

AG
T

*Algebraic & Geometric
Topology*

Volume 25 (2025)

Cubulating a free-product-by-cyclic group

FRANÇOIS DAHMANI

SURAJ KRISHNA MEDA SATISH

Cubulating a free-product-by-cyclic group

FRANÇOIS DAHMANI
SURAJ KRISHNA MEDA SATISH

Let $G = H_1 * \cdots * H_k * F_r$ be a finitely generated torsion-free group and ϕ an automorphism of G that preserves this free factor system. We show that when ϕ is fully irreducible and atoroidal relative to this free factor system, the mapping torus $\Gamma = G \rtimes_{\phi} \mathbb{Z}$ acts relatively geometrically on a hyperbolic CAT(0) cube complex. This is a generalisation of a result of Hagen and Wise for hyperbolic free-by-cyclic groups.

20E08, 20E36, 20F65, 20F67

1. Introduction	1561
2. The flow space	1565
3. Relative cubulation	1576
4. Walls in the flow space	1577
5. Cutting principal flow lines and geodesics by wall saturations	1589
6. Separating pairs of points in the Bowditch boundary	1593
References	1596

1 Introduction

Consider a finitely generated free group F and an automorphism $\phi: F \rightarrow F$. Hagen and Wise showed in [HW16] that if ϕ is atoroidal and fully irreducible, then the mapping torus $F \rtimes_{\phi} \mathbb{Z}$ acts properly and cocompactly on a hyperbolic CAT(0) cube complex. It often happens that automorphisms of free groups are neither atoroidal nor fully irreducible, suggesting various directions of generalisation. In [HW15], Hagen and Wise relaxed the requirement of full irreducibility and, using the sophisticated machinery of relative train track maps, showed that the mapping torus still acts geometrically on a CAT(0) cube complex. They asked (see the discussion around [HW15, Problem B]) if a systematic answer to which free-by-cyclic groups admit cubulations is possible, especially in the presence of polynomially growing subgroups.

Here, by investigating *relative cubulations* instead of usual cubulations, we provide an answer in great generality as to when such groups are relatively cubulated. Let ϕ be an automorphism of F and let $F = H_1 * H_2 * \cdots * H_k * F_r$ be a free decomposition that is preserved by ϕ (up to taking conjugates of the factors). Such a free decomposition always exists for any ϕ . In particular, when ϕ is not fully irreducible, there exists a free decomposition preserved by ϕ , relative to which ϕ is fully irreducible. Let

us also assume that ϕ is atoroidal relative to the free decomposition. We will clarify the meanings of both these terms more precisely, using the notion of free factor systems, in [Section 2](#). In our setting, we allow, for instance, elements to have polynomial growth under ϕ as long as they are elliptic in the free product. Under a mild complexity condition, the first author and Li showed in [\[DL22\]](#) that the mapping torus $F \rtimes_{\phi} \mathbb{Z}$ is hyperbolic relative to the suspensions of the free factors H_i . A particular case of our main result shows that such a mapping torus acts relatively geometrically on a hyperbolic CAT(0) cube complex.

1.1 Main result

Let G be a finitely generated group, ϕ an automorphism of G and H a subgroup of G whose conjugacy class is preserved by a power of ϕ . Let n be the minimal positive power of ϕ such that $\phi^n(H) = g^{-1}Hg$. Then we say that the suspension of H by ϕ in the semidirect product $G \rtimes_{\phi} \langle t \rangle$ is the group $H \rtimes_{\text{ad}_g \circ \phi^n} \langle t^n g^{-1} \rangle$, where $\text{ad}_g : G \rightarrow G$ denotes the inner automorphism given by $h \mapsto ghg^{-1}$.

To state our main result for automorphisms of free products of groups, we need the notions of full irreducibility and atoroidality *relative to a free factor system*. These are analogous to the corresponding notions for automorphisms of free groups, with the condition that the free decomposition is preserved, but within a free factor there are no restrictions. We also need a technical notion of *no twinned subgroups*.¹ We refer the reader to [Section 2](#) for the definitions.

We also recall the notion of relative cubulation introduced by Einstein and Groves in [\[EG20\]](#): a relatively hyperbolic group (Γ, \mathcal{P}) is *relatively cubulated* if it acts cocompactly on a CAT(0) cube complex with cell stabilisers either trivial or conjugate to a finite index subgroup in \mathcal{P} .

Theorem 1.1 *Let G be a finitely generated torsion-free group and let $G \cong H_1 * \cdots * H_k * F_r$ be a free decomposition such that each H_i is nontrivial. Let ϕ be an automorphism that preserves the associated free factor system. Assume that $k + r \geq 3$, that ϕ is fully irreducible relative to the free factor system and atoroidal relative to the free factor system, and that there exist no twinned subgroups. Then the mapping torus $\Gamma = G \rtimes_{\phi} \mathbb{Z}$ admits a relative cubulation for the peripheral structure of the suspensions of the free factors H_i .*

We recover the result of [\[HW16\]](#) when $G = F_r$ above,² as well as some cases of [\[HW15\]](#), through a telescopic argument of Groves and Manning [\[GM23, Theorem D\]](#). In general, the relatively geometric action on CAT(0) cube complexes that we obtain otherwise still has interesting consequences.

Note that any group acting on a CAT(0) cube complex admits an action on an ℓ^2 -space built using characteristic functions of hyperplane-halfspaces; see Niblo and Reeves [\[NR97\]](#). We thus have:

Corollary 1.2 *Let G and ϕ be as in [Theorem 1.1](#). Then the mapping torus Γ acts on a Hilbert space with unbounded orbits, with no global fixed point for G .*

¹This requirement can in fact be removed; see work of the authors and Mutanguha [\[DMM25, Lemma 3.3\]](#) or [Remark 2.5](#).

²It is stated there in terms of irreducible atoroidal automorphisms, which are fully irreducible by Dowdall, Kapovich, and Leininger [\[DKL15, Corollary B.4\]](#).

In addition, when G is residually finite, using a generalisation by Einstein and Groves [EG22, Theorem 1.6] of a well-known result of Haglund and Wise [HW08], we have:

Corollary 1.3 *Let G and ϕ be as in Theorem 1.1. Further, let G be residually finite. Then every full relatively quasiconvex subgroup of the mapping torus Γ is separable.*

Another consequence of Theorem 1.1 can be seen in our recent work with Mutanguha [DMM25], where we showed that all hyperbolic hyperbolic-by-cyclic groups are virtually special.

1.2 Method

Our procedure to cubulate follows the scheme laid out by Hagen and Wise [HW16]. The goal is to obtain a collection of codimension-1 subgroups of Γ and then apply Sageev’s dual cube complex construction [Sag95]. The codimension-1 subgroups we build will be stabilised by full relatively quasiconvex subgroups in the relatively hyperbolic group Γ . We then apply the boundary criterion of Einstein and Groves [EG20].

In order to do this, an important tool is Francaviglia and Martino’s absolute train tracks for free products [FM15]. Given G and an automorphism ϕ satisfying the hypotheses of the main result, there exists a G -tree T and a train track map $f: T \rightarrow T$ representing ϕ . Taking mapping cylinders for f , one then obtains a *flow space* on which Γ acts. The flow space has the structure of a tree of spaces, where the underlying graph is a line and vertex and edge spaces are copies of the tree T . The map f “flows” a point on any tree to its image in the next tree. We describe the flow space and various properties we need in order to define walls in the flow space in Section 2.

Before explaining how we build walls in our setup, let us motivate our construction in the surface case. Let M_f be the mapping torus of a closed hyperbolic surface S_g under a pseudo-Anosov map f . Cooper, Long and Reid showed in [CLR94] that in this case there exists an immersed quasi-Fuchsian surface in M_f (and hence a quasiconvex wall in the universal cover). First, take a simple closed curve C in S_g which is disjoint from its f -image. Such a simple closed curve exists up to taking a finite cover of M_f . The required immersed surface in M_f is then obtained by cutting S_g along C and $f(C)$ and gluing C_{\pm} to $f(C_{\mp})$ (the cut-and-cross-join technique).

Hagen and Wise mimicked this construction for the setup of hyperbolic free-by-cyclic groups in the fully irreducible case. A surface with a pseudo-Anosov map is now replaced by a graph with a train track map. The analogue of cutting along a simple closed curve is cutting along a point in the graph. However, the situation here is more complicated as a train track map is only a homotopy equivalence and not a homeomorphism. A point often has multiple preimages and the cut-and-cross-join operation is performed along all points with the same image.

We use the same operation, but now our train track representative is defined not on a finite graph but on the G -tree T . The lack of local finiteness of the tree T gives rise to additional difficulties, but we were able to manage them because of the behaviour of train track maps in this setting, and considerations of angles at vertices under relative hyperbolicity.

While cocompact cubulations require walls to be (relatively) quasiconvex, relative cubulations require walls to also be full. This forced us to introduce *saturations* of our walls in order to ensure fullness. The construction of walls and their saturations can be found in [Section 4](#).

Finally, in order to use the boundary criterion, we need sufficiently many wall saturations to not only cut biinfinite geodesics in the flow space, but also to cut pairs of *principal flow lines* that are stabilised by maximal parabolic subgroups and pairs consisting of a geodesic ray and a principal flow line. We show this in [Section 5](#), ensuring the separation of every pair of points in the Bowditch boundary. We verify the latter in [Section 6](#), where we also give a proof of [Theorem 1.1](#).

1.3 Questions

We end this introduction with three questions arising from this work.

Question 1.4 Let $G = A * B$ be a torsion-free group and ϕ be an automorphism that is fully irreducible and atoroidal relative to the above free decomposition. Does the mapping torus of G admit a relative cubulation?

The above makes a case for a combination theorem of relatively cubulated groups, which is as yet a largely unexplored area of research.

Question 1.5 Let G be a free product and ϕ be an automorphism that is fully irreducible but not necessarily atoroidal, relative to the given free decomposition. Let \mathcal{P} be the peripheral structure of suspensions of maximal subgroups of G on which iterations of ϕ make lengths of conjugacy classes grow at most polynomially. Does the mapping torus of G admit a relative cubulation?

Motivated by work of the first author and Li [\[DL22\]](#), such a construction would apply to free group automorphisms, refining the cartography of possible cubulations of free-by-cyclic groups.

In [\[DM23\]](#), we showed that the mapping torus of a torsion-free hyperbolic group is hyperbolic relative to the suspensions of the maximal polynomially growing subgroups. This leads to a natural question:

Question 1.6 Let G be a torsion-free hyperbolic group and ϕ be an automorphism of G . Does the mapping torus of G admit a relative cubulation?

The answer is yes when ϕ is atoroidal; see our work with Mutanguha [\[DMM25\]](#).

Acknowledgements Meda Satish was partially supported during this work by CEFIPRA grant 5801-1, *Interactions between dynamical systems, geometry and number theory* at TIFR and by ISF grant 1226/19 at the Technion. Dahmani is supported by ANR-22-CE40-0004 GoFR. Both acknowledge support by LabEx CARMIN, ANR-10-LABX-59-01. Different institutions must be thanked for hosting different stages of this work: the IHP (UAR 839 CNRS-Sorbonne Université) hosting the trimester program *Groups acting on fractals* in 2022, the Institut Fourier, the Math department of the Technion, for mutual visits, and the IHES.

We would like to thank Mark Hagen, Thomas Ng and Dani Wise for helpful conversations, and the referee for numerous comments which made the paper more pleasant.

2 The flow space

2.1 Free G -trees relative to \mathcal{H} , and train track maps

Let us fix G to be a finitely generated group for the rest of the paper.

A free factor system for G is a tuple (H_1, \dots, H_k) of subgroups such that there exists a free subgroup $F < G$ for which $G = H_1 * H_2 * \dots * H_k * F$. Another free factor system (J_1, \dots, J_ℓ) of G is strictly larger if each H_i is conjugate into some J_r , and one inclusion is strict.

A G -tree is a metric tree endowed with an isometric action of G .

A *free G -tree relative to $\mathcal{H} = \{H_1, \dots, H_k\}$* is a G -tree which is minimal for its G -action, its edge stabilisers are trivial, and its nontrivial elliptic subgroups are exactly the conjugates of the $\{H_1, \dots, H_k\}$ in G . We may as well require that there is no vertex of valence 2 that has trivial stabiliser.

A vertex is *singular* if its stabiliser is nontrivial.

Observe that, because G is finitely generated, any such G -tree has finite quotient by G . There is a whole space of such free G -trees relative to $\{H_1, \dots, H_k\}$, as studied in [GL07].

It is convenient to have a notion of *angle* in these nonlocally finite trees. Let T be a free G -tree relative to \mathcal{H} . Let us choose a word metric for each H_j , which is finitely generated as G is. Let $v_j \in T$ be the unique vertex fixed by H_j , and finally choose a finite set of orbit representatives for the H_j -action of edges issuing from v_j . Define the angle between two edges e and e' issuing from the vertex fixed by H_j to be the word length of the element hh'^{-1} given by h and h' sending e and e' , respectively, into our set of representatives. If e and e' are distinct edges in the finite set of orbit representatives, the angle between them is one. It is clear that the $\text{Stab}(v)$ -action on edges adjacent to v preserves angles. We hence complete the definition by G -equivariance: it defines the angle between two edges adjacent to singular vertices stabilised by conjugates of the H_j . In other (locally finite) vertices, we may say that angles are 0 (if edges are equal) or 1 (if they are different). Observe that only finitely many edges make a given angle with a given edge, and that angles (around a given vertex) satisfy the triangle inequality.

Let ϕ be an automorphism of G that preserves the conjugacy class of each H_i . Consider a free G -tree T relative to $\{H_1, \dots, H_k\}$. One says that a continuous map $f: T \rightarrow T$ *realises the automorphism ϕ* if it is equivariant in the sense that for all $p \in T$ and all $g \in G$, $f(gp) = \phi(g)f(p)$. Such maps realising ϕ exist in our context [FM15, Lemma 4.2]. Such a map is also, by equivariance, a quasiisometry of T .

A *turn* is a pair of edges e_1 and e_2 in T starting at a common vertex v . The pair is a *proper turn* if e_1 and e_2 are distinct. We say that e_1 and e_2 make a *legal turn* if the two paths $f(e_1)$ and $f(e_2)$ starting at $f(v)$ share no proper common subpath. In other words, by definition, f sends legal turns to proper turns. We say that f is a *train track map* if it sends edges to reduced paths without nonlegal turns, and if moreover, for all such e_1 and e_2 in T making a legal turn, their images $f(e_1)$ and $f(e_2)$ are two paths starting at $f(v)$ by two edges that themselves make a legal turn. In other words, by definition, f is a train track map if it sends legal turns to legal turns.

It is far from obvious that train track maps representing automorphisms exist. A theorem of Francaviglia and Martino [FM15] ensures that if ϕ is fully irreducible relative to $\{H_1, \dots, H_k\}$, then there exists a free G -tree relative to $\{H_1, \dots, H_k\}$, which we denote by T , and there exists $f: T \rightarrow T$ continuous, with constant speed on edges, that realises ϕ and is a train track map.

Recall that, if (H_1, \dots, H_k) is a free factor system of G , and if ϕ is an automorphism permuting it (ie preserving the set of conjugacy classes $\{[H_1], \dots, [H_k]\}$), we say that ϕ is *fully irreducible relative to* $\{H_1, \dots, H_k\}$ [FM15, Definition 8.2] if no positive power of ϕ preserves any larger free factor system.

We recall for later use an equivalent formulation of irreducibility (see [FM15, Definition 8.1 and Lemma 8.3]). Let $f: T \rightarrow T$ be a map realising the automorphism ϕ . We say f is *irreducible* if for every proper subgraph W of T that is f -invariant and G -invariant, the quotient of W by G is a forest such that each component subtree contains at most one nonfree vertex. The map f is *fully irreducible* if f^i is irreducible for all $i > 0$. We say ϕ is (fully) irreducible relative to $\{H_1, \dots, H_k\}$ if every f realising ϕ is (fully) irreducible.

2.2 The flow space of an automorphism

From now on T and f are thus chosen so that f is a train track map on T representing the relatively fully irreducible ϕ .

Define for all $i \in \mathbb{Z}$ a tree T_i equivariantly isomorphic to T . Let us denote by $(p \mapsto p_i)$ the identification $T \rightarrow T_i$. Define the action of G on T_i as $g.p_i = (\phi^i(g)p)_i$. Observe that this makes T_i a free G -tree relative to $\{H_1, \dots, H_k\}$. Define $f_i: T_i \rightarrow T_{i+1}$ by $f_i(p_i) = (f(p))_{i+1}$. Observe, by the property of train tracks, how f_i sends a turn: if it is legal when seen in T , its image is a legal turn when seen in T . We will make the abuse of language that any composition of the form $f_{i+k} \circ f_{i+k-1} \circ \dots \circ f_{i+1} \circ f_i$ (from T_i to T_{i+k+1}) is called a (positive) iteration of f .

Start with the disjoint union of all the T_i for $i \in \mathbb{Z}$. For each i , and each (unoriented) edge e_i in T_i , we choose an orientation and glue a rectangle $R_{e_i} = [0, 1] \times [0, 1]$ so that $\{0\} \times [0, 1]$ is glued to e_i and $\{1\} \times [0, 1]$ is glued on the path $f_i(e_i)$ in T_{i+1} (respecting orientation). See Figure 1 for an illustration. Finally, for each vertex v_i , we glue together the sides $[0, 1] \times \{0\}$ of the rectangles R_{e_i} for all e_i starting at this vertex, and the sides $[0, 1] \times \{1\}$ of the rectangles R_{e_i} for all e_i terminating at this vertex. The resulting space, with a natural structure of a cell complex, is denoted by \tilde{X} . We call this space the *flow space of ϕ* (with respect to \mathcal{H} , T and f). Thus the flow space is a *tree of spaces*, where the underlying tree is a biinfinite combinatorial line and vertex spaces are the trees T_i .

In \tilde{X} , we will call any edge in some tree T_i a vertical edge, and the image of a side $[0, 1] \times \{1\}$ or $[0, 1] \times \{0\}$ of the rectangle $R(e_i)$ a horizontal edge. We call a path horizontal if it intersects rectangles in horizontal segments $[0, 1] \times \{h\}$.

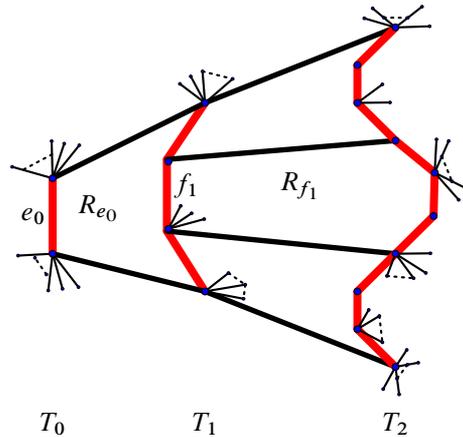


Figure 1: 2-cells in the flow space. The 2-cell R_{e_0} is bounded by the vertical edge e_0 on the left, two horizontal edges, and three vertical edges (including f_1) in T_1 .

If $f : T \rightarrow T$ is a train track map realising ϕ , then for all $L \geq 1$, the map f^L is a train track map realising the automorphism ϕ^L . Keeping the same T and f we may thus produce the space \tilde{X}_L by using the map f^L realising ϕ^L . We will prefer to index the trees by $L\mathbb{Z}$ though.

Observe that \tilde{X}_L need not be isomorphic to \tilde{X} . However we have a natural map $\varrho_L : \tilde{X}_L \rightarrow \tilde{X}$: it is obviously defined on the trees T_{Li} , and on a rectangle $R(e_{Li})$, one sends a horizontal edge onto the unique reduced concatenation of horizontal edges of \tilde{X} with the same endpoints.

2.3 Action on the flow space

Consider the semidirect product $G \rtimes_{\phi} \mathbb{Z}$, and write t for the generator of \mathbb{Z} that induces the automorphism ϕ by conjugation on G : $G \rtimes_{\phi} \mathbb{Z} = \langle G, t \mid t^{-1}gt = \phi(g), \forall g \in G \rangle$. Then $G \rtimes_{\phi} \mathbb{Z}$ acts cofinitely on \tilde{X} by defining the action of G on each T_i as above, and defining the action of t to be the shift of indices: t maps T_i on T_{i+1} isometrically, through the identification with T .

The group G preserves each of the trees T_i , and in each of them it induces the action of G on T precomposed by ϕ^i . So the orbits of G on each of the T_i are the same, after identification to T . In particular:

Lemma 2.1 *If $x, y \in T$ are such that there is $g \in G$ for which $gx = y$, and if $i_1, i_2 \in \mathbb{Z}$, then there exists $\gamma \in G \rtimes_{\phi} \mathbb{Z}$ for which $\gamma x_{i_1} = y_{i_2}$. □*

2.4 Forward flow, backward flow and principal flow lines

For each i , from each point x_i of T_i , there is a unique horizontal ray starting at x_i , containing all its images by positive iterates of f . We denote it by $\sigma(x_i)$, and $\sigma_k(x_i)$ is its initial subsegment of length k . Its endpoint is denoted (slightly abusively) by $f^k(x_i)$ as already mentioned. The segment $\sigma_1(x_i)$ of this ray whose endpoints are x_i and $f_1(x_i) \in T_{i+1}$ is called the midsegment at x_i . We call the ray $\sigma(x_i)$ the *forward flow (ray)* from x_i (forward path in [HW16]). It is sometimes useful to use the forward flow

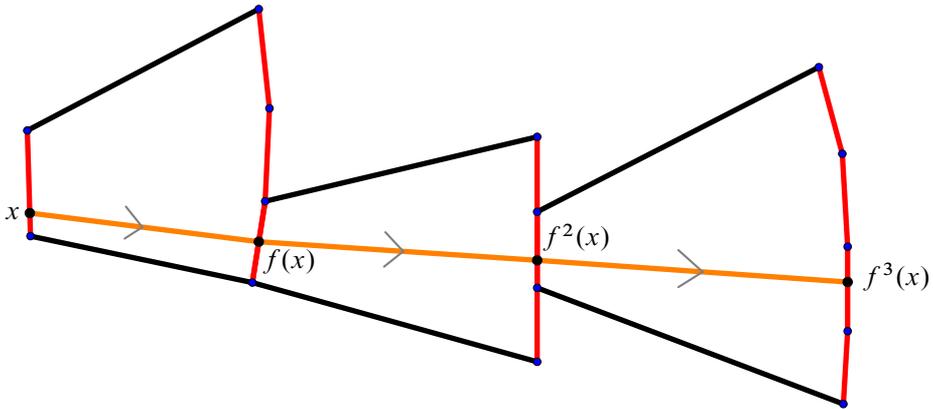


Figure 2: The orange segment between x and $f^3(x)$ is the forward flow of length 3 from x . Its forward ladder is the union of the 2-cells in the picture.

of length L , which is the initial subsegment of length L of the forward flow ray (see Figure 2 for an illustration). The *forward ladder* of a forward flow (ray or segment) σ , denoted by $N(\sigma)$, is the smallest subcomplex of the cell complex \tilde{X} containing σ .

We emphasise that the forward flow is different from the action of t .

In the backward direction, we note the following:

Lemma 2.2 *For any $x_i \in T_i$ with x_i not a singular vertex, its preimage $f_{i-1}^{-1}(x_i)$ in T_{i-1} is finite. In particular, every vertical edge of \tilde{X} is contained in finitely many 2-cells.*

Proof It suffices to prove that each edge of T is contained in finitely many images of edges of T under f . Assume on the contrary that there are infinitely many such edges e_k , whose images meet x_i . Since there are finitely many G -orbits of edges in T , we may assume that all edges e_k are images of e_0 by some elements $g_k \in G$, all different. The map f being a quasiisometry of T , all edges e_k are at bounded distances, so the g_k have bounded displacement. This forces that for large k , the g_k and their images by ϕ have a contribution in some of the free factors H_i that is larger than the maximal angle of the path $f(e_0)$. However $f(e_k) = f(g_k(e_0)) = \phi(g_k)f(e_0)$. This forces $f(e_k)$ and $f(e_0)$ to be disjoint, a contradiction. \square

In the case of the preimage of an infinite-valence vertex of T_i , we have an even clearer picture. For $j \in \mathbb{N}$, let us denote by $f^{-j}(v)$ the set $(f^j)^{-1}(v)$.

Lemma 2.3 *Let $v \in T_i$ be a singular vertex of T_i . Then for each $j \in \mathbb{N}$, the set $f^{-j}(v)$ is finite and contains a unique singular vertex.*

Proof We may assume that $v \in T$. The set $f^{-j}(v)$ lies in T . We will first show that there is a unique singular vertex in $f^{-j}(v)$ and then show that the set is finite.

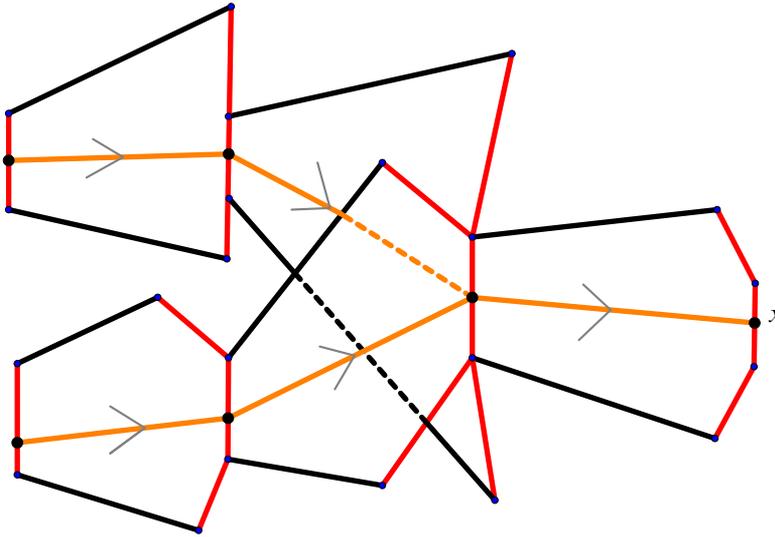


Figure 3: The backward flow of length 3 from x (in orange).

Assume that it contains two different singular vertices w_1 and w_2 such that $f^j(w_1) = f^j(w_2) = v$. By equivariance of f , the stabilisers G_{w_1} and G_{w_2} of w_1 and w_2 are sent by ϕ^j inside the stabiliser of G_v . But they are also sent onto stabilisers of vertices, since ϕ preserves the conjugacy classes of the H_k . Thus $\phi^j(G_{w_1}) = \phi^j(G_{w_2})$, and therefore $G_{w_1} = G_{w_2}$. In particular, the tree-geodesic between the vertices w_1 and w_2 is pointwise fixed by G_{w_1} , and since $w_1 \neq w_2$, this contradicts the triviality of stabilisers of edges of the tree.

Assume now that the preimage $f^{-j}(v)$ has infinitely many nonsingular points. Denote by $(e_i)_{i \in \mathbb{N}}$ the collection of edges containing a preimage. Since f is a quasiisometry, they are all at a bounded distance from each other. Up to extracting a subsequence, we may assume they are in the same G -orbit: write $e_i = g_i e_0$. Being at bounded distance from each other, the displacement of g_i remains bounded. Consider the segment s_i between e_0 and $e_i = g_i e_0$ containing both. Taking a subsequence if necessary, we may assume that there is a vertex w for which the s_i have the same prefix until w and then start having angles going to infinity at this vertex. Consider $f(e_i) = \phi(g_i)f(e_0)$. The displacement of $\phi(g_i)$ also remains bounded, and the angle of $f(s_i)$ at the image $f(w)$ tends to infinity. After w , all these images of the $f(s_i)$ are thus disjoint. It follows that all the images e_i are either disjoint, or share possibly only $f(w)$. In other words, the $f(e_i \setminus \{w\})$ are disjoint. Since at w , angles are arbitrarily large, it is a singular point, contradicting the initial assumption that the e_i each contain a nonsingular preimage of a point. \square

We may now define the backward flow. For a point $x_i \in T_i$ the backward flow (level in [HW16]) $\tau_L(x_i)$ from x_i of length L is the union of length L forward flows from each point in $f^{-L}(x_i)$; see Figure 3. Note that $\tau_L(x_i)$ is a rooted tree, rooted at x_i . When $x \notin \tilde{X}^0$, by Lemma 2.2, $\tau_L(x)$ is a finite rooted tree. In particular, [HW16, Proposition 2.5] gives the following observation:

Proposition 2.4 *Let $x \notin \tilde{X}^0$. Then for any $L \geq 0$, there exists a topological embedding $\tau_L(x) \times [-1, 1] \rightarrow \tilde{X}$ such that $\tau_L(x) \times \{0\}$ maps isomorphically onto $\tau_L(x)$.* \square

Finally, by [Lemma 2.3](#), there is a unique singular vertex v_j in each T_j for $j < i$ that is in the backward flow of a singular vertex v_i of T_i . Hence we can construct the *principal flow line* of v_i to be the direct limit of the forward flow rays of v_j , for $j \rightarrow -\infty$. It is well defined, biinfinite, as all v_j as above, and all images of v_i by positive iterates of f have the same principal flow line.

2.5 Geometry of the flow space

In order to have relative hyperbolicity for the group $\Gamma := G \rtimes_{\phi} \mathbb{Z}$, following [\[DL22\]](#), we will need the following three additional properties, which we will assume to hold from now on:

- Any fundamental domain of T contains at least two edges, ie $k + r \geq 3$ or $r \geq 2$.
- The automorphism ϕ is *atoroidal relative to* $\{H_1, \dots, H_k\}$: given any element $g \in G$ such that g is not contained in any conjugate of any H_i , then for all $n \in \mathbb{N}$, $[\phi^n(g)] \neq [g]$.
- The automorphism ϕ has *no twinned subgroups*: given two subgroups $H \neq K$ such that $[H], [K] \in \mathcal{H}$, then given any $g \in G$ and any $m \in \mathbb{N}$, $\phi^m(H) \neq gHg^{-1}$ whenever $\phi^m(K) = gKg^{-1}$.

Remark 2.5 We in fact do not need to assume the no twinning property, as by [\[DMM25, Lemma 3.3\]](#), it automatically holds whenever ϕ is relatively fully irreducible and $k + r \geq 3$. We clarify that the case $r \geq 2$ is redundant if $k = 0$ as there are no atoroidal maps of F_2 .

Let us recall that the *cone-off* of a graph Z over a family of subgraphs \mathcal{L} , which we denote here by \hat{Z} , is the graph obtained by adding to Z a vertex v_L for each $L \in \mathcal{L}$ and an edge between each vertex of L and v_L . Usually, each such edge is assigned to have length $\frac{1}{2}$.

It can help to picture T as equivariantly quasiisometric (by a collapse map) to a cone-off of the Cayley graph of G over the left cosets of the free factors, and to picture \tilde{X} as equivariantly quasiisometric to a cone-off of the Cayley graph of Γ over the left cosets of the same free factors of $G < \Gamma$. Finally, one can see the cone-off of \tilde{X} over principal flow lines as equivariantly quasiisometric to the cone-off of the Cayley graph of Γ over the cosets of the suspensions of the free factors of G .

Theorem 2.6 *The 1-skeleton \tilde{X}^1 is δ -hyperbolic for some $\delta > 0$. Moreover, the cone-off of \tilde{X}^1 over the principal flow lines is also hyperbolic.*

This is essentially proved in [\[DL22\]](#). Let us cover how to obtain it. We want to use [\[MR08, Theorem 4.5\]](#). However, a little care is in order. The assumption of this theorem is that one has a tree (denoted by T in [\[MR08\]](#), and temporarily denoted by T_{MR} here) of relatively hyperbolic spaces S_v , for $v \in T_{MR}^{(0)}$, with some properties, and the conclusion is that the whole space, denoted by X in [\[MR08\]](#) and temporarily

denoted by X_{MR} , is itself hyperbolic relative to maximal cone subtrees. Here the tree T_{MR} is just a biinfinite line \mathbb{Z} indexing our trees T_i , and the spaces S_i for $i \in \mathbb{Z}$ are indeed the trees T_i , all isometric to T , on which G acts through ϕ^i , and the total space X_{MR} is \tilde{X} . The different assumptions of the theorem were checked in [DL22, Section 2.3], and the conclusion is that \tilde{X} is hyperbolic relative to a collection of quasiconvex lines (the principal flow lines of singular vertices). It follows that \tilde{X} is hyperbolic itself. It also implies that its cone-off over those lines is hyperbolic as well.

Recall that a graph is said to be *fine* if for each $n \in \mathbb{N}$, every edge of the graph belongs to only finitely many simple cycles of length n . Recall also that a finitely generated group Γ is hyperbolic relative to a finite collection of finitely generated subgroups \mathcal{P} if the cone-off of a Cayley graph of Γ over the left cosets of elements of \mathcal{P} is hyperbolic and fine [Bow12].

The main result of [DL22] (Theorem 0.2) was actually a related statement, that the group $G \rtimes_{\phi} \langle t \rangle$ is itself relatively hyperbolic with respect to the collection \mathcal{P} consisting of mapping tori of the subgroups H_1, \dots, H_k . Since the stabiliser of a singular vertex in T_0 is a conjugate of some H_i , each principal flow line is in fact cocompactly stabilised by the suspension of a conjugate of some H_i (this can, for instance, be deduced from Lemma 2.15(2)).

2.6 Quasiconvexity and divergence of forward flow rays

We first observe, following [HW16], that forward flow rays are uniformly quasiconvex (Proposition 2.7).

From $x_i \in T_i$ to $f^n(x_i)$ any path intersects each tree T_{i+k} and therefore has at least n horizontal edges. However, the forward-path has exactly n horizontal edges and no vertical contribution. If D is an upper bound to the diameter of 2-cells, we see that the intersection of the forward flow ray with $\bigcup_i T_i$ is a D -bilipschitz embedding of \mathbb{N} in the 1-skeleton of \tilde{X} . Since the latter is hyperbolic, we have quasiconvexity. We thus have:

Proposition 2.7 [HW16, Proposition 2.3] *There exists $\lambda \geq 0$ such that the 1-skeleton $N(\sigma)^1$ of any forward ladder is λ -quasiconvex in \tilde{X}^1 .* □

Second, two flow rays starting at different points in the same edge of T_i diverge from each other:

Lemma 2.8 *Given $\epsilon > 0$, there exists $N > 0$ such that for any two points x and y contained in any vertical edge $e_i \subset T_i$ with $d(x, y) \geq \epsilon$, the distance in T_{i+N} between $f^N(x)$ and $f^N(y)$ is at least $e^{100(\delta+\lambda)}$ (and thus the forward rays from x and y diverge).*

Proof The forward images of a single edge are all legal paths, and by [DL22, Lemma 1.11] f applies a uniform stretching factor > 1 on all legal paths. □

Finally, two flow rays starting at different points in the star of a vertex either uniformly diverge from each other or fellow-travel forever. More precisely we have the following:

Proposition 2.9 For all $\epsilon > 0$, there exists $N > 0$ such that, for all singular vertices $v \in T_i$, for all edges $e_1 \neq e_2$ in T_i starting at v , and for all x_1 and x_2 in e_1 and e_2 , respectively, at distance at least ϵ from v , either the distance in T_{i+N} between $f^N(x_1)$ and $f^N(x_2)$ is at least $e^{100(\delta+\lambda)}$, or the forward rays fellow-travel until infinity.

Proof We fix ϵ and v . For each edge e issued at v , and $x \in e$ at distance at least ϵ from v , we know by Lemma 2.8 that there is N_ϵ such that, for all $n \geq N_\epsilon$, $f^n(x)$ is at distance (in T_{i+n}) at least $e^{1000(\delta+\lambda)}$ from the principal flow line of v .

By [DL22, Lemma 2.6] f induces a quasiisometry on angles: there exists $\theta_1 > 1$ such that, if e_1 and e_2 issued at v make an angle θ , then the paths $f(e_1)$ and $f(e_2)$ make an angle at least $\theta/\theta_1 - \theta_1$ at their common initial point $f(v)$ — while it is slightly stronger than the stated lemma, the claim is still true with same proof, however, the stated lemma is also sufficient to be used in a similar way. Call $\rho(\theta) = \theta/\theta_1 - \theta_1$. Calculus ensures that it is possible to find θ_0 such that $\rho^{N_0}(\theta_0) > 0$. If e_1 and e_2 issued at v make an angle greater than θ_0 , the paths $f^{N_0}(e_1)$ and $f^{N_0}(e_2)$ make a positive angle at their initial common point $f^{N_0}(v)$. In particular, they do not overlap. It follows that, measured in T_{i+N_0} , the distance $d(f^{N_0}(x_1), f^{N_0}(x_2))$ is equal to $d(f^{N_0}(x_1), f^{N_0}(v)) + d(f^{N_0}(v), f^{N_0}(x_2))$. It is then greater than $2e^{1000(\delta+\lambda)}$. Since the two rays are quasiconvex in the hyperbolic space \tilde{X} and have started to diverge after N_0 edges, they will diverge onward after that point.

It remains to treat the case of two edges making an angle less than θ_0 . There are finitely many $\text{Stab}(v)$ -orbits of such pairs of edges. Partition them into classes of those whose forward rays fellow-travel, and those whose forward rays exponentially diverge. Since there are finitely many, one can choose N that is suitable for all of those in the second class, and larger than N_0 . \square

Observe that this does not prevent some points in different edges from having fellow-travelling forward rays, with the same endpoint at infinity. However, if the origins of the flow rays are translates of each other by an elliptic element g in T_i , they diverge, provided that g is not a torsion element:

Lemma 2.10 For any point x_i in T_i , and for each $g \in G$ elliptic in T_i such that $g^n x_i \neq x_i$ for any $n \in \mathbb{N}$, the flow rays from $g x_i$ and x_i do not fellow-travel until infinity.

Proof Note that the midpoint of the segment of T_i between x_i and $g x_i$ is fixed by g , and is hence a singular vertex v_i . If the flow rays from x_i and $g x_i$ fellow-travel until infinity, then by equivariance, so do the flow rays from $g^n x_i$ and $g^{n'} x_i$, for all $n, n' \in \mathbb{Z}$. But for large enough powers, the segments from v_i to $g^n x_i$ and x_i make arbitrarily large angles. By arguing as in Proposition 2.9, we can conclude that the flow rays from x_i and $g^n x_i$ cannot fellow-travel for arbitrarily large powers, a contradiction. \square

Henceforth, we will assume that the group G is torsion-free, so the above result always holds for elliptic elements.

Let us also record a basic observation about principal flow lines:

Lemma 2.11 *Let $v_1 \neq v_2$ be two singular vertices in \tilde{X} . Then either the principal flow lines of v_1 and v_2 coincide or uniformly diverge from each other in both the forward and backward directions, ie for each $R > 0$ there exists $B_R > 0$ such that for any pair of singular vertices $v_1 \neq v_2$, either the principal flow lines Λ_{v_i} through v_i coincide, or $N_R(\Lambda_{v_1}) \cap \Lambda_{v_2}$ has diameter at most B_R .*

Proof Observe that by Lemma 2.3, for each singular vertex there exists a unique principal flow line going through this vertex. Since the automorphism ϕ has no twinned subgroups for the (finite) collection \mathcal{H} , if two principal flow lines are different, they must uniformly diverge as required. \square

2.7 Periodic points and forward flow rays

In the case of certain special points called periodic points, we have no fellow-travelling of flow rays between any translates, which we prove below. We will also show that periodic points with flow rays (lines, in fact) diverging from every principal flow line are dense in edges of T (Lemma 2.18).

A point $x \in T$ is a *periodic point* if there exist $g = g_x \in G$ and $n = n_x > 0$ such that $f^n(x) = gx$. If x is periodic, then we will also call each $x_i \in T_i$ a periodic point. We say that the periodic point $x \in T$ has *period* n if $f^n(x) = gx$ as above and for all $0 < k < n$, $f^k(x) \neq hx$ for any $h \in G$.

Proposition 2.12 *For any periodic point x_i in T_i , and for each $g \in G$ with $gx_i \neq x_i$, the flow rays from gx_i and x_i diverge.*

Proof Observe that by Lemma 2.10, the statement has to be proved only for loxodromic elements g . Since x_i is periodic, there is a minimal positive k_i and $g_i \in G$ such that $g_i x_{i+k_i} = (f^{k_i}(x_i))$ (which is in T_{i+k_i}). Thus the element $g_i t^{k_i}$ sends x_i to $f^{k_i}(x_i)$ in the flow space. This flow space is hyperbolic, and this element is loxodromic, since for any $x \in T_0$, $d_{\tilde{X}^1}((g_i t^{k_i})^r x, x) \geq rk_i$ (the distance separating T_0 from T_{rk_i}). It follows that the forward flow from x_i remains at bounded distance from $((g_i t^{k_i})^n x_i)_{n \in \mathbb{N}}$, and hence from the axis of $(g_i t^{k_i})$. Now, from gx_i , the forward flow remains close to the axis of $g(g_i t^{k_i})g^{-1}$.

Since the element g is loxodromic on T_i , g and $g_i t^{k_i}$ generate a nonelementary subgroup of isometries of \tilde{X} , because their axes eventually diverge. It follows that the limit points of $g_i t^{k_i}$ and of g in the boundary of \tilde{X}^1 cannot be the same. One concludes that g does not fix the limit points of $g_i t^{k_i}$, and $g(g_i t^{k_i})g^{-1}$ and $(g_i t^{k_i})$ thus have divergent axes. \square

Lemma 2.13 *For all edges e and e' in T , there exist $n \geq 0$ and $g \in G$ such that $f^n(e)$ contains ge' .*

Proof Assume that for some e and e' , for every n , the path $f^n(e)$ does not contain any G -translate of e' . Note that this means that no G -translate of $f^n(e)$ contains any G -translate of e' . Let W be the union of all G -translates of all paths $f^n(e)$. Then W is a proper subgraph which is both f -invariant

and G -invariant. By the irreducibility of f , the quotient of W by G is a (bounded) forest with at most one nonfree vertex in each component. Thus the quotient of each $f^n(e)$ is a segment with at most one nonfree vertex. This is not possible as the lengths of $f^n(e)$ are arbitrarily long. \square

Lemma 2.14 *For each $\epsilon > 0$ and each closed subinterval d of any edge e in T of length ϵ , there exists a periodic point x in d .*

Proof Lemma 1.11 of [DL22] ensures that there exists a growth factor $\lambda > 1$ such that every subsegment of every edge of T expands under f by λ . Thus, given $d \subset e$, there exists n such that $f^n(d)$ contains an edge. The result then follows from Lemma 2.13, by applying Brouwer's fixed-point theorem. \square

Let us now observe the following useful facts about periodic points. Given a periodic point x of period n , the preimage $f^{-n}(x)$ contains a G -translate of x (Lemma 2.15(1)). One can then define a *periodic flow line* through x as the direct limit of forward flow rays from the G -translates of x at the various $f^{-nk}(x)$. We will show in Lemma 2.15(2) that a periodic flow line is really *periodic*, ie there is an infinite subgroup of Γ that stabilises the line and acts cocompactly on the line.

Lemma 2.15 *Let $x \in T$ be a periodic point of period n . Then:*

- (1) *There exists $g' \in G$ such that $g'x \in f^{-n}(x)$.*
- (2) *The periodic flow line through x is periodic. In particular, principal flow lines are periodic and stabilised by suspensions of the relevant free factors of G .*

Proof (1) Let $g \in G$ be such that $f^n(x) = gx$. Observe that, by equivariance of f , for $g' = \phi^{-n}(g^{-1})$, $g'x \in f^{-n}(x)$. Indeed, $f^n(\phi^{-n}(g^{-1})x) = \phi^n(\phi^{-n}(g^{-1}))f^n(x) = g^{-1}f^n(x) = x$.

(2) Consider the flow segment $\sigma_n(x)$ from $x \in T_0$ to $f_{n-1} \circ \dots \circ f_0(x) \in T_n$. Recall that the element t acts by shifting indices, and therefore, $f_{n-1} \circ \dots \circ f_0(x) = t^n f^n(x)$. This implies that $t^n g(\sigma_n(x))$ is the flow segment from $t^n gx \in T_n$ to $t^n gt^n f^n(x) \in T_{2n}$. Using the facts that $gt^n = t^n \phi^n(g)$ and $gx = f^n(x)$, we conclude that $t^n g(\sigma_n(x))$ is the flow segment from $t^n f^n(x)$ to $t^{2n} \phi^n(g) f^n(x)$. Since $\phi^n(g) f^n(x) = f^n(gx) = f^{2n}(x)$, we have that $\sigma_n(x) \cdot t^n g(\sigma_n(x)) = \sigma_{2n}(x)$. It is also easy to check that the flow segment of length n from $t^{-n} g'(x)$ is equal to $(t^n g)^{-1}(\sigma_n(x))$. Thus the periodic flow line through x is the union over all $k \in \mathbb{Z}$ of $(t^n g)^k(\sigma_n(x))$. \square

Lemma 2.16 *If a periodic flow line is asymptotic to another periodic flow line in one direction, they are asymptotic in both directions.*

Proof We will prove the contrapositive. Assume Λ' and Λ'' are periodic flow lines that are nonasymptotic in the forward direction (the other case is similar). Since they are periodic, there are elements γ' and γ'' in the group Γ (Lemma 2.15(2)) that fix the endpoints of Λ' and Λ'' , respectively, in the boundary $\partial \tilde{X}^1$ of

the hyperbolic space \tilde{X}^1 . In particular, γ' and γ'' cannot be elements of the same parabolic subgroup of Γ . Assume that the backward directions of Λ' and Λ'' converge to the same point of $\partial\tilde{X}^1$. It would follow that the commutator $[\gamma', \gamma'']$ is in G and has small displacement on infinitely many T_{-i} for $i \in \mathbb{N}$, realised near the ray Λ' . Since γ' and γ'' have different axes, they do not commute. However, by hyperbolicity of the automorphism, $[\gamma', \gamma'']$ must be elliptic, and hence it is in a unique conjugate of a free factor of the free product G . It hence fixes nonfree vertices in each T_{-i} , at bounded distance from Λ' , and all in the same principal flow line. It then follows that both Λ' and Λ'' are asymptotic to the same principal flow line, and therefore that the elements γ' and γ'' preserving them fix a point fixed by a parabolic group, thus ensuring that they are in the same parabolic subgroup, contradicting what we previously had. \square

Given two flow lines that intersect the tree T_0 in x and y , we say that the two flow lines are separated by the segment $[x, y]$ in T_0 . The next statement says that a periodic flow line either diverges from every principal flow line, or it is asymptotic to a principal flow line which is Hausdorff-close to it in terms of both distance and angle.

Lemma 2.17 *There exist constants δ_0 and θ_0 such that the following holds: Let Λ be a periodic flow line in \tilde{X} and x be its intersection with T_0 . If a principal flow line Λ' , with intersection y at T_0 , is asymptotic to Λ , then $d_{T_0}(x, y) \leq \delta_0$. Further, there exists a principal flow line Λ'' , asymptotic to Λ and Λ' , such that if $z \in \Lambda'' \cap T_0$, then for every vertex v in the interior of the segment $[x, z]_{T_0}$ the angle subtended by $[x, z]_{T_0}$ at v is less than θ_0 .*

Proof Consider a geodesic $[x, y]_{\tilde{X}^1}$ in \tilde{X}^1 . By asymptoticity of the biinfinite lines Λ and Λ' in the hyperbolic space \tilde{X}^1 , its length is at most δ . Denote by e_1, \dots, e_r its consecutive vertical edges, let T_{k_i} be the tree containing e_i and write $e_i = (v_i^{k_i}, w_i^{k_i})$. Observe that $r \leq \delta$ and $|k_i| \leq \delta$ for all i .

Since e_1 is the first vertical edge, $x \in f^{-k_1}(v_1^{k_1})$, and similarly, $y \in f^{-k_r}(w_r^{k_r})$. Take, for all $i < r$, $x_i \in f^{-k_i}(w_i^{k_i})$. Such a point x_i is in T_0 and is in $f^{-k_{i+1}}(v_{i+1}^{k_{i+1}})$. Therefore, for all $i < r$, $d_{T_0}(x_i, x_{i+1}) \leq K^{|k_{i+1}|}$, where K is the Lipschitz constant of f . Since for all i , $|k_{i+1}| \leq \delta$, and $r \leq \delta$, it follows that $d_{T_0}(x, y) \leq \delta K^\delta$.

We turn to the second part of the statement. By Lemma 2.10, for any edge e incident to a vertex w , and any $g \in \text{Stab}(w)$, e and ge form a legal turn and so their iterates under f subtend a positive angle at all $f^n(w)$. Since there are finitely many orbits of edges incident to w , there are only finitely many edges incident to w that make an illegal turn with e . Since there are finitely many orbits of vertices, there exists a maximal angle, which we call $\theta_0 - 1$, such that any illegal turn between any pair of edges at any vertex subtends an angle of at most $\theta_0 - 1$.

Assume now that there exists $v \in [x, y]_{T_0}$ such that $\text{Ang}_v[x, y]_{T_0} \geq \theta_0$. We choose v to be closest to x with this property. Note that, after changing θ_0 if necessary, v is a singular vertex. Let Λ_v be the associated principal flow line.

For all $n \in \mathbb{N}$, let us denote by x_{-n} the point $f^{-n}(x) \cap \Lambda$ and by y_{-n} the point $f^{-n}(y) \cap \Lambda'$. Let us also write $x_n = f^n(x)$, $y_n = f^n(y)$, $v_n = f^n(v)$ and v_{-n} the point $f^{-n}(v) \cap \Lambda_v$.

Since $\theta_0 - 1$ is the maximal angle at an illegal turn, for all $n \in \mathbb{N}$, $\text{Ang}_{v_n}[x_n, y_n]_{T_n} \geq 1$. In particular, v_n is between x_n and y_n in T_n for all $n \in \mathbb{N}$, and so the principal flow line Λ_v is asymptotic to Λ in one and hence both directions. \square

Lemma 2.18 *For each $\epsilon > 0$ and each closed subinterval d of any edge e in T of length ϵ , there exists a periodic point x in d such that the periodic flow line through x diverges from every principal flow line in both the forward and backward directions.*

Proof There are finitely many orbits of vertices, and hence finitely many orbits of periodic lines containing a vertex. Let n_0 be the maximum of their period. If a subinterval d of e is given, one may take a sufficiently small subinterval d' such that its images by f^k for $k = 1, \dots, n_0$ do not meet any vertex. This, together with the periodicity of periodic points, guarantees that any periodic point in this subinterval will have a periodic flow line that misses all vertices.

There are only finitely many vertices in the tree T_0 that are at distance $\leq \delta K^\delta + 1$ from d with a path that makes no angle greater than θ_1 . Let N be this number. This produces N principal flow lines.

Consider $2N + 1$ periodic points in the subinterval d' . Since the interval between any two of them is a legal path, their forward flow rays diverge from each other. Therefore at least one of these periodic points, say x , belongs to a periodic flow line that diverges from the N principal flow lines in both directions. By [Lemma 2.17](#) it diverges from all principal flow lines. \square

3 Relative cubulation

In this section, we will recall the notion of relative cubulation and the boundary criterion for relative cubulation. We will construct walls in [Section 4](#) for the flow space \tilde{X} and show in [Sections 5 and 6](#) that the stabilisers in $\Gamma = G \rtimes_\phi \mathbb{Z}$ of “wall saturations” satisfy the hypotheses required for the boundary criterion to hold.

Definition 3.1 [[EG20](#)] Let Γ be hyperbolic relative to a collection of subgroups \mathcal{P} . Then (Γ, \mathcal{P}) is *relatively cubulated* if there exists a CAT(0) cube complex C such that

- (1) there is a cubical cocompact action of Γ on C ,
- (2) each element of \mathcal{P} is elliptic, and
- (3) the stabiliser of any cube of C is either finite or conjugate to a finite-index subgroup of an element of \mathcal{P} .

We call such an action of Γ on C a *relatively geometric* action.

Let us recall a few definitions before stating the boundary criterion for relative cubulation. A subgroup $H \leq \Gamma$ is *relatively quasiconvex* (with respect to \mathcal{P}) if there exists a constant K such that given $h_1, h_2 \in H$, any geodesic path γ between h_1 and h_2 in the cone-off $\hat{\Gamma}$ (with respect to \mathcal{P}) is such that every noncone vertex of γ is at Γ -distance at most K from a vertex of H ; see [Hru10; Osi06].

A subgroup H of (Γ, \mathcal{P}) is *full* if the intersection of any conjugate of H with any element $P \in \mathcal{P}$ is either finite or is of finite index in P .

A subgroup H of Γ is a *codimension-1 subgroup* if for some $r \geq 0$, $\Gamma \setminus \mathcal{N}_r(H)$ contains at least two orbits of components that are not contained in any finite neighbourhood of H .

Let us also recall a definition (of the many equivalent ones) of the *Bowditch boundary* of a relatively hyperbolic group (Γ, \mathcal{P}) : it is a compact metrisable space $\partial_B(\Gamma, \mathcal{P})$ such that

- (1) Γ acts properly discontinuously on the space of distinct triples of $\partial_B(\Gamma, \mathcal{P})$,
- (2) each point of $\partial_B(\Gamma, \mathcal{P})$ is either a conical point or a bounded parabolic point, and
- (3) the stabiliser of a bounded parabolic point is a conjugate of an element of \mathcal{P} .

As a set, $\partial_B(\Gamma, \mathcal{P})$ is the union of the Gromov boundary of the cone-off $\hat{\Gamma}$ relative to \mathcal{P} and the set of cone vertices of $\hat{\Gamma}$. We refer to [Bow12, Section 9] for more details, in particular for the definitions of conical and bounded parabolic points, which will not be needed here.

The following theorem is implicit in the work of Bergeron and Wise [BW12] and explicitly stated and proved by Einstein and Groves [EG20].

Theorem 3.2 (boundary criterion for relative cubulation) *Let (Γ, \mathcal{P}) be relatively hyperbolic with one-ended parabolics. Suppose that for each pair of distinct points $u, v \in \partial_B(\Gamma, \mathcal{P})$ there exists a full relatively quasiconvex codimension-1 subgroup H of Γ such that u and v lie in H -distinct components of $\partial_B(\Gamma, \mathcal{P}) \setminus \Lambda H$. Then there exists a finite collection of full relatively quasiconvex codimension-1 subgroups such that the action of Γ on the dual cube complex is relatively geometric.*

We note that since our parabolics are suspensions of the infinite free factors of G , they are always one-ended.

4 Walls in the flow space

4.1 Constructing immersed walls

Our construction of walls starts with that of Hagen and Wise in [HW16], although they work in a locally finite setup. However, the stabilisers of walls thus obtained would not be full subgroups in general, so we later “saturate” these walls to obtain full codimension-1 subgroups.

Let us explain the construction of Hagen and Wise. Fix a fundamental domain $D \subset T$ for the G -action on T , ie a smallest subtree that contains exactly one edge from each orbit of edges. In order to construct immersed walls in \tilde{X} , we need the following data: a choice of a *tunnel length* $L \in \mathbb{N}$ and a choice of a subinterval of each open edge of D , called a *primary bust*.

Choose the tunnel length $L \in \mathbb{N}$. The immersed wall will be first constructed in \tilde{X}_L and then pushed to \tilde{X} via $q_L: \tilde{X}_L \rightarrow \tilde{X}$.

Choose a closed subinterval in the interior of each edge of D . This choice extends equivariantly to a choice of a subinterval d_k in every edge e_k of T . For each $i \in \mathbb{Z}$, the copy $d_{k,Li}$ in T_{Li} of each $d_k \subset e_k$ is a *primary bust*. Note that $f^{-L}(d_k)$ is a union of finitely many subintervals $\{d_{kj}\} \subset T$, by [Lemma 2.2](#). For each $i \in \mathbb{Z}$, the copy $d_{kj,Li}$ in T_{Li} of each $\{d_{kj}\}$ is a *secondary bust*. We will always choose primary busts satisfying the following:

Lemma 4.1 *Let $L \in \mathbb{N}$ and $\epsilon > 0$. Let $\{x_k\}$ be a collection of points in D , with exactly one point in any edge of D . The collection of primary busts in \tilde{X}_L can be chosen so that*

- (1) every primary bust is disjoint from every secondary bust in the collection,
- (2) the primary busts in D lie in the ϵ -neighbourhood of $\{x_k\}$,
- (3) the flow rays from the endpoints of the primary busts do not meet any vertex of \tilde{X}_L ,
- (4) if the flow ray from x_k does not meet a vertex, then one of the endpoints of the primary bust d_k in D can be chosen to be x_k ,
- (5) the f^L -image of any primary bust does not contain two points in the same G -orbit,
- (6) if $f^L(x_k) \neq f^L(x_j)$ whenever $k \neq j$, then $f^L(d_k) \cap f^L(d_j) = \emptyset$.

The lemma is a reproduction of [[HW16](#), Lemma 3.5], whose proof can be replicated in our case. Since there are several properties to check, the proof is long, but not difficult, based on Brouwer's fixed-point theorem. We thus refer the reader to [[HW16](#)] for the proof. In fact, (3) can be obtained either by the proof of Hagen and Wise [[HW16](#)], or by applying [Lemma 2.18](#).

Denote by $T_{Li+1/2} = T'_{Li} \subset \tilde{X}_L$ the parallel copy of T at distance $\frac{1}{2}$ from T_{Li} and distance $L - \frac{1}{2}$ from $T_{(i+1)L}$. We will denote by $d'_{k,Li}$ and $d'_{kj,Li}$, respectively, the copies of $d_{k,Li}$ and $d_{kj,Li}$ in T'_{Li} . We refer the reader to [Figure 4](#) for an illustration of what follows.

- (1) Consider the subspace C_{Li} of T'_{Li} obtained by removing each open primary bust and each open secondary bust in T'_{Li} . A component N_{Li} of C_{Li} is a *nucleus*. Note that each nucleus is either a subinterval of some edge or contained in the star of a vertex.
- (2) The endpoints of the nuclei are endpoints of either primary busts or of secondary busts. To each endpoint of each secondary bust $d'_{kj,L(i-1)}$ in a nucleus, glue a forward flow segment of length $L - \frac{1}{2}$. Note that this flow segment ends at the primary bust $d_{k,Li}$. The union of all flow segments which end at a common endpoint (an endpoint of some $d_{k,Li}$) is a *level L* . Such an endpoint is a *forward endpoint*.

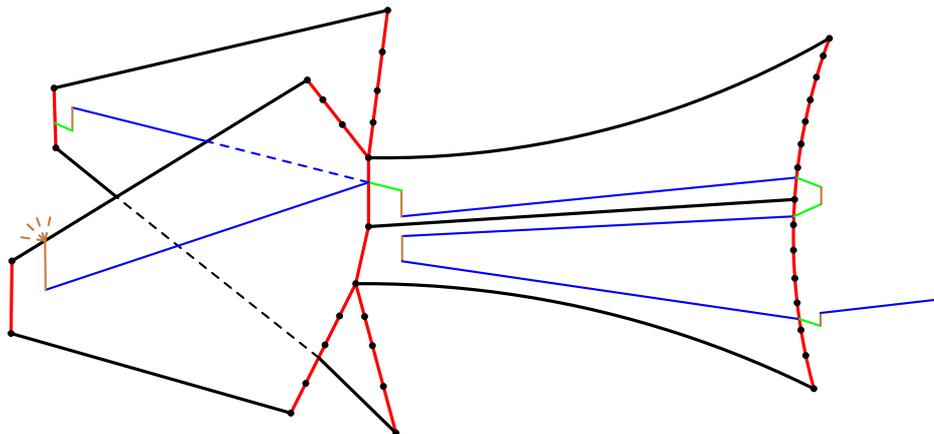


Figure 4: A part of a wall in \tilde{X}_L . Brown paths are nuclei, blue paths are levels and green paths are slopes.

(3) Join the endpoints $d_{k,Li}^\pm$ of each primary bust $d_{k,Li}$ to the endpoints $d_{k,Li}'^\mp$ (note the change in order) of its parallel copy $d_{k,Li}'$ by segments called *slopes*. A slope is denoted by S .

The graph \bar{W}_L is constructed as a quotient of the disjoint union of the nuclei, levels and slopes, with the gluing performed in a natural way: an endpoint of a primary bust of a nucleus is glued to an endpoint of the corresponding slope, an endpoint of a secondary bust in a nucleus is glued to the initial point of the relevant flow segment of a level, while a forward endpoint is glued to the relevant endpoint of a slope. There exists a noncombinatorial immersion of the graph \bar{W}_L in \tilde{X}_L , with loss of injectivity at midpoints of slopes. Let \bar{W} be the graph obtained from \bar{W}_L by “folding” its levels according to q_L so that the following diagram commutes:

$$\begin{array}{ccc} \bar{W}_L & \longrightarrow & \tilde{X}_L \\ \downarrow & & \downarrow q_L \\ \bar{W} & \longrightarrow & \tilde{X} \end{array}$$

A component W_u of \bar{W} is an *immersed wall*.³

Remark 4.2 The immersion $W_u \rightarrow \tilde{X}$ extends to a local homeomorphism $W_u \times [-1, 1] \rightarrow \tilde{X}$, with W_u identified with $W_u \times \{0\}$. Indeed, this is because the embedding of each nucleus, level and slope of W_u in \tilde{X} extends to a local homeomorphism of the above type.

We will denote by $N(W_u)^1 \subset \tilde{X}^1$ the 1-skeleton of the smallest subcomplex of \tilde{X} that contains W_u . When the tunnel length L is sufficiently large, it turns out that $N(W_u)^1 \hookrightarrow \tilde{X}^1$ is a quasiconvex embedding (Proposition 4.6) and that W_u separates \tilde{X} into exactly two components (Proposition 4.8). Thus the local homeomorphism of Remark 4.2 is in fact a homeomorphism in that case.

³The subscript u in W_u denotes “unsaturated”, as we will soon saturate W_u in order to obtain a full codimension-1 subgroup.

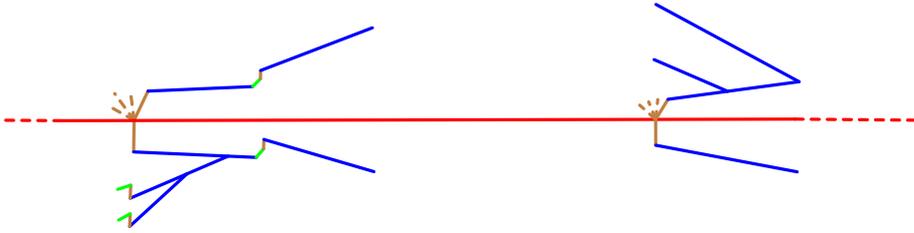


Figure 5: A part of a wall saturation. The red line is a principal flow line, with two unsaturated components (on the left and right) attached.

4.2 Wall saturations

Let W_u be an immersed wall as above, which is a component of the graph \overline{W} . Fix a *saturation length* $M \in L\mathbb{N}$. We will define below the M -saturation W of W_u as a union of components of \overline{W} and principal flow lines. We start with a definition:

Let $v \in T_{Li}$ be a singular vertex. The M -saturation of v is the set $\{f^{Mk}(v) \in T_{Mk+Li}\}_{k \in \mathbb{Z}}$ of singular vertices. (Recall that for each $k \in \mathbb{Z}$, there is a unique singular vertex $f^{Mk}(v)$, by Lemma 2.3.) By abuse of notation, for a vertex v in a “fractional” copy $T_{1/2+Li}$, we will denote by the M -saturation of v the intersection points of the principal flow line of v in the trees $T_{1/2+Mk+Li}$.

Let $v \in T_{1/2+Li}$ be a singular vertex in a nucleus of \overline{W} . Denote by $W(v) \subset \overline{W}$ the smallest union of components of \overline{W} such that each vertex in the M -saturation of v is contained in the nucleus of some component of $W(v)$.

Let W_u be an immersed wall. Let $W'_u \subset \overline{W}$ be the smallest subgraph of \overline{W} satisfying

- $W_u \subset W'_u$, and
- for each singular vertex $v \in W'_u$, either the principal flow line of v intersects a component infinitely many times, or $W(v) \subset W'_u$ (but not both).

The M -saturation W of W_u is defined as the union of W'_u with the principal flow lines of all singular vertices of W'_u . We refer the reader to Figure 5 for an illustration.

Lemma 4.3 *Let W_u be an immersed wall. For any $M \in L\mathbb{N}$, the M -saturation W of W_u is a connected graph.*

Proof We will prove the lemma by showing that a certain connected subgraph containing W_u in W is in fact equal to W . Consider the subgraph W' of W built inductively as the ascending union of subgraphs W^n as follows: $W^1 = W_u$. W^2 is the union of each principal flow line intersecting W_u with the union of all $W(v)$, where v is a singular vertex in W^1 . W^n is the union of each principal flow line intersecting W^{n-1} along with the union of all $W(v)$, where v is a singular vertex in W^{n-1} with the property that

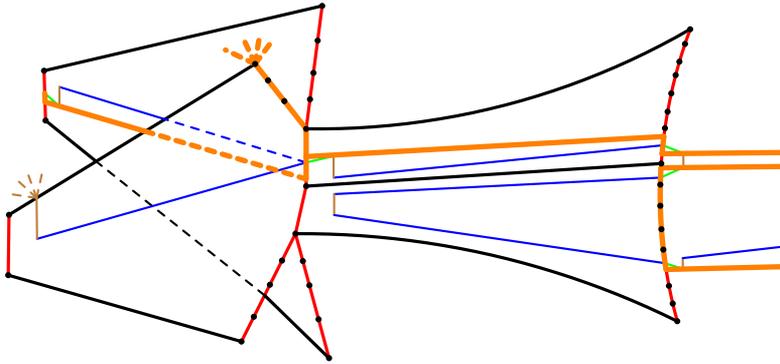


Figure 6: The approximation of the part of the wall from Figure 4 is indicated in thick orange.

$W(v) \subset W$. Note that W' is a connected subgraph of W , by construction. Further, the subgraph W'_u (introduced in the definition of the M -saturation W of W_u) is contained in W' . This gives the reverse containment $W' \supset W$. □

4.3 Approximations of walls

In order to show that families of immersed walls are uniformly quasiconvex, Hagen and Wise use a technical construction called approximations. Approximations are not walls in \tilde{X} (they are walls in \tilde{X}_L , as it turns out), but have the advantage that, unlike immersed walls, they do not have backward flows that can have long fellow-travelling subpaths.

We will first state the definition of approximations given in [HW16] and then extend their definition to saturations. Let $W_u \rightarrow \tilde{X}$ be an immersed wall of tunnel length L . The *approximation* $A: W_u \rightarrow \tilde{X}$ is defined as below. We refer the reader to Figure 6 for an illustration.

- (1) For each $x' \in W_u$ such that x' lies in a nucleus $N_{Li} \subset T_{Li+1/2}$, let $x \in T_{Li}$ denote the parallel copy of x' behind it at distance $\frac{1}{2}$. Then $A(x') := f^L(x) \in T_{Li+L} \subset \tilde{X}$.
- (2) For each x in a level of W_u , $A(x)$ is the unique forward endpoint of the level.
- (3) Let S be a slope in W_u associated to the primary bust d . Recall that S is a segment from an endpoint d^+ (say) of d to the other endpoint d'^- of d' , where d' is the parallel copy of d at distance $\frac{1}{2}$. A maps S homeomorphically to the concatenation of d and the forward flow segment of length L from d^- .

Let $M \in L\mathbb{N}$ and let $W \rightarrow \tilde{X}$ be an M -saturation of W_u . The *approximation of the M -saturation* $A: W \rightarrow \tilde{X}$ is defined the same way on nuclei, slopes and levels. On each point x of a principal flow line, by abuse of notation, $A(x) := f^{L-1/2}(x)$.

We will denote by $N(A(W_u))^1$ and $N(A(W))^1$ the 1-skeletons of the smallest subcomplexes of \tilde{X} containing $A(W_u)$ and $A(W)$, respectively.

4.4 Large tunnel length and quasiconvexity

In this subsection, we recall results from [HW16, Section 4] towards quasiconvexity of immersed walls. We note that their methods do not require local finiteness of \tilde{X}^1 and thus work in our case as well, both for immersed walls and their saturations.

The main lemma that we will repeatedly use is the following. Recall that $\mathcal{N}_r(Y)$ denotes the r -neighbourhood of Y .

Lemma 4.4 [HW16, Lemma 4.3] *Let Z be a δ -hyperbolic space and let $P = \alpha_0\beta_1\alpha_1 \dots \beta_k\alpha_k$ be a path such that each α_i is a (λ_1, λ_2) -quasigeodesic and each β_i is a (μ_1, μ_2) -quasigeodesic. Suppose that for each $R \geq 0$ there exists a $B_R \geq 0$ such that for all i each intersection below has diameter $\leq B_R$:*

$$\mathcal{N}_{3\delta+R}(\beta_i) \cap \beta_{i+1}, \quad \mathcal{N}_{3\delta+R}(\beta_i) \cap \alpha_i, \quad \mathcal{N}_{3\delta+R}(\beta_i) \cap \alpha_{i-1}.$$

Then there exists L_0 such that if each β_i is of length at least L_0 , then the path P is a $(4\lambda_1\mu_1, \frac{1}{2}\mu_2)$ -quasigeodesic.

Let $\mathcal{W} := \{W_u \rightarrow \tilde{X}\}$ be a family of immersed walls in \tilde{X} . In order to use Lemma 4.4, we need the following property to be satisfied:

Definition 4.5 [HW16, Definition 3.15] *The family of immersed walls \mathcal{W} has the ladder overlap property if there exists $B \geq 0$ such that for all $W_u \in \mathcal{W}$ and all distinct slopes $\mathcal{S}_1, \mathcal{S}_2 \subset W_u$,*

$$\text{diam}(\mathcal{N}_{3\delta+2\lambda}(N(A(\mathcal{S}_1)))^1 \cap \mathcal{N}_{3\delta+2\lambda}(N(A(\mathcal{S}_2))))^1 \leq B,$$

where λ is the quasiconvexity constant for forward ladders (Proposition 2.7).

In practice, any family of immersed walls is constructed in the following way: We will first choose a finite set of periodic points in the fundamental domain D of T such that the flow rays from these points diverge and do not meet any vertex. The various immersed walls of the family will then be constructed by choosing tunnel lengths and by choosing primary busts in small neighbourhoods of these periodic points, with the size of the neighbourhoods depending on the chosen tunnel lengths. The ladder overlap property will then hold because the finitely many flow rays from the periodic points have bounded overlaps.

Proposition 4.6 *Let \mathcal{W} be a family of immersed walls satisfying the ladder overlap property. Then there exist L_0, κ_1 and κ_2 such that for all $W_u \in \mathcal{W}$ with tunnel length at least L_0 , the inclusion $N(A(W_u))^1 \hookrightarrow \tilde{X}^1$ is a (κ_1, κ_2) -quasiisometric embedding.*

We refer to [HW16, Proposition 4.1] for a proof. They show that when tunnels are long enough, the hypotheses of Lemma 4.4 are satisfied. The same proof works in our case, as Lemma 4.4 works for any hyperbolic space, not necessarily proper. The main point is that geodesic paths in $N(A(W_u))^1$ can be written as alternating paths $P = \alpha_0\beta_1\alpha_1 \dots \beta_k\alpha_k$ satisfying the conditions of Lemma 4.4, where each β_i is a geodesic fellow-travelling with a forward flow segment of length $L \geq L_0$.

We note that the constants depend only on the following data: L_0 depends on the quasiconvexity constant μ of nuclei of walls (there is a uniform quasiconvexity constant μ as nuclei are either subintervals of edges or stars of vertices) and the constant B of the ladder overlap property. The constants κ_1 and κ_2 depend on L_0 and μ , but not on the tunnel length $L \geq L_0$ of individual immersed walls in \mathcal{W} .

Proposition 4.7 *Let \mathcal{W} be a family of immersed walls with the ladder overlap property. Then there exists $L_1 \geq L_0$ such that for each $W_u \in \mathcal{W}$ with tunnel length at least L_1 , $A(W_u)$ is a tree.*

This is a consequence of Proposition 4.6: Indeed, any closed path in $A(W_u)$ has length at most $\kappa_1\kappa_2$. Thus if $L_1 \geq L_0, \kappa_1\kappa_2 + 1$, then no closed path can contain a forward flow segment from a slope approximation. This implies that the closed path has to be contained in the intersection of a tree T_i in the flow space with the approximation of W_u , which is impossible. See [HW16, Proposition 4.4] for a precise proof.

Recall that a wall is a subspace $Y \subset \tilde{X}$ such that $\tilde{X} \setminus Y$ has exactly two components, neither of which is contained in any finite neighbourhood of Y (ie components are deep). We now state the main result of this subsection:

Proposition 4.8 *Let \mathcal{W} be a family of immersed walls with the ladder overlap property and let L_1 be the constant from Proposition 4.7. Then for each $W_u \in \mathcal{W}$ with tunnel length at least L_1 , W_u is a wall.*

Again, a detailed proof is given in [HW16, Proposition 4.6]. Since the flow space \tilde{X} is simply connected, it is enough to show that W_u locally separates (a small neighbourhood) into two components. Using Remark 4.2, the only place where this can fail is when two distinct slopes intersect in their interior. But when the tunnel length L is at least L_1 , such a scenario would contradict the fact that $A(W_u)$ is a tree.

Proposition 4.9 (approximations of saturations are quasiconvex) *Let \mathcal{W} be a family of immersed walls satisfying the ladder overlap property. Then there exists $L_2 \geq L_1$ such that for all $W_u \in \mathcal{W}$ with tunnel length $L \geq L_2$, there exist M_0, Θ_1 and Θ_2 such that for any M -saturation W of W_u with $M \geq M_0L$, the inclusion $N(A(W)) \hookrightarrow \tilde{X}^1$ is a (Θ_1, Θ_2) -quasiisometric embedding.*

Proof We will prove the statement by again using Lemma 4.4. Let P be a path in $N(A(W)) \hookrightarrow \tilde{X}^1$ such that P is a concatenation $\alpha_0\beta_1\alpha_1 \dots \beta_k\alpha_k$, with each α_i a geodesic in the approximation of a component subgraph of \bar{W} and each β_i a geodesic segment in a principal flow line. By Proposition 4.6, each α_i is a (κ_1, κ_2) -quasigeodesic, while Proposition 2.7 ensures that each β_i is a (λ_1, λ_2) -quasigeodesic. We choose tunnel length L large enough so that, in length smaller than L , α_i and β_j diverge from each other: recall that primary bust endpoints are chosen so that their flow rays are not asymptotic to any principal flow line (Lemma 2.18). By choosing L larger than the minimal divergence distance between the flow ray of any primary bust endpoint and a principal flow line, we obtain the required bounds. \square

We also observe that since any immersed wall W_u (respectively its M -saturation W) of tunnel length L is at Hausdorff distance L from its approximation $A(W_u)$ (respectively $A(W)$), we have:

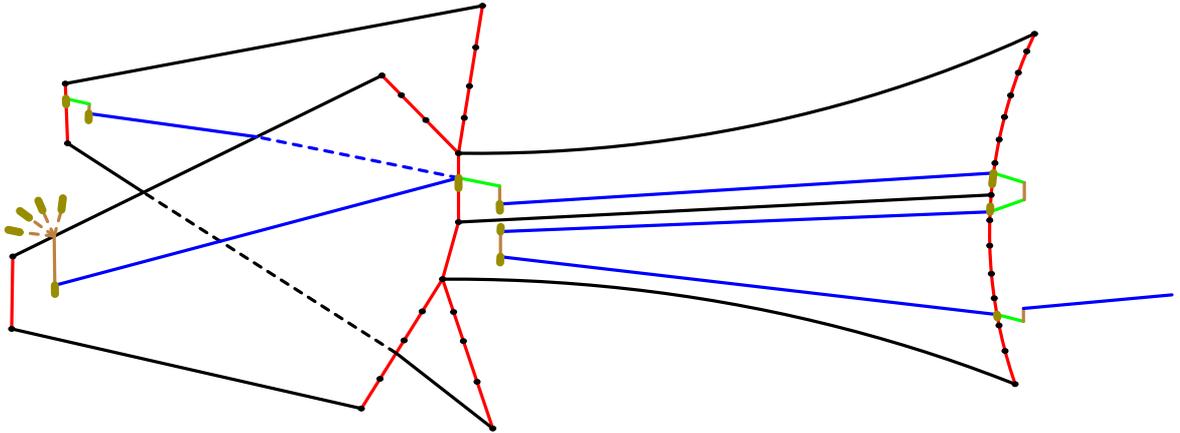


Figure 7: The busts in olive are all in a single complementary component.

Lemma 4.10 For any wall W_u (with M -saturation W), the limit sets in $\partial \tilde{X}^1$ of $N(W_u)^1$ and $N(A(W_u))^1$ coincide, as do those of $N(W)^1$ and $N(A(W))^1$. □

Lemma 4.11 Let d and d' be primary busts such that for some $m, m' \in \mathbb{Z}$, we have $d \subset T_{mL}$ and $d' \subset T_{m'L}$, and one of the endpoints of each of d and d' is contained in W_u . Then the interiors of d and d' lie in the same complementary component of W_u .

Proof For the purposes of this argument, we endow W_u with a graph structure where the vertices are vertices of nuclei and the endpoints of various primary and secondary busts. Edges consist of forward flow segments in levels, slopes and segments in nuclei between vertices. Let us assume that $d \subset T_0$ and let P be a (geodesic) path in W_u between the endpoints of d and d' . The edge of P with endpoint d^+ (say), which is either a slope starting at d^+ or a forward flow segment ending at d^+ , leads to a nucleus N in a “fractional” tree $T_{1/2}$ or $T_{-L+1/2}$. Except for the central vertex, every vertex of this nucleus is the endpoint of either a primary bust or of a secondary bust. In the latter case, note that each such secondary bust (in this fractional copy of T) meets the same complementary component C as d (see Figure 7), and thus flowing such a secondary bust to a primary bust in T_L or T_0 (along an edge of W_u without crossing it) will not lead to a change of components. In the former case, the primary bust in the fractional copy, with an endpoint in N , is not in C , but travelling back along the slope will lead to its parallel copy (in either T_0 or T_{-L}), and this primary bust is contained in C . One can now proceed by induction on the distance between d and d' to obtain the desired result. □

Proposition 4.12 Let $W_u \in \mathcal{W}$ be an immersed wall with tunnel length at least L_2 (from Proposition 4.8). Then $\text{Stab}(W_u) < \Gamma$ is a codimension-1 subgroup that stabilises each complementary component of W_u and acts cocompactly on W_u .

Proof Complementary components are not flipped: as a consequence of Lemma 4.11, if $\gamma \in \text{Stab}(W_u)$ and C is a complementary component of W_u , we have that $\gamma C = C$.

Let us show cocompactness. Recall (from Section 4.1) that W_u is a component of the graph \overline{W} . For a point $x \in \tilde{X}$ such that $x \in \overline{W}$ (the image of \overline{W} , to be precise), denote by $W(x)$ the component of \overline{W} that contains x . Observe that for two points $x, y \in \overline{W}$, $W(x) = W(y)$ if and only if $y \in W(x)$ (and $x \in W(y)$).

Let L be the tunnel length of W_u . Denote by Γ_L the subgroup of Γ generated by G and t^L . We first note that, by the equivariant choice of primary busts, and the fact that no primary bust intersects any secondary bust (Lemma 4.1(1)), the action of Γ_L on the flow space restricts to an action on \overline{W} .

Let $x \in W_u$. Let us denote by $H < \Gamma$ the stabiliser of W_u . For an element $\gamma \in \Gamma_L$, we observe that $\gamma \in H$ if and only if $\gamma x \in W_u$. Indeed, $\gamma x \in W_u$ if and only if $W_u = W(x) = W(\gamma x)$, with the latter being equal to $\gamma W(x)$. Let us denote by X the compact quotient of \tilde{X} by the action of Γ . Observe that X is the mapping torus of the compact graph $G \setminus T$ by the map induced by f . There are finitely many images in X of busts and nuclei of W_u , and hence the image of W_u in X is compact. This implies that H acts cocompactly on W_u .

Finally, we show that $\text{Stab}(W_u)$ has codimension 1. There is a continuous equivariant collapse map from a Cayley graph of Γ to \tilde{X}^1 which crushes free factors of G to points. Since W_u separates \tilde{X} into deep components, so does its preimage in Γ , giving us the desired result. □

Note that if W_u is a wall, then any M -saturation of W_u separates \tilde{X}^1 into several deep components and therefore separates the boundary $\partial\tilde{X}^1$. In fact:

Proposition 4.13 *Let $W_u \in \mathcal{W}$ be a wall with tunnel length at least L_2 (from Proposition 4.9). Then there exists $L_3 \geq L_2$ such that for all $L > L_3$ and for all M large enough, the stabiliser of the M -saturation W of W_u is a relatively quasiconvex full subgroup. Further, W separates the flow space into at least two deep components.*

Proof To prove separation with deep components, since W_u itself is of codimension 1 by Proposition 4.12, it suffices to find, for each deep component of $\tilde{X} \setminus W_u$, a ray starting at a point of W_u , entering the given component, and never encountering W again. To do that, it suffices to find a ray avoiding all the principal flow lines issued from W_u , and diverging from each of them within $\frac{1}{10}M$ of distance from W_u . Since \tilde{X} is nonelementary hyperbolic and principal flow lines are disjoint and diverge from each other in bounded time, it is possible to find such a ray, provided M is sufficiently large compared to δ .

Since W is quasiconvex, an application of [EGN21, Theorem 1.2] implies that the stabiliser of W is relatively quasiconvex in Γ . Let us show that $\text{Stab}(W)$ is also full. To this end, consider a principal flow line Λ whose associated parabolic group (a suspension of a vertex stabiliser in T_0 by an element translating along Λ) intersects $\text{Stab}(W)$. We want to show that $\text{Stab}(W)$ contains a finite-index subgroup of the parabolic group. For this it is sufficient to show that Λ intersects W , since in that case, the saturation property of W ensures that W is stabilised by such a subgroup.

Thus, let $\text{Stab}(W)$ intersect $\text{Stab}(\Lambda)$. We distinguish whether the intersection contains an element translating along Λ or an element fixing a vertex in Λ .

First assume that $\text{Stab}(W)$ contains an element of Γ which acts loxodromically on Λ . This implies that a finite neighbourhood of W contains Λ . Since every pair of principal flow lines diverge, either a vertex of Λ is contained in W (and we are done, by the property of saturations), or there exists a periodic union of tunnels, slopes and nuclei of W that fellow-travels with Λ . Assume the latter. If there exists a singular vertex (of some nucleus) in this periodic union, then the no twinning property ensures a contradiction. Let γ be translating along Λ , and such that γ^k stabilises W (for some large k). Let $x \in W$ be near Λ , and consider a shortest path p in W from x to $\gamma^k x \in W$. Since W is assumed not to cross Λ , the path p does not contain an arc of Λ . Thus p consists of nuclei, slopes and levels, and is a quasigeodesic that must remain close to Λ . Note that p does not contain any piece of a principal flow line by the no twinning assumption. Since p moves from a tree T_r to a tree T_{r+k} for some large k , there must be a level in p . As a subpath of p , it travels close to Λ . But a level is a piece of a periodic flow line, and by our choice of primary bust endpoints (Lemma 2.18) and large tunnel length, levels diverge from any principal flow line Λ . Hence p does not stay close to Λ , as promised.

Assume now that the stabiliser of W contains an elliptic element g , fixing some vertex v of T_0 . Let $x \in T_0$ be in W so that for each $n \in \mathbb{Z}$, $g^n x \in W$. Since the T_0 -distance between $g^m x$ and $g^n x$ is bounded for all m and n (it equals twice the T_0 -distance between x and v), there exists $K \geq 0$ such that any geodesic path $\alpha_{m,n}$ in the quasiconvex space W between $g^m x$ and $g^n x$ does not meet T_i for $i > K$. We flow each such geodesic to T_K to obtain a path in T_K between $f^K(g^m x)$ and $f^K(g^n x)$. Observe that there is a uniform bound on the length of each such path. Thus, up to taking a subsequence, there is an infinite valence vertex of T_K contained in each of the flowed paths between $f^K(g^m x)$ and $f^K(g^n x)$. By equivariance of f , this vertex w is fixed by $\phi^K(g^n)$. By uniqueness of vertex stabilisers under the G -action on T_K , we have $w = f^K(v)$. We claim that this implies that infinitely many of the paths $\alpha_{m,n}$ intersect the principal flow line Λ_v of v . Indeed, by Lemma 2.3 the intersection of the backward flow of w with the trees T_i intersects the infinite collection $\alpha_{m,n}$ in a finite set. Let $w' \in T_i$ be one such point that lies in infinitely many $\alpha_{m,n}$. If w' is in Λ_v , we are done. If not, then by equivariance, we again have that w' is stabilised by $\phi^i(g^k)$ for infinitely many k , a contradiction. \square

4.5 Many effective walls

Recall that $D \subset T$ denotes the fixed fundamental domain.

Definition 4.14 Using the terminology of [HW16], we say \tilde{X} has *many effective walls* if the following two conditions are satisfied:

- (1) For each $y \in D \subset T_0$ such that the flow ray from y does not meet a vertex, there exists a set \mathcal{W} of immersed walls with the ladder overlap property such that for every $\epsilon > 0$, there exists $W_u \in \mathcal{W}$ of arbitrarily large tunnel length L and a primary bust in the ϵ -neighbourhood of y . We will call such a set of walls a *regular effective set of walls*.

(2) For each periodic point $a \in D$, there exists $k = k(a)$ and a set of immersed walls \mathcal{W}_a with the ladder overlap property such that for each primary bust $d' \subset T' = T_{1/2}$ that is joined to $a' = a_{1/2} \in T'$ by a path in W_u disjoint from the slopes of W_u , we have that $d_{\tilde{X}^1}(f^n(a), f^n(d)) \geq 3\delta + 2\lambda$, for all $n \geq k$. We will call such a set of walls a *periodic effective set of walls*.

Let us comment on why one needs the property of many effective walls. The full details are available in [HW16, Section 5]. For a hyperbolic group to admit a proper cocompact cubulation, the boundary criterion of [BW12] stipulates that every pair of distinct points in the boundary of the group (equivalently, the limit set of every biinfinite geodesic of \tilde{X}^1 in the setup of [HW16]) should be separated by the limit set of a quasiconvex codimension-1 subgroup (equivalently, a quasiconvex wall in \tilde{X}^1). Having many effective walls assures that this can be done in \tilde{X}^1 . Indeed, there are two types of biinfinite geodesics in \tilde{X}^1 , “horizontal” geodesics and “nonhorizontal” geodesics, which we define below. If \tilde{X}^1 has many effective walls, then each horizontal geodesic is cut by a wall from a periodic effective set, while each nonhorizontal geodesic is cut by a wall from a regular effective set.

Definition 4.15 (geodesic classification) Let $N > 0$. A biinfinite geodesic γ in \tilde{X}^1 is *N -horizontal* (N -ladderlike in [HW16]) if there exists a forward flow segment σ of length N such that a geodesic of the forward ladder $N(\sigma)$ joining the endpoints of σ fellow-travels with a subpath of γ at distance at most $2\delta + \lambda$. Here δ is the hyperbolicity constant of \tilde{X}^1 and λ is the quasiconvexity constant of forward ladders (Proposition 2.7). Otherwise, γ is *N -nonhorizontal* (N -deviating in [HW16]).

Theorem 4.16 *The flow space \tilde{X} has many effective walls.*

This is proved in [HW16, Theorem 6.16], to which we refer the reader for a detailed proof. The main ingredients of their proof, which are available in our case as well, are the facts that periodic points are dense in T (Lemma 2.14) and that flow rays starting from translates of a periodic point diverge (Proposition 2.12). With these ingredients at hand, here is a brief sketch of the proof. The first condition of Definition 4.14 can be verified by first choosing a periodic point in an ϵ -neighbourhood of the given point y and then choosing one periodic point in each edge of D so that pairwise, their flow lines diverge. This will ensure that the ladder overlap property holds when primary busts are chosen in small enough neighbourhoods of these points, depending on the tunnel length. The family of walls is now constructed by choosing tunnel lengths to be large common multiples of the periods of the chosen periodic points, and choosing primary busts as above. To verify the second condition, we choose one periodic point per edge of D as above, including one for the edge containing a , while ensuring that pairwise, the flow lines of the periodic points and of a diverge. The family of walls is then constructed as done for the first condition.

Let us conclude this section with another property that will be useful in the next section. Following [HW16], we say \tilde{X} is *level-separated* if for each $N > 0$ and $K \geq 0$, and each N -nonhorizontal geodesic γ , there exists a point $y \in \tilde{X}$ such that the backward flow of length n meets γ in an odd-cardinality set for all large n , and the intersection is at distance at least $N + K$ from both y and the leaves of the backward flow.

Lemma 4.17 *The flow space \tilde{X} is level-separated.*

In order to prove this statement, we will make use of an \mathbb{R} -tree T_∞ obtained as a limit of the trees T_m in the flow space with metrics $d_m = d|_{T_m}/\lambda^m$, where λ is the (maximal) stretching factor of the train track map f . We will not go into details, but refer the reader to [HW16, Section 6; Hor17]. Formally, T_∞ is an ultralimit; its points are equivalence classes of sequences of points in the T_m . In particular, to each point in T_0 (and in \tilde{X}), its flow ray defines a point in T_∞ . This defines a tautological map $\rho: \tilde{X} \rightarrow T_\infty$.

The main points we will use are easy facts: that the map $\rho: \tilde{X} \rightarrow T_\infty$ is continuous, restricts to an isometric embedding on any vertical edge (as f is a train track map) and that the preimage of any point in T_∞ consists of points whose forward flow rays either intersect or remain at bounded Hausdorff distances from each other.

Proof Let $\gamma: \mathbb{R} \rightarrow \tilde{X}^1$ be an N -nonhorizontal geodesic. Observe that the flow-like parts of γ are horizontal segments starting and ending at vertices and that the vertical parts are paths in trees T_k . Since γ is N -nonhorizontal, it maps infinitely many unit subintervals of \mathbb{R} to vertical edges, which we call “vertical” subintervals.

First, we make a slight change. Let γ_{vert} be the union of paths defined on the vertical subintervals of γ , but now reparametrised on \mathbb{R} .

Let $\gamma^\infty = \rho \circ \gamma_{\text{vert}}: \mathbb{R} \rightarrow T_\infty$ be the image of γ_{vert} (and also of γ , in fact) in T_∞ . Call $\gamma_{\text{vert}}(s)$ a “flat” point if s is a point of discontinuity for γ_{vert} . Observe that this happens only when $\gamma(s)$ is the starting point in γ of a horizontal segment, and that there are only countably many flat points. Since all points in any horizontal segment have the same flow, all such points have the same image in T_∞ . Therefore the map $\gamma^\infty: \mathbb{R} \rightarrow T_\infty$ is continuous, and in fact, 1-Lipschitz.

Further, there exists D such that the preimage of a point in γ^∞ is of diameter $\leq D$. Indeed, if s_1 and s_2 have the same γ^∞ -image, observe that $\gamma_{\text{vert}}(s_1)$ and $\gamma_{\text{vert}}(s_2)$ correspond to points in the image of γ which fellow-travel a forward flow segment. Since γ is N -nonhorizontal, there exists a D as required.

In fact, the same argument shows that $\gamma^\infty(\mathbb{R}_+)$ and $\gamma^\infty(\mathbb{R}_-)$ have infinite diameter and go to ends γ_+^∞ and γ_-^∞ of T_∞ , respectively. Note that, by the bound on the diameter of preimages, $\gamma_+^\infty \neq \gamma_-^\infty$. In particular, there are uncountably many points that cut $\gamma^\infty(\mathbb{R})$ into two unbounded components, with each such point contained in $\gamma^\infty(\mathbb{R})$ (as T_∞ is an \mathbb{R} -tree).

Let us now show that there exists D' such that the preimage of a point under γ^∞ has cardinality at most D' . Indeed, if there are many points in the preimage, there must be many in the same tree T_k , and their diameter being bounded, either they accumulate or have to be placed so that there is a vertex and a huge angle to reach them. Accumulation is not permitted since all edges are stretched by the scaling factor (more precisely, we are guaranteed that two points in the same edge are flown to different points of T_∞). Also for two points separated by a huge angle, we already noticed that their flow lines cannot

fellow-travel (see Proposition 2.9), and hence they must diverge with speed defined by the stretch factor λ , thus defining different points in the limit tree T_∞ .

Call $\gamma^\infty(s)$ a “backtrack” point if there are two intervals $[s_-, s)$ and $(s, s_+]$ that are sent by γ^∞ to the same component of $T_\infty \setminus \{\gamma^\infty(s)\}$.

Observe that since the map ρ is injective on edges, if $\gamma^\infty(s)$ is a backtrack point, then $\gamma_{\text{vert}}(s)$ is a vertex. Therefore there are countably many backtrack points in γ^∞ .

There are two kinds of backtrack points: those such that the component of the intervals $[s_-, s)$ and $(s, s_+]$ contains both the ends γ_+^∞ and γ_-^∞ , and those that don't. For the former, take the centre of the tripod $(\gamma_+^\infty, \gamma_-^\infty, \gamma^\infty(s))$, and call it a “crossroad” point.

There are countably many backtrack points, so the union of flat points, backtrack points and crossroad points is a countable set.

Pick a point ξ in $\gamma^\infty(\mathbb{R})$ that is not a backtrack point nor a crossroad point nor a flat point, and such that ξ separates $\gamma^\infty(\mathbb{R})$ in two infinite components. Since it is not a backtrack point, there is a well-defined direction of crossing ξ at each preimage s (a small interval around s such that (s_-, s) and (s, s_+) are sent on different sides of ξ , each containing only one of the ends γ_+^∞ and γ_-^∞).

Since preimages in γ_{vert} are finite, for ξ there can only be an odd number of preimages in γ_{vert} . Since ξ is not a flat point, it has the same (odd) number of preimages in both γ and γ_{vert} . Then at least one of the points x in the preimage is such that for a large n , if $y = f^n(x)$, then the backward flow of large enough length from x meets γ in an odd-cardinality set, with the intersection at a large distance from y and the leaves of the backward flow, as required. □

5 Cutting principal flow lines and geodesics by wall saturations

In this section, we will show that three types of subsets of \tilde{X}^1 are separated by the saturations of walls. In [HW16], they show that all biinfinite geodesics are separated by (unsaturated) walls. For relative cubulation (see Section 3), we will need more: that saturations of walls also separate principal flow lines from geodesic rays and pairs of principal flow lines.

Let us begin with a tautological observation in preparation for the future coning-off of principal flow lines:

Lemma 5.1 *Let W_u be a quasiconvex wall in \tilde{X} and W an M -saturation of W_u . Then for each principal flow line Λ of a singular vertex, either $\Lambda \subset W$ or the two points of $\partial\Lambda$ lie in a single component of $\partial\tilde{X}^1 \setminus \partial W$.* □

We will also need the notion of “cut” in order to disconnect points in the Bowditch boundary:

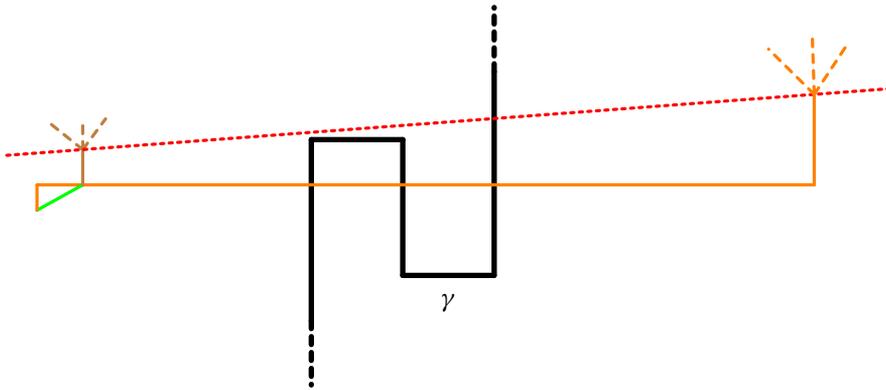


Figure 8: Proof of Proposition 5.3. The geodesic γ hits the approximation of a slope of W_u . The red dotted line indicates a flow line in the saturation of W_u .

Definition 5.2 Let W_u be a quasiconvex wall in \tilde{X} and let W be an M -saturation of W_u . Let A be a quasiconvex subspace of \tilde{X}^1 such that ∂A has at least two points. We say that W cuts A if

- (1) $\partial A \cap \partial W = \emptyset$ (as subsets of $\partial \tilde{X}^1$), and
- (2) ∂W separates ∂A , ie ∂A nontrivially intersects at least two components of $\partial \tilde{X}^1 \setminus \partial W$.

5.1 Cutting geodesics

Recall that there are two types of geodesics (see Definition 4.15) in \tilde{X}^1 . In this subsection, we will show that for each geodesic line that is not a principal flow line, there exists a wall whose saturation cuts the geodesic.

Proposition 5.3 Let $\gamma: \mathbb{R} \rightarrow \tilde{X}^1$ be an N -nonhorizontal biinfinite geodesic in the flow space. Then there exists a quasiconvex wall $W_u \rightarrow \tilde{X}$ and an M -saturation W of W_u such that W cuts γ .

Proof Proposition 5.18 of [HW16] shows that in this case, there exists a quasiconvex wall W_u which separates $\partial \gamma$. Let us quickly explain the idea behind their proof. We refer the reader to Figure 8 for an illustration of what follows. We choose a wall W_u from a regular effective set of walls (Definition 4.14(1)) with tunnel length L sufficiently larger than N , while ensuring that the approximation $A(S)$ of some slope S of W_u intersects a vertical segment of γ in an odd number of points (in the interior of some vertical edges), and far away from the endpoints of $A(S)$ (arguing exactly as in [HW16] while using Lemma 4.17 for level separatedness). Such a W_u exists because of Theorem 4.16. The fact that γ is N -nonhorizontal will ensure that γ does not have long subpaths which $(3\delta + 2\lambda)$ -fellow-travel with W_u (as the tunnel length L of W_u is much larger than N). This implies, by Lemma 4.4, that the union of γ and $N(A(W_u))^1$ quasiisometrically embeds in \tilde{X}^1 and thus ∂W_u is disjoint from, and separates, $\partial \gamma$ in $\partial \tilde{X}^1$.

We now choose $M \in L\mathbb{N}$ sufficiently large, and let W be the M -saturation of W_u . Since ∂W_u separates $\partial \gamma$, and $W_u \subset W$, in order to show that ∂W also separates $\partial \gamma$, it suffices to show that $\partial \gamma$ is disjoint from ∂W ,

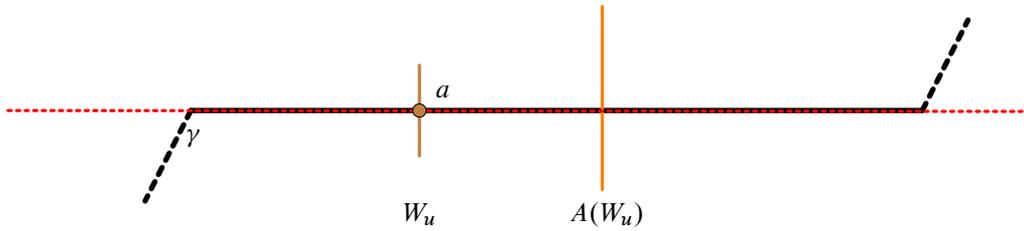


Figure 9: Proof of Proposition 5.4. The geodesic γ is cut along a nucleus of W_u at a . The red dotted line indicates the principal flow line Λ_a .

and for that it suffices to show that the union of γ and $N(A(W))$ ¹ quasiisometrically embeds in \tilde{X} ¹. Since we already know that the union of γ and $N(A(W_u))$ ¹ quasiisometrically embeds in \tilde{X} ¹ and that the saturation consists of translates of W_u that are at distance larger than M from each other, and of principal flow lines between them, a failure of quasiconvexity for the union of γ and $N(A(W))$ ¹ would mean that γ fellow-travels with a principal flow line for a long distance (at least M by definition of saturation). This contradicts the N -nonhorizontality of γ . □

Proposition 5.4 *Let $\gamma: \mathbb{R} \rightarrow \tilde{X}$ ¹ be a biinfinite geodesic in the flow space such that*

- (1) γ is N -horizontal for all large N , and
- (2) no ray of γ is asymptotic to any ray of a principal flow line.

Then there exists a quasiconvex wall $W_u \rightarrow \tilde{X}$ and an M -saturation W of W_u such that W cuts γ .

Proof Proposition 5.19 of [HW16] ensures that there exists a quasiconvex wall W_u which separates $\partial\gamma$; the wall W_u is chosen from a periodic effective set of walls \mathcal{W}_a (Definition 4.14(2)) so that, up to translation, the point a in the nucleus of W_u intersects γ roughly in the middle of a long horizontal subpath of γ , whose length is much bigger than the tunnel length of W_u (Figure 9). Such a W_u exists, again, by Theorem 4.16. By the way W_u was chosen, slope approximations of W_u do not have long $(3\delta+2\lambda)$ -fellow-travelling subpaths with the flow line of a , and therefore with γ . Lemma 4.4 now ensures that the union of γ and $N(A(W_u))$ ¹ quasiisometrically embeds in \tilde{X} ¹. Therefore ∂W_u is disjoint from, and separates, $\partial\gamma$ in $\partial\tilde{X}$ ¹.

In order to prove the proposition, we will now choose an M -saturation W of W_u with M much larger than the length of the maximal subpath of γ that $(3\delta+2\lambda)$ -fellow-travels with the principal flow line Λ_a through a . Observe that this maximal subpath is bounded as $\partial\gamma$ is disjoint from $\partial\Lambda_a$. An application of Lemma 4.4 now shows that the union of $N(A(W))$ ¹ with γ is a quasiisometric embedding and thus W cuts γ . □

5.2 Separating geodesics from principal flow lines

In this subsection, we will show that the union of a principal flow line and a geodesic ray is cut by a wall. More precisely:

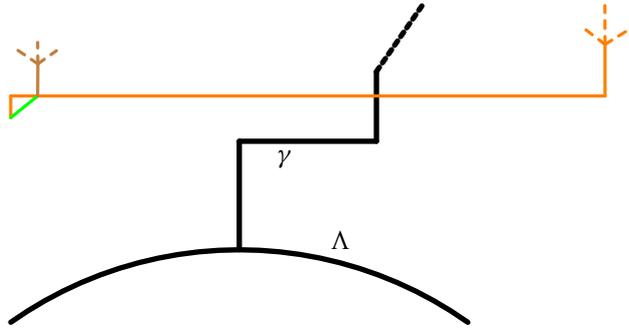


Figure 10: The wall W_u is chosen so as to hit the geodesic ray γ far from the principal flow line Λ .

Proposition 5.5 *Let Λ be a principal flow line and γ be a geodesic ray starting at some point on Λ . Suppose that γ and Λ are not asymptotic. Then there exists a quasiconvex wall $W_u \rightarrow \tilde{X}$ and an M -saturation W of W_u , such that*

- (1) W does not separate $\partial\Lambda$, and
- (2) W cuts $\Lambda \cup \gamma$.

Proof Let $p \in \Lambda$ be the starting point of γ . Denote by Λ_+ and Λ_- the two geodesic subrays of Λ starting at p . Denote the endpoints at infinity of these rays by $\Lambda_+(\infty)$, $\Lambda_-(\infty)$ and $\gamma(\infty)$. Let ϵ be the maximal Gromov product at $p = \gamma(0)$ of $\gamma(\infty)$ and $\Lambda_{\pm}(\infty)$. Then both the concatenations $\gamma \cup \Lambda_{\pm}$ are $(1, 2\epsilon + 10\delta)$ -quasigeodesic lines. We will now choose a quasiconvex wall W_u (see Figure 10) such that

- (1) $W_u \cap (\gamma \cup \Lambda_-) = W_u \cap (\gamma \cup \Lambda_+) \subset \gamma$,
- (2) the union of the quasigeodesic lines $\gamma \cup \Lambda_{\pm}$ with $N(A(W_u))^1$ embeds quasiisometrically, and
- (3) ∂W_u separates $\gamma(\infty)$ from $\Lambda_+(\infty)$ and $\Lambda_-(\infty)$.

The choice of W_u that will work for us is a wall given by either Proposition 5.3 or 5.4 depending on whether or not γ is N -horizontal for all N .

Let λ be the quasiconvexity constant from Proposition 2.7, and ϵ as before. Let κ and L be the quasiconvexity constants for $N(A(W_u))^1$ and the tunnel length of W_u , respectively, for such walls, as guaranteed by Proposition 4.6. We choose W_u so that W_u intersects γ at points at distance much larger than $2\epsilon + 20\delta + 2\kappa + \lambda + L$ from p .

Thanks to the large distance between the point p and $W_u \cap \gamma$, Lemma 4.4 says that the union of $N(A(W_u))^1$ with $\gamma \cup \Lambda$ embeds quasiisometrically and both W_u and $A(W_u)$ are disjoint from Λ . This implies that ∂W_u does not separate $\partial\Lambda$. We now choose an M -saturation W of W_u with M much larger than the maximal $(3\delta + 2\lambda)$ -fellow-travelling length between pairs of principal flow lines. This choice ensures, yet again by Lemma 4.4, that W is disjoint from Λ and therefore cuts $\gamma \cup \Lambda$ but does not separate $\partial\Lambda$. \square

5.3 Cutting pairs of principal flow lines

Proposition 5.6 *Let $\Lambda_1 \neq \Lambda_2$ be two principal flow lines. Then there exists a quasiconvex wall $W_u \rightarrow \tilde{X}$ and an M -saturation W of W_u such that*

- (1) W cuts $\Lambda_1 \cup \Lambda_2$, and
- (2) ∂W does not separate $\partial\Lambda_i$.

Proof Let γ be a geodesic segment between Λ_1 and Λ_2 . Note that γ contains at least one vertical edge e , since $\Lambda_1 \neq \Lambda_2$. By Lemma 2.18, there exists a periodic point x in the edge e such that the periodic flow line Λ_x is disjoint from the lines Λ_i and diverges from Λ_i in both the forward and backward directions. We now choose an immersed wall W_u from a regular effective set (Definition 4.14(1)) and satisfying the following:

- (1) W_u has tunnel length L much larger than the $(3\delta+2\lambda)$ -fellow-travelling length of Λ_x with each Λ_i ,
- (2) the intersection of Λ_x with the approximation $A(W_u)$ is a segment of length L with centre x ,
- (3) the approximation $A(W_u)$ intersects the edge e at the singleton $\{x\}$ and the geodesic γ in an odd cardinality set, and
- (4) the intersection of the union $\Lambda_1 \cup \Lambda_2 \cup \gamma$ with $A(W_u)$ coincides with the intersection of γ with $A(W_u)$.

Such a W_u exists as \tilde{X} is level-separated (Lemma 4.17) and has many effective walls (Theorem 4.16). Property (1) above allows us to apply Lemma 4.4 to the union of $\Lambda_1 \cup \Lambda_2 \cup \gamma$ with $N(A(W))$ ¹ and conclude that it is a quasiisometric embedding in \tilde{X}^1 . Therefore ∂W_u is disjoint from, and separates, $\partial(\Lambda_1 \cup \Lambda_2)$. Further, by (4), it does not separate $\partial\Lambda_i$. We now choose an M -saturation W of W_u for M much larger than L . By applying Lemma 4.4 again and arguing as above for W_u , we obtain the desired result. \square

6 Separating pairs of points in the Bowditch boundary

The goal of this section is to show that the mapping torus $\Gamma = G \rtimes_{\phi} \mathbb{Z}$ satisfies the hypothesis of Theorem 3.2. Recall that Γ is hyperbolic relative to the collection \mathcal{P} consisting of the suspensions of the free factors H_i of G .

We first make two observations connecting the Gromov boundary of the flow space \tilde{X}^1 with the Bowditch boundary $\partial_B(\Gamma, \mathcal{P})$. We denote by $\widehat{\tilde{X}^1}$ the cone-off of the (1-skeleton of the) flow space \tilde{X}^1 over the set of principal flow lines. Recall also that $\widehat{\Gamma}$ denotes the cone-off of Γ relative to \mathcal{P} .

As a consequence of the relative Švarc–Milnor lemma [CC07, Theorem 5.1], we obtain:

Proposition 6.1 *The cone-off $\widehat{\tilde{X}^1}$ is Γ -equivariantly quasiisometric to the cone-off $\widehat{\Gamma}$.*

Let $\beta: \tilde{X}^1 \rightarrow \widehat{\tilde{X}^1}$ denote the inclusion map from the flow space to its cone-off. We will denote by $\partial\beta$ the induced map from $\partial\tilde{X}^1$ to the Bowditch boundary $\partial_B(\Gamma, \mathcal{P})$. Observe that $\partial\beta$ maps the limit points of any principal flow line Λ to the corresponding cone vertex v_{Λ} . In fact:

Corollary 6.2 *The map $\partial\beta: \partial\tilde{X}^1 \rightarrow \partial_B(\Gamma, \mathcal{P})$ is continuous and surjective. Further, for all $\xi \in \partial_B(\Gamma, \mathcal{P})$, $\partial\beta^{-1}(\xi)$ is either a singleton or contains two points. The latter arises if and only if ξ is a cone vertex. \square*

Proposition 6.3 *If ξ and ζ are points in $\partial\tilde{X}^1$ that are preimages of conical limit points of the Bowditch boundary, then there exists a geodesic line in \tilde{X}^1 joining them.*

If $\xi \in \partial\tilde{X}^1$ is the preimage of a conical limit point of the Bowditch boundary, and z_0 a vertex in \tilde{X} , then there exists a geodesic ray from z_0 that converges to ξ .

Proof Let (x_n) and (y_n) be sequences of vertices in \tilde{X} going to ξ and ζ , respectively, and let $[x_n, y_n]$ be a geodesic in \tilde{X}^1 . Let $z_0 \in T_0$ minimise the Gromov product $(\xi, \zeta)_{z_0} \in \mathbb{N}$, and let R_0 be the minimum.

Let $B(z_0, R)$ be the ball of \tilde{X}^1 of radius R centred at z_0 . We aim to prove that for all $R > R_0$, there exists $m \in \mathbb{N}$ such that $\{B(z_0, R) \cap [x_n, y_n] \mid n > m\}$ is a finite collection of segments of length $\geq 2R - 2R_0 - 10\delta$.

First, for n large enough, $B(z_0, R) \cap [x_n, y_n]$ is indeed such a segment, otherwise one easily gets that $(\xi, \zeta)_{z_0} > R_0$.

If for some R the collection is infinite, let k be such that in T_k the intersection $T_k \cap B(z_0, R) \cap [x_n, y_n]$ is an infinite collection of segments as n varies (since horizontal segments are determined by their endpoints, such a k exists). Let $z_k \in T_k$, and let u_n be in $T_k \cap B(z_0, R) \cap [x_n, y_n]$. Assume that the maximal angle of $[z_k, u_n]_{T_k}$ goes to infinity with n . Let v_k be the closest vertex to z_k at which the angle goes to infinity. Then the same is true for the initial vertical segment of either $[z_k, x_n]$ or $[z_k, y_n]$. It follows that either (x_n) or (y_n) converges to a point that is an image of the parabolic point of the principal flow line of v_k , contrary to our assumption.

The number of subsegments of $[x_n, y_n]$ around z_0 of a given length is therefore bounded when n varies. One can thus diagonally extract subsequences such that, inductively, for each given length, the subsegments of this length of $[x_n, y_n]$ around z_0 make a constant sequence. We thus obtain a biinfinite geodesic between ξ and ζ in \tilde{X} .

The second assertion is obtained with a similar argument. \square

Proposition 6.4 *Let ξ and ζ be two distinct points in the Bowditch boundary, and let W be the M -saturation of a quasiconvex wall W_u such that W cuts a geodesic between $\beta^{-1}(\xi)$ and $\beta^{-1}(\zeta)$. Then $\partial\beta(\partial W) = \partial \text{Stab}(W)$ separates ξ and ζ and there exists a subgroup K of index at most 2 of $\text{Stab}(W)$ such that ξ and ζ are in K -distinct components of $\partial_B(\Gamma, \mathcal{P}) \setminus \partial\beta(\partial W)$.*

Let us mention a word of caution here. Though $\text{Stab}(W)$ is a codimension-1 subgroup (Proposition 4.13) of Γ , the statement above is necessary because, in general, a codimension-1 subgroup need not separate the Bowditch boundary. As an example, take the cyclic subgroup $P = \langle aba^{-1}b^{-1} \rangle$ of the free group $F(a, b)$ with peripheral structure $\mathcal{P} = P$. Here P is a maximal parabolic subgroup that is full, relatively quasiconvex and codimension-1, but does not separate the Bowditch boundary, a circle.

Proof Let ξ and ζ , and W_u be as in the statement, and $\tilde{\xi}$ and $\tilde{\zeta}$ be two points in \tilde{X}^1 in the preimages $\partial\beta^{-1}(\xi)$ and $\partial\beta^{-1}(\zeta)$, respectively. Let \mathcal{U}_ξ and \mathcal{U}_ζ be a clopen partition of $\partial\tilde{X}^1 \setminus \partial W$ containing $\tilde{\xi}$ and $\tilde{\zeta}$, respectively. Consider points in $\partial\tilde{X}^1$ identified in $\partial_B(\Gamma, \mathcal{P})$. By [Corollary 6.2](#) they are at the end of the same principal flow line, and by [Lemma 5.1](#), if one is in \mathcal{U}_ξ , so is the other one. It follows that the map $\partial\beta$ sends \mathcal{U}_ξ and \mathcal{U}_ζ on disjoint subsets $\partial\beta(\mathcal{U}_\xi)$ and $\partial\beta(\mathcal{U}_\zeta)$ of $\partial_B(\Gamma, \mathcal{P}) \setminus \partial\beta(\partial W)$. By surjectivity, they provide a partition of $\partial_B(\Gamma, \mathcal{P}) \setminus \partial\beta(\partial W)$, each containing ξ and ζ , respectively.

We need to check that both $\partial\beta(\mathcal{U}_\xi)$ and $\partial\beta(\mathcal{U}_\zeta)$ are closed. By symmetry, it is enough to check that one is closed. If a sequence (x_n) in $\partial\beta(\mathcal{U}_\xi)$ converges in $\partial_B(\Gamma, \mathcal{P}) \setminus \partial\beta(\partial W)$, lift it to a sequence (x'_n) in $\partial\tilde{X}^1 \setminus \partial W$. Since $\partial\tilde{X}^1$ is not compact, the sequence (x'_n) may or may not have an accumulation point. If it does, then the fact that \mathcal{U}_ξ is closed, along with the continuity of $\partial\beta$, gives us the necessary conclusion. Now consider the case where it has no accumulation point in $\partial\tilde{X}^1$. If x'_n is a point represented by a ray ρ'_n in $\tilde{X}^1 \setminus W$ (from a given basepoint) that lives in the component of which ξ is adherent, it means that, up to taking a subsequence, all the rays ρ'_n pass through a singular vertex, and make an angle θ_n at that vertex such that $\theta_n \rightarrow \infty$ with n . Take v to be the closest singular vertex to the basepoint with this property. The limit of (x_n) in $\partial_B(\Gamma, \mathcal{P}) \setminus \partial\beta(\partial W)$ is then the parabolic point associated to the principal flow line of this vertex v . Since the rays live in the component of $\tilde{X}^1 \setminus W$ adherent to ξ , this principal flow line has at least one endpoint not in \mathcal{U}_ξ . By the property of saturations, either both endpoints are in \mathcal{U}_ξ , or both are in ∂W , and in that case the limit of (x_n) is in $\partial\beta(\partial W)$, which we assumed otherwise. So both points of the line are in \mathcal{U}_ξ , and their images, which are the limit of (x_n) , are in $\partial\beta\mathcal{U}_\xi$. Therefore this set is closed.

Therefore ξ and ζ are separated by $\partial\beta(\partial W)$, which is $\partial \text{Stab}(W)$, since the action of the latter is cocompact on W .

Fix a basepoint in the complement of W . Let A be the union of components C of $\tilde{X} \setminus W$ such that there is a path from the basepoint to C that crosses the set W'_u an even number of times. Recall that W'_u is the complement in W of the set of its principal flow lines. We note that $\tilde{X} \setminus W$ is thus partitioned into two subsets: A and its complement, which we denote here by B . Since the stabiliser of W sends complementary components of W to complementary components, it preserves this partition of W into A and B . Therefore there exists a subgroup of index at most 2 which preserves A (and hence also B).

By the way the saturation W of W_u cuts the preimages of ξ and ζ (as in [Proposition 5.3](#), [5.4](#), [5.5](#) or [5.6](#)), it is easy to see that ξ and ζ do not both lie in the limit set of A , or in the limit set of B , which gives us the desired result. □

We now prove the main result [Theorem 1.1](#), namely, in our current notation, that the relatively hyperbolic group (Γ, \mathcal{P}) is relatively cubulated:

Proof of Theorem 1.1 We will show that the boundary criterion [Theorem 3.2](#) holds. Let $\xi \neq \zeta$ be two points in the Bowditch boundary $\partial_B(\Gamma, \mathcal{P})$. We have three cases:

Case 1 (both ξ and ζ are conical points) By abuse of notation, we denote by ξ and ζ the unique preimages in $\partial\tilde{X}^1$ of ξ and ζ (Corollary 6.2). Let γ be a geodesic in \tilde{X}^1 joining ξ and ζ as given by Proposition 6.3. Then Propositions 5.3 and 5.4 ensure that there exists a wall saturation W that cuts γ . An application of Proposition 6.4 gives the desired result.

Case 2 (without loss of generality, ξ is conical while ζ is parabolic) We consider a geodesic ray γ in \tilde{X}^1 , between a vertex of the principal flow line Λ associated to ζ and the preimage of ξ , as given by Proposition 6.3. Proposition 5.5 ensures that the union of γ and Λ is cut by a wall saturation W . Proposition 6.4 then gives the result.

Case 3 (both ξ and ζ are parabolic points) Proposition 5.6 ensures that the principal flow lines associated to ξ and η are cut by a wall saturation W . The result then follows from Proposition 6.4. \square

References

- [Bow12] B H Bowditch, *Relatively hyperbolic groups*, Int. J. Algebra Comput. 22 (2012) art. id. 1250016 MR Zbl
- [BW12] N Bergeron, D T Wise, *A boundary criterion for cubulation*, Amer. J. Math. 134 (2012) 843–859 MR Zbl
- [CC07] R Charney, J Crisp, *Relative hyperbolicity and Artin groups*, Geom. Dedicata 129 (2007) 1–13 MR Zbl
- [CLR94] D Cooper, D D Long, A W Reid, *Bundles and finite foliations*, Invent. Math. 118 (1994) 255–283 MR Zbl
- [DKL15] S Dowdall, I Kapovich, C J Leininger, *Dynamics on free-by-cyclic groups*, Geom. Topol. 19 (2015) 2801–2899 MR Zbl
- [DL22] F Dahmani, R Li, *Relative hyperbolicity for automorphisms of free products and free groups*, J. Topol. Anal. 14 (2022) 55–92 MR Zbl
- [DM23] F Dahmani, S Krishna M S, *Relative hyperbolicity of hyperbolic-by-cyclic groups*, Groups Geom. Dyn. 17 (2023) 403–426 MR Zbl
- [DMM25] F Dahmani, S K Meda Satish, J P Mutanguha, *Hyperbolic hyperbolic-by-cyclic groups are cubulable*, Geom. Topol. 29 (2025) 259–268 MR Zbl
- [EG20] E Einstein, D Groves, *Relative cubulations and groups with a 2-sphere boundary*, Compos. Math. 156 (2020) 862–867 MR Zbl
- [EG22] E Einstein, D Groves, *Relatively geometric actions on CAT(0) cube complexes*, J. Lond. Math. Soc. 105 (2022) 691–708 MR Zbl
- [EGN21] E Einstein, D Groves, T Ng, *Separation and relative quasiconvexity criteria for relatively geometric actions*, Groups Geom. Dyn. 18 (2024) 649–676 MR Zbl
- [FM15] S Francaviglia, A Martino, *Stretching factors, metrics and train tracks for free products*, Illinois J. Math. 59 (2015) 859–899 MR
- [GL07] V Guirardel, G Levitt, *The outer space of a free product*, Proc. Lond. Math. Soc. 94 (2007) 695–714 MR Zbl

- [GM23] **D Groves, J F Manning**, *Hyperbolic groups acting improperly*, *Geom. Topol.* 27 (2023) 3387–3460 [MR](#) [Zbl](#)
- [Hor17] **C Horbez**, *The boundary of the outer space of a free product*, *Israel J. Math.* 221 (2017) 179–234 [MR](#) [Zbl](#)
- [Hru10] **G C Hruska**, *Relative hyperbolicity and relative quasiconvexity for countable groups*, *Algebr. Geom. Topol.* 10 (2010) 1807–1856 [MR](#) [Zbl](#)
- [HW08] **F Haglund, D T Wise**, *Special cube complexes*, *Geom. Funct. Anal.* 17 (2008) 1551–1620 [MR](#) [Zbl](#)
- [HW15] **M F Hagen, D T Wise**, *Cubulating hyperbolic free-by-cyclic groups: the general case*, *Geom. Funct. Anal.* 25 (2015) 134–179 [MR](#) [Zbl](#)
- [HW16] **M F Hagen, D T Wise**, *Cubulating hyperbolic free-by-cyclic groups: the irreducible case*, *Duke Math. J.* 165 (2016) 1753–1813 [MR](#) [Zbl](#)
- [MR08] **M Mj, L Reeves**, *A combination theorem for strong relative hyperbolicity*, *Geom. Topol.* 12 (2008) 1777–1798 [MR](#) [Zbl](#)
- [NR97] **G Niblo, L Reeves**, *Groups acting on CAT(0) cube complexes*, *Geom. Topol.* 1 (1997) 1–7 [MR](#) [Zbl](#)
- [Osi06] **D V Osin**, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, *Mem. Amer. Math. Soc.* 843, Amer. Math. Soc., Providence, RI (2006) [MR](#) [Zbl](#)
- [Sag95] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, *Proc. Lond. Math. Soc.* 71 (1995) 585–617 [MR](#) [Zbl](#)

*Institut Fourier, University Grenoble Alpes
Grenoble, France*

*Department of Mathematics, Ashoka University
Haryana, India*

francois.dahmani@univ-grenoble-alpes.fr, suraj.meda@ashoka.edu.in

Received: 9 April 2023 Revised: 1 December 2023

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Kathryn Hess
kathryn.hess@epfl.ch
École Polytechnique Fédérale de Lausanne

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Thomas Koberda	University of Virginia thomas.koberda@virginia.edu
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Markus Land	LMU München markus.land@math.lmu.de
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Norihiko Minami	Yamato University minami.norihiko@yamato-u.ac.jp
Octav Cornea	Université de Montreal cornea@dms.umontreal.ca	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Robert Oliver	Université Paris 13 bobol@math.univ-paris13.fr
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
David Futer	Temple University dfuter@temple.edu	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Vesna Stojanoska	Univ. of Illinois at Urbana-Champaign vesna@illinois.edu
Ian Hambleton	McMaster University ian@math.mcmaster.ca	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Kristen Hendricks	Rutgers University kristen.hendricks@rutgers.edu	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com
Daniel Isaksen	Wayne State University isaksen@math.wayne.edu		

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2025 is US \$760/year for the electronic version, and \$1110/year (+\$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by [Mathematical Reviews](#), [Zentralblatt MATH](#), [Current Mathematical Publications](#) and the [Science Citation Index](#).

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing
<https://msp.org/>

© 2025 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 25 Issue 3 (pages 1265–1915) 2025

A-polynomials, Ptolemy equations and Dehn filling	1265
JOSHUA A HOWIE, DANIEL V MATHEWS and JESSICA S PURCELL	
The Alexandrov theorem for $2 + 1$ flat radiant spacetimes	1321
LÉO MAXIME BRUNSWIC	
Real algebraic overtwisted contact structures on 3-spheres	1377
ŞEYMA KARADERELI and FERİT ÖZTÜRK	
Fully augmented links in the thickened torus	1411
ALICE KWON	
Unbounded \mathfrak{sl}_3 -laminations and their shear coordinates	1433
TSUKASA ISHIBASHI and SHUNSUKE KANO	
Bridge trisections and Seifert solids	1501
JASON JOSEPH, JEFFREY MEIER, MAGGIE MILLER and ALEXANDER ZUPAN	
Random Artin groups	1523
ANTOINE GOLDSBOROUGH and NICOLAS VASKOU	
A deformation of Asaeda–Przytycki–Sikora homology	1545
ZHENKUN LI, YI XIE and BOYU ZHANG	
Cubulating a free-product-by-cyclic group	1561
FRANÇOIS DAHMANI and SURAJ KRISHNA MEDA SATISH	
Virtual domination of 3-manifolds, III	1599
HONGBIN SUN	
The Kakimizu complex for genus one hyperbolic knots in the 3-sphere	1667
LUIS G VALDEZ-SÁNCHEZ	
Band diagrams of immersed surfaces in 4-manifolds	1731
MARK HUGHES, SEUNGWON KIM and MAGGIE MILLER	
Anosov flows and Liouville pairs in dimension three	1793
THOMAS MASSONI	
Hamiltonian classification of toric fibres and symmetric probes	1839
JOÉ BRENDEL	
An example of higher-dimensional Heegaard Floer homology	1877
YIN TIAN and TIANYU YUAN	
Fibered 3-manifolds and Veech groups	1897
CHRISTOPHER J LEININGER, KASRA RAFI, NICHOLAS ROUSE, EMILY SHINKLE and YVON VERBERNE	