

Nielsen Realization by Gluing: Limit Groups and Free Products

SEBASTIAN HENSEL & DAWID KIELAK

ABSTRACT. We generalize the Karrass–Pietrowski–Solitar and the Nielsen realization 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 and Mosher and on the outer space of a free product of Guirardel and Levitt, and also a relative version of the Nielsen realization theorem, which, in the case of free groups, answers a question of Karen Vogtmann. We also prove Nielsen realization for limit groups and, as a byproduct, obtain a new proof that limit groups are $\text{CAT}(0)$.

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

1. Introduction

In its original form, the *Nielsen realization problem* asks which finite subgroups of the mapping class group of a surface can be realized as groups of homeomorphisms of the surface. A celebrated result of Kerckhoff [Ker1; Ker2] answers this positively for all finite subgroups and even allows for realizations by isometries of a suitable hyperbolic metric.

Subsequently, similar realization results were found in other contexts, perhaps most notably for realizing finite groups in $\text{Out}(F_n)$ by isometries of a suitable graph (independently by [Cul], [Khr], and [Zim]; compare [HOP] for a different approach).

In this article, we begin to develop a *relative* approach to Nielsen realization problems. The philosophy here is that if a group G allows for a natural decomposition into pieces, then Nielsen realization for $\text{Out}(G)$ may be reduced to realization in the pieces and to a *gluing problem*. In addition to just solving Nielsen realization for finite subgroups of $\text{Out}(G)$, such an approach yields more explicit realizations, which also exhibit the structure of pieces for G .

We demonstrate this strategy for two classes of groups, free products and limit groups. In another article, we use the results presented here, together with the philosophy of relative Nielsen realization, to prove Nielsen realization for certain right-angled Artin groups [HK].

Received September 19, 2016. Revision received August 22, 2017.

The second-named author was supported by the SFB 701.

The early proofs of Nielsen realization for free groups rely in a fundamental way on a result of Karrass, Pietrowski, and Solitar [KPS], which states that every finitely generated virtually free group acts on a tree with finite edge and vertex stabilizers. In the language of Bass–Serre theory, it amounts to saying that such a virtually free group is a fundamental group of a finite graph of groups with finite edge and vertex groups.

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

In the first part of this article, we generalize these results to the setting of a free product

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

in which we (usually) require the factors A_i to be finitely generated, and B to be a finitely generated free group. Consider any finite group H acting on A by outer automorphisms in a way preserving the given free-product decomposition, by which we mean that each element of H sends each subgroup A_i to some A_j (up to conjugation); note that we do not require the action of H to preserve B in any way. We then obtain the corresponding group extension

$$1 \rightarrow A \rightarrow \overline{A} \rightarrow H \rightarrow 1.$$

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 4.1) \overline{A} acts on a tree with finite edge stabilizers, with each A_i fixing a vertex of the tree, and with, informally speaking, all other vertex groups finite.

Relative Nielsen realization theorem (Theorem 7.5) Suppose that we are given complete nonpositively curved (i.e. locally CAT(0)) spaces X_i realizing the induced actions of H on the factors A_i . Then the action of H can be realized by a complete nonpositively curved space X ; in fact, X can be chosen to contain the X_i in an equivariant manner.

We emphasize that such a relative Nielsen realization is new even if all A_i are free groups, in which case it answers a question of Karen Vogtmann.

The classical Nielsen realization for graphs immediately implies that a finite subgroup $H < \text{Out}(F_n)$ fixes points in the Culler–Vogtmann outer space (defined in [CV]) and in the complex of free splittings of F_n (which is a simplicial closure of outer space). As an application of the work in this article, we similarly obtain fixed point statements (Corollaries 5.1 and 6.1) for the graph of relative free splittings defined by Handel and Mosher [HM] and for the outer space of a free product defined by Guirardel and Levitt [GL].

In the last section of the paper, we prove the following:

THEOREM 8.11. *Let A be a limit group, and let*

$$A \rightarrow \overline{A} \rightarrow H$$

be an extension of A by a finite group H . Then there exists a complete locally CAT(κ) space X realizing the extension \overline{A} , where $\kappa = -1$ when A is hyperbolic and $\kappa = 0$ otherwise.

This theorem is obtained by combining the classical Nielsen realization theorems (for free, free-Abelian, and surface groups; see Theorems 8.1, 8.2 and 8.3) with the existence of an invariant JSJ decomposition shown by Bumagin, Kharlampovich, and Myasnikov [BKM].

Note that, in general, having a graph of groups decomposition for a group G with CAT(0) vertex groups and virtually cyclic edge groups does not allow us to build a CAT(0) space for G to act on and thus to conclude that G is itself CAT(0); the JSJ decompositions of limit groups are however special in this respect, and the extra structure allows for the conclusion. This has been observed by Brown [Bro], where he developed techniques for building up a CAT(0) space for G to act on.

Observe that we obtain optimal curvature bounds for our space X : it has been proved by Alibegović and Bestvina [AB] that limit groups are CAT(0), and by Brown [Bro] that a limit group is CAT(-1) if and only if it is hyperbolic.

Also, taking H to be the trivial group gives a new (more direct) proof of the fact that limit groups are CAT(0).

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

2. Relative Stallings' Theorem

In this section, we prove a 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.

CONVENTION. When talking about free factor decompositions

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

of some group A , we always assume that at least two of the factors $\{A_1, \dots, A_n, B\}$ are nontrivial.

DEFINITION 2.1. Suppose that $\phi : H \rightarrow \text{Out}(A)$ is a homomorphism with 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 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 $j \geq i$.

REMARK 2.2. Note that when the decomposition is preserved, we obtain an induced action $H \rightarrow \text{Sym}(n)$ on the indices $1, \dots, n$. We may thus speak of the stabilizers $\text{Stab}_H(i)$ inside H . Furthermore, we obtain an induced action

$$\text{Stab}_H(i) \rightarrow \text{Out}(A_i).$$

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

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

$$1 \rightarrow A \rightarrow \overline{A} \rightarrow H \rightarrow 1,$$

and the left action of H by outer automorphisms agrees with the left conjugation action inside the extension \overline{A} .

Observe that then, for each i , we also obtain an extension

$$1 \rightarrow A_i \rightarrow \overline{A}_i \rightarrow \text{Stab}_H(i) \rightarrow 1,$$

where \overline{A}_i is the normalizer of A_i in \overline{A} .

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

We now begin the proof of the relative version of Stallings' theorem. It uses ideas from both Dunwoody's proof [Dun1] 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 , then we write $\Theta - E$ to mean the graph obtained from Θ by removing the interiors of edges in E .

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^* need not be cuts.

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

¹We warn the reader that later parts of Krön's paper are not entirely correct; we only rely on the early, correct sections.

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

Sketch of proof. To prove this, we recall the following terminology, roughly following Dunwoody. We say that C is a \mathcal{P} -cut if and only if its boundary lies in \mathcal{P} . We say that a \mathcal{P} -cut is \mathcal{P} -narrow if and only 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 $e \in \mathcal{P}$, there are finitely many \mathcal{P} -narrow cuts that contain e in its boundary. This is shown by Dunwoody [Dun1, Section 2.5] for narrow cuts, and the proof carries over to the \mathcal{P} -narrow case. Alternatively, Krön [Krö, Lemma 2.1] shows this for sets of cutting edges that cut the graph into exactly two connected components, and \mathcal{P} -narrow cuts have this property.

Now, for each \mathcal{P} -narrow cut C , consider the number $m(C)$ of \mathcal{P} -narrow cuts that are not nested with C (this is finite by [Dun1, Section 2.6]). We call a \mathcal{P} -narrow cut *optimally nested* if $m(C)$ is smallest amongst all \mathcal{P} -narrow cuts. The proof of Theorem 3.3 of [Krö] now shows that optimally nested \mathcal{P} -cuts are all nested with each other. This shows Theorem 2.5. \square

To use that theorem, recall the following:

THEOREM 2.6 ([Dun1, Theorem 4.1]). *Let G be a group acting on a graph Θ . Suppose that there exists a cut C such that*

- (1) C^* is also a cut,
- (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 $h \in G$.

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

Now we are ready for our main splitting result.

THEOREM 2.7 (Relative Stallings' theorem). *Let $\phi : H \rightarrow \text{Out}(A)$ be a monomorphism with 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 \bar{A} be the preimage of $H = \text{im } \phi$ in $\text{Aut}(A)$. Then \bar{A} acts on a tree with finite quotient so that each A_i fixes a vertex, and no nontrivial subgroup of A fixes any edge.*

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

Proof of Theorem 2.7. Before we begin the proof in earnest, we give a brief outline of the strategy. First, we will define a variant of the Cayley graph for \bar{A} in




Figure 1 A local picture of the graph Θ

which the free product structure of A will be visible (in fact, a subgraph will collapse to the Bass–Serre tree of the free product decomposition of A). This graph will contain the different copies of A_i disjointly, separated by edges labeled with a certain label. We will then aim to show that there is a set of cutting edges just using edges with that label, which, using Theorem 2.6, will yield the desired action on a tree.

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

We define Θ to be a variation of the (right) Cayley graph of \bar{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 Figure 1). More formally, the vertex set of Θ is

$$V(\Theta) = \bar{A} \sqcup (\bar{A} \times \{0, \dots, n\}).$$

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

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

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

The group \bar{A} acts on Θ on the left, preserving the labels. The action is free and cocompact. The graph Θ retracts via a quasi-isometry onto a usual Cayley graph

of \bar{A} by collapsing edges connecting g to g_i . Also note that there are copies of the Cayley graphs of the A_i with respect to the generating set \mathcal{A}_i in Θ , where each edge has the label 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 specific quasi-isometric retractions of Θ onto Ω . These will be used later to modify paths to avoid edges with certain labels.

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

Such a choice defines a retraction $\rho_i : \Theta \rightarrow \Omega$ 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 labeled 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 labeled by i is sent A -equivariantly to a path connecting the images of its endpoints such that the path contains edges labeled only by some j (where j is determined by the coset of A the endpoints lie in); such a path exists 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 $\kappa_i \geq 1$; 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 (all of whose edges are labeled by 0). Since A is a nontrivial free product, the identity element disconnects the Cayley graph into at least two infinite components. Hence, $B_\Omega(1, 1)$ disconnects Ω also into at least two infinite components; let us take two vertices of Ω , x and y , lying in distinct infinite components of $\Omega - B_\Omega(1, 1)$, and such that

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

Now let E denote the set of all edges lying in the ball $B_\Theta(1, \kappa^2 + 4)$ labeled by 0. We claim that E disconnects Θ into at least two infinite components. Note that $\Theta - E$ has finitely many components, since E is finite. By possibly choosing x, y even further from each other, it therefore suffices to show that E disconnects x from y (viewed as vertices of Θ).

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 $B_{\Omega}(1, 1)$. Note that if $\gamma' \subset \gamma$ is a subpath lying completely in $B_{\Theta}(1, \kappa^2 + 4)$, then γ' only traverses edges with the same label (as γ does not intersect E). Thus, we can write γ as a concatenation

$$\gamma = \gamma_1 * \cdots * \gamma_m,$$

where each γ_i intersects $B_{\Theta}(1, \kappa^2 + 4)$ only at edges of one label, and its endpoints lie outside of $B_{\Theta}(1, \kappa^2 + 4)$. 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 labeled by a single label.

Now we construct a new path γ' as follows. Suppose that k_i is such that each edge in $\gamma_i \cap B_{\Theta}(1, \kappa^2 + 1)$ has label k_i . We put

$$\gamma'_i = \rho_{k_i}(\gamma_i).$$

Note that as ρ_{k_i} is a retraction onto Ω and the endpoints of γ_i are in Ω , the path γ'_i has the same endpoints as γ_i . Put

$$\gamma' = \gamma'_1 * \cdots * \gamma'_m.$$

This is now a path joining x to y in Ω and thus contains an edge

$$e \in B_{\Omega}(1, 1).$$

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

We now apply Theorem 2.5 taking \mathcal{P} to be the set of edges labeled 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 labeled 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 labeled 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.

It remains to show the triviality of edge stabilizers in A . In fact, we will show that no nontrivial subgroup $G < A$ fixes a narrow cut in Θ with boundary consisting only of edges labeled by 0. To this end, let C be such a cut, and let be F the set of edges forming the boundary of C .

We begin by considering the subgraph Ω . Let Γ be an infinite component of $\Omega - F$, and let $h \in H$ be arbitrary. There are infinitely many vertices v in Γ such that no edge emanating from v lies in F (as the latter is finite). Take one such

vertex and consider an edge e in its star that corresponds to right multiplication with h . Since h normalizes A , it in fact connects Ω to $h.\Omega$. On the other hand, there can be only a single component of $h.\Omega - F$ that is connected to Γ as the cut C is narrow: otherwise, the components of $h.\Omega - F$ would lie in the same component of $\Theta - F$, and F would fail the definition of a boundary of a cut.

In summary, we have shown that, for each h , each infinite component Γ of $\Omega - F$ is connected (via an edge corresponding to right multiplication by h) to a unique infinite component of $h.\Omega - F$. In other words, infinite components of $\Omega - F$ and $h.\Omega - F$ are in bijection to each other, where the bijection identifies components that are connected in $\Theta - F$.

Now, we can think of Ω as the Bass–Serre tree for the splitting of A , whose vertices have been “blown up” to Cayley graphs of the subgroups A_i . In particular, each edge labeled by 0 disconnects Ω . This implies that $\Omega - F$, and hence each $h.\Omega - F$, has exactly two components, both of which are infinite. Namely, if $\Omega - F$ would have more than two infinite components, or just a single one, the same would be true for $\Theta - F$, violating the narrowness of the cut F . It also implies that $F \cap h.\Omega$ consists of exactly one edge for each h . Since A acts freely on Ω , this implies the final claim of the theorem. \square

3. Blow-Ups

We make the convention that graphs of groups are always connected unless explicitly stated otherwise.

PROPOSITION 3.1 (Blow-up with finite edge groups). *Let G be a graph of groups with finite edge groups. For each vertex v , suppose that the associated vertex group G_v acts on a connected space X_v in such a way that each finite subgroup of G_v fixes a point of X_v . Then there exists a connected space Y on which $\pi_1(G)$ acts satisfying the following:*

- (1) *there is a $\pi_1(G)$ -equivariant map $\pi : Y \rightarrow \tilde{G}$;*
- (2) *if w is a vertex of \tilde{G} fixed by G_v , then $\pi^{-1}(w)$ is G_v -equivariantly isometric to X_v ; and*
- (3) *every finite subgroup of G fixes a point of Y .*

Moreover, when the spaces X_v are complete and $\text{CAT}(0)$, then Y is a complete $\text{CAT}(0)$ space.

Proof. Recall that the vertices of \tilde{G} are left cosets of the vertex groups G_v of G ; for each vertex w , we pick an element $z_w \in G$ to be a coset representative of such a coset.

We will build the space Y in two steps. First, we construct the preimage under π of the vertices of \tilde{G} and call it V . We define V to be the disjoint union of spaces X_w , where w runs over the vertices of \tilde{G} , and X_w is an isometric copy of X_v , where v is the image of w under the quotient map $\tilde{G} \rightarrow G$. We construct $\pi : V \rightarrow \tilde{G}$ by declaring $\pi(X_w) = \{w\}$.

We now construct an action of $A = \pi_1(G)$ on V . Let us take $X_w \subset V$, and let $a \in A$. Let $u = a.w$, and note that its image in G is still v . The action of a on V will take X_w to X_u ; using the identifications $X_w \simeq X_v \simeq X_u$, we only need to say how a is supposed to act on X_v , and here it acts as $z_w^{-1}a$.

We now construct the space Y by adding edges to V .

Let e be an edge of \tilde{G} with terminal endpoint w and initial endpoint u . Let X_e denote a copy of the unit interval. Now G_e is a finite subgroup of G_w and so fixes a point in X_w seen as a subset of V . We glue the endpoint 1 of X_e to this point. Analogously, we glue the endpoint 0 to a point in X_u . Now, using the action of A , we equivariantly glue all the endpoints of the edges in the A -orbit of e . We proceed this way for all (geometric) edges. Note that this construction allows us to extend the definition of π .

When all the vertex spaces are complete CAT(0), it is clear that so is Y . \square

REMARK 3.2. Suppose that the spaces X_v in Proposition 3.1 are trees. Then the resulting space Y is a tree, and the quotient graph of groups is obtained from G by replacing v by the quotient graph of groups $X//G_v$.

We will refer to this construction as *blowing up* G by the spaces X_v . We warn the reader that our notion of a blow-up is not standard terminology (and has nothing to do with blow-ups in other fields).

When dealing with limit groups, we will need a more powerful version of a blow-up. We will use a method by Brown, essentially following [Bro, Theorem 3.1]; to this end, let us start with a number of definitions and standard facts.

DEFINITION 3.3. An n -simplex of type M_κ is the convex hull of $n + 1$ points in general position lying in the n -dimensional model space M_κ of curvature κ , as defined in [BH].

An M_κ -simplicial complex K is a simplicial complex in which each simplex is endowed with the metric of a simplex of type M_κ and the face inclusions are isometries.

Note that we are interested in the case of $n = 2$ and negative κ , where the model space M_κ is just a suitably rescaled hyperbolic plane.

DEFINITION 3.4. Let K be an M_κ -simplicial complex of dimension at most 2. The *link* of a vertex v is a metric graph whose vertices are edges of K incident at v , and edges are 2-simplices of K containing v . Inclusion of edges into simplices in X induces the inclusion of vertices into edges in the link. The length of an edge in the link is equal to the angle the edges corresponding to its endpoints make in the simplex.

Let us state a version of Gromov's link condition adapted to our setting.

THEOREM 3.5 (Gromov's link condition [BH, Theorem II.5.2]). *Let K be an M_κ -simplicial complex of dimension at most 2, endowed with a cocompact simplicial*

isometric action. Then K is a locally $\text{CAT}(\kappa)$ space if and only if the link of each vertex in K is $\text{CAT}(1)$.

Of course, for a graph, being $\text{CAT}(1)$ is equivalent to having no nontrivial simple loop of length less than 2π .

LEMMA 3.6 ([Bro, Lemma 2.29]). *For any $0 < \theta < \pi$ and any A, C with $C > A > 0$, there exist $k < 0$ and a locally $\text{CAT}(k)$ M_k -simplicial annulus with one locally geodesic boundary component of length A , and one boundary component of length C , which is locally geodesic everywhere except for one point, where it subtends an angle greater than θ .*

LEMMA 3.7. *Let Z be an infinite virtually cyclic group. Any two cocompact isometric actions on \mathbb{R} have the same kernel, and the quotient of Z by the kernel is isomorphic to either \mathbb{Z} or the infinite dihedral group D_∞ .*

Proof. Clearly, both actions on \mathbb{R} can be made into actions on 2-regular trees with a single edge orbit and no edge inversions; each such action gives us a decomposition of \mathbb{Z} into a graph of finite groups, where the kernel of the action is the unique edge group, and the quotient is as claimed. Let G_1 and G_2 denote the graphs of groups, and let K_1 and K_2 denote the respective edge groups.

Suppose that one of the graphs, say G_1 , has only one vertex. Then K_1 is also equal to the vertex group, and we have $K_2 \leq K_1$, since any finite group acting on a tree has a fixed point. If G_2 also has a single vertex, then $K_1 \leq K_2$ by the same argument, and we are done. Otherwise, $Z/K_1 \simeq \mathbb{Z}$ is a quotient of $Z/K_2 \simeq D_\infty$, which is impossible.

Now suppose that both G_1 and G_2 have two vertices each. Let G_v be a vertex group of G_1 . Arguing as before, we see that it fixes a point in the action of Z on G_2 , and so some index 2 subgroup of G_v fixes an edge. Thus $K_1 \cap K_2$ is a subgroup of K_1 of index at most two. If the index is two, then the image of K_1 in $Z/K_2 \simeq D_\infty$ is a normal subgroup of cardinality 2, but D_∞ has no such subgroups, and so $K_1 \leq K_2$. By symmetry $K_2 \leq K_1$, and we are done. \square

Let us record the following standard fact.

LEMMA 3.8. *Let Z be an infinite virtually cyclic group acting properly by semisimple isometries on a complete $\text{CAT}(0)$ space X . Then Z fixes an image of a geodesic in X (called an axis).*

For the next proposition, let us introduce some notation.

DEFINITION 3.9. A $\text{CAT}(-1)$ M_{-1} -simplicial complex of dimension at most 2 with finitely many isometry classes of simplices will be called *useful*.

PROPOSITION 3.10 (Blow-up with virtually cyclic edge groups). *Suppose that $\kappa \in \{0, -1\}$. Let G be a finite graph of groups with virtually cyclic edge groups. For each vertex v , suppose that the associated vertex group G_v acts properly on*

a connected complete $\text{CAT}(\kappa)$ simplicial complex X_v by semisimple isometries. Suppose further that

- (A1) there exists an orientation of geometric edges of G such that the initial vertex of every edge e is useful: it is a vertex u with useful X_u ; and
- (A2) when X_u is useful and e_1, \dots, e_n are all the edges of G incident at u carrying an infinite edge group, then the axes preserved by $g^{-1}X_{e_i}g$ with $i \in \{1, \dots, n\}$ and $g_i \notin X_{e_i}$ can be taken to be simplicial and pairwise transverse.

Then there exists a connected complete $\text{CAT}(\kappa)$ space Y on which $\pi_1(G)$ acts satisfying the following:

- (1) there is a $\pi_1(G)$ -equivariant map $\pi : Y \rightarrow \tilde{G}$;
- (2) if w is a vertex of \tilde{G} fixed by G_v , then $\pi^{-1}(w)$ is G_v -equivariantly isometric to X_v .

Proof. We proceed exactly as in the proof of Proposition 3.1, with two exceptions: firstly, we rescale the spaces X_v before we start the construction; secondly, we need to deal with infinite virtually cyclic edge groups. Let us first explain how to deal with the infinite edge groups, and then it will become apparent how we need to rescale the useful spaces.

Let e be an oriented edge of \tilde{G} with infinite stabilizer G_e (note that this is a slight abuse of notation, as we usually reserve G_e to be an edge group in G rather than a stabilizer in \tilde{G}). The group is virtually cyclic and so, by Lemma 3.8, fixes an axis in each of the vertex spaces corresponding to the endpoints of e (it could of course be two axes in a single space if e is a loop). The actions on these axes are equivariant by Lemma 3.7, and the only difference is the length of the quotient of the axis by G_e ; we denote the two lengths by λ_e^+ and λ_e^- ; λ_e^+ and λ_e^- are the amounts by which G_e translate the axes corresponding to the terminus and origin of e , respectively.

We claim that we can rescale the spaces X_v and orient the geometric edges so that, for any edge e with infinite stabilizer, we have the initial vertex of e useful, and $\lambda_e^+ \leq \lambda_e^-$. Let us assume that we have already performed a suitable rescaling; we will come back to it at the end of the proof.

Let u denote the initial (useful) endpoint of e , and let w denote the other endpoint of e . We replace each two-dimensional simplex in X_u by the comparison simplex of type $M_{-1/2}$; note that, in particular, this does not affect the metric on the 1-skeleton of X_u and hence does not affect the constant λ_e^- . Let \hat{X}_u denote the resulting space.

In \hat{X}_u we have generated, in Brown's terminology, an *excess angle* δ (depending on u), that is, in the link of any vertex x in \hat{X}_u the distance between any two points that were of distance at least π in the link of x in X_u is at least $\pi + 2\delta$ in the link in \hat{X}_u . By possibly decreasing δ we may assume that $\delta < \frac{\pi}{3}$ and that the distance between any two distinct vertices in a link of a vertex in \hat{X}_u is at least δ (this is possible since there are only finitely many different isometry types of

simplices in X_u , and so in \widehat{X}_u). We still have G_u acting on \widehat{X}_u simplicially and isometrically.

Suppose that $\lambda_e^+ = \lambda_e^-$. Then we take X_e to be a flat strip $[0, 1] \times \mathbb{R}$ on which G_e acts by translating the \mathbb{R} factor so that the quotient is isometric to $[0, 1] \times \mathbb{R}/\lambda_e^+\mathbb{Z}$.

If $\lambda_e^+ \neq \lambda_e^-$, then we take X_e to be the universal cover of an annulus from Lemma 3.6 with boundary curves of length λ_e^+ and λ_e^- , and $\theta = \pi - \delta$. The space X_e is a $\text{CAT}(k_e)$ M_{k_e} -simplicial complex for some $k_e < 0$.

We glue the preimage (in X_e) of each of the boundary curves to the corresponding axis of G_e , so that the gluing is a G_e -equivariant isometry. The gluing along the preimage of the shorter curve (or both curves if they are of equal length) proceeds along convex subspaces, and so if the vertex space was $\text{CAT}(\mu)$ with $\mu \leq 0$, then the glued-up space is still locally $\text{CAT}(\mu)$ along the axis of G_e .

The situation is different at the useful end: here we glue in along a nonconvex curve. We claim that the resulting space is still locally $\text{CAT}(k_e)$ along this geodesic. This follows from Gromov’s link condition (Theorem 3.5) and from the observation that, in the link of any vertex of \widehat{X}_u , we introduced a single path (a *shortcut*) of length at least $\pi - \delta$ between vertices whose distance before the introduction of the shortcut was at least $\pi + 2\delta$. A simple closed curve that traverses both endpoints of the shortcut therefore had length at least $2\pi + 4\delta$ before introducing the shortcut and thus still has length $\geq 2\pi + \delta$ afterward. Thus there is still no nontrivial simple loop shorter than 2π .

We now use the action of $A = \pi_1(G)$ to equivariantly glue in copies of X_e for all edges in the orbit of e . We proceed in the same way for all the other (geometric) edges.

Now we need to look at the curvature. The useful spaces have all been altered to be $M_{-1/2}$ -simplicial complexes, and so they are now $\text{CAT}(-\frac{1}{2})$. If we had any $\text{CAT}(0)$ vertex spaces, then they would remain $\text{CAT}(0)$. The universal covers X_e of annuli are $\text{CAT}(k_e)$ with $k_e < 0$; the infinite strips are $\text{CAT}(0)$. The gluing into the nonuseful spaces did not disturb the curvature. A single gluing into a useful space did not disturb the curvature either, but the situation is more complicated when we glue more than one space X_e into a single \widehat{X}_u , since we could introduce multiple shortcuts of length at least $\pi - \delta$ into a link of a single vertex. If a curve traverses one (or no) shortcut, then the previous argument shows that it has length at least 2π . If it traverses more more than 2, then (as $\delta < \frac{\pi}{3}$) it also has length $\geq 2\pi$. In the final case where it goes through exactly two, note that the endpoints of the shortcuts are all distinct by the transversality assumption (A2). Hence, by the choice of δ , any path connecting these endpoints has length $> \delta$, and so the total path has length $> 2(\pi - \delta) + 2\delta$ as well.

We conclude that our space Y is complete and $\text{CAT}(k)$, where k is the maximum of the values k_e, κ , and $-\frac{1}{2}$. When $\kappa = 0$, we have $k = 0$, and we are done. Otherwise, observing that we had only finitely many edges in G , we have $k < 0$, and so we can rescale Y to obtain a $\text{CAT}(-1)$ space, as claimed.

We still need to explain how to rescale the vertex spaces. We order the vertices of the graph of groups G in some way, obtaining a list v_1, \dots, v_m . We do not

rescale the space X_{v_1} . Up to reorienting the geometric edges running from v_1 to itself, we see that the constants λ_e^+ and λ_e^- for such edges satisfy $\lambda_e^+ \leq \lambda_e^-$.

We look at the full subgraph Γ of G spanned by the vertices v_1, \dots, v_i . Inductively, we assume that the spaces corresponding to vertices in Γ have already been rescaled as required. Now we attach v_{i+1} to Γ , together with all edges connecting v_{i+1} to itself or Γ . If $X_{v_{i+1}}$ is not useful, then we have no edges of the latter type, and all edges connecting v_{i+1} to Γ are oriented toward v_{i+1} . Clearly, we can rescale X_v to be sufficiently small so that the desired inequalities are satisfied (note that there are only finitely many edges to consider).

If $X_{v_{i+1}}$ is useful, then we can reorient all edges connecting v_{i+1} to Γ so that they run away from v_{i+1} . Now we can make $X_{v_{i+1}}$ sufficiently big to satisfy the desired inequalities. We also reorient the edges connecting v_{i+1} to itself in a suitable manner. \square

4. Relative Karrass–Pietrowski–Solitar Theorem

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

THEOREM 4.1 (Relative Karrass–Pietrowski–Solitar theorem). *Let*

$$\phi : H \rightarrow \text{Out}(A)$$

be a monomorphism with finite domain, and let

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

be a decomposition preserved by H , with each A_i finitely generated, nontrivial, and B a (possibly trivial) finitely generated free group. Let A_1, \dots, A_m be the minimal factors. Then the associated extension \bar{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, \dots, v_m , such that the vertex group associated with v_i is a conjugate of the extension \bar{A}_i of A_i by $\text{Stab}_H(i)$, and vertex groups associated with 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.

We will prove the result by an induction on a complexity (n, f) , where n is the number of factors A_i , and f is the rank of the free group B in the decomposition. We order the complexity lexicographically. The cases of complexity $(0, f)$ follow from the usual Nielsen realization theorem for free groups (see Theorem 8.1).

Thus, for the inductive step, we assume a complexity (m, f) with $m > 0$. We begin by applying Theorem 2.7 to the finite extension \bar{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 and no non-trivial subgroup of any factor A_i fixes an edge.

Let v be any vertex of \tilde{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.

Consider the graph of groups associated with the product $A_1 * \dots * A_n * B$ 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 $P_v \cap x A_i x^{-1}$ for some $x \in A$ and of a free group B' .

Let us suppose that the intersection $P_v \cap x A_i x^{-1}$ is nontrivial for some i and $x \in A$. This implies that a nontrivial subgroup G of A_i fixes the vertex $x^{-1}.v$. We also know that A_i fixes some vertex v_i in \tilde{P} by construction, and thus so does G . If $x^{-1}.v \neq v_i$, then this would imply that G fixes an edge, which is impossible. Hence $v_i = x^{-1}.v$, and in particular we have that $x A_i x^{-1} \leq P_v$.

Now suppose that $P_v \cap y A_i y^{-1}$ is nontrivial 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 $P_v \cap x A_i x^{-1}$ and $P_v \cap y A_i y^{-1}$ of $P_v \cap A$ are conjugate inside the group, and so they must coincide.

We consider the action of A on the tree \tilde{P} and conclude that A is equal to the fundamental group of the graph of groups $\tilde{P} // A$. Our discussion shows that:

- (i) The stabilizer of a vertex $v \in \tilde{P}$ has the structure

$$P_v \cap A = x_{i(v,1)} A_i x_{i(v,1)}^{-1} * \dots * x_{i(v,k)} A_i x_{i(v,k)}^{-1} * B',$$

where the indices $i(v, k)$ are all distinct, and B' is some free group.

- (ii) If a conjugate of A_i intersects some stabilizer of v nontrivially, then it stabilizes v .
- (iii) For each i , there is exactly one vertex v such that a conjugate of A_i appears as $A_{i(v,l)}$ in the description above.
- (iv) The edge groups in $\tilde{P} // A$ are trivial.

Since the splitting defined by P is nontrivial, the index of $P_v \cap A$ in \bar{A} is infinite, and thus A is not a subgroup of P_v for any v .

Next, we aim to show that the complexity of each $P_v \cap A$ is strictly smaller than that of A . To begin, note that the only way that this could fail is if there is some vertex w such that

$$P_w \cap A = x_1 A_1 x_1^{-1} * \dots * x_m A_m x_m^{-1} * B'$$

for a free group B' . Since all edge groups in $\tilde{P} // A$ are trivial, A is obtained from $P_w \cap A$ by a free product with a free group. Such an operation cannot decrease the rank of B' and in fact increases it unless the free product is trivial. However, in the latter case, we would have $P_w \cap A = A$, which is impossible.

We have thus shown that each P_v is an extension

$$P_v \cap A \rightarrow P_v \rightarrow H_v,$$

where H_v is a subgroup of H , the group $P_v \cap A$ decomposes in a way 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 3.1 (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 \bar{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 \bar{A} but not in P_v . Let $h \in \bar{A}$ 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 $h \in P_v$, 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 $\text{Stab}_H(i)$. This would happen if $\text{Stab}_H(i)$ is not a subgroup of H_v . Let us take $h \in \bar{A}$ in the preimage of $\text{Stab}_H(i)$ such that $hA_ih^{-1} = A_i$. Then in the action on \tilde{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

5. Fixed Points in the Graph of Relative Free Splittings

Consider a free product decomposition

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

with a finitely generated free group B . Handel and Mosher [HM] (see also the work of Horbez [Hor]) defined a *graph of relative free splittings* $\mathcal{FS}(A, \{A_1, \dots, A_n\})$ associated with 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 $\mathcal{FS}(A, \{A_1, \dots, A_n\})$ is connected and Gromov hyperbolic [HM, Theorem 1.1].

Observe that the subgroup $\text{Out}(A, \{A_1, \dots, A_n\})$ of $\text{Out}(A)$ consisting of those outer automorphisms of A that 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 $\mathcal{FS}(A, \{A_1, \dots, A_n\})$.

COROLLARY 5.1. *Let $H \leq \text{Out}(A, \{A_1, \dots, A_n\})$ be a finite subgroup, and suppose that the factors A_i are finitely generated. Then H fixes a point in the free-splitting graph $\mathcal{FS}(A, \{A_1, \dots, A_n\})$.*

Proof. Theorem 4.1 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 $\mathcal{FS}(A, \{A_1, \dots, A_n\})$. Since the whole of \overline{A} acts on T , every outer automorphism in H fixes this vertex. \square

6. Fixed Points in the Outer Space of a Free Product

Take any finitely generated group A and consider its *Grushko decomposition*, that is, a free splitting

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

where B is a finitely generated free group, and each group A_i is finitely generated and freely indecomposable, that, is it cannot act on a tree without a global fixed point (note that \mathbb{Z} is not freely indecomposable in this sense).

Grushko's theorem [Gru] tells us that such a decomposition is essentially unique; more precisely, if

$$A = A'_1 * \dots * A'_m * B'$$

is another such decomposition, then $B \cong B'$, $m = n$, and there is a permutation β of the set $\{1, \dots, n\}$ such that A_i is conjugate to $A'_{\beta(i)}$. In particular, this implies that the decomposition

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

is preserved in our sense by every outer automorphism of A .

Guirardel and Levitt [GL] introduced PO , the (projectivized) outer space of a free product. It is a simplicial complex whose vertices are equivalence classes of pairs (G, ι) , where:

- (1) G is a finite graph of groups with trivial edge groups;
- (2) edges of G are given positive lengths;
- (3) for every $i \in \{1, \dots, n\}$, there is a unique vertex v_i in G such that the vertex group G_{v_i} is conjugate to A_i ;
- (4) all other vertices have trivial vertex groups;
- (5) every leaf of G is one of the vertices $\{v_1, \dots, v_n\}$;
- (6) $\iota : \pi_1(G) \rightarrow A$ is an isomorphism.

The equivalence relation is given by postcomposing ι with an inner automorphism of A and by multiplying the lengths of all edges of G by a positive constant. We also consider two pairs G, ι and G', ι' equivalent if there exists an isometry $\psi : G \rightarrow G'$ such that $\iota = \iota' \circ \psi$.

Because of the essential uniqueness of the Grushko decomposition, the group $\text{Out}(A)$ acts on PO by postcomposing the marking ι . We offer the following result for this action.

COROLLARY 6.1. *Let A be a finitely generated group, and let $H \leq \text{Out}(A)$ be a finite subgroup. Then H fixes a vertex in PO .*

Proof. Theorem 4.1 gives us an action of the extension \overline{A} on a tree T , and we may assume that this action is minimal; in particular, A acts on this tree, and this action

satisfies the definition of a vertex in \mathcal{PO} (with all edge lengths equal to 1). Since the whole of \overline{A} acts on T , every outer automorphism in H fixes this vertex. \square

Note that \mathcal{PO} has been shown in [GL, Theorem 4.2, Corollary 4.4] to be contractible.

7. Relative Nielsen Realization

In this section, we use Theorem 4.1 to prove relative Nielsen realization for free products. To do this, we need to formalize the notion of a marking of a space.

DEFINITION 7.1. We say that a path-connected topological space X with a universal covering \tilde{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 \tilde{X} .

REMARK 7.2. Given a space X marked by a group A , we obtain an isomorphism $A \cong \pi_1(X, p)$ by choosing a basepoint $\tilde{p} \in \tilde{X}$ (where p denotes its projection in X).

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

DEFINITION 7.3. 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 $\iota : X \rightarrow Y$ is said to *respect the markings via the map $\tilde{\iota}$* if and only if $\tilde{\iota} : \tilde{X} \rightarrow \tilde{Y}$ is $\pi_1(X)$ -equivariant (with respect to the given embedding $\pi_1(X) \hookrightarrow \pi_1(Y)$) and satisfies the commutative diagram

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

We say that ι *respects the markings* if and only if such an $\tilde{\iota}$ 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 $H \rightarrow \text{Out}(A)$, but in fact it does more: it gives us an extension

$$1 \rightarrow A \rightarrow \overline{A} \rightarrow H \rightarrow 1,$$

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

DEFINITION 7.4. Suppose that we are given a group extension

$$A \rightarrow \overline{A} \rightarrow H.$$

We say that an action $\phi : H \rightarrow \text{Isom}(X)$ of H on a metric space X *realizes the extension \overline{A}* if and only if X is marked by A , and the extension

$$\pi_1(X) \rightarrow G \rightarrow H$$

induced by ϕ fits into the commutative diagram

$$\begin{array}{ccccc}
 A & \longrightarrow & \overline{A} & \longrightarrow & H \\
 \downarrow \simeq & & \downarrow \simeq & & \parallel \\
 \pi_1(X) & \longrightarrow & G & \longrightarrow & H
 \end{array}$$

When A is center-free and we are given an embedding $H \leq \text{Out}(A)$, we say that an action ϕ as before realizes the action $H \rightarrow \text{Out}(A)$ if and only if it realizes the corresponding extension.

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

THEOREM 7.5 (Relative Nielsen realization). *Let $\phi : H \rightarrow \text{Out}(A)$ be a homomorphism with finite domain, and let*

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

be a decomposition preserved by H , with each A_i finitely generated, and a (possibly trivial) finitely generated free group B . 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 $\text{Stab}_i(H)$ acts in such a way that the associated extension of A_i by $\text{Stab}_H(i)$ is isomorphic (as an extension) to the extension \overline{A}_i coming from \overline{A} . Then there exists a complete NPC space X realizing the action ϕ and such that, for each $i \in \{1, \dots, m\}$, we have a $\text{Stab}_H(i)$ -equivariant embedding $\iota_i : X_i \rightarrow X$ that 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 realization, with our new relative Stallings’ and Karrass–Pietrowski–Solitar theorems in place of the classical ones.

Proof. When ϕ is injective, we first apply Theorem 4.1 to obtain a graph of groups \overline{G} and then use Proposition 3.1 and blow up each vertex of \overline{G} by the appropriate \tilde{X}_i ; we call the resulting space \tilde{X} . The space X is obtained by taking the quotient of the action of A on \tilde{X} .

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

$$H / \ker \phi \rightarrow \text{Out}(A),$$

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

REMARK 7.6. In the theorem the hypothesis on the spaces X_i being complete and NPC can be replaced by the condition that they are semilocally simply connected, and any finite group acting on their universal covering fixes at least one point.

REMARK 7.7. 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), 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 7.8. In [HOP] Osajda, Przytycki, and the first-named author develop a more topological approach to Nielsen realization and the Karrass–Pietrowski–Solitar theorem. In that article, Nielsen realization 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 realization 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 that have dismantling properties analogous to those shown in [HOP] to hold for arc, sphere, and disk graphs.

8. Nielsen Realization for Limit Groups

In the last section, we prove a Nielsen realization statement for limit groups. It relies on the following three classical Nielsen realization theorems.

THEOREM 8.1 ([Cul; Khr; Zim]). *Let H be a finite subgroup of $\text{Out}(F_n)$, where F_n denotes the free group of rank n . There exists a finite graph X realizing the given action $H < \text{Out}(F_n)$.*

THEOREM 8.2. *Let*

$$\mathbb{Z}^n \rightarrow \overline{\mathbb{Z}^n} \rightarrow H$$

be a finite extension of \mathbb{Z}^n . There exists a metric n -torus X realizing this extension.

THEOREM 8.3 (Kerckhoff [Ker1; Ker2]). *Let H be a finite subgroup of $\text{Out}(\pi_1(\Sigma))$ where Σ is a closed surface of genus at least 2. There exists a hyperbolic metric on Σ such that Σ endowed with this metric realizes the given action $H < \text{Out}(\pi_1(\Sigma))$.*

Now we are ready to proceed with limit groups.

DEFINITION 8.4. A group A is called *fully residually free* if and only if for any finite subset $\{a_1, \dots, a_n\} \subseteq A \setminus \{1\}$, there exists a free quotient $q : A \rightarrow F$ such that $q(a_i) \neq 1$ for each i .

A finitely generated fully residually free group is called a *limit group*.

Note that the definition immediately implies that limit groups are torsion free.

The crucial property of one-ended limit groups is that they admit JSJ-decompositions invariant under automorphisms.

THEOREM 8.5 (Bumagin–Kharlampovich–Myasnikov [BKM, Theorem 3.13 and Lemma 3.16]). *Let A be a one-ended limit group. Then there exists a finite graph of groups G with all edge groups cyclic, each vertex group being finitely generated free, finitely generated free abelian, or the fundamental group of a closed surface, such that $\pi_1(G) = A$ and such that any automorphism ϕ of A induces an A -equivariant isometry ψ of \tilde{G} such that the following diagram commutes:*

$$\begin{array}{ccc} A & \longrightarrow & \text{Isom}(\tilde{G}) \\ \downarrow \phi & & \downarrow c_\psi \\ A & \longrightarrow & \text{Isom}(\tilde{G}) \end{array}$$

where c_ψ denotes conjugation by ψ .

Moreover, every maximal Abelian subgroup of A is conjugate to a vertex group of G , and every edge in G connects a vertex carrying a maximal Abelian subgroup to a vertex carrying a non-Abelian free group or a surface group.

We will refer to the graph of groups G as the *canonical JSJ-decomposition*.

DEFINITION 8.6. Recall that a subgroup $G \leq A$ is *malnormal* if and only if $a^{-1}Ga \cap G \neq \{1\}$ implies that $a \in G$ for every $a \in A$.

Following Brown, we say that a family of subgroups G_1, \dots, G_n of A is *malnormal* if and only if for every $a \in A$, we have that $a^{-1}G_i a \cap G_j \neq \{1\}$ implies that $i = j$ and $a \in G_i$.

We will use another property of limit groups and their canonical JSJ-decompositions.

PROPOSITION 8.7 ([BKM, Theorem 3.1(3), (4)]). *Let A be a limit group. Every nontrivial Abelian subgroup of A lies in a unique maximal Abelian subgroup, and every maximal Abelian subgroup is malnormal.*

COROLLARY 8.8. *Let G_v be a non-Abelian vertex group in a canonical JSJ-decomposition of a one-ended limit group A . Then the edge groups carried by edges incident at v form a malnormal family in G_v .*

Proof. Let Z_1 and Z_2 be two edge groups carried by distinct edges, e and e' say, incident at v . Without loss of generality, we may assume that each of these groups is infinite cyclic. Suppose that there exist $g \in G_v$ and a nontrivial $z \in g^{-1}Z_1g \cap Z_2$. Each Z_i lies in a unique maximal subgroup M_i of A . But then the Abelian subgroup generated by z lies in both M_1 and M_2 , which forces $M_1 = M_2$ by uniqueness. Now

$$g^{-1}M_1g \cap M_1 \neq \{1\},$$

which implies that $g \in M_1$ (since M_1 is malnormal), and so $g^{-1}Z_1g = Z_1$, which in turn implies that $z \in Z_1 \cap Z_2$.

The edges e and e' form a loop in G , and so there is the corresponding element t in $A = \pi_1(G)$. Observe that t commutes with z , and so the group $\langle t, z \rangle$ must lie in M_1 . But this is a contradiction, as t does not fix any vertices in \tilde{G} . \square

We are now going to use [Bro, Lemma 2.31]; we are however going to break the argument of this lemma into two parts.

LEMMA 8.9 (Brown). *Let X be a connected M_{-1} -simplicial complex of dimension at most 2. Let $A = \pi_1(X)$, and suppose that we are given a malnormal family $\{G_1, \dots, G_n\}$ of infinite cyclic subgroups of A . Then, after possibly subdividing X , each group G_i fixes a (simplicial) axis a_i in the universal cover of X , and the images in X of axes a_i and a_j for $i \neq j$ are distinct.*

In the second part of [Bro, Lemma 2.31], we need to introduce an extra component, namely a simplicial action of a finite group H on X , which permutes the groups G_i up to conjugation.

LEMMA 8.10 (Brown). *Let X be a locally CAT(-1) connected finite M_{-1} -simplicial complex of dimension at most 2. Let $A = \pi_1(X)$, and suppose that we are given a family $\{c_1, \dots, c_n\}$ of locally geodesic simplicial closed curves with images pairwise distinct. Suppose that we have a finite group H acting simplicially on X in a way preserving the images of the curves c_1, \dots, c_n setwise. Then there exists a locally CAT(k) two-dimensional finite simplicial complex X' of curvature k , with $k < 0$, with a transverse family of locally geodesic simplicial closed curves $\{c'_1, \dots, c'_n\}$, such that X' is H -equivariantly homotopic to X , and the homotopy takes c'_i to c_i for each i .*

Sketch of proof. The proof of [Bro, Lemma 2.31] goes through verbatim, with a slight modification; to explain the modification, let us first briefly recount Brown's proof.

We start by finding two local geodesics, say c_1 and c_2 , that contain segments whose union is a tripod—one arm of the tripod is shared by both segments. We glue in a *fin*, that is, a two-dimensional M_k -simplex, so that one side of the simplex is glued to the shared segment of the tripod, and another side is glued to another arm (the intersection of the two sides goes to the central vertex of the tripod). This way one of the curves, say c_1 , is no longer locally geodesic, and we replace it by a locally geodesic curve identical to c_1 except that, instead of travelling along two sides of the fin, it goes along the third side.

The problem is that, after the gluing of a fin, our space will usually not be locally CAT(-1) (the third side of the fin introduces a shortcut in the link of the central vertex of the tripod). To deal with this, we first replace simplices in X by the corresponding M_k -simplices, and this creates an excess angle δ (compare also the proof of Proposition 3.10). Then gluing in the fin does not affect the property of being locally CAT(k).

We glue such fins multiple times, until all local geodesics intersect transversely; after each gluing, we perform a replacement of simplices to generate the excess angle.

Now let us describe what changes in our argument. When gluing in a fin, we need to do it H -equivariantly in the following sense: a fin is glued along two consecutive edges, say (e, e') , and H acts on pairs of consecutive edges. We thus glue in one fin for each coset of the stabilizer of (e, e') in H . This way, when we introduce shortcuts in a link of a vertex, no two points are joined by more than one shortcut. Since we are gluing multiple fins simultaneously, we need to make the angle $\pi - \delta$ sufficiently close to π . \square

Note that when we say that the family $\{c'_1, \dots, c'_n\}$ is transverse, we mean that each curve c'_i intersects transversely with the other curves and itself.

THEOREM 8.11. *Let A be a limit group, and let*

$$A \rightarrow \bar{A} \rightarrow H$$

be an extension of A by a finite group H . Then there exists a complete locally CAT(κ) space X realizing the extension \bar{A} , where $\kappa = -1$ when A is hyperbolic and $\kappa = 0$ otherwise.

Proof. We first assume that A , and so \bar{A} , are one-ended. We apply Theorem 8.5 and obtain a connected graph of groups G with

$$\pi_1(G) = A,$$

for which we can extend the natural action of A on \tilde{G} to an action of \bar{A} . Taking the quotient by \bar{A} , we obtain a new graph of groups Γ with

$$\pi_1(\Gamma) = \bar{A}.$$

The edge groups of Γ are virtually cyclic, and vertices are finite extensions of finitely generated free or free-Abelian groups, or finite extensions of fundamental groups of closed surfaces.

Using Theorems 8.1, 8.2 and 8.3, for each vertex group Γ_v we construct a complete NPC space X_v marked by $A_v = A \cap \Gamma_v$, on which Γ_v/A_v acts in such a way that the induced extension is isomorphic to Γ_v . The space \tilde{X}_v is isometric either to a Euclidean space, the hyperbolic plane, or a tree, and the group A_v acts by deck transformations upon it.

When \tilde{X}_u is the hyperbolic space, we can triangulate it Γ_u -equivariantly, and so \tilde{X}_u and X_u have the structure of two-dimensional finite M_{-1} -simplicial complexes. Moreover, we can triangulate it in such a way that each axis fixed by an infinite cyclic group carried by an edge incident at u is also simplicial. Observe that Γ_v/A_v permutes these axes, and so each of the corresponding edge groups in Γ preserves such an axis as well.

Now we apply Lemma 8.9 and conclude that distinct axes do not coincide. Thus we may use Lemma 8.10 and replace X_u by a new CAT(-1) M_{-1} -simplicial complex (after rescaling) of dimension at most 2, which has only finitely many

isometry classes of simplices, and in which our axes intersect each other and themselves transversely.

We argue in the analogous manner for spaces X_μ that are trees.

Observing that each infinite edge group preserves an axis in each of the relevant vertex spaces by Lemma 3.8, we apply Proposition 3.10 and take the resulting space to be X .

Let us now consider a limit group A that is not one-ended. In this case, we apply the classical version of the Stallings theorem to \bar{A} and split it over a finite group. We will in fact apply the theorem multiple times, so that we obtain a finite graph of groups B with finite edge groups, with all vertex groups finitely generated and one-ended, and $\pi_1(G) = A$; the fact that we only have to apply the theorem finitely many times follows from finite presentability of A (see [BKM, Theorem 3.1(5)]) and Dunwoody's accessibility [Dun2].

The one-ended vertex groups are themselves finite extensions of limit groups, and so for each of them, we have a connected metric space to act on by the first part of the current proof. We finish the argument by an application of Proposition 3.1: the assumption on finite groups fixing points is satisfied since the vertex spaces are complete and CAT(0). \square

ACKNOWLEDGMENTS. The authors would like to thank Karen Vogtmann for discussions and suggesting the statement of relative Nielsen realization for free groups, Stefan Witzel for pointing out the work of Sam Brown, and the referee for extremely valuable comments.

References

- [AB] E. Alibegović and M. Bestvina, *Limit groups are CAT(0)*, J. Lond. Math. Soc. (2) 74 (2006), 259–272.
- [BH] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Grundlehren Math. Wiss., 319, Springer-Verlag, Berlin, 1999.
- [Bro] S. Brown, *A gluing theorem for negatively curved complexes*, J. Lond. Math. Soc. (2) 93 (2016), 741–762.
- [BKM] I. Bumagin, O. Kharlampovich, and A. Miasnikov, *The isomorphism problem for finitely generated fully residually free groups*, J. Pure Appl. Algebra 208 (2007), 961–977.
- [Cul] M. Culler, *Finite groups of outer automorphisms of a free group*, Contributions to group theory, Contemp. Math., 33, pp. 197–207, Amer. Math. Soc., Providence, RI, 1984.
- [CV] M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. 84 (1986), 91–119.
- [Dun1] M. J. Dunwoody, *Cutting up graphs*, Combinatorica 2 (1982), 15–23.
- [Dun2] ———, *The accessibility of finitely presented groups*, Invent. Math. 81 (1985), 449–457.
- [Gru] I. Gruschko, *Über die Basen eines freien Produktes von Gruppen*, Rec. Math. [Mat. Sbornik] N.S. 8(50) (1940), 169–182.
- [GL] V. Guirardel and G. Levitt, *The outer space of a free product*, Proc. Lond. Math. Soc. (3) 94 (2007), 695–714.

- [HM] M. Handel and L. Mosher, *Relative free splitting and free factor complexes I: Hyperbolicity*, [arXiv:1407.3508](#).
- [HK] S. Hensel and D. Kielak, *Nielsen realization for untwisted right-angled Artin groups*, [arXiv:1410.1618](#).
- [HOP] S. Hensel, D. Osajda, and P. Przytycki, *Realisation and dismantlability*, *Geom. Topol.* 18 (2014), 2079–2126.
- [Hor] C. 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. Aust. Math. Soc.* 16 (1973), 458–466, Collection of articles dedicated to the memory of Hanna Neumann, IV.
- [Ker1] S. P. Kerckhoff, *The Nielsen realization problem*, *Bull. Amer. Math. Soc. (N.S.)* 2 (1980), 452–454.
- [Ker2] ———, *The Nielsen realization problem*, *Ann. of Math. (2)* 117 (1983), 235–265.
- [Khr] D. G. Khramtsov, *Finite groups of automorphisms of free groups*, *Mat. Zametki* 38 (1985), 386–392, 476.
- [Krö] B. Krön, *Cutting up graphs revisited—A short proof of Stallings’ structure theorem*, *Groups Complex. Cryptol.* 2 (2010), 213–221.
- [Ser] J.-P. Serre, *Trees*, Springer Monogr. Math., Springer-Verlag, Berlin, 2003, translated from the French original by John Stillwell, corrected 2nd printing of the 1980 English translation.
- [Sta1] J. R. Stallings, *On torsion-free groups with infinitely many ends*, *Ann. of Math. (2)* 88 (1968), 312–334.
- [Sta2] ———, *Group theory and three-dimensional manifolds*, Yale University Press, New Haven, CT, 1971.
- [Zim] B. Zimmermann, *Über Homöomorphismen n -dimensionaler Henkelkörper und endliche Erweiterungen von Schottky-Gruppen*, *Comment. Math. Helv.* 56 (1981), 474–486.

S. Hensel
Mathematisches Institut
Universität Bonn
Endenicher Allee 60
D-53115 Bonn
Germany

hensel@math.uni-bonn.de

D. Kielak
Fakultät für Mathematik
Universität Bielefeld
Postfach 100131
D-33501 Bielefeld
Germany

dkielak@math.uni-bielefeld.de