RELATIVE NIELSEN REALISATION FOR FREE PRODUCTS

SEBASTIAN HENSEL AND DAWID KIELAK

Abstract. We generalise the Karrass–Pietrowski–Solitar and the Nielsen realisation theorems from the setting of free groups to that of free products. As a result, we obtain a fixed point theorem for finite groups of outer automorphisms acting on the relative free splitting complex of Handel–Mosher, as well as a relative version of the Nielsen realisation theorem, which is new even for free groups.

The proofs rely on a new version of Stallings' theorem on groups with at least two ends, in which some control over the behaviour of virtual free factors is gained.

1. Introduction

In the 1980's Marc Culler [Cul], Dmitry Khramtsov [Khr], and Bruno Zimmermann [Zim] independently proved the Nielsen Realisation theorem for free groups. It states that every finite subgroup $H < \mathrm{Out}(F_n)$ can be realised as a group of automorphisms of a graph with fundamental group F_n.

All three proofs rely in a fundamental way on a result of Karrass–Pietrowski–Solitar [KPS], which states that every finitely generated virtually free group acts on a tree with finite edge and vertex stabilisers. In the language of Bass–Serre theory, it amounts to saying that such a virtually free group is a fundamental group of a graph of groups with finite edge and vertex groups (compare [HOP] for a different approach to Nielsen realisation).

This result of Karrass–Pietrowski–Solitar in turn relies on the celebrated theorem of Stallings on groups with at least two ends [Sta1, Sta2], which states that any finitely generated group with at least two ends splits over a finite group, that is it acts on a tree with a single edge orbit and finite edge stabilisers, or equivalently, that it is a fundamental group of a graph of groups with a single edge and a finite edge group.

The purpose of this article is to generalise these three results to the setting of a free product

$$A = A_1 * \dots A_n * B$$

in which we (usually) require the factors A_i to be finitely generated torsion-free, and B to be a finitely generated free group. Consider any finite group

The second author is supported by the SFB 701.

H acting on A by outer automorphisms in a way preserving the given free-product decomposition (up to conjugation). We then obtain a corresponding group extension

$$1 \to A \to \overline{A} \to H \to 1$$

In this setting we prove (for formal statements, see the appropriate sections)

Relative Stallings' theorem (Theorem 2.7): $\overline{A}$ splits over a finite group, in such a way that each A_i fixes a vertex in the associated action on a tree.

Relative Karrass–Pietrowski–Solitar theorem (Theorem 3.4): $\overline{A}$ acts on a tree with finite edge stabilisers, and with each A_i fixing a vertex of the tree, and with, informally speaking, all other vertex groups finite.

Relative Nielsen realisation theorem (Theorem 5.4): Suppose that we are given complete non-positively curved (i.e. locally CAT(0)) spaces X_i realising the induced actions of H on the factors A_i. Then the action of H can be realised by a complete non-positively curved space X; in fact X can be chosen to contain the X_i in an equivariant manner.

We emphasise that such a relative Nielsen realisation is new even if all A_i are free groups. It is used as a crucial ingredient in [HK] by the same authors, where Nielsen realisation for some classes of right-angled Artin groups is proven.

The classical Nielsen realisation for graphs immediately implies that a finite subgroup $H < \mathrm{Out}(F_n)$ fixes points in the Culler–Vogtmann Outer Space (defined in [CV]), as well as in the complex of free splittings of F_n (which is a simplicial closure of Outer Space).

As another application of the work in this article, we similarly obtain a fixed point statement (Corollary 4.1) for the graph of relative free splittings defined by Handel and Mosher [HM].

Throughout the paper, we are going to make liberal use of the standard terminology of graphs of groups. The reader may find all the necessary information in Serre's book [Ser]. We are also going to make use of standard facts about CAT(0) and non-positively curved (NPC) spaces; the standard reference here is the book by Bridson–Haefliger [BH].

Acknowledgements. The authors would like to thank Karen Vogtmann for discussions and suggesting the statement of relative Nielsen realisation for free groups.

2. Relative Stallings' theorem

In this section we will prove the relative version of Stallings' theorem. Before we can begin with the proof, we need a number of definitions to formalize the notion of a free splitting that is preserved by a finite group action.

When talking about free factor decompositions $A = A_1 * \cdots * A_n * B$ of some group A, we will always assume that at least two of the factors (including B) are non-trivial.

Definition 2.1. Suppose that $\phi\colon H \to \mathrm{Out}(A)$ is a homomorphism with a finite domain. Let $A = A_1 * \cdots * A_n * B$ be a free factor decomposition of A. We say that this decomposition is *preserved by* H if and only if for every i and every $h \in H$, there is some j such that $h(A_i)$ is conjugate to A_j.

We say that a factor A_i is *minimal* if and only if for any $h \in H$ the fact that $h(A_i)$ is conjugate to A_j implies that $j \geqslant i$.

Remark 2.2. Note that when the decomposition is preserved, we obtain an induced action $H \to \mathrm{Sym}(n)$ on the indices $1, \ldots, n$. We may thus speak of the stabilisers $\mathrm{Stab}_H(i)$ inside H. Furthermore, we obtain an induced action

$$\mathrm{Stab}_H(i) \to \mathrm{Out}(A_i)$$

The minimality of factors is merely a way of choosing a representative of each H orbit in the action $H \to \mathrm{Sym}(n)$.

Remark 2.3. Given an action $\phi\colon H \to \mathrm{Out}(A)$, with ϕ injective and A with trivial centre, we can define $\overline{A} \leqslant \mathrm{Aut}(A)$ to be the preimage of $H = \mathrm{im}\,\phi$ under the natural map $\mathrm{Aut}(A) \to \mathrm{Out}(A)$. We then note that $\overline{A}$ is an extension of A by H:

$$1 \to A \to \overline{A} \to H \to 1$$

and the left action of H as outer automorphism agrees with the left conjugation action inside the extension $\overline{A}$.

Observe that then for each i we also obtain an extension

$$1 \to A_i \to \overline{A_i} \to \mathrm{Stab}_H(i) \to 1$$

where $\overline{A_i}$ is the subgroup of $\overline{A}$ generated by A_i and a set of elements in $\overline{A}$ which bijectively surject to $\mathrm{Stab}_H(i)$ and are contained in the normaliser of $A_i < A$. Note that as the normaliser of a free factor in a nontrivial free product is that free factor, the subgroup $\overline{A_i}$ does not depend on the choices.

We emphasise that this construction works even when A_i itself is not centre-free. In this case it carries more information than the induced action $\mathrm{Stab}_H(i) \to \mathrm{Out}(A_i)$ (e.g. consider the case of $A_i = \mathbb{Z}$ – there are many different extensions corresponding to the same map to $\mathrm{Out}(\mathbb{Z})$).

We will now begin the proof of the relative version of Stallings' theorem. It will use ideas from both Dunwoody's proof [Dun] and Krön's proof [Krö][1] of Stallings' theorem, which we now recall.

Convention. If E is a set of edges in a graph Θ, we write $\Theta - E$ to mean the graph obtained from Θ by removing the interiors of edges in E.

[1] We warn the reader that Krön's paper contains some arguments which are not entirely correct; we will indicate what changes we make below.

Definition 2.4. Let Θ be a graph. A finite subset E of the edge set of Θ is called a set of *cutting edges* if and only if $\Theta - E$ is disconnected and has at least two infinite components.

A *cut* C is the union of all vertices contained in an infinite connected complementary component of some set of cutting edges. The *boundary* of C consists of all edges with exactly one endpoint in C.

Given two cuts C and D, we call them *nested* if and only if C or its complement C^* is contained in D or its complement D^*. Note that C^* and D^* do not need to be cuts.

We first aim to show the following theorem which is implicit in [Krö].

Theorem 2.5 ([Krö]). *Suppose that Θ is a connected graph on which a group G acts. Let $\mathcal{P}$ be a property of subsets of the edge set of Θ, which is stable under the G-action, taking subsets and unions. If there exists a set of cutting edges with $\mathcal{P}$, then there exists a cut C whose boundary has $\mathcal{P}$, such that the cuts C and $g.C$ are nested for any $g \in G$, and such that C^* is also a cut.*

Sketch of proof. In order to prove this, we recall the following terminology, roughly following Dunwoody. We say that C is a $\mathcal{P}$-cut, if its boundary has $\mathcal{P}$. Say that a $\mathcal{P}$-cut is *$\mathcal{P}$-narrow*, if its boundary contains the minimal number of elements among all $\mathcal{P}$-cuts. Note that for each $\mathcal{P}$-narrow cut C, the complement C^* is also a cut, as otherwise we could remove some edges from the boundary of C and get another $\mathcal{P}$-cut.

Given any edge e with $\mathcal{P}$, there are finitely many $\mathcal{P}$-narrow cuts which contain e in its boundary. This is shown by Dunwoody [Dun, 2.5] for narrow cuts, and the proof carries over to the $\mathcal{P}$-narrow case. Similarly, Krön [Krö, Lemma 2.1] shows this for sets of cutting edges which cut the graph into exactly two connected components, and $\mathcal{P}$-narrow cuts have this property.

Now, consider for each $\mathcal{P}$-cut C the number $m(C)$ of $\mathcal{P}$-cuts which are not nested with C (this is finite by the remark above). Call a $\mathcal{P}$-cut *optimally nested* if $m(C)$ is smallest amongst all $\mathcal{P}$-cuts. The proof of Theorem 3.3 of [Krö] now shows that optimally nested $\mathcal{P}$-cuts are all nested with each other[2]. This shows Theorem 2.5. $\square$

To use that theorem, recall

Theorem 2.6 ([Dun, Theorem 4.1]). *Suppose that there exists a cut C, such that*

(1) C^* is also a cut; and
(2) there exists $g \in G$ such that $g.C$ is properly contained in C or C^*; and
(3) C and $h.C$ are nested for any $h \in G$.

[2] Krön's proof involves intersections of cuts, which by Krön's definition need not be cuts (he assumes that the a cut and its complement is connected) – this does not actually pose a major problem; and does not appear when our definition of a cut is used.

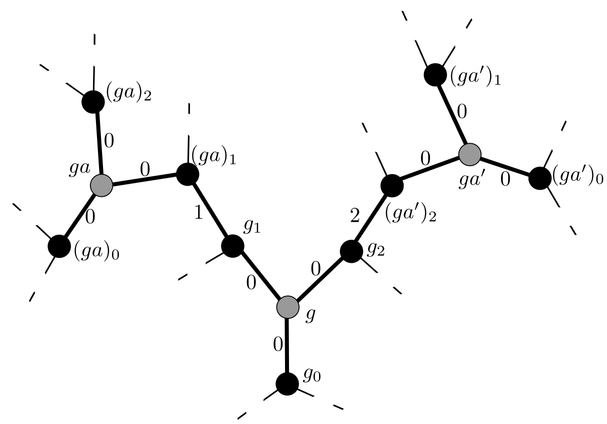

FIGURE 1. A local picture of the graph Θ.

Let E be the boundary of C. Then G splits over the stabiliser of E (which is a finite group), and the stabiliser of any component of $\Theta - G.E$ is contained in a conjugate of a vertex group.

Now we are ready for our splitting result.

Theorem 2.7 (Relative Stallings' Theorem)**.** *Let $\phi\colon H \to \mathrm{Out}(A)$ be a monomorphism with a finite domain. Let $A = A_1 * \cdots * A_n * B$ be a free product decomposition with each A_i and B finitely generated, and suppose that it is preserved by H. Let $\overline{A}$ be the preimage of $H = \operatorname{im}\phi$ in $\mathrm{Aut}(A)$. Then $\overline{A}$ splits over a finite group in such a way that each A_i fixes a vertex in the associated action on a tree.*

Note in particular that the quotient of the associated tree by $\overline{A}$ has a single edge.

Proof. Let $\mathcal{A}_i$ and $\mathcal{B}$ be finite generating sets of A_i and B, respectively (for all $i \leqslant n$). We also choose a finite set $\mathcal{H} \subset \overline{A}$ which maps onto H under the natural epimorphism $\overline{A} \to H$. Note that $\bigcup_i \mathcal{A}_i \cup \mathcal{B} \cup \mathcal{H}$ is a generating set of $\overline{A}$.

We define Θ to be a variation of the (right) Cayley graph of $\overline{A}$ with respect to the generating set $\bigcup_i \mathcal{A}_i \cup \mathcal{B} \cup \mathcal{H}$. Intuitively, every vertex of the Cayley graph will be "blown up" to a finite tree (see Fig. 1). More formally, the vertex set of Θ is

$$V(\Theta) = \overline{A} \sqcup \overline{A} \times \{0, \ldots, n\}$$

We adopt the notation that a vertex corresponding to an element in $\overline{A}$ will simply be denoted by g, whereas a vertex (g, i) in the second part will be denoted by g_i.

We now define the edge set, together with a labelling of the edges by integers $0, 1, \ldots, n$, as follows:

- for each $g \in \overline{A}$ and each $i \in \{0, \ldots, n\}$ we have an edge labelled by 0 connecting g to g_i;
- for each $g \in \overline{A}$, each $i \geqslant 1$ and each $a \in \mathcal{A}_i$, we have an edge labelled by i from g_i to $(ga)_i$;
- for each $g \in \overline{A}$, and each $b \in \mathcal{B} \cup \mathcal{H}$, we have an edge labelled by 0 from g_0 to $(gb)_0$.

The group $\overline{A}$ acts on Θ on the left, preserving the labels. The action is free and co-compact. The graph Θ retracts via a quasi-isometry onto a usual Cayley graph of $\overline{A}$ by collapsing edges connecting g to g_i.

Let Ω denote a graph constructed in the same way for the group A with respect to the generating set $\bigcup \mathcal{A}_i \cup \mathcal{B}$. There is a natural embedding of Ω into Θ, and hence we will consider Ω as a subgraph of Θ. Note that this embedding is also a quasi-isometry.

We will now construct n quasi-isometric retractions of Θ onto Ω.

Let us fix $i \in \{1, \ldots, n\}$. For each $h \in H$ we pick a representative $h_i \in \overline{A}$ thereof, such that $h_i A_i {h_i}^{-1} = A_j$ for a suitable (and unique) j; for $1 \in H$ we pick $1 \in \overline{A}$ as a representative. These elements h_i are coset representatives of the normal subgroup A of $\overline{A}$.

Such a choice defines a retraction $\rho_i \colon \Theta \to \Omega$ in the following way: each vertex g is mapped to the unique vertex g' where $g' \in A$ and $g' h_i = g$ for some h_i; the vertex g_k is then mapped to $(g')_k$. An edge labelled by 0 connecting g to g_k is sent to the edge connecting g' to ${g'}_k$. The remaining edges with label 0 are sent in an A-equivariant fashion to paths connecting the image of their endpoints; the lengths of such paths are uniformly bounded, since (up to the A-action) there are only finitely many edges with label 0.

Similarly, the edges of label $k \notin \{0, i\}$ are mapped in an A-equivariant manner to paths connecting the images of their endpoints; again, their length is uniformly bounded.

Each edge labelled by i is sent A-equivariantly to a path connecting the images of its endpoints, such that the path contains edges labelled only by some j (where j is determined by the coset of A the endpoints lie in); such a path exist by the choice of the representatives h_i.

Note that each such retraction ρ_i is a (κ_i, κ_i)-quasi-isometry for some $\kappa_i \geqslant 1$; we set $\kappa = \max_i \kappa_i$.

Now we are ready to construct a set of cutting edges in Θ.

Consider the ball $B_\Omega(1, 1)$ of radius 1 around the vertex 1 in Ω and let E' denote the set of all edges in $B_\Omega(1, 1)$ labelled by 0. This set disconnects Ω into at least two infinite components; let us take two vertices of Ω, x and y, lying in distinct infinite components of $\Omega - E'$, and such that

$$d_\Omega(1, x) = d_\Omega(1, y) \geqslant \kappa^2 + 4$$

Now let E denote the set of all edges lying in the ball $B_\Theta(1, \kappa^2 + 4)$ labelled by 0. We claim that E disconnects Θ into at least two infinite components. It is enough to show that it disconnects x from y (viewed as vertices of Θ), since we may take x and y to be arbitrarily far from 1 in Ω, and thus in Θ (as Ω and Θ are quasi isometric), and $\Theta - E$ has finitely many components, since E is finite.

Suppose for a contradiction that there exists a path γ in $\Theta - E$ connecting x to y. Using any of the quasi-isometries ρ_i we immediately see that γ has to go through $B_\Theta(1, \kappa^2 + 4)$, since $\rho_i(\gamma)$ must intersect $E' \subseteq B_\Omega(1, 1)$. We write γ as a concatenation of paths $\gamma_1, \ldots, \gamma_m$, such that each γ_i intersects $B_\Theta(1, \kappa^2 + 4)$ only at edges of one label, and its endpoints lie outside of $B_\Theta(1, \kappa^2 + 4)$ (this is possible since γ does not intersect E). We modify each γ_i by pre- and post-concatenating it with a path of length at most 4 (note that all the elements of $\mathcal{H}$ correspond to edges), so that it now starts and ends at Ω. Still, the new path (which we will continue to call γ_i) intersects $B_\Theta(1, \kappa^2 + 1)$ only at edges labelled by a single label.

Now we construct a new path γ' as follows: to each γ_i we apply the retraction ρ_k, where k is the label of edges of γ_i inside $B_\Theta(1, \kappa^2 + 1)$; we now define γ' to be the concatenation of these paths. Such a construction is possible, since the maps ρ_i are retractions, and so in particular they preserve the endpoints of the paths γ_j for all j. Also, γ' runs from x to y. By construction it does so in Ω, and thus it contains an edge of E'; let us denote it by e.

There exists an edge f in some γ_i, such that e lies in the image of f under the map ρ_k that we applied to γ_i. Since ρ_k is an (κ, κ)-quasi-isometry, the edge f lies within $B_\Theta(1, \kappa^2 + 1)$. But then $\rho_k(f)$ is a path the edges of whom are never labelled by 0, and so in particular $e \notin E'$, a contradiction.

We successively remove edges from E until the newly obtained set satisfies the definition of a set of cutting edges. We now apply Theorem 2.5, taking $\mathcal{P}$ to be the property of having all edges labelled by 0. Let C denote the cut we obtain, and let F denote its boundary.

To apply Theorem 2.6 we need to only show that for some $g \in \overline{A}$ we have $g.C$ properly contained in C or C^*. Since C^* is infinite, it contains an element $g \in \overline{A}$ such that $g.F \neq F$. Taking such a g, we see that either $g.C$ is properly contained in C^* (in which case we are done), or C is properly contained in $g.C$. In the latter case we have $g^{-1}.C \subset C$. We have thus verified all the hypotheses of Theorem 2.6.

Since the boundary F of the final cut C is labelled by 0, upon removal of the open edges in $\overline{A}.F$, the connected component containing 1_i contains the entire subgroup A_i, since vertices corresponding to elements of this subgroup are connected to 1_i by paths labelled by i. Thus A_i is a subgroup of a conjugate of a vertex group, and so it fixes a vertex in the associated action on a tree. □

3. Relative Karrass–Pietrowski–Solitar theorem

Definition 3.1. Let T be a metric space, and $v \in T$ a point which admits a neighbourhood isometric to the neighbourhood of a vertex in a tree. Let X

be a connected metric space. We say that the metric space Y is a *blow-up* of T at v by X if and only if X embeds into Y, and collapsing X to a point yields an isometry onto T which collapses X onto v.

We warn the reader that our notion of blow-up is not standard terminology (and has nothing to do with blow-ups in other fields).

Proposition 3.2. *Let G be a graph of groups with finite edge groups. Let G_v be the vertex group associated to v, and suppose that G_v acts on a complete CAT(0) space X. Then there exists a complete CAT(0) space Y on which $\pi_1(G)$ acts, satisfying the following:*

(1) Y is obtained from the universal cover $\widetilde{G}$ by blowing up each preimage u of v by $X_u = X$;
(2) the restricted action of G_v on Y preserves X_w, where w is the vertex in $\widetilde{G}$ fixed by G_v, and the induced action is the given action of G_v on $X_w = X$;
(3) collapsing each X_u individually to a point is $\pi_1(G)$-equivariant, and the resulting tree with the $\pi_1(G)$-action is equivariantly isomorphic to $\widetilde{G}$.

Proof. Let w be the vertex defined in (2). We start by blowing $\widetilde{G}$ up at w by X; such a blow-up will be defined by the way edges emanating from w are attached to X: let e be such an edge. Its stabiliser is a finite subgroup of G_v by assumption, and hence there is a point $p_e \in X$ fixed by the given action $G_v \curvearrowright X$ (since X is a complete CAT(0) space). We attach the edge e to this point p_e. Let e' be another edge in the G_v-orbit of e. There exists $x \in G_v$ taking e to e', and we attach e' at $x.p_e$. This way we attach all edges in the orbit of e, and then we proceed to attach edges in the remaining orbits in the same way.

Now we are going to blow up the other vertices in the preimage of v. Let u be such a vertex. Its stabiliser is a conjugate of G_v; pick once and for all a conjugating element x. We now blow up u by $X_u = X$, and attach the edges in the following way: each edge f emanating form u is the image under x of some e emanating from w; we attach f to $p_e \in X = X_w$. The space Y we constructed this way certainly satisfies (1).

Now we are going to construct an action of $\pi_1(G)$ on Y. Let us take $z \in \pi_1(G)$ and $p \in Y$. If p lies outside any of the X_u, then $z.p$ is defined to be the unique point in Y mapping onto $z.p \in \widetilde{G}$ under the map collapsing each X_u individually to a point.

Now let us suppose that $p \in X_u$ for some u. Let $u' = z.u \in \widetilde{G}$. We have the identification $X_u = X = X_{u'}$, and when constructing Y we picked elements $x_1, x_2 \in \pi_1(G)$ such that $\mathrm{Stab}(u) = x_1 G_v {x_1}^{-1}$ and $\mathrm{Stab}(u') = x_2 G_v {x_2}^{-1}$. We now declare $z.p$ to be the image in $X_{u'} = X$ of

$$x_1 {x_2}^{-1} z.p \in X_w = X$$

(observing that $x_1 {x_2}^{-1} x \in G_v$).

We have thus defined the action, and it is clear that it satisfies (2) and (3). □

Remark 3.3. Suppose that the space X in the above proposition is a tree. Then the resulting space is a tree, and the quotient graph of groups is obtained from G by replacing v by the quotient graph of groups $X/\!\!/G_v$.

The following theorem is a generalisation of a theorem of Karrass–Pietrowski–Solitar [KPS], which lies behind the Nielsen realisation theorem for free groups.

Theorem 3.4 (Relative Karrass–Pietrowski–Solitar theorem)**.** *Let*

$$\phi\colon H \to \mathrm{Out}(A)$$

be a monomorphism with a finite domain, and let

$$A = A_1 * \cdots * A_n * B$$

be a decomposition preserved by H, with each A_i finitely generated, non-trivial and torsion-free, and B a (possibly trivial) finitely generated free group. Let $A_1, \ldots, A_m$ be the minimal factors. Then the associated extension $\overline{A}$ of A by H is isomorphic to the fundamental group of a finite graph of groups with finite edge groups, with m distinguished vertices $v_1, \ldots, v_m$, such that the vertex group associated to v_i is a conjugate of the extension $\overline{A_i}$ of A_i by $\mathrm{Stab}_H(i)$, and vertex groups associated to other vertices are finite.

Proof. The proof goes along precisely the same lines as the original proof of Karrass–Pietrowski–Solitar [KPS], with the exception that we use Relative Stallings' Theorem (Theorem 2.7) instead of the classical one.

Formally, the proof is an induction on the *complexity* $n + \mathrm{rk}(B)$, where n is the number of factors A_i in A, and $\mathrm{rk}(B)$ denotes the usual rank of the free group B. When the complexity of A is 0 the result trivially follows by looking at the graph of groups with a single vertex and no edges.

In the general case, we apply Theorem 2.7 to the finite extension $\overline{A}$. We obtain a graph of groups P with one edge and a finite edge group, such that each A_i lies up to conjugation in a vertex group.

Let v be a vertex of $\widetilde{P}$. The group P_v is a finite extension of $A \cap P_v$ by a subgroup H_v of H.

Let us look at the structure of $P_v \cap A$ more closely. To this end, consider the graph of groups associated to the product $A_1 * \ldots A_n * B$ and apply Kurosh's theorem [Ser, Theorem I.14] to the subgroup $P_v \cap A$. We obtain that $P_v \cap A$ is a free product of groups of the form $P_v \cap xA_ix^{-1}$ for some $x \in A$, and a free group.

Let us suppose that the intersection $P_v \cap xA_ix^{-1}$ is non-trivial for some i and $x \in A$. This implies that a non-trivial subgroup of A_i fixes the vertex $x^{-1}.v$. Since A_i is torsion-free, this subgroup is infinite. We also know that A_i fixes some vertex, say v_i, in $\widetilde{P}$, and thus so does the infinite subgroup we are discussing. But edge stabilisers are finite, and so $v_i = x^{-1}.v$.

Now suppose that $P_v \cap yA_iy^{-1}$ is non-trivial for some other element $y \in A$. Then $x^{-1}.v = v_i = y^{-1}.v$, and so $xy^{-1} \in A \cap P_v$. This implies that the two free factors $P_v \cap xA_ix^{-1}$ and $P_v \cap yA_iy^{-1}$ of $P_v \cap A$ are conjugate inside the group, and so they must coincide.

Note also that $P_v \cap yA_iy^{-1}$ being non-trivial forces $yA_iy^{-1} \leqslant P_v$.

This discussion shows that $P_v \cap A$ is is a free product of at most n non-trivial factors of the form xA_ix^{-1} (at most one for each i), and a free group.

Kurosh's theorem applied to $A \leqslant \pi_1(P) = \overline{A}$ tells us that A is a free product of conjugates of its intersections with the vertex groups and a free group. In particular $P_v \cap A$ is a free factor of A, and hence it has at most the same complexity (by the discussion above), and the equality of complexitites is equivalent to $P_v \cap A = A$. Since the splitting P defines is non-trivial, the index of $P_v \cap A$ in $\overline{A}$ is infinite, and thus A is not a subgroup of P_v. We immediately conclude that the complexity of $A \cap P_v$ is strictly lower than that of A.

We have thus shown that P_v is an extension

$$P_v \cap A \to P_v \to H_v$$

where H_v is a subgroup of H, the group $P_v \cap A$ decomposes in a way which is preserved by H_v, and its complexity is smaller than that of A. Therefore the group P_v satisfies the assumption of the inductive hypothesis.

We now use Proposition 3.2 (together with the remark following it) to construct a new graph of groups Q, by blowing P up at u by the result of the theorem applied to P_u, with u varying over some chosen lifts of the vertices of P.

By construction, Q is a finite graph of groups with finite edge groups, and the fundamental group of Q is indeed $\overline{A}$. Also, Q inherits distinguished vertices from the graphs of groups we blew up with. Thus, Q is as required in the assertion of our theorem, with two possible exceptions.

Firstly, it might have too many distinguished vertices. This would happen if for some i and j we have A_i and A_j both being subgroups of, say, P_v, which are conjugate in $\overline{A}$ but not in P_v. Let $h \in \overline{A}$ be an element such that $hA_ih^{-1} = A_j$. Since both A_i and A_j fix only one vertex, and this vertex is v, we must have $h \in \mathcal{P}_v$, and so A_i and A_j are conjugate inside P_v.

Secondly, it could be that the finite extensions of A_i we obtain as vertex groups are not extensions by $\mathrm{Stab}_H(i)$. This would happen if $\mathrm{Stab}_H(i)$ is not a subgroup of H_v. Let us take $h \in \overline{A}$ in the preimage of $\mathrm{Stab}_H(i)$, such that $hA_ih^{-1} = A_i$. Then in the action on $\widetilde{P}$ the element h takes a vertex fixed by A_i to another such; if these were different, then A_i would fix an edge, which is impossible. Thus h fixes the same vertex as A_i. This finishes the proof. $\square$

4. Fixed points in the graph of relative free splittings

Consider a free product decomposition

$$A = A_1 * \cdots * A_n * B$$

with B a finitely generated free group. Handel and Mosher [HM] (see also the work of Horbez [Hor]) defined a *graph of relative free splittings* $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$ associated to such a decomposition. Its vertices are finite non-trivial graphs of groups with trivial edge groups, and such that each A_i is contained in a conjugate of a vertex group; two such graphs of groups define the same vertex when the associated universal covers are A-equivariantly isometric. Two vertices are connected by an edge if and only if the graphs of groups admit a common refinement.

In their article Handel and Mosher prove that $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$ is connected and Gromov hyperbolic [HM, Theorem 1.1].

Observe that the subgroup $\mathrm{Out}(A, \{A_1, \ldots, A_n\})$ of $\mathrm{Out}(A)$ consisting of those outer automorphisms of A which preserve the decomposition

$$A = A_1 * \cdots * A_n * B$$

acts on this graph. We offer the following fixed point theorem for this action on $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$.

Corollary 4.1. *Let $H \leqslant \mathrm{Out}(A, \{A_1, \ldots, A_n\})$ be a finite subgroup, and suppose that the factors A_i are finitely generated and torsion-free. Then H fixes a point in the free-splitting graph $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$.*

Proof. Theorem 3.4 gives us an action of the extension $\overline{A}$ on a tree T; in particular A acts on this tree, and this action satisfies the definition of a vertex in $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$. Since the whole of $\overline{A}$ acts on T, every outer automorphism in H fixes this vertex. □

5. Relative Nielsen realisation

In this section we use Theorem 3.4 to prove relative Nielsen Realisation for free products. To do this we need to formalise the notion of a marking of a space.

Definition 5.1. We say that a path-connected topological space X with a universal covering $\widetilde{X}$ is *marked* by a group A if and only if it comes equipped with an isomorphism between A and the group of deck transformations of $\widetilde{X}$.

Remark 5.2. Given a space X marked by a group A, we obtain an isomorphism $A \cong \pi_1(X, p)$ by choosing a basepoint $\widetilde{p} \in \widetilde{X}$ (where p denotes its projection in X).

Conversely, an isomorphism $A \cong \pi_1(X, p)$ together with a choice of a lift $\widetilde{p} \in \widetilde{X}$ of p determines the marking in the sense of the previous definition.

Definition 5.3. Suppose that we are given an embedding $\pi_1(X) \hookrightarrow \pi_1(Y)$ of fundamental groups of two path-connected spaces X and Y, both marked. A map $\iota\colon X \to Y$ is said to *respect the markings via the map* $\widetilde{\iota}$ if and only if $\widetilde{\iota}\colon \widetilde{X} \to \widetilde{Y}$ is $\pi_1(X)$-equivariant (with respect to the given embedding $\pi_1(X) \hookrightarrow \pi_1(Y)$), and satisfies the commutative diagram

$$\begin{array}{ccc} \widetilde{X} & \xrightarrow{\widetilde{\iota}} & \widetilde{Y} \\ \downarrow & & \downarrow \\ X & \xrightarrow{\iota} & Y \end{array}$$

We say that ι *respects the markings* if and only if such an $\widetilde{\iota}$ exists.

Suppose that we have a metric space X marked by a group A, and a group H acting on X. Of course such a setup yields the induced action $H \to \mathrm{Out}(A)$, but in fact it does more: it gives us an extension

$$1 \to A \to \overline{A} \to H \to 1$$

where $\overline{A}$ is the group of all lifts of elements of H to automorphisms of the universal covering $\widetilde{X}$ of X.

Now we are ready to state the relative Nielsen Realisation theorem for free products.

Theorem 5.4 (Relative Nielsen Realisation)**.** *Let $\phi\colon H \to \mathrm{Out}(A)$ be a homomorphism with a finite domain, and let*

$$A = A_1 * \cdots * A_n * B$$

be a decomposition preserved by H, with each A_i finitely generated, and B a (possibly trivial) finitely generated free group. Let $A_1, \dots, A_m$ be the minimal factors.

Suppose that for each $i \in \{1, \dots, m\}$ we are given a complete NPC space X_i marked by A_i, on which $\mathrm{Stab}_i(H)$ acts in such a way that the associated extension of A_i by $\mathrm{Stab}_H(i)$ is isomorphic (as an extension) to the extension $\overline{A}_i$ coming from $\overline{A}$. Then there exists a complete NPC space X realising the action ϕ, and such that for each $i \in \{1, \dots, m\}$ we have a $\mathrm{Stab}_H(i)$-equivariant embedding $\iota_i\colon X_i \to X$ which preserves the marking.

Moreover, the images of the spaces X_i are disjoint, and collapsing each X_i and its images under the action of H individually to a point yields a graph with fundamental group abstractly isomorphic to the free group B.

As outlined in the introduction, the proof is very similar to the classical proof of Nielsen realisation, with our new relative Stallings' and Karrass–Pietrowski–Solitar theorems in place of the classical ones.

Proof. Note that the groups A_i are torsion-free, since they are fundamental groups of complete NPC spaces.

When ϕ is injective we first apply Theorem 3.4 to obtain a graph of groups G, and then use Proposition 3.2 and blow up each vertex of $\widetilde{G}$ by the appropriate $\widetilde{X_i}$; we call the resulting space $\widetilde{X}$. The space X is obtained by taking the quotient of the action of A on $\widetilde{X}$.

If ϕ is not injective, then we consider the induced map

$$H/\ker\phi \to \mathrm{Out}(A)$$

apply the previous paragraph, and declare H to act on the resulting space with $\ker\phi$ in the kernel. □

Remark 5.5. In the above theorem the hypothesis on the spaces X_i being complete and NPC can be replaced by the condition that they are semi-locally simply connected, and any finite group acting on their universal covering fixes at least one point.

Remark 5.6. On the other hand, when we strengthen the hypothesis and require the spaces X_i to be NPC cube complexes (with the actions of our finite groups preserving the combinatorial structure), then we may arrange for X to also be a cube complex. When constructing the blow ups, we may always take the fixed points of the finite groups to be midpoints of cubes, and then X is naturally a cube complex, when we take the cubical barycentric subdivisions of the complexes X_i instead of the original cube complexes X_i.

Remark 5.7. In [HOP] Osajda, Przytycki and the first-named author develop a more topological approach to Nielsen realisation and the Karrass–Pietrowski–Solitar theorem. In that article, Nielsen realisation is shown first, using *dismantlability* of the sphere graph (or free splitting graph) of a free group, and the Karrass–Pietrowski–Solitar theorem then follows as a consequence.

The relative Nielsen realisation theorem with all free factors A_i being finitely generated free groups is a fairly quick consequence of the methods developed in [HOP] – however, the more general version proved here cannot at the current time be shown using the methods of [HOP]: to the authors knowledge no analogue of the sphere graph exhibits suitable properties. It would be an interesting problem to find a "splitting graph" for free products which has dismantling properties analogous to the ones shown in [HOP] to hold for arc, sphere and disk graphs.

Sebastian Hensel
Mathematisches Institut
Universität Bonn
Endenicher Allee 60
D-53115 Bonn
Germany
hensel@math.uni-bonn.de

Dawid Kielak
Fakultät für Mathematik
Universität Bielefeld
Postfach 100131
D-33501 Bielefeld
Germany
dkielak@math.uni-bielefeld.de

References

[BH] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[Cul] Marc Culler. Finite groups of outer automorphisms of a free group. In *Contributions to group theory*, volume 33 of *Contemp. Math.*, pages 197–207. Amer. Math. Soc., Providence, RI, 1984.

[CV] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.* **84**(1986), 91–119.

[Dun] M. J. Dunwoody. Cutting up graphs. *Combinatorica* **2**(1982), 15–23.

[HM] Michael Handel and Lee Mosher. Relative free splitting and free factor complexes I: Hyperbolicity. *arXiv:1407.3508*.

[HK] S. Hensel and D. Kielak. Nielsen realisation for untwisted right-angled Artin groups. *arXiv:1410.1618*.

[HOP] Sebastian Hensel, Damian Osajda, and Piotr Przytycki. Realisation and dismantlability. *Geom. Topol.* **18**(2014), 2079–2126.

[Hor] Camille Horbez. The boundary of the outer space of a free product. *arXiv:1408.0543*.

[KPS] A. Karrass, A. Pietrowski, and D. Solitar. Finite and infinite cyclic extensions of free groups. *J. Austral. Math. Soc.* **16**(1973), 458–466. Collection of articles dedicated to the memory of Hanna Neumann, IV.

[Khr] D. G. Khramtsov. Finite groups of automorphisms of free groups. *Mat. Zametki* **38**(1985), 386–392, 476.

[Krö] Bernhard Krön. Cutting up graphs revisited—a short proof of Stallings' structure theorem. *Groups Complex. Cryptol.* **2**(2010), 213–221.

[Ser] Jean-Pierre Serre. *Trees.* Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[Sta1] John R. Stallings. On Torsion-Free Groups with Infinitely Many Ends. *Annals of Mathematics* **Second Series, 88**(September 1968), 312–334.

[Sta2] John R. Stallings. *Group theory and three-dimensional manifolds.* Yale University Press, New Haven,, 1971.

[Zim] Bruno Zimmermann. Über Homöomorphismen n-dimensionaler Henkelkörper und endliche Erweiterungen von Schottky-Gruppen. *Comment. Math. Helv.* **56**(1981), 474–486.