathrm{rr}(Z, Y)$ is a group invariant, we see that Item 5 implies Item 1 and we are done. $\square$

4.2. **Proof of Theorem 1.1.** It is clear that G is finitely generated if $\mathbb{F}$ admits such a free product decomposition so we only focus on proving the other direction.

Let R be a rose graph with $\pi_1(R, v_R) \cong \mathbb{F}$. Let (Z, Y, X) be a finite tight minimal ψ-bi-invariant graph triple for $\mathbb{F} \rtimes_\psi \mathbb{Z}$, with $f_Z\colon (Z, v_Z) \to (R, v_R)$ the associated graph map. Applying Theorem 4.1 we obtain that

$$\mathbb{F} = \left(\ast_{i=0}^{-\infty}\psi^i(D)\right) * E * \left(\ast_{i=0}^{\infty}\psi^i(C)\right).$$

where $E = (X \cap Y)^{\#}$. For each $j \geqslant 0$, denote by

$$\mathbb{F}_j = \left(\ast_{i=0}^{-j}\psi^i(D)\right) * E * \left(\ast_{i=0}^{\infty}\psi^i(C)\right).$$

For each $j \geqslant 0$, we have

$$\psi^{-j}(\mathbb{F}_0) = \mathbb{F}_j = \psi^{-j}(E) * \left(\ast_{i=-j}^{\infty}\psi^i(C)\right).$$

Since $\mathbb{F}_j$ is a free factor of $\mathbb{F}$, we see that

$$\mathop{\ast}_{i \geqslant -j} \psi^i(C)$$

is a free factor of $\mathbb{F}$ for all $j \geqslant 0$ and so

$$\mathop{\ast}_{i \in \mathbb{Z}} \psi^i(C)$$

is a free factor of $\mathbb{F}$. Letting $A \leqslant \mathbb{F}$ be a free factor such that $\mathbb{F} = A * \left(\ast_{i\in\mathbb{Z}}\,\psi^i(C)\right)$, it suffices to show that A is finitely generated to complete the proof as we put $C_0 = C$.

Let A' be any finitely generated free factor of A. Then there is some $j \geqslant 0$ such that $A' \leqslant \mathbb{F}_j$. Since $A' * \ast_{i \geqslant -j}\,\psi^i(C)$ is a free factor of $\mathbb{F}$, it is also a free factor of $\mathbb{F}_j$. Since $\mathbb{F}_j/\langle\!\langle \ast_{i\geqslant -j}\,\psi^i(C)\rangle\!\rangle \cong \psi^{-j}(E)$, and since the projection map $\mathbb{F}_j \to \psi^{-j}(E)$ sends the free factor A' injectively to a free factor of $\psi^{-j}(E)$, it follows that $\mathrm{rk}(A') \leqslant \mathrm{rk}(E)$. Since A' was an arbitrary finitely generated free factor of A, we see that $\mathrm{rk}(A) \leqslant \mathrm{rk}(E)$ and so A is finitely generated as claimed.

Remark 4.2. Note that, using the notation of Theorem 1.1, $\chi(\mathbb{F} \rtimes_\psi \mathbb{Z}) = -\mathrm{rk}(C_0)$. In particular, the rank of C_0 is an invariant of the free-by-cyclic group. The rank of the group A is certainly not an invariant since there are groups with several epimorphisms to $\mathbb{Z}$ with finitely generated free kernels of different ranks. However, we do not know whether $\mathrm{rk}(A)$ is an invariant of the automorphism ψ. When $C_0 = 1$, then this is trivially the case as $\mathbb{F} = A$.

4.3. **The geometry of free-by-cyclic groups.** Combining Theorem 1.1 with the main result from [Lin25] (see also [Lin25, Theorem 5.1] and the discussion preceding it) we may also state a geometric refinement. The reader is invited to consult Hruska's article [Hru10] for more information on relatively hyperbolic groups and relatively quasi-convex subgroups.

Theorem 4.3. *Let $\mathbb{F}$ be a free group, let $\psi \in \mathrm{Aut}(\mathbb{F})$ be an automorphism and suppose that $G = \mathbb{F} \rtimes_\psi \mathbb{Z}$ is finitely generated. Then $\mathbb{F}$ admits a free product decomposition*

$$\mathbb{F} = A_1 * \ldots * A_n * B * (\ast_{i\in\mathbb{Z}} C_i)$$

with the following properties:

(1) *$A_1, \dots, A_n, B, C_0$ are each finitely generated and $C_i = \psi^i(C_0)$ for each $i \in \mathbb{Z}$.*
(2) *There is a constant k so that for any finitely generated subgroup $H \leqslant \mathbb{F}$, if $\psi^m(H)$ intersects a conjugate of H non-trivially for some $m \geqslant k$, then H is conjugate into some A_i.*
(3) *There is an element $\sigma \in \mathrm{Sym}(n)$ and elements $f_1, \dots, f_n \in \mathbb{F}$ so that $\psi(A_i)^{f_i} = A_{\sigma(i)}$.*
(4) *For each i, if $m_i \geqslant 1$ is minimal so that $\sigma^{m_i}(i) = i$, then the group*
$$H_i = \langle A_i, tf_i tf_{\sigma(i)} \dots tf_{\sigma^{m_i-1}(i)} \rangle$$
is {fg free}-by-cyclic.
(5) *If $\mathcal{P}$ is a collection of conjugacy class representatives of the subgroups H_i, then any subgroup of G isomorphic to a {fg free}-by-cyclic group is conjugate into a unique $P \in \mathcal{P}$.*

Moreover, $(G, \mathcal{P})$ is relatively hyperbolic and locally relatively quasi-convex.

4.4. **Proof of Theorem 1.2.** Let $\mathbb{F} = A * (*_{i \in \mathbb{Z}} C_i)$ be the free product decomposition from Theorem 1.1. Let D be a finitely generated free group of rank $\mathrm{rk}(C_0)$ and let $\phi \in \mathrm{Aut}(D)$ be an atoroidal automorphism so that $D \rtimes_\phi \mathbb{Z}$ is hyperbolic by a result of Brinkmann [Bri00]. Choose any free generating set $\mathcal{D} \subset D$ for D so that $\mathcal{D} \cup \{t\}$ is a generating set for $D \rtimes_\phi \mathbb{Z}$. For some $n \gg 1$, if $\mathcal{D}_n$ denotes the set obtained from $\mathcal{D}$ by elevating each generator to the power n, the subgroup $\langle \mathcal{D}_n, t^n \rangle$ is a quasi-convex free group of rank $|\mathcal{D}_n| + 1$ by a result of Arzhantseva [Arz01, Theorem 1]. Hence, after replacing ϕ with ϕ^n, we may assume that $\langle \mathcal{D}_n, t \rangle$ is a quasi-convex free subgroup and hence that
$$\langle \mathcal{D}_n, t \rangle \cap D \cong *_{i \in \mathbb{Z}} \langle \phi^i(\mathcal{D}_n) \rangle.$$
Choose any isomorphism $\lambda_0 \colon C_0 \to \langle \mathcal{D}_n \rangle$ and define $\lambda_i \colon C_i \to \langle \phi^i(\mathcal{D}_n) \rangle$ as $\lambda_i = \phi^i \circ \lambda_0 \circ \psi^{-i} \mid C_i$. This then extends to an isomorphism
$$\lambda \colon *_{i \in \mathbb{Z}} C_i \to *_{i \in \mathbb{Z}} \langle \phi^i(\mathcal{D}_n) \rangle$$
so that $\lambda \circ \psi = \phi \circ \lambda$ when restricted to $*_{i \in \mathbb{Z}} C_i$.

Now let $\iota \colon \mathbb{F} \to F = A * D$ be the homomorphism given by extending the inclusion of A and λ. Note that $\iota(\mathbb{F}) \cap D = *_{i \in \mathbb{Z}} C_i$. By construction, the homomorphism $\theta \colon A * D \to A * D$ given by $\theta \mid A = \iota \circ \psi \mid A$ and $\theta \mid D = \phi$ is an automorphism which restricts to ψ on $\iota(\mathbb{F})$. Hence, $\mathbb{F} \rtimes_\psi \mathbb{Z}$ embeds into $(A * D) \rtimes_\theta \mathbb{Z}$ via the homomorphism extending ι and the identity map on $\mathbb{Z}$.

Now suppose that G is hyperbolic. Let $\mathbb{F} = (*_{i=1}^n A_i) * B * (*_{i \in \mathbb{Z}} C_i)$ be the free product decomposition from Theorem 4.3. By Theorem 4.3, G is hyperbolic relative to a finite collection of {finitely generated free}-by-$\mathbb{Z}$ subgroups which each intersect $\mathbb{F}$ in conjugates of some A_i. In particular, the subgroup $*_{i \in \mathbb{Z}} C_i$ intersects each (conjugate of each) peripheral trivially. Hence, the subgroup $\langle C_0, t \rangle$ intersects each (conjugate of each) peripheral in a subgroup isomorphic to $\mathbb{Z}$ or the trivial group. By Theorem 4.3, $\langle C_0, t \rangle$ is relatively quasi-convex in G. Since cyclic subgroups of hyperbolic groups are quasi-convex and quasi-convex subgroups of hyperbolic groups are precisely the undistorted subgroups, a result of Hruska [Hru10, Theorem 1.4] implies that $\langle C_0, t \rangle$ is quasi-convex in G. We claim further that $\langle C_0, t \rangle$ is malnormal in G. Indeed, let $t^k f \in G$ be any element with $f \in \mathbb{F}$ (every element is of this form). Then $\langle C_0, t\rangle^{t^k} = \langle C_0, t\rangle$ and $\langle C_0, t\rangle^f \cap \langle C_0, t\rangle = 1$ if $f \notin \ast_{i\in\mathbb{Z}} C_i$. Hence, $\langle C_0, t\rangle^{t^k f} \cap \langle C_0, t\rangle = 1$ if $f \notin \ast_{i\in\mathbb{Z}} C_i$. Since $t^k f \notin \langle C_0, t\rangle$ only if $f \notin \ast_{i\in\mathbb{Z}} C_i$, it follows that $\langle C_0, t\rangle$ is malnormal in G. By our construction of $F \rtimes_\theta \mathbb{Z}$ above, we see that

$$F \rtimes_\theta \mathbb{Z} = \mathbb{F} \rtimes_\psi \mathbb{Z} \underset{\langle C_0, t\rangle = \langle \mathcal{D}_n, t\rangle}{\ast} D \rtimes_\phi \mathbb{Z}.$$

Since $\langle C_0, t\rangle$ is malnormal and quasi-convex in G and since $\langle \mathcal{D}_n, t\rangle$ is quasi-convex in $D \rtimes_\phi \mathbb{Z}$, we may use a result of Kapovich [Kap01] to see that $\mathbb{F} \rtimes_\psi \mathbb{Z}$ is quasi-convex in $F \rtimes_\theta \mathbb{Z}$ as required.

Remark 4.4. Replacing the use of Theorem 1.1 with [Lin25, Corollary 4.7] and making other minor modifications to the proof of Theorem 1.2, one may obtain the analogous result for mapping tori of free groups. This was already obtained by Chong–Wise in [CW24].

References

[Ago13] Ian Agol, *The virtual Haken conjecture*, Doc. Math. **18** (2013), 1045–1087, With an appendix by Agol, Daniel Groves, and Jason Manning. MR 3104553

[Arz01] G. N. Arzhantseva, *On quasiconvex subgroups of word hyperbolic groups*, Geom. Dedicata **87** (2001), no. 1-3, 191–208. MR 1866849

[Bri00] P. Brinkmann, *Hyperbolic automorphisms of free groups*, Geom. Funct. Anal. **10** (2000), no. 5, 1071–1089. MR 1800064

[CW24] Hip Kuen Chong and Daniel T. Wise, *Embedding partial HNN extensions in ascending HNN extensions*, 2024, arXiv:2408.00453.

[FH99] Mark Feighn and Michael Handel, *Mapping tori of free group automorphisms are coherent*, Ann. of Math. (2) **149** (1999), no. 3, 1061–1077. MR 1709311

[Ger96] S. M. Gersten, *Subgroups of word hyperbolic groups in dimension 2*, J. Lond. Math. Soc., II. Ser. **54** (1996), no. 2, 261–283 (English). MR 1405055

[Hru10] G. Christopher Hruska, *Relative hyperbolicity and relative quasiconvexity for countable groups*, Algebr. Geom. Topol. **10** (2010), no. 3, 1807–1856. MR 2684983

[HW08] Frédéric Haglund and Daniel T. Wise, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), no. 5, 1551–1620. MR 2377497

[HW15] Mark F. Hagen and Daniel T. Wise, *Cubulating hyperbolic free-by-cyclic groups: the general case*, Geom. Funct. Anal. **25** (2015), no. 1, 134–179. MR 3320891

[Kap01] Ilya Kapovich, *The combination theorem and quasiconvexity*, Internat. J. Algebra Comput. **11** (2001), no. 2, 185–216. MR 1829050

[Lin25] Marco Linton, *The geometry of subgroups of mapping tori of free groups*, 2025, arXiv:2510.03145.

[Sta83] John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565. MR 695906

[WY25] Xiaolei Wu and Shengkui Ye, *Some questions related to free-by-cyclic groups and tubular groups*, 2025, arXiv:2504.14192.

Instituto de Ciencias Matemáticas, CSIC-UAM-UC3M-UCM, Madrid, Spain
Email address: `marco.linton@icmat.es`