

Hyperbolic extensions of free groups

SPENCER DOWDALL

SAMUEL J TAYLOR

Given a finitely generated subgroup $\Gamma \leq \text{Out}(\mathbb{F})$ of the outer automorphism group of the rank- r free group $\mathbb{F} = F_r$, there is a corresponding free group extension $1 \rightarrow \mathbb{F} \rightarrow E_\Gamma \rightarrow \Gamma \rightarrow 1$. We give sufficient conditions for when the extension E_Γ is hyperbolic. In particular, we show that if all infinite-order elements of Γ are atoroidal and the action of Γ on the free factor complex of \mathbb{F} has a quasi-isometric orbit map, then E_Γ is hyperbolic. As an application, we produce examples of hyperbolic \mathbb{F} -extensions E_Γ for which Γ has torsion and is not virtually cyclic. The proof of our main theorem involves a detailed study of quasigeodesics in Outer space that make progress in the free factor complex. This may be of independent interest.

20F28, 20F67; 20E06, 57M07

1 Introduction

Let $\mathbb{F} = F_r$ denote the free group of rank $r \geq 3$ and consider its group $\text{Out}(\mathbb{F})$ of outer automorphisms. These groups fit into the short exact sequence

$$1 \longrightarrow \mathbb{F} \xrightarrow{i} \text{Aut}(\mathbb{F}) \xrightarrow{p} \text{Out}(\mathbb{F}) \longrightarrow 1,$$

where $a \in \mathbb{F}$ is mapped to its corresponding inner automorphism i_a , which is defined by $x \mapsto axa^{-1}$ for $x \in \mathbb{F}$. Hence, for any $\Gamma \leq \text{Out}(\mathbb{F})$ we obtain the extension of \mathbb{F}

$$1 \longrightarrow \mathbb{F} \xrightarrow{i} E_\Gamma \xrightarrow{p} \Gamma \longrightarrow 1,$$

where E_Γ is equal to the preimage $p^{-1}(\Gamma) \leq \text{Aut}(\mathbb{F})$. In fact, any extension of \mathbb{F} induces a homomorphism to $\text{Out}(\mathbb{F})$ and thereby produces an extension of the above form (see Section 2.5 for details). This paper will address the following question:

What conditions on $\Gamma \leq \text{Out}(\mathbb{F})$ imply that the extension E_Γ is a hyperbolic group?

This question fits in to a long history of understanding hyperbolic group extensions that goes back to Thurston's work on the hyperbolization of fibered 3-manifolds. From a group-theoretic perspective, the combination theorem of Bestvina and Feighn [4] provides a combinatorial framework to understand the structure of more general hyper-

bolic group extensions. Using this, Farb and Mosher's influential work [16] initiated the systematic study of hyperbolic extensions of surface groups (see Section 1.2). Our answer to the question above continues this investigation in the setting of free group extensions.

1.1 Statements of results

To state our main theorem, we briefly recall the relevant definitions and refer the reader to Section 2 for additional details. First, an outer automorphism $\phi \in \text{Out}(\mathbb{F})$ is *atoroidal*, or *hyperbolic*, if no power of ϕ fixes any nontrivial conjugacy class in \mathbb{F} . Similarly, $\phi \in \text{Out}(\mathbb{F})$ is *fully irreducible* if no power of ϕ preserves the conjugacy class of any proper free factor of \mathbb{F} . The (*free*) *factor complex* \mathcal{F} for the free group \mathbb{F} is the simplicial complex in which each k -simplex corresponds to a set $[A_0], \dots, [A_k]$ of $k+1$ conjugacy classes of proper free factors of \mathbb{F} with properly nested representatives: $A_0 < \dots < A_k$. Note that there is an obvious simplicial action $\text{Out}(\mathbb{F}) \curvearrowright \mathcal{F}$. We prove the following:

Theorem 1.1 *Suppose that each infinite-order element of a finitely generated subgroup $\Gamma \leq \text{Out}(\mathbb{F})$ is atoroidal and that some orbit map $\Gamma \rightarrow \mathcal{F}$ is a quasi-isometric embedding. Then the free group extension E_Γ is hyperbolic.*

Remark Bestvina and Feighn [5] have proven that the factor complex \mathcal{F} is hyperbolic. Hence, the hypotheses of Theorem 1.1 additionally imply that the subgroup Γ is itself hyperbolic and that all infinite-order elements of Γ are fully irreducible. See Section 2.8 for details.

Theorem 1.1 provides combinatorial conditions on a subgroup $\Gamma \leq \text{Out}(\mathbb{F})$ which guarantee that the corresponding extension E_Γ is hyperbolic. This is similar to the better understood situation of hyperbolic extensions of surface groups. For surface group extensions, it follows from work of Farb and Mosher [16], Kent and Leininger [25] and Hamenstädt [18], that a subgroup H of the mapping class group induces a hyperbolic extension of the surface group if and only if H admits a quasi-isometric embedding into the curve complex of the surface. See Section 1.2 for details.

Remark Unlike the surface group case (see Theorem 1.2), the converse to Theorem 1.1 does not hold: there exists subgroups $\Gamma \leq \text{Out}(\mathbb{F})$ for which E_Γ is hyperbolic but Γ does not quasi-isometrically embed into \mathcal{F} . For example, Theorem 2.13, which is due to Brinkmann [11], shows that any $\phi \in \text{Out}(\mathbb{F})$ that is atoroidal but not fully irreducible generates a cyclic subgroup of this form.

The proof of Theorem 1.1 requires several steps and is completed in Section 8 (see Corollary 8.3). The first of these steps is to show that the assumption that the orbit

map $\Gamma \rightarrow \mathcal{F}$ is a quasi-isometric embedding implies a strong quasiconvexity property for the orbit of Γ in Outer space \mathcal{X} , the space of \mathbb{F} -marked metric graphs. This follows from our next main result, Theorem 4.1, which says that quasigeodesics in Outer space that make definite progress in the factor complex are stable. For the statement, the injectivity radius of $G \in \mathcal{X}$ is the length of the shortest loop in the marked metric graph G , and the ϵ -thick part \mathcal{X}_ϵ is the set of points with injectivity radius at least ϵ . Additionally, $\pi: \mathcal{X} \rightarrow \mathcal{F}$ denotes the (coarse) map that associates to each marked graph $G \in \mathcal{X}$ the collection $\pi(G)$ of nontrivial free factors that arise as the fundamental group of a proper subgraph of G .

Theorem 4.1 *Let $\gamma: I \rightarrow \mathcal{X}$ be a K -quasigeodesic whose projection $\pi \circ \gamma: I \rightarrow \mathcal{F}$ is also a K -quasigeodesic. Then there exist constants $A, \epsilon > 0$ and $K' \geq 1$ depending only on K (and the injectivity radius of the terminal endpoint $\gamma(I_+)$ when $I_+ < \infty$) with the following property: if $\rho: J \rightarrow \mathcal{X}$ is any geodesic with the same endpoints as γ , then*

- (i) $\gamma(I), \rho(J) \subset \mathcal{X}_\epsilon$,
- (ii) $d_{\text{Haus}}(\gamma(I), \rho(J)) < A$, and
- (iii) $\pi \circ \rho: J \rightarrow \mathcal{F}$ is a (parametrized) K' -quasigeodesic.

In the statement of Theorem 4.1, γ and ρ are directed (quasi)geodesics with respect to the asymmetric Lipschitz metric $d_{\mathcal{X}}$ on Outer space, and d_{Haus} denotes the Hausdorff distance with respect to the symmetrized Lipschitz distance; see Section 2 for a more detailed discussion of this terminology. Theorem 4.1 is analogous to the stability theorem of Hamenstädt [19] for quasigeodesics in Teichmüller space that make definite progress in the curve complex.

Theorem 1.1 allows one to easily construct hyperbolic extensions of free groups using ping-pong arguments on hyperbolic $\text{Out}(\mathbb{F})$ -graphs. For example, we can recover (Theorem 9.3) the theorem of Bestvina, Feighn and Handel [7] which states that if f_1, \dots, f_k are atoroidal, fully irreducible elements of $\text{Out}(\mathbb{F})$, then for all sufficiently large $N \geq 1$ the extension E_Γ is hyperbolic for $\Gamma = \langle f_1^N, \dots, f_k^N \rangle \leq \text{Out}(\mathbb{F})$. (In [7], this is proven for $k = 2$.) Further, we use Theorem 1.1 to construct the first examples of hyperbolic free group extensions E_Γ for which $\Gamma \leq \text{Out}(\mathbb{F})$ has torsion and is not virtually cyclic. First, say that $f \in \text{Out}(\mathbb{F})$ is independent for a finite subgroup $H \leq \text{Out}(\mathbb{F})$ if f and hfh^{-1} have no common powers for each $h \in H \setminus 1$. We prove the following:

Theorem 9.4 *Let H be a finite subgroup of $\text{Out}(\mathbb{F})$ and let $f \in \text{Out}(\mathbb{F})$ be a hyperbolic, fully irreducible outer automorphism that is independent for H . Then for all*

sufficiently large $N \geq 1$, the subgroup

$$\Gamma = \langle H, f^N \rangle$$

is isomorphic to $H * \mathbb{Z}$ and the \mathbb{F} -by- $(H * \mathbb{Z})$ extension E_Γ is hyperbolic.

1.2 Motivation from surface group extensions and some previous results

In [16], Farb and Mosher introduced convex cocompact subgroups of $\text{Mod}(S)$, the mapping class group of an orientable surface S . We will discuss the case where S is further assumed to be closed. A finitely generated subgroup $\Gamma \leq \text{Mod}(S)$ is *convex cocompact* if for some (any) $x \in \text{Teich}(S)$, the Teichmüller space of the surface S , the orbit $\Gamma \cdot x \subset \text{Teich}(S)$ is quasiconvex with respect to the Teichmüller metric. (See the papers of Farb and Mosher [16] and Kent and Leininger [25; 24] for definitions and details.) Similar to the situation described above, a subgroup $\Gamma \leq \text{Mod}(S)$ gives rise to a surface group extension

$$1 \rightarrow \pi_1(S) \rightarrow E_\Gamma \rightarrow \Gamma \rightarrow 1.$$

Farb and Mosher show that if E_Γ is hyperbolic then Γ is convex cocompact. Moreover, they prove that if Γ is assumed to be free, then convex cocompactness of Γ implies that the extension E_Γ is hyperbolic [16]. The assumption that Γ is free was later removed by Hamenstädt in [18]. Hence, the surface group extension E_Γ is hyperbolic exactly when $\Gamma \leq \text{Mod}(S)$ is convex cocompact. We note that the first examples of hyperbolic surface group extensions follow from work of Thurston, whose geometrization theorem for fibered 3-manifolds produces examples of hyperbolic surface-by-cyclic groups. Later, Mosher [29] constructed more general hyperbolic surface-by-free groups using the combination theorem of Bestvina and Feighn [4].

Since their introduction, there have been several additional characterizations of convex cocompact subgroups of $\text{Mod}(S)$. A particularly useful characterization of convex cocompactness is the following theorem of Kent and Leininger and Hamenstädt. In the statement, $\mathcal{C}(S)$ denotes the curve complex for the closed surface S .

Theorem 1.2 (Kent and Leininger [25]; Hamenstädt [18]) *A finitely generated subgroup $\Gamma \leq \text{Mod}(S)$ is convex cocompact if and only if some (any) orbit map $\Gamma \rightarrow \mathcal{C}(S)$ is a quasi-isometric embedding.*

From this we see that the surface group extension E_Γ is hyperbolic if the orbit map from $\Gamma \leq \text{Mod}(S)$ into the curve complex is a quasi-isometric embedding. Hence, strong geometric features of surface group extensions arise from combinatorial conditions on their corresponding subgroups of $\text{Mod}(S)$. With Theorem 1.1,

we provide analogous conditions under which combinatorial information about a subgroup $\Gamma \leq \text{Out}(\mathbb{F})$ implies geometric information about the corresponding free group extension E_Γ .

Remark The condition that every infinite-order element of Γ is atoroidal is necessary for E_Γ to be hyperbolic, but this condition is not implied by having a quasi-isometric orbit map into the factor complex \mathcal{F} . This contrasts with the surface group situation (see Theorem 1.2), where having a quasi-isometric orbit map $\Gamma \rightarrow \mathcal{C}(S)$ automatically implies every infinite-order element of Γ is pseudo-Anosov. Indeed, there are elements of $\text{Out}(\mathbb{F})$ that act with positive translation length on \mathcal{F} but are not atoroidal. By Bestvina and Handel [9], these all arise as pseudo-Anosov mapping classes on surfaces with a single puncture. Since such outer automorphisms each fix a conjugacy class in \mathbb{F} (corresponding to the loop enclosing the puncture), they cannot be contained in a subgroup Γ for which E_Γ is hyperbolic.

We conclude this section with a brief review of previous examples of hyperbolic extensions of free groups. In [4], Bestvina and Feighn produce examples of hyperbolic free-by-cyclic groups (ie $\Gamma \cong \mathbb{Z}$) using automorphisms assumed to satisfy the Bestvina–Feighn flaring conditions. Later, Brinkmann [11] showed that any atoroidal automorphism induces a hyperbolic free-by-cyclic group by showing that all such automorphisms satisfy these flaring conditions. This is recorded in Theorem 2.13.

The first examples where $\Gamma \leq \text{Out}(\mathbb{F})$ is not cyclic are given in [7]. There, Bestvina, Feighn and Handel show that if one starts with fully irreducible and atoroidal elements $\phi, \psi \in \text{Out}(\mathbb{F})$ that do not have a common power, then there is an $N \geq 1$ such that $\Gamma = \langle \phi^N, \psi^N \rangle$ is a rank-2 free group and the corresponding extension E_Γ is hyperbolic. A different proof of this fact (still using the Bestvina–Feighn combination theorem) is given by Kapovich and Lustig [23], who additionally show that for large N each nonidentity element of Γ is fully irreducible.

1.3 Outline of proof

To show that the extension E_Γ is hyperbolic, we use the combination theorem of Mj and Sardar [28], which is recalled in Section 2.4. Their theorem states that if a metric graph bundle satisfies a certain *flaring property* (terminology coming from the Bestvina–Feighn combination theorem), then the bundle is hyperbolic. Using the map between the Cayley graphs of E_Γ and Γ as our graph bundle, we show in Section 8 that this flaring property is implied by the following conjugacy flaring property of $\Gamma \leq \text{Out}(\mathbb{F})$. First let S be a finite symmetric generating set for Γ with associated word norm $|\cdot|_S$. Also

fix a basis X for \mathbb{F} . We say that Γ has (λ, M) -conjugacy flaring for the given $\lambda > 1$ and positive integer $M \in \mathbb{N}$ if the following condition is satisfied:

For all $\alpha \in \mathbb{F}$ and $g_1, g_2 \in \Gamma$ with $|g_i|_S \geq M$ and $|g_1 g_2|_S = |g_1|_S + |g_2|_S$, we have

$$\lambda \|\alpha\|_X \leq \max\{\|g_1(\alpha)\|_X, \|g_2^{-1}(\alpha)\|_X\},$$

where $\|\cdot\|_X$ denotes conjugacy length (ie the shortest word length with respect to X of any element in the given conjugacy class).

Proposition 8.1 shows that if $\Gamma \leq \text{Out}(\mathbb{F})$ has conjugacy flaring, then E_Γ has the Mj–Sardar flaring property, and hence, E_Γ is hyperbolic. Thus it suffices to show that any $\Gamma \leq \text{Out}(\mathbb{F})$ satisfying the hypotheses of Theorem 1.1 has conjugacy flaring. This is accomplished by using the geometry of Outer space.

First, Theorem 4.1 is used to show that geodesic words in $(\Gamma, |\cdot|_S)$ are sent via the orbit map $\Gamma \rightarrow \mathcal{X}$ to quasigeodesics that fellow-travel a special class of paths in \mathcal{X} , called *folding paths*. Therefore, by the definition of distance in \mathcal{X} (Proposition 2.5), the conjugacy length of $\alpha \in \mathbb{F}$ along the quasigeodesic in Γ is proportional to the conjugacy length of α along the nearby folding path. Thus it suffices to show that the length of every conjugacy class “flares” along any folding path that remains close to the orbit of Γ in \mathcal{X} , meaning that the length grows at a definite exponential rate in either the forward direction or the backward direction. (See Section 6 for details.) Proposition 6.11 proves exactly this type of flaring for folding paths that remain close to the orbit of any group Γ that satisfies the hypotheses of Theorem 1.1.

To summarize: If the orbit map $\Gamma \rightarrow \mathcal{F}$ is a quasi-isometric embedding and every infinite-order element of Γ is atoroidal then folding paths between points in the orbit $\Gamma \cdot R$ (for $R \in \mathcal{X}$) have the flaring property (Section 6). This, together with the fact that these folding paths fellow-travel the image of geodesics in the group Γ (Theorem 4.1), implies that Γ has conjugacy flaring (Theorem 6.5). Finally, Proposition 8.1 shows that conjugacy flaring of Γ implies that the hypotheses of the Mj–Sardar theorem are satisfied and that Γ is hyperbolic.

Acknowledgements We would like to thank Chris Leininger, Alan Reid, Patrick Reynolds and Caglar Uyanik for helpful conversations, as well as Ilya Kapovich for valuable comments on an earlier version of this paper. We also thank the referee for a careful reading of the paper and for several helpful suggestions. We are grateful to the University of Illinois at Urbana-Champaign and the University of Texas at Austin for hosting the authors during different stages of this project, and to the GEAR Network for supporting a trip during which this paper was completed.

Dowdall was partially supported by the NSF postdoctoral fellowship, NSF MSPRF no. 1204814. Taylor was partially supported by a department fellowship from the University of Texas at Austin and by the NSF postdoctoral fellowship, NSF MSPRF no. 1400498. Both authors acknowledge support from U.S. National Science Foundation grants DMS 1107452, 1107263, 1107367 “RNMS: GEometric structures And Representation varieties” (the GEAR Network).

2 Preliminaries

2.1 Paths

Throughout this paper, the notation I (or sometimes J) will be used to denote a closed, connected interval $I \subseteq \mathbb{R}$. We write $I_{\pm} \in \mathbb{R} \cup \{\pm\infty\}$ for the positive and negative endpoints of I , respectively, and correspondingly write $I = [I_-, I_+]$. By a *discrete interval*, we simply mean the integer points $I \cap \mathbb{Z}$ of an interval $I \subset \mathbb{R}$.

A *path* in a topological space Y is a map $\gamma: I \rightarrow Y$. If Y is a metric space, then the path γ is said to be a *geodesic* if $d_Y(\gamma(a), \gamma(b)) = |a - b|$ for all $a, b \in I$ (that is, if γ is an isometric embedding of I into Y). A *discrete geodesic* is similarly a map $\gamma: (I \cap \mathbb{Z}) \rightarrow Y$ of a discrete interval into Y such that $d_Y(\gamma(a), \gamma(b)) = |a - b|$ for all $a, b \in I \cap \mathbb{Z}$. The space Y is said to be a *geodesic metric space* if it is a metric space and for any points $y_+, y_- \in Y$ there exists a finite geodesic $\gamma: I \rightarrow Y$ with $\gamma(I_{\pm}) = y_{\pm}$.

2.2 Coarse geometry

Suppose that X and Y are metric spaces. Given a constant $K \geq 1$, a map $f: X \rightarrow Y$ is said to be a *K -quasi-isometric embedding* if for all $a, b \in X$ we have

$$\frac{1}{K}d_X(a, b) - K \leq d_Y(f(a), f(b)) \leq Kd_X(a, b) + K.$$

More generally, the map is said to be *coarsely K -Lipschitz* if the rightmost inequality above holds. A *K -quasi-isometry* is a K -quasi-isometric embedding $f: X \rightarrow Y$ whose image $f(X)$ is D -dense for some $D \geq 0$. (This the equivalent to the existence of a K' -quasi-isometric embedding $g: Y \rightarrow X$ for which $f \circ g$ and $g \circ f$ are within bounded distance of Id_Y and Id_X , respectively.)

A *K -quasigeodesic* in a metric space Y is a K -quasi-isometric embedding $\gamma: I \rightarrow Y$ of an interval $I \subset \mathbb{R}$ into Y . Similarly, a *discrete K -quasigeodesic* is a K -quasi-isometric embedding $\gamma: (I \cap \mathbb{Z}) \rightarrow Y$ of a discrete interval into Y .

For $A \geq 0$, the A -neighborhood of a subset Z of a metric space Y will be denoted by

$$\mathcal{N}_A(Z) := \{y \in Y : \inf\{d(z, y) : z \in Z\} < A\}.$$

The Hausdorff distance between two subsets $Z, Z' \subset Y$ is then defined to be

$$d_{\text{Haus}}(Z, Z') := \inf\{\epsilon > 0 : Z \subset \mathcal{N}_\epsilon(Z') \text{ and } Z' \subset \mathcal{N}_\epsilon(Z)\}.$$

Finally, when Y is a geodesic metric space, a subset $Z \subset Y$ is said to be A -quasiconvex if every (finite) geodesic with endpoints in Z is contained $\mathcal{N}_A(Z)$.

2.3 Gromov hyperbolicity

Given $\delta \geq 0$, a geodesic metric space Y is δ -hyperbolic if every geodesic triangle Δ in Y is δ -thin, meaning that each side of Δ lies in the δ -neighborhood of the union of the other two sides. A metric space is hyperbolic if it is δ -hyperbolic for some $\delta \geq 0$. It is a fact (see [10, Proposition III.H.1.17]) that if X is a δ -hyperbolic space then there is a constant $\delta' = \delta'(\delta)$ such that every triangle Δ has a δ' -barycenter, meaning a point $c \in X$ that lies within δ' of each side of Δ .

Every hyperbolic metric space Y has a Gromov boundary ∂Y defined to be the set of equivalence classes of admissible sequences in Y , where a sequence $\{y_n\}$ is admissible if $\lim_{n,m} (y_n | y_m)_x = \infty$ and two sequences $\{y_n\}$ and $\{z_n\}$ are equivalent if $\lim_{n,m} (y_n | z_m)_x = \infty$ for some basepoint $x \in Y$. (Here $(a | b)_x$ denotes the Gromov product $(d(a, x) + d(b, x) - d(a, b))/2$ of $a, b \in Y$ with respect to $x \in Y$.) Notice that by the triangle inequality, the notions of “admissible” and “equivalent” do not depend on the point $x \in Y$. One says that the admissible sequence $y_1, y_2, \dots \in Y$ converges to the point $\{y_n\} \in \partial Y$. In particular, every quasigeodesic ray $\gamma: [0, \infty) \rightarrow Y$ converges to a well-defined endpoint at infinity $\gamma(\infty) := \{\gamma(n)\}_{n=1}^\infty \in \partial Y$, and we note that any two quasigeodesic rays whose images have finite Hausdorff distance converge to the same endpoint at infinity. We refer the reader to [12, Section 2.2] for additional details.

Consequently, to each quasigeodesic $\gamma: \mathbf{I} \rightarrow Y$ we may associate two well-defined endpoints $\gamma(\mathbf{I}_+), \gamma(\mathbf{I}_-) \in Y \cup \partial Y$, where $\gamma(\mathbf{I}_\pm)$ is understood to be a point of ∂Y when $\mathbf{I}_\pm = \pm\infty$ and is a point of Y when $\mathbf{I}_\pm \in \mathbb{R}$. With this terminology, we have the following well-known consequence of hyperbolicity; see [10, Theorem III.H.1.7] for a proof.

Proposition 2.1 (stability of quasigeodesics) *For any given $K, \delta > 0$, there exists a stability constant $R_0 = R_0(\delta, K) > 0$ with the following property: Let Y be a δ -hyperbolic space. If $\gamma: \mathbf{I} \rightarrow Y$ and $\rho: \mathbf{J} \rightarrow Y$ are K -quasigeodesics with the same endpoints, then $\gamma(\mathbf{I})$ and $\rho(\mathbf{J})$ have Hausdorff distance at most R_0 from each other.*

Thinness of triangles in a hyperbolic spaces extends to ideal triangles: that is, given $\delta \geq 0$ there is a constant δ'' such that every geodesic triangle with vertices in $X \cup \partial X$ is δ'' -thin, and there exists a barycenter point $c \in X$ that lies within δ'' of each side of the triangle [35, Theorem 6.24].

Hyperbolic groups Let Γ be a finitely generated group. For any finite generating set S , we may build the corresponding Cayley graph $\text{Cay}(\Gamma, S)$ and equip it with the path metric in which all edges have length one. The group Γ is then given the subspace metric, which is equal to the word metric for the given generating set S . Up to quasi-isometry, this metric is independent of the choice of generating set. Since the inclusion $\Gamma \hookrightarrow \text{Cay}(\Gamma, S)$ is a 1-quasi-isometry, we often blur the distinction between Γ and its Cayley graph when considering Γ as a metric space. Accordingly, the group Γ is said to be δ -hyperbolic if there is a finite generating set whose Cayley graph is δ -hyperbolic. In this case, boundary $\partial\Gamma$ of Γ is defined to be the Gromov boundary of the Cayley graph. Equivalently $\partial\Gamma$ is the set of equivalence classes of discrete quasigeodesic rays $\gamma: \mathbb{N} \rightarrow \Gamma$.

2.4 Metric bundles

We will make use of the concept of metric graph bundles introduced by Mj and Sardar in [28]. Let X and B be connected graphs equipped their respective path metrics (in which each edge has length 1), and let $p: X \rightarrow B$ be a simplicial surjection. Write $V(B)$ for the vertex set of the graph B . We say that X is a *metric graph bundle over B* if there is a function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that:

- For each vertex $b \in V(B)$, the fiber $F_b = p^{-1}(b)$ is a connected subgraph of X and the induced path metric d_b on F_b satisfies $d_b(x, y) \leq f(d_X(x, y))$ for all vertices x and y of F_b .
- For any adjacent vertices $b_1, b_2 \in V(B)$ and any vertex $x_1 \in F_{b_1}$, there is a vertex $x_2 \in F_{b_2}$ that is adjacent to x_1 .

Suppose that $p: X \rightarrow B$ is a metric graph bundle. By a k -qi lift of a geodesic $\gamma: I \rightarrow B$ (where $k \geq 1$) we mean any k -quasigeodesic $\tilde{\gamma}: I \rightarrow X$ such that $p(\tilde{\gamma}(n)) = \gamma(n)$ for all $n \in I \cap \mathbb{Z}$. We then say that the metric bundle $p: X \rightarrow B$ satisfies the *flaring condition* if for all $k \geq 1$ there exists $\lambda_k > 1$ and $n_k, M_k \in \mathbb{N}$ such that the following holds: for any geodesic $\gamma: [-n_k, n_k] \rightarrow B$ and any two k -qi lifts $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ satisfying $d_{\gamma(0)}(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) \geq M_k$ we have

$$\lambda_k \cdot d_{\gamma(0)}(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) \leq \max\{d_{\gamma(n_k)}(\tilde{\gamma}_1(n_k), \tilde{\gamma}_2(n_k)), d_{\gamma(-n_k)}(\tilde{\gamma}_1(-n_k), \tilde{\gamma}_2(-n_k))\}.$$

The following combination theorem of Mj and Sardar [28] is the key tool that allows us to prove hyperbolicity of group extensions. It builds on the original Bestvina–Feighn

combination theorem [4] (in the case where B is a tree) and is also related to a combination theorem of Hamenstädt [18].

Theorem 2.2 (Mj and Sardar [28]) *Suppose that a metric graph bundle $p: X \rightarrow B$ satisfies:*

- (1) B is δ -hyperbolic, and each fiber $F_b = p^{-1}(b)$, for $b \in V(B)$, is δ -hyperbolic with respect to the path metric d_b induced by X .
- (2) For each $b \in V(B)$, the set of barycenters of ideal triangles in F_b is D -dense.
- (3) The flaring condition holds.

Then X is a hyperbolic metric space.

2.5 Free group extensions

In general, an \mathbb{F} -extension is any group E that fits into a short exact sequence of the form

$$(1) \quad 1 \rightarrow \mathbb{F} \rightarrow E \rightarrow Q \rightarrow 1.$$

We often blur the distinction between the group E and the short exact sequence itself. Every such extension gives rise to a homomorphism $\chi: Q \rightarrow \text{Out}(\mathbb{F})$ by sending $q \in Q$ to the outer automorphism class of $(\alpha \mapsto \tilde{q}\alpha\tilde{q}^{-1}) \in \text{Aut}(\mathbb{F})$, where $\tilde{q} \in E$ is any lift of q . Since different choices of lift give automorphisms that differ by conjugation by an element of \mathbb{F} , this gives a well defined homomorphism to $\text{Out}(\mathbb{F})$. Conversely, any homomorphism $\chi: Q \rightarrow \text{Out}(\mathbb{F})$ gives rise to an \mathbb{F} -extension E_χ via the fiber product construction:

$$E_\chi := \{(t, q) \in \text{Aut}(\mathbb{F}) \times Q : p(t) = \chi(q)\}.$$

Indeed, if E is the extension in (1) with corresponding homomorphism $\chi: Q \rightarrow \text{Out}(\mathbb{F})$, then $E \cong E_\chi$. In the case of a subgroup $\Gamma \leq \text{Out}(\mathbb{F})$, we write E_Γ for the \mathbb{F} -extension induced by the inclusion $\Gamma \hookrightarrow \text{Out}(\mathbb{F})$.

As in the introduction, there is a canonical short exact sequence

$$1 \longrightarrow \mathbb{F} \xrightarrow{i} \text{Aut}(\mathbb{F}) \xrightarrow{p} \text{Out}(\mathbb{F}) \longrightarrow 1.$$

This sequence is natural for \mathbb{F} -extensions in the sense that any extension E as in (1) with corresponding homomorphism $\chi: Q \rightarrow \text{Out}(\mathbb{F})$ fits into a commutative diagram

$$\begin{array}{ccccccc} 1 & \longrightarrow & \mathbb{F} & \longrightarrow & E \cong E_\chi & \longrightarrow & Q \longrightarrow 1 \\ & & \parallel & & \downarrow \hat{\chi} & & \downarrow \chi \\ 1 & \longrightarrow & \mathbb{F} & \xrightarrow{i} & \text{Aut}(\mathbb{F}) & \xrightarrow{p} & \text{Out}(\mathbb{F}) \longrightarrow 1 \end{array}$$

in which $\hat{\chi}$ is the restriction of the projection $\text{Aut}(\mathbb{F}) \times Q \rightarrow \text{Aut}(\mathbb{F})$ to E_χ . Therefore $\hat{\chi}$ surjects E_χ onto the preimage of $\chi(Q) \leq \text{Out}(\mathbb{F})$ in $\text{Aut}(\mathbb{F})$. From this we note that the $\ker(\hat{\chi}) = 1 \times \ker(\chi) \leq \text{Aut}(\mathbb{F}) \times Q$; thus $\hat{\chi}$ and χ have isomorphic kernels. Moreover, we see that in the case of a subgroup $\Gamma \leq \text{Out}(\mathbb{F})$, the extension E_Γ mentioned above is isomorphic to the preimage $E_\Gamma = p^{-1}(\Gamma) \leq \text{Aut}(\mathbb{F})$.

Note that in order for the extension E_χ to be hyperbolic, it is necessary that the map $\chi: Q \rightarrow \text{Out}(\mathbb{F})$ have finite kernel and for its image to be purely hyperbolic; in which case the above shows E_χ is quasi-isometric to $E_{\chi(Q)}$. Otherwise, it is easily seen that E_χ contains a $\mathbb{Z} \oplus \mathbb{Z}$ and thus that E_χ cannot be hyperbolic. Hence, to address the question of hyperbolicity of \mathbb{F} -extensions, it suffices to focus on the case of extensions E_Γ associated to subgroups $\Gamma \leq \text{Out}(\mathbb{F})$. With this perspective, we only consider such extensions E_Γ throughout the rest of this paper.

2.6 Metric properties of Outer space

Outer space Let \mathbb{F} denote the free group of rank $r = \text{rk}(\mathbb{F})$. Since \mathbb{F} is fixed throughout our discussion, its rank r will often be suppressed from the notation. Letting \mathcal{R} denote the r -petal rose (that is, a wedge of r circles) with vertex $v \in \mathcal{R}$, we fix once and for all an isomorphism $\mathbb{F} \cong \pi_1(\mathcal{R}, v)$. A *graph* is a 1-dimensional CW complex, and a connected, simply connected graph is a *tree*. A *core graph* is a graph all of whose vertices have valence at least 2. Any connected graph G with nontrivial, finitely generated fundamental group has a unique core subgraph whose inclusion into G is a homotopy equivalence. This subgraph is called the *core* of G .

Culler and Vogtmann's [14] Outer space \mathcal{X} of marked metric graphs will play a central role in our discussion. A *marked graph* (G, g) is a core graph G together with a homotopy equivalence $g: \mathcal{R} \rightarrow G$, called a marking. A *metric* on G is a function $\ell: E(G) \rightarrow \mathbb{R}_{>0}$ from the set of edges of G to the positive real numbers; we say that an edge $e \in E(G)$ of G has length $\ell(e)$. The *volume* of G is defined to be $\sum_{e \in E(G)} \ell(e)$. We view the metric ℓ as making G into a path metric space in which each edge e has length $\ell(e)$. A *marked metric graph* is then defined to be the triple (G, g, ℓ) , and we say that two triples (G_1, g_1, ℓ_1) and (G_2, g_2, ℓ_2) are *equivalent* if there is a graph isometry $\phi: G_1 \rightarrow G_2$ that preserves the markings in the sense that $\phi \circ g_1$ is homotopic to g_2 . *Outer space* \mathcal{X} is the set of equivalence classes of marked metric graphs of volume 1. We use the notation $\hat{\mathcal{X}}$ to denote *unprojectivized Outer space*, which is the space of marked metric graphs with no restriction on volume. When discussing points in \mathcal{X} or $\hat{\mathcal{X}}$ we typically suppress the marking/metric and just write the core graph.

Conjugacy classes The marking $\mathcal{R} \rightarrow G$ attached to a point $G \in \mathcal{X}$ allows us to view any nontrivial conjugacy class α in \mathbb{F} as a homotopy class of loops in the core graph G .

Following the notation of [5], we denote the unique immersed loop in this homotopy class by $\alpha|G$, which we view as an immersion of S^1 into G . We use $\ell(\alpha|G)$ to denote the *length of α in $G \in \mathcal{X}$* , that is, the sum of the lengths of the edges crossed by $\alpha|G$, counted with multiplicities. Note that if $X = \{x_1, \dots, x_r\}$ is a free basis of \mathbb{F} and $G \in \mathcal{X}$ is the rose whose edges each have length $1/r$ and are consecutively labeled by the elements x_1, \dots, x_r , then $r \cdot \ell(\alpha|G)$ is simply the *conjugacy length* $\|\alpha\|_X$ of α with respect to the free basis X . That is, $r \cdot \ell(\alpha|G) = \|\alpha\|_X$ is the length of the shortest word in the letters x_1^\pm, \dots, x_r^\pm that represents an element of the conjugacy class α . We often blur the distinction between an element of \mathbb{F} and its conjugacy class.

The *standard topology on \mathcal{X}* is defined to be the coarsest topology such that all of the length functions $\ell(\alpha|\cdot) : \mathcal{X} \rightarrow \mathbb{R}_+$ are continuous [14]. Though we will not discuss it, this topology may also be obtained as a simplicial complex with missing faces, or as the equivariant Gromov–Hausdorff topology (see [14] and [31]). For $\epsilon > 0$, we additionally define the ϵ -*thick part of \mathcal{X}* to be the subset

$$\mathcal{X}_\epsilon := \{G \in \mathcal{X} : \ell(\alpha|G) \geq \epsilon \text{ for every nontrivial conjugacy class } \alpha \text{ in } \mathbb{F}\}.$$

Lipschitz metric A *difference of markings* from $G \in \mathcal{X}$ to $H \in \mathcal{X}$ is any (not necessarily cellular) map $\phi : G \rightarrow H$ that is homotopic to $h \circ g^{-1}$, where g and h are the markings on G and H , respectively. The *Lipschitz distance* from G to H is then defined to be

$$d_{\mathcal{X}}(G, H) := \inf\{\log(\text{Lip}(\phi)) : \phi \simeq h \circ g^{-1}\},$$

where $\text{Lip}(\phi)$ denotes the Lipschitz constant of the difference of markings ϕ . While $d_{\mathcal{X}}$ is in general asymmetric (that is, $d_{\mathcal{X}}(G, H) \neq d_{\mathcal{X}}(H, G)$), we often regard it as a metric since it satisfies definiteness ($d_{\mathcal{X}}(G, H) = 0$ if and only if $G = H$) and the ordered triangle inequality ($d_{\mathcal{X}}(E, H) \leq d_{\mathcal{X}}(E, G) + d_{\mathcal{X}}(G, H)$). Its symmetrization

$$d_{\mathcal{X}}^{\text{sym}}(G, H) := d_{\mathcal{X}}(G, H) + d_{\mathcal{X}}(H, G)$$

is therefore an honest metric on \mathcal{X} , which we note induces the standard topology [17]. The preference to work with the asymmetric metric $d_{\mathcal{X}}$ comes from the fact, discussed below, that folding paths are directed geodesics, whereas the symmetrized metric on \mathcal{X} is not a geodesic metric.

Note that for any $\alpha \in \mathbb{F}$ and any difference of marking $\phi : G \rightarrow H$, by definition we have $\ell(\alpha|H) \leq \text{Lip}(\phi) \cdot \ell(\alpha|G)$. This implies that

$$(2) \quad \log\left(\sup_{\alpha \in \mathbb{F}} \frac{\ell(\alpha|H)}{\ell(\alpha|G)}\right) \leq \log(\inf_{\phi} \text{Lip}(\phi)) = d_{\mathcal{X}}(G, H).$$

We will see below that this is, in fact, an equality; see also [17] and [3]. It follows that for any free basis $X = \{x_1, \dots, x_r\}$ of \mathbb{F} and any point $G \in \mathcal{X}$, there is a constant $K = K(X, G) \geq 1$ such that $(1/K)\|\alpha\|_X \leq \ell(\alpha|G) \leq K\|\alpha\|_X$ for every conjugacy class α in \mathbb{F} .

Coping with asymmetry Since the Lipschitz metric $d_{\mathcal{X}}$ is not symmetric, some care must be taken when discussing distances in \mathcal{X} . Thankfully, the difficulty is somewhat mitigated in the thick part \mathcal{X}_ϵ .

Lemma 2.3 (Handel and Mosher [20]; Algom-Kfir and Bestvina [2]) *For any $\epsilon > 0$, there exists $M_\epsilon \geq 1$ such that for all $G, H \in \mathcal{X}_\epsilon$ we have*

$$d_{\mathcal{X}}(H, G) \leq d_{\mathcal{X}}^{\text{sym}}(H, G) = d_{\mathcal{X}}^{\text{sym}}(G, H) \leq M_\epsilon \cdot d_{\mathcal{X}}(G, H).$$

Nevertheless, whenever discussing neighborhoods, we always use the symmetrized distance $d_{\mathcal{X}}^{\text{sym}}$. That is, the A -neighborhood of a subset $Z \subset \mathcal{X}$ is defined to be

$$\mathcal{N}_A(Z) := \{G \in \mathcal{X} : \inf\{d_{\mathcal{X}}^{\text{sym}}(G, H) : H \in Z\} < A\}.$$

In particular, if $G \in \mathcal{N}_A(Z)$, then there exists some $H \in Z$ such that $d_{\mathcal{X}}(G, H)$ and $d_{\mathcal{X}}(H, G)$ are both less than A . Note that by [17], if $Z \subset \mathcal{X}$ is compact, then so is the closed neighborhood $\overline{\mathcal{N}_A(Z)}$. The Hausdorff distance between two subsets of \mathcal{X} is then defined as usual using these symmetrized neighborhoods.

We will say that two paths $\gamma: I \rightarrow \mathcal{X}$ and $\gamma': I' \rightarrow \mathcal{X}$ have the *same terminal endpoint* either if $I_+, I'_+ < \infty$ and $\gamma(I_+) = \gamma'(I'_+)$ or if $I_+ = I'_+ = \infty$ and the sets $\gamma([t, \infty))$ and $\gamma'([t', \infty))$ have finite Hausdorff distance for all $t \in I$ and $t' \in I'$. Having the *same initial endpoint* is defined similarly. Accordingly, γ and γ' are said to *have the same endpoints* if their initial and terminal endpoints agree.

By a *geodesic in \mathcal{X}* we always mean a *directed geodesic*, that is, a path $\gamma: I \rightarrow \mathcal{X}$ such that $d_{\mathcal{X}}(\gamma(s), \gamma(t)) = t - s$ for all $s < t$. Similarly a K -*quasigeodesic in \mathcal{X}* means a path $\gamma: I \rightarrow \mathcal{X}$ such that

$$\frac{1}{K}(t - s) - K \leq d_{\mathcal{X}}(\gamma(s), \gamma(t)) \leq K(t - s) + K$$

for all $s < t$. Note that a K -quasigeodesic typically will not be a K -quasigeodesic when traversed in reverse.

Convention 2.4 Our default metric on \mathcal{X} is the Lipschitz metric and geodesics are directed geodesics with respect to this metric. When discussing neighborhoods and Hausdorff distance, however, we make use of the symmetrized metric as discussed above.

2.7 Navigating Outer space

Optimal maps For any $G, H \in \mathcal{X}$, there exists a (nonunique) difference of markings $\phi: G \rightarrow H$ that realizes the infimum in the definition of $d_{\mathcal{X}}(G, H)$ [17; 3]. Such a map is called *optimal*. Here, we describe some structure of optimal maps and refer to the references above for details. Firstly, we say that a difference of markings $\phi: G \rightarrow H$ is *linear on edges* if ϕ has a constant slope $\sigma(e)$ on each edge e of G , meaning that ϕ is a local $\sigma(e)$ -homothety on e with respect to the local path metrics on G and H . In this case $\text{Lip}(\phi) = \max_e \{\sigma(e)\}$. We define the *tension subgraph* Δ_ϕ to be the subgraph of G consisting of maximally stretched edges, that is, the edges e of G with $\sigma(e) = \text{Lip}(\phi)$. Since every difference of markings is homotopic rel vertices to a map that is linear on edges and whose Lipschitz constant is no greater than the original, we may always suppose optimal maps are linear on edges.

Train tracks Let us define a *segment* $[p, q]$ between points $p, q \in G$ to be a locally isometric immersion $[0, L] \rightarrow G$ of an interval $[0, L] \subset \mathbb{R}$ sending $0 \mapsto p$ and $L \mapsto q$. A *direction* at $p \in G$ is a germ of nondegenerate segments $[p, q]$ with $p \neq q$. A map $\phi: G \rightarrow H$ that is linear on edges with slope $\sigma(e) \neq 0$ for all edges e of G then induces a derivative map D_ϕ which sends a direction at p to a direction at $\phi(p)$. We say that two directions at $p \in G$ are in the same *gate* if the directions are identified by D_ϕ . The gates form an equivalence relation on the set of directions in G .

An unordered pair $\{d, d'\}$ of distinct directions at a vertex v of G is called a *turn*. The turn $\{d, d'\}$ is *illegal* (with respect to ϕ) if d and d' belong to the same gate and is *legal* otherwise. Accordingly, the set of gates in G is also called the *illegal turn structure* on G induced by ϕ . An illegal turn structure is moreover a *train track structure* if there are at least two gates at each $p \in G$. This is equivalent to requiring that ϕ is locally injective on (the interior of) each edge of G and that every vertex has at least two gates.

For any $G, H \in \mathcal{X}$ there is an optimal map $\phi: G \rightarrow H$ such that Δ_ϕ is a core graph and the illegal turn structure induced by the restriction of ϕ to Δ_ϕ is a train track structure [17; 3]. Hence, the tension subgraph Δ_ϕ contains an immersed loop that is legal (crosses only legal turns). If α denotes the conjugacy class represented by a legal loop contained in Δ_ϕ , it follows that $\ell(\alpha|H) = \text{Lip}(\phi) \cdot \ell(\alpha|G)$. Conversely, any difference of markings $\phi: G \rightarrow H$ satisfying $\ell(\alpha|H) = \text{Lip}(\phi) \cdot \ell(\alpha|G)$ for some conjugacy class α is necessarily optimal. The existence of optimal maps thus shows that the inequality in (2) is in fact an equality. We collect these facts into the following proposition:

Proposition 2.5 (see Francaviglia and Martino [17] or Algom-Kfir [1]) *There exists, for every $G \in \mathcal{X}$, a finite set \mathcal{C}_G of primitive conjugacy classes, called **candidates**,*

whose immersed representatives in G cross each edge at most twice and such that for any $H \in \mathcal{X}$,

$$d_{\mathcal{X}}(G, H) = \max_{\alpha \in \mathcal{C}_G} \log \frac{\ell(\alpha|H)}{\ell(\alpha|G)} = \sup_{\alpha \in \mathbb{F}} \log \frac{\ell(\alpha|H)}{\ell(\alpha|G)}.$$

Folding For a linear difference of markings $\phi: G \rightarrow H$, if $\Delta_{\phi} = G$ and ϕ induces a train track structure on G , then ϕ induces a unique *folding path* $\gamma = \gamma^{\phi}: [0, L] \rightarrow \mathcal{X}$ with $\gamma^{\phi}(0) = G$ and $\gamma^{\phi}(L) = H$ such that $d_{\mathcal{X}}(\gamma^{\phi}(s), \gamma^{\phi}(t)) = t - s$ for $0 \leq s \leq t \leq L$. Thus γ^{ϕ} is a (directed) geodesic from G to H . The path γ^{ϕ} is obtained by “folding all illegal turns at unit speed”, as follows: Fix some sufficiently small $\epsilon > 0$. Then for $0 \leq s \leq \epsilon$, form the quotient graph \widehat{G}_s obtained by declaring two points $x, y \in G$ to be equivalent if $\phi(x) = \phi(y)$ and $d(x, v) = d(y, v) \leq s$ for some vertex v of G . Then ϕ factors through the quotient map $G \rightarrow \widehat{G}_s$, and \widehat{G}_s inherits a natural metric so that this quotient map is a local isometry on each edge of G . If we let $L_s = \text{vol}(\widehat{G}_s)$, then the rescaled graph $\overline{G}_s = (1/L_s)\widehat{G}_s$ lies in \mathcal{X} (ie has volume 1), and ϕ factors as a composition $G \rightarrow \overline{G}_s \rightarrow H$ of two optimal maps with Lipschitz constants L_s and L/L_s , respectively. Accordingly, we set $\gamma^{\phi}(\log(L_s)) = \overline{G}_s$. This defines $\gamma^{\phi}(t)$ for all sufficiently small $t \geq 0$. Staring now with the optimal map $\overline{G}_{\epsilon} \rightarrow H$, we may repeat this procedure to define $\gamma^{\phi}(t)$ for more values of t . While it is not obvious, after finitely many iterations we will eventually arrive at $\gamma^{\phi}(L) = H$. See [5, Proposition 2.2] for a justification of this claim and a more detailed construction of γ^{ϕ} .

Remark The folding paths used this paper, as defined above, are sometimes called “greedy folding paths” [5] or “fast folding lines” [17] in the literature. These are a special case of the more flexible “liberal folding paths” that are sometimes considered (see the appendix of [6]) and which include the original Stallings paths [32].

If $\gamma^{\phi}: [0, L] \rightarrow \mathcal{X}$ is a folding path, as described above, we often use $G_t, t \in [0, L]$, to denote $\gamma^{\phi}(t)$. Observe that for all $0 \leq s \leq t \leq L$, the construction of γ^{ϕ} provides induced optimal maps $\gamma_{st}^{\phi}: G_s \rightarrow G_t$, which we refer to as *folding maps*. These maps compose so that $\gamma_{rt}^{\phi} = \gamma_{st}^{\phi} \circ \gamma_{rs}^{\phi}$ for $r \leq s \leq t$, and we additionally have $\gamma_{0L}^{\phi} = \phi$ and $\gamma_{tt}^{\phi} = \text{Id}_{G_t}$ for all t . Furthermore, for all $t > s$, the maps $\gamma_{st}^{\phi}: G_s \rightarrow G_t$ (i) induce the same train track structure on G_s (independent of t), (ii) send legal segments (segments crossing only legal turns) to legal segments, and (iii) have associated folding paths exactly given by the restrictions $\gamma^{\phi}|_{[s,t]}$.

Lastly, we note that it is also possible to construct *bi-infinite* folding paths, by which we mean a directed geodesic $\gamma: \mathbb{R} \rightarrow \mathcal{X}$ together with maps $\gamma_{st}: G_s \rightarrow G_t$ (where $G_t = \gamma(t)$) for all $s \leq t$ satisfying the above properties.

Standard geodesics It is not true that any two points of $G, H \in \mathcal{X}$ may be connected by a folding path. There is, however, a nonunique *standard geodesic* from G to H [17]. In [5, Proposition 2.5], Bestvina and Feighn give a detailed construction of such a standard geodesic, which we summarize here: First, take an optimal map $\phi: G \rightarrow H$ that is linear on edges and consider the tension subgraph Δ_ϕ of G . Let $\Sigma_G \subset \mathcal{X}$ denote the simplex of all (volume-1) length functions on the marked graph G . By shortening some of the edges outside of Δ_ϕ (and rescaling to maintain volume 1), one may then find a point $G' \in \overline{\Sigma_G}$ in the closed simplex together with an optimal difference of markings $\phi': G' \rightarrow H$ whose tension graph $\Delta_{\phi'}$ is all of G' and such that

$$d_{\mathcal{X}}(G, H) = d_{\mathcal{X}}(G, G') + d_{\mathcal{X}}(G', H).$$

If γ_1 denotes the linear path in $\overline{\Sigma_G}$ from G to G' (which when parametrized by arclength is a directed geodesic) and $\gamma_2 = \gamma^{\phi'}$ denotes the folding path from G' to H induced by ϕ' , it follows from the equation above that the concatenation $\gamma_1\gamma_2$ is a directed geodesic from G to H .

Let us introduce the following terminology. By a *rescaling path* we mean a linear path $I \rightarrow \mathcal{X}$ in a closed simplex $\overline{\Sigma_G}$ parametrized by arclength. While such a path can in principle have infinite length in the negative direction (if the volume of a core subgraph tends to 0 as $t \rightarrow -\infty$), every rescaling path has finite length in the forward direction since a subgraph can only stretch until its volume is equal to 1. More specifically:

Lemma 2.6 *If $G_t, t \in [0, L]$, is a rescaling path with $G_0 \in \mathcal{X}_\epsilon$, then $L \leq \log(2/\epsilon)$.*

Proof Let α be any candidate for G_0 , so the immersed loop representing α in G_0 crosses each edge at most twice. Since G_0 and G_L represent the same marked graphs up to collapsing some edges of G_0 , the loop $\alpha|_{G_L}$ crosses no edge more than twice. Thus we have $\ell(\alpha|_{G_L}) \leq 2$. On the other hand $\ell(\alpha|_{G_0}) \geq \epsilon$ by assumption. Thus $\ell(\alpha|_{G_L})/\ell(\alpha|_{G_0}) \leq 2/\epsilon$. Since this holds for every candidate of G , Proposition 2.5 implies that $L = d_{\mathcal{X}}(G_0, G_L) \leq \log(2/\epsilon)$, as claimed. \square

In general, by a *standard geodesic* we mean a (directed) geodesic $\gamma: I \rightarrow \mathcal{X}$ that is either a folding path, a rescaling path, or a concatenation $\gamma: I \rightarrow \mathcal{X}$ of a rescaling path $\gamma: I^s \rightarrow \mathcal{X}$ and a folding path $\gamma: I^f \rightarrow \mathcal{X}$, where in the latter case we require that $I^s_+ = I^f_- \in \mathbb{R}$, that $I = I^s \cup I^f$, and that the concatenation is a directed geodesic. In this latter case the *folding image* of the standard geodesic is denoted by $\text{Im}^f(\gamma) = \gamma(I^f)$, and the *scaling image* is similarly denoted by $\text{Im}^{\text{sc}}(\gamma) = \gamma(I^s)$. For notational convenience, when the standard geodesic $\gamma: I \rightarrow \mathcal{X}$ is simply a rescaling path, we define $\text{Im}^{\text{sc}}(\gamma) = \gamma(I)$ and $\text{Im}^f(\gamma) = \gamma(I_+)$ (recall that $I_+ < \infty$ for rescaling paths); when γ is simply a folding path we define $\text{Im}^f(\gamma) = \gamma(I)$ and either $\text{Im}^{\text{sc}}(\gamma) = \gamma(I_-)$

or $\text{Im}^{\text{sc}}(\gamma) = \emptyset$ depending on whether $I_- \in \mathbb{R}$ or $I_- = -\infty$. In particular, note that the $\text{Im}^f(\gamma)$ is nonempty for every standard geodesic.

Folding and unfolding In Section 5 of [5], Bestvina and Feighn give a detailed account of what happens to an immersed path in the graph G_t under folding and unfolding. We review the basics here, as they will be needed in Section 6. For additional details and examples, see [5].

Fix a folding path $\gamma(t) = G_t$ with $t \in [a, b]$, and let p_b be an immersed path in G_b . It is always possible to lift (or unfold) p_b to an immersed path p_t in G_t with the property that p_t maps to a path in G_b whose immersed representative, rel endpoints, is p_b . (Recall that the folding path γ comes equipped with folding maps $\gamma_{tb}: G_t \rightarrow G_b$.) These lifts are not necessarily unique, but Bestvina and Feighn show that we can remove segments from the ends of p_b to obtain unique lifts. This is their unfolding principle, which we state as the following lemma:

Lemma 2.7 (unfolding principle [5]) *With the setup above, lifting p_b to G_t is unique between the first and last illegal turns of p_b , including the germs of directions beyond these turns.*

The process of lifting (uniquely) an immersed segment p_b whose endpoints are at illegal turns is called *unfolding*. Note that the unfolding principle applies to an illegal turn itself. In particular, if α is a conjugacy class of \mathbb{F} and p_b is either a subpath of $\alpha|G_t$ with endpoints illegal turns or an illegal turn of $\alpha|G_b$, then p_b unfolds to a path (or an illegal turn) p_t that is contained in $\alpha|G_t$. Moreover, multiple occurrences of p_b in $\alpha|G_b$ all unfold to p_t as a subpath of G_t . This all follows from the unfolding principle.

Similarly, we can understand the image of certain subpaths p_a of $\alpha|G_a$ under the folding path G_t . Note that the image of p_a in G_t is not necessarily contained in the image of $\alpha|G_t$, even after tightening (ie passing to the immersed representative). However, if there is a subpath p_b of $\alpha|G_b$ with endpoints at illegal turns which unfolds to p_a in $\alpha|G_a$, then unfolding gives a unique path p_t of G_t whose endpoints are at illegal turns of G_t . By the above paragraph, these unfolded paths have the property that p_t is a subsegment of $\alpha|G_t$ for all $t \in [a, b]$.

Projecting to standard geodesics In [5, Definition 6.3] Bestvina and Feighn define for any folding path $\gamma: I \rightarrow \mathcal{X}$ a *projection* $\text{Pr}_\gamma: \mathcal{X} \rightarrow \gamma(I)$ onto the image of the folding path. (One could alternately think of the projection as landing in the domain interval I .) As the definition of $\text{Pr}_\gamma(H)$ is rather technical — in short it involves looking at the infimum of times t for which a certain cover of $\gamma(t)$ contains an immersed legal segment of length 3 — we delay a careful discussion until Section 4.1 where a precise

construction of the projection $\text{Pr}_\gamma: \mathcal{X} \rightarrow \gamma(\mathbf{I})$ is given in Definition 4.9. However, although Pr_γ does coarsely agree with the closest-point projection to $\gamma(\mathbf{I})$ in special circumstances (see Lemma 4.11), we caution that Pr_γ is generally *unrelated* to the closest-point projection onto $\gamma(\mathbf{I})$.

Taking the existence of this projection for granted for the time being, we presently extend this construction in the natural way to any standard geodesic $\gamma: \mathbf{I} \rightarrow \mathcal{X}$ by declaring $\text{Pr}_\gamma := \text{Pr}_{\gamma^f}: \mathcal{X} \rightarrow \text{Im}^f(\gamma)$, where $\mathbf{I} = \mathbf{I}^s \cup \mathbf{I}^f$ and $\gamma^f = \gamma|_{\mathbf{I}^f}$ is the folding portion of γ . (Recall that $\text{Im}^f(\gamma) \neq \emptyset$ for every standard geodesic γ .)

2.8 The free factor complex

The (free) factor complex \mathcal{F} of \mathbb{F} is the simplicial complex whose vertices are conjugacy classes of nontrivial, proper free factors of \mathbb{F} . A collection of vertices $\{[A_0], \dots, [A_k]\}$ determines a k -simplex if, after reordering and choosing conjugacy representatives, we have $A_0 < \dots < A_k$. The free factor complex was first introduced by Hatcher and Vogtmann in [21]. When it should cause no confusion to do so, we will usually drop the conjugacy symbol from the notation and denote a conjugacy class of free factors by $A \in \mathcal{F}^0$.

We equip the factor complex \mathcal{F} with its simplicial path metric. That is, we geometrically view \mathcal{F} as the simplicial graph \mathcal{F}^1 equipped with the path metric in which each edge has length 1. For our purposes, the significance of the factor complex stems from the following foundational result of Bestvina and Feighn:

Theorem 2.8 (Bestvina and Feighn [5]) *The factor complex \mathcal{F} is Gromov-hyperbolic.*

There is a natural (coarse) projection $\pi: \mathcal{X} \rightarrow \mathcal{F}$ defined by sending $G \in \mathcal{X}$ to the set of free factors corresponding to proper subgraphs of G . That is,

$$\pi(G) = \{\pi_1(G') : G' \text{ is a proper, connected, noncontractible subgraph of } G\} \subset \mathcal{F}^0,$$

where $\pi_1(G') \leq \pi_1(G)$ is identified with a free factor of \mathbb{F} via the marking $\mathcal{R} \rightarrow G$. This projection is a key tool in the proof of Theorem 2.8 above. For $G \in \mathcal{X}$, it is quickly verified that $\text{diam}_{\mathcal{F}}(\pi(G)) \leq 4$ [5, Lemma 3.1].

Let us define the *factor distance* between two points $G, H \in \mathcal{X}$ to be

$$d_{\mathcal{F}}(G, H) = \text{diam}_{\mathcal{F}}(\pi(G) \cup \pi(H)).$$

Corollary 3.5 of [5] shows that $d_{\mathcal{F}}(G, H) \leq 12e^{d_{\mathcal{X}}(G, H)} + 32$. In fact, as indicated in [5], this may easily be strengthened to show that π is coarsely 80-Lipschitz:

Lemma 2.9 *For any $G, H \in \mathcal{X}$ we have $d_{\mathcal{F}}(G, H) \leq 80 d_{\mathcal{X}}(G, H) + 80$.*

Proof Let $L = d_{\mathcal{X}}(G, H)$, and let $\gamma: [0, L] \rightarrow \mathcal{X}$ be a standard geodesic from G to H . If $k = \lceil L \rceil$, then we may find times $0 = t_0 < \dots < t_k = L$ so that $d_{\mathcal{X}}(\gamma(t_i), \gamma(t_{i+1})) \leq 1$ for all $0 \leq i < k$. By Corollary 3.5 of [5], we have $d_{\mathcal{F}}(\gamma(t_i), \gamma(t_{i+1})) \leq 12e + 32$ for each i , and thus that $d_{\mathcal{F}}(G, H) \leq (12e + 32)k \leq 80L + 80$ by the triangle inequality. \square

Similarly, we will use the following easy lemma:

Lemma 2.10 *Suppose that there is a nontrivial conjugacy class α which has length less than 1 on both $G, H \in \mathcal{X}$. Then $d_{\mathcal{F}}(G, H) \leq 10$.*

Proof Such a conjugacy class would determine an immersed loop contained in a proper core subgraph of each graph. Hence, α is simultaneously contained in free factors A and B appearing in the diameter-4 projections of G and H , respectively. In this case, $d_{\mathcal{F}}(A, B) \leq 2$ [33, Section 3.2], showing that the union $\pi(G) \cup \pi(H)$ has diameter at most 10. \square

In the process of showing that \mathcal{F} is hyperbolic, Bestvina and Feighn also prove the following very useful result; it essentially says that the projection onto a folding path is strongly contracting when viewed from the factor complex.

Proposition 2.11 (Bestvina and Feighn [5, Proposition 7.2]) *There exists a universal constant B (depending only on $\text{rk}(\mathbb{F})$) such that the following holds. If $H, H' \in \mathcal{X}$ satisfy $d_{\mathcal{X}}(H, H') \leq M$ and $\gamma: I \rightarrow \mathcal{X}$ is a standard geodesic with $d_{\mathcal{X}}(H, \gamma(t)) \geq M$ for all t , then $d_{\mathcal{F}}(\text{Pr}_{\gamma}(H), \text{Pr}_{\gamma}(H')) \leq B$.*

Remark While Proposition 7.2 of [5] is only stated for the projection to a finite-length folding path, it clearly holds for our generalized projection to a finite-length standard geodesic. By considering an exhaustion by finite-length subpaths, the result is also seen to hold for infinite-length standard geodesics.

2.9 Out(\mathbb{F}) basics

We recall some of the structure of automorphisms of \mathbb{F} and the dynamics of their actions on \mathcal{X} and \mathcal{F} . The group $\text{Out}(\mathbb{F})$ acts naturally on \mathcal{X} by changing the marking: $\phi \cdot (G, g, \ell) = (G, g \circ \hat{\phi}^{-1}, \ell)$, where $\hat{\phi}^{-1}: \mathcal{R} \rightarrow \mathcal{R}$ is any homotopy equivalence whose induced map on $\mathbb{F} \cong \pi_1(\mathcal{R})$ is in the outer automorphism class $\phi^{-1} \in \text{Out}(\mathbb{F})$. One may easily verify that $G \mapsto \phi \cdot G$ defines an isometry of $(\mathcal{X}, d_{\mathcal{X}})$. Each outer automorphism $\phi \in \text{Out}(\mathbb{F})$ permutes the set \mathcal{F}^0 of conjugacy classes of free factors via $\phi \cdot [A] = [\phi(A)]$, and this extends to a simplicial (and hence isometric) action of $\text{Out}(\mathbb{F})$ on \mathcal{F} . The actions of $\text{Out}(\mathbb{F})$ on \mathcal{X} and \mathcal{F} are equivariant with respect to the projection $\pi: \mathcal{X} \rightarrow \mathcal{F}$: for each $G \in \mathcal{X}$, we have, as subsets of \mathcal{F} ,

$$\pi(\phi \cdot G) = \phi \cdot \pi(G).$$

Full irreducibility We are primarily interested in elements $\phi \in \text{Out}(\mathbb{F})$ that are *fully irreducible*, meaning that no positive power of ϕ fixes the conjugacy class of any free factor of \mathbb{F} . Hence ϕ is fully irreducible if and only if its action on \mathcal{F} has no periodic vertices. In fact, Bestvina and Feighn have shown the following:

Theorem 2.12 (Bestvina and Feighn [5]) *An element $\phi \in \text{Out}(\mathbb{F})$ acts with positive translation length on the free factor complex \mathcal{F} if and only if ϕ is fully irreducible.*

Recall that the (stable) translation length of $\phi \in \text{Out}(\mathbb{F})$ acting on \mathcal{F} is by definition

$$\ell_{\mathcal{F}}(\phi) = \lim_{n \rightarrow \infty} \frac{1}{n} d_{\mathcal{F}}(A, \phi^n A),$$

for any $A \in \mathcal{F}^0$. It is well known (and easily verified) that $\ell_{\mathcal{F}}(\phi)$ does not depend on the choice of A and that $\ell_{\mathcal{F}}(\phi^n) = n \cdot \ell_{\mathcal{F}}(\phi)$. Having positive translation length implies that for any $A \in \mathcal{F}^0$, the orbit map $\mathbb{Z} \rightarrow \mathcal{F}$ defined by $n \mapsto \phi^n \cdot A$ is a quasigeodesic in \mathcal{F} . In Section 9 we also discuss translation lengths of elements of $\text{Out}(\mathbb{F})$ acting on a different hyperbolic complex. Regardless of the context, we call an isometry of a hyperbolic space *loxodromic* if it acts with positive translation length.

Hyperbolicity An element $\phi \in \text{Out}(\mathbb{F})$ is said to be *hyperbolic* or *atoroidal* if $\phi^i(\alpha) \neq \alpha$ for every nontrivial conjugacy class α in \mathbb{F} and every $i \geq 1$. While neither hyperbolicity nor full irreducibility implies the other, there are many automorphisms of \mathbb{F} that have both these properties. Hyperbolic elements of $\text{Out}(\mathbb{F})$ are essential to our discussion because of the following theorem of Brinkmann.

Theorem 2.13 (Brinkmann [11]) *The outer automorphism class of $\Phi \in \text{Aut}(\mathbb{F})$ is hyperbolic if and only if the semidirect product $\mathbb{F} \rtimes_{\Phi} \mathbb{Z}$ is a Gromov-hyperbolic group.*

We say that $\Gamma \leq \text{Out}(\mathbb{F})$ is *purely hyperbolic* if every infinite-order element of Γ is hyperbolic. Before concluding this section, we observe that when Γ is purely hyperbolic there is a uniform upper bound (depending only on $\text{rk}(\mathbb{F})$) on the number of elements of Γ that fix any given conjugacy class. To this end, for α a conjugacy class in \mathbb{F} set

$$\Gamma_{\alpha} = \{\phi \in \Gamma : \phi(\alpha) = \alpha\}.$$

Lemma 2.14 *There is a constant e_r depending only on the rank $r = \text{rk}(\mathbb{F})$ such that for any purely hyperbolic $\Gamma \leq \text{Out}(\mathbb{F})$ we have $|\Gamma_{\alpha}| \leq e_r$ for each nontrivial conjugacy class α of \mathbb{F} .*

Proof Since Γ is purely hyperbolic, Γ_{α} is a torsion subgroup of $\text{Out}(\mathbb{F})$. It is known that any torsion element survives in the quotient $\text{Out}(\mathbb{F}) \rightarrow \text{GL}_r(\mathbb{Z}/3\mathbb{Z})$ [14] and so Γ_{α} injects into $\text{GL}_r(\mathbb{Z}/3\mathbb{Z})$. Hence, we may take $e_r = |\text{GL}_r(\mathbb{Z}/3\mathbb{Z})|$. □

3 Quasiconvexity and folding paths

For the main results of Section 4 we will need to know that outgoing balls in the Lipschitz metric are quasiconvex with respect to folding paths. This is proven in Corollary 3.3. We first show in Proposition 3.2 that the length of every conjugacy class is quasiconvex along folding paths.

We begin by recalling some notation from [5]. For a folding path G_t , $t \in I$, define the *illegality* $m(G_{t_0})$ of G_{t_0} at time t_0 to be

$$m(G_{t_0}) = \sum_v \sum_{\Omega_v} (|\Omega_v| - 1),$$

where v varies over the vertices of G_t and Ω_v varies over all gates of G_{t_0} at the vertex v (so each Ω_v is an equivalence class of directions at v). Note that if we set $M = 6 \operatorname{rk}(\mathbb{F}) - 6$, which bounds twice the number of edges of any graph in \mathcal{X} , then $1 \leq m(G_t) \leq M$ for all t . We often write m_t for $m(G_t)$ when the folding path is understood. For any conjugacy class α , we additionally let $k_t = k(\alpha|G_t)$ denote the number of illegal turns in $\alpha|G_t$.

In Corollary 4.5 and Lemma 4.4 of [5], Bestvina and Feighn show that the function $t \mapsto \ell(\alpha|G_t)$ is piecewise exponential and that its right derivative at time t_0 is given by

$$\ell(\alpha|G_{t_0}) - 2 \frac{k(\alpha|G_{t_0})}{m(G_{t_0})}.$$

Using this, they prove the following estimate:

Lemma 3.1 (Bestvina and Feighn [5, Lemma 4.10]) *Suppose that G_t , $t \in [0, L]$, is a folding path and that α is any conjugacy class in \mathbb{F} . Then for all $t \in [0, L]$ we have*

$$\ell(\alpha|G_t) \leq \max\{2k(\alpha|G_0), \ell(\alpha|G_L)\}.$$

Proposition 3.2 (quasiconvexity of lengths along folding paths) *Let G_t , $t \in [0, L]$, be a folding path, and let α be any conjugacy class in \mathbb{F} . Then*

$$\ell(\alpha|G_t) \leq 6 \operatorname{rk}(\mathbb{F}) \cdot \max\{\ell(\alpha|G_0), \ell(\alpha|G_L)\}.$$

Proof Let $l \in [0, L]$ be the supremum of times for which the piecewise exponential function $t \mapsto \ell(\alpha|G_t)$ is decreasing on $[0, l]$. Hence the right derivative of $\ell(\alpha|G_t)$ at time l is nonnegative. If $l = L$, then we are done. Otherwise, by the derivative formula above we have $\ell(\alpha|G_l) \geq 2(k_l/m_l) \geq 2(k_l/6r)$, where $r = \operatorname{rk}(\mathbb{F})$. Hence

$$k_l \leq 3r \cdot \ell(\alpha|G_l) \leq 3r \cdot \ell(\alpha|G_0)$$

by the choice of l . Applying Lemma 3.1, we see that for all $t \in [l, L]$,

$$\begin{aligned} \ell(\alpha|G_t) &\leq \max\{2k_l, \ell(\alpha|G_L)\} \\ &\leq \max\{6r \cdot \ell(\alpha|G_0), \ell(\alpha|G_L)\}. \end{aligned}$$

Since $\ell(\alpha|G_t) \leq \ell(\alpha|G_0)$ for all $t \in [0, l]$, this completes the proof. □

Corollary 3.3 (outgoing balls are folding-path-quasiconvex) *There exists a universal constant A (depending only on $\text{rk}(\mathbb{F})$) such that the following holds. For any $H \in \mathcal{X}$ and $R > 0$, if $\gamma: [0, L] \rightarrow \mathcal{X}$ is a folding path $\gamma(t) = G_t$ with $d_{\mathcal{X}}(H, G_0) \leq R$ and $d_{\mathcal{X}}(H, G_L) \leq R$, then for all $t \in [0, L]$ we have*

$$d_{\mathcal{X}}(H, G_t) \leq R + A.$$

Proof Applying Proposition 3.2, for any time $t \in [0, L]$ we have

$$\begin{aligned} d_{\mathcal{X}}(H, G_t) &= \sup_{c \in \mathbb{F}} \log \frac{\ell(c|G_t)}{\ell(c|H)} \\ &\leq \sup_{c \in \mathbb{F}} \log \left(6 \text{rk}(\mathbb{F}) \cdot \frac{\max\{\ell(c|G_0), \ell(c|G_L)\}}{\ell(c|H)} \right) \\ &\leq \log(6 \text{rk}(\mathbb{F})) + \max \left\{ \sup_{c \in \mathbb{F}} \log \frac{\ell(c|G_0)}{\ell(c|H)}, \sup_{c \in \mathbb{F}} \log \frac{\ell(c|G_L)}{\ell(c|H)} \right\} \\ &\leq \log(6 \text{rk}(\mathbb{F})) + R. \end{aligned} \quad \square$$

4 Stability for \mathcal{F} -progressing quasigeodesics

In this section we explore the structure of quasigeodesics in Outer space that project to parametrized quasigeodesics in the factor complex. We show that, as in a hyperbolic space, such quasigeodesics are stable in the sense that they fellow-travel any geodesic with the same endpoints. More specifically, we prove the following.

Theorem 4.1 (\mathcal{F} -progressing quasigeodesics are stable) *Let $\gamma: I \rightarrow \mathcal{X}$ be a K -quasigeodesic whose projection $\pi \circ \gamma: I \rightarrow \mathcal{F}$ is also a K -quasigeodesic. Then there exist constants $A, \epsilon > 0$ and $K' \geq 1$ depending only on K (and the injectivity radius of the terminal endpoint $\gamma(I_+)$ when $I_+ < \infty$) with the following property: if $\rho: J \rightarrow \mathcal{X}$ is any geodesic with the same endpoints as γ , then*

- (i) $\gamma(I), \rho(J) \subset \mathcal{X}_\epsilon$,
- (ii) $d_{\text{Haus}}(\gamma(I), \rho(J)) < A$, and
- (iii) $\pi \circ \rho: J \rightarrow \mathcal{F}$ is a (parametrized) K' -quasigeodesic.

Our proof relies crucially on the projection $\text{Pr}_\gamma: \mathcal{X} \rightarrow \gamma(\mathbf{I})$ from Outer space to the image of any standard geodesic $\gamma: \mathbf{I} \rightarrow \mathcal{X}$. As recorded in Proposition 2.11, Bestvina and Feighn prove that this projection is strongly contracting when viewed in the factor complex, and they use this to show that \mathcal{F} is δ -hyperbolic [5]. The projection $\pi \circ \gamma$ of γ to the factor complex is also shown to be an unparametrized K_f -quasigeodesic, where K_f depends only on $\text{rk}(\mathbb{F})$ [5]. As a quasigeodesic, the nearest-point retraction $n_{\pi \circ \gamma}: \mathcal{F} \rightarrow \pi(\gamma(\mathbf{I}))$ onto the image $\pi(\gamma(\mathbf{I}))$ is coarsely L_0 -Lipschitz for some L_0 that depends only on δ and K_f . The next lemma verifies that $n_{\pi \circ \gamma}: \mathcal{F} \rightarrow \pi(\gamma(\mathbf{I}))$ agrees with the Bestvina–Feighn projection $\pi \circ \text{Pr}_\gamma: \mathcal{F} \rightarrow \pi(\gamma(\mathbf{I}))$ up to uniformly bounded error.

Lemma 4.2 *There is a constant $D_1 \geq 0$, depending only on $\text{rk}(\mathbb{F})$, such that for any $H \in \mathcal{X}$ and any standard geodesic $\gamma: \mathbf{I} \rightarrow \mathcal{X}$ we have*

$$d_{\mathcal{F}}(\pi(\text{Pr}_\gamma(H)), n_{\pi \circ \gamma}(\pi(H))) \leq D_1.$$

Proof To simplify notation, set $\hat{C} = \pi(\text{Pr}_\gamma(H))$, $C = \pi(H)$ and $A' = n_{\pi \circ \gamma}(C)$; both of these points lie on the unparametrized K_f -quasigeodesic $\pi(\gamma(\mathbf{I}))$. Now let $\hat{\rho}$ and ρ' be folding paths whose images in \mathcal{F} join C to \hat{C} and A' , respectively. We are now in the situation of [5, Proposition 9.1], which states that there is a Q' on $\pi(\rho')$ whose distance from \hat{C} is no greater than B_1 , where B_1 is a uniform constant.

Since $\pi(\rho')$ is an unparametrized K_f -quasigeodesic, any geodesic $[C, A']$ in \mathcal{F} joining C and A' contains a point Q_0 with $d_{\mathcal{F}}(Q', Q_0) \leq R_0$, where $R_0 = R_0(\delta, K_f)$ is the constant from Proposition 2.1. Hence, $d_{\mathcal{F}}(Q_0, \hat{C}) \leq B_1 + R_0$. Since no factor on $\pi(\gamma)$ is closer to C than A' , we must have $d_{\mathcal{F}}(Q_0, A') \leq B_1 + R_0$. Hence, we conclude that $d_{\mathcal{F}}(A', \hat{C}) \leq 2(B_1 + R_0)$. Thus the lemma holds with $D_1 = 2(B_1 + R_0)$. \square

The proof of Theorem 4.1 will take the rest of the section and require several lemmas. In fact, we first prove the theorem in the special case that ρ is a standard geodesic (Proposition 4.7) and complete the general proof in Section 4.1. We note that only the special case is needed for the proof of our main result.

First, we observe that quasigeodesics that make definite progress in the factor graph cannot become arbitrarily thin.

Lemma 4.3 *Let $\gamma: \mathbf{I} \rightarrow \mathcal{X}$ be a K -quasigeodesic whose projection $\pi \circ \gamma: \mathbf{I} \rightarrow \mathcal{F}$ is also a K -quasigeodesic. Then there is an $\epsilon > 0$ depending only on K such that $\gamma(i) \in \mathcal{X}_\epsilon$ for all $i \in \mathbf{I}$ with $i + K(K + 11) \in \mathbf{I}$.*

Furthermore, for any $i \in \mathbf{I}$ with $i + K(K + 11) \notin \mathbf{I}$ (so that necessarily $\mathbf{I}_+ < \infty$), we have $\gamma(i) \in \mathcal{X}_{\epsilon'}$ for some $\epsilon' > 0$ depending only on K and the injectivity radius of $\gamma(\mathbf{I}_+)$.

Proof Since γ is a K -quasigeodesic in \mathcal{F} , we have $d_{\mathcal{F}}(\gamma(i), \gamma(j)) \geq (1/K)|j-i| - K$. If $b \geq 0$ is chosen to equal $K(K + 11)$, then $\gamma(i)$ and $\gamma(i + b)$ have distance in the factor complex at least 11. By Lemma 2.10, this implies that there is no nontrivial conjugacy class having length less than 1 in each of $\gamma(i), \gamma(i + b) \in \mathcal{X}$. If there exists $\alpha \in \mathbb{F}$ with $\ell(\alpha|\gamma(i)) = \epsilon \leq 1$, our choice of b thus forces $\ell(\alpha|\gamma(i + b)) \geq 1$. Hence we find that

$$Kb + K \geq d_{\mathcal{X}}(\gamma(i), \gamma(i + b)) \geq \log \frac{\ell(\alpha|\gamma(i + b))}{\ell(\alpha|\gamma(i))} \geq \log \frac{1}{\epsilon}.$$

This ensures $\epsilon \geq e^{-(Kb+K)}$, and so we conclude $\gamma(i) \in \mathcal{X}_{e^{-(Kb+K)}}$ for all $i \in I$ with $i + b \in I$.

Finally suppose $I_+ < \infty$ and that $\gamma(I_+) \in \mathcal{X}_{\epsilon_0}$. If $i \in I$ fails to satisfy $i + b \in I$, then for any nontrivial $\alpha \in \mathbb{F}$ we similarly have

$$Kb + K \geq d_{\mathcal{X}}(\gamma(i), \gamma(I_+)) \geq \log \frac{\ell(\alpha|\gamma(I_+))}{\ell(\alpha|\gamma(i))} \geq \log \frac{\epsilon_0}{\ell(\alpha|\gamma(i))}.$$

Thus $\ell(\alpha|\gamma(i)) \geq \epsilon_0 e^{-(Kb+K)}$ for every nontrivial $\alpha \in \mathbb{F}$, which proves the claim. \square

Proposition 4.4 *Let $\gamma: I \rightarrow \mathcal{X}$ be a K -quasigeodesic whose projection $\pi \circ \gamma: I \rightarrow \mathcal{F}$ is also a K -quasigeodesic, and let $\rho: J \rightarrow \mathcal{X}$ be a standard geodesic with the same endpoints as γ . Then there exists a constant $D_0 \geq 0$ depending only on K such that:*

- (i) *For each $i \in I$ there is a $t_i \in J$ such that $d_{\mathcal{X}}(\gamma(i), \rho(t_i)) \leq D_0$.*
- (ii) *For each $j \in J$ there is a $s_j \in I$ such that $d_{\mathcal{X}}(\gamma(s_j), \rho(j)) \leq D_0$.*

Proof We first prove (i). Let B be the universal constant from Proposition 2.11, and let $\tau' = R_0(\delta, \max\{K, K_f\})$ be the fellow-traveling constant (Proposition 2.1) for $\max\{K, K_f\}$ -quasigeodesics in a δ -hyperbolic space (recall that \mathcal{F} is δ -hyperbolic). Set $\tau = \tau' + D_1$, where D_1 is the constant appearing in Lemma 4.2. Define

$$M = 2K^2B + K, \quad C_0 = \frac{M-K}{K} = 2KB \quad \text{and} \quad L_0 = 2K(B + 2\tau + K).$$

Suppose that $[a, b] \subset I$ is a subinterval such that $d_{\mathcal{X}}(\gamma(t), \text{Im}^f(\rho)) \geq M$ for all $t \in [a, b]$. Setting $n = \lceil (b-a)/C_0 \rceil$, we then have the sequence of points $q_j = \gamma(a + ((b-a)/n)j)$ for $j = 0, \dots, n$. Notice that these points enjoy

$$d_{\mathcal{X}}(q_j, q_{j+1}) \leq K \left(\frac{b-a}{n} \right) + K \leq K(C_0) + K = M$$

for each $j = 0, \dots, n - 1$. Since $d_{\mathcal{X}}(q_j, \text{Im}^f(\rho)) \geq M$ by assumption, Proposition 2.11 now implies that

$$d_{\mathcal{F}}(\text{Pr}_{\rho}(q_j), \text{Pr}_{\rho}(q_{j+1})) \leq B$$

for all $j = 0, \dots, n - 1$.

Recall that by Lemma 4.2, $\pi(\text{Pr}_\rho(q_j)) \in \pi(\text{Im}^f(\rho))$ is within distance D_1 from the closest-point projection of $\pi(q_j) \in \mathcal{F}$ to the path $\pi(\rho(\mathbf{J}))$. (Note that $\pi(\rho(\mathbf{J})) = \pi(\text{Im}^f(\rho))$ since rescaling paths have constant projection in \mathcal{F} by construction.) Since $\pi \circ \gamma$ and $\pi \circ \rho$ are both (unparametrized) $\max\{K, K_f\}$ -quasigeodesics and \mathcal{F} is δ -hyperbolic, these paths have Hausdorff distance at most τ' in \mathcal{F} . It follows that for each $j = 0, \dots, n$ we also have $d_{\mathcal{F}}(\pi(q_j), \pi(\text{Pr}_\rho(q_j))) \leq \tau' + D_1 = \tau$. By the triangle inequality, we now have

$$\begin{aligned} d_{\mathcal{F}}(\gamma(a), \gamma(b)) &\leq d_{\mathcal{F}}(q_0, \text{Pr}_\rho(q_0)) + \sum_{j=0}^{n-1} d_{\mathcal{F}}(\text{Pr}_\rho(q_j), \text{Pr}_\rho(q_{j+1})) + d_{\mathcal{F}}(\text{Pr}_\rho(q_n), q_n) \\ &\leq nB + 2\tau \\ &\leq \left(\frac{b-a}{C_0} + 1\right)B + 2\tau \\ &= \frac{b-a}{2K} + B + 2\tau. \end{aligned}$$

On the other hand, by hypothesis we also have $d_{\mathcal{F}}(\gamma(a), \gamma(b)) \geq (b-a)/K - K$. Combining these, we find that

$$b - a \leq 2K(B + 2\tau + K) = L_0.$$

That is, L_0 is an upper bound for the length of any subinterval of \mathbf{I} on which γ stays at least distance M from $\text{Im}^f(\rho)$. Said differently, for any $t \in \mathbf{I}$, there exists $0 \leq t' \leq L_0$ such that $d_{\mathcal{X}}(\gamma(t+t'), \text{Im}^f(\rho)) < M$. (When $t+L_0 \in \mathbf{I}$, this is clear. When $t+L_0 \notin \mathbf{I}$, then we necessarily have $\mathbf{I}_+ < \infty$ and the assumption that γ and ρ have the same ends ensures $\gamma(\mathbf{I}_+) \in \text{Im}^f(\rho)$.) In particular, we conclude that

$$d_{\mathcal{X}}(\gamma(t), \text{Im}^f(\rho)) \leq d_{\mathcal{X}}(\gamma(t), \gamma(t+t')) + d_{\mathcal{X}}(\gamma(t+t'), \text{Im}^f(\rho)) \leq KL_0 + K + M.$$

This proves (i) with $D_0 = KL_0 + K + M$.

We now prove (ii). Let E_0 denote the maximum value of $D_0 = KL_0 + K + M$ and of the quasiconvexity constant A provided by Corollary 3.3. Note that $E_0 \geq K$. For each point $i \in \mathbf{I}$, let

$$U_i = \{y \in \text{Im}^f(\rho) : d_{\mathcal{X}}(\gamma(i), y) \leq 4E_0\}.$$

By the proof of (i), we know that there exists a point $y_i \in U_i$ with $d_{\mathcal{X}}(\gamma(i), y_i) \leq 2E_0$; in particular U_i contains the length- $2E_0$ subinterval of $\text{Im}^f(\rho)$ starting at y_i . Let $W_i \subset \text{Im}^f(\rho)$ denote the smallest connected interval containing U_i . It follows that each interval W_i with $\rho(\mathbf{J}_+) \notin W_i$ has length at least $2E_0$. By Corollary 3.3 we additionally know that $d_{\mathcal{X}}(\gamma(i), w) \leq 4E_0 + A \leq 5E_0$ for all $w \in W_i$.

Using that the projection $\pi: \mathcal{X} \rightarrow \mathcal{F}$ is coarsely 80-Lipschitz (Lemma 2.9) we see that $\text{diam}_{\mathcal{F}} \pi(\{\gamma(i)\} \cup W_i) \leq 80(10E_0)$. In particular, if $i, j \in \mathbf{I}$ satisfy the

inequality $|i - j| \geq 2 \cdot 80(10KE_0)$, then $d_{\mathcal{F}}(\gamma(i), \gamma(j)) \geq 2 \cdot 80(10E_0)$, ensuring that $\pi(W_i)$ and $\pi(W_j)$ are disjoint. In particular, this implies $W_i \cap W_j = \emptyset$. On the other hand, if $i, j \in \mathbf{I}$ satisfy $i < j$ and $(j - i) \leq 1$, then

$$d_{\mathcal{X}}(\gamma(i), y_j) \leq d_{\mathcal{X}}(\gamma(i), \gamma(j)) + d_{\mathcal{X}}(\gamma(j), y_j) \leq K(1) + K + 2E_0 \leq 4E_0,$$

showing that $y_j \in U_i$ by definition. Thus W_i and W_j intersect whenever $|i - j| \leq 1$. This implies that the union

$$W = \bigcup_{i \in \mathbf{I}} W_i$$

is a connected subinterval of $\text{Im}^f(\rho)$. We claim that in fact $W = \text{Im}^f(\rho)$.

To see this, first suppose that $\mathbf{I}_+ < \infty$, in which case we also have $\mathbf{J}_+ < \infty$ and $\gamma(\mathbf{I}_+) = \rho(\mathbf{J}_+) \in \text{Im}^f(\rho)$ by assumption. In particular, $\rho(\mathbf{J}_+) \in W_{\mathbf{I}_+} \subset W$ by definition. If instead $\mathbf{I}_+ = \mathbf{J}_+ = \infty$, then the above shows that for any $t \in \mathbf{J}$ we can find infinitely many disjoint intervals $W_i \subset \rho([t, \infty))$ that each have length at least $2E_0$. Thus $W \cap \rho([t, \infty))$ is an infinite-length interval and so covers the positive end of $\text{Im}^f(\rho)$.

Now suppose $\mathbf{I}_- = \mathbf{J}_- = -\infty$. In this case, we claim ρ cannot have an initial rescaling segment (ie that $\text{Im}^{\text{sc}}(\rho) = \emptyset$ and consequently that $\rho(\mathbf{J}) = \text{Im}^f(\rho)$). Indeed, if $\text{Im}^{\text{sc}}(\rho)$ were nonempty then it must have infinite length in the negative direction. Since it is a rescaling path, this implies $\text{Im}^{\text{sc}}(\rho)$ contains arbitrarily thin points (Lemma 2.6). However this contradicts the fact that $\gamma(\mathbf{I})$ is contained in some thick part \mathcal{X}_ϵ (by Lemma 4.3) and that the initial rays of γ and ρ have finite Hausdorff distance. Therefore, $\text{Im}^f(\rho)$ has infinite length in the negative direction and the same argument as above shows that $W \cap \rho((-\infty, t])$ has infinite length for any $t \in \mathbf{J}$. Whence $W = \text{Im}^f(\rho)$ as claimed.

Finally suppose $\mathbf{I}_- \neq -\infty$. Let $t \in \mathbf{J}$ be such that $\rho(t) = y_{\mathbf{I}_-} \in U_{\mathbf{I}_-} \subset \text{Im}^f(\rho)$. Then $d_{\mathcal{X}}(\gamma(\mathbf{I}_-), \rho(t)) \leq 4E_0$ by definition, and since ρ is a geodesic, it follows that

$$d_{\mathcal{X}}(\gamma(\mathbf{I}_-), \rho(s)) = d_{\mathcal{X}}(\rho(\mathbf{J}_-), \rho(s)) \leq 4E_0$$

for all $s \in [\mathbf{J}_-, t]$. In particular, $U_{\mathbf{I}_-} \subset W$ contains the left endpoint of $\text{Im}^f(\rho)$, which proves the desired equality $W = \text{Im}^f(\rho)$. Moreover, the above equation shows that any point $y \in \text{Im}^{\text{sc}}(\rho)$ satisfies $d_{\mathcal{X}}(\gamma(\mathbf{I}_-), y) \leq 4E_0$. Therefore we conclude that for every $s \in \mathbf{J}$ the point $\rho(s) \in \text{Im}^{\text{sc}}(\rho) \cup W$ satisfies $d_{\mathcal{X}}(\gamma(\mathbf{I}), \rho(s)) \leq 5E_0$. Hence (ii) holds with $D_0 = 5E_0$. □

Lemma 4.5 (thinness prevents factor progress) *Suppose that $\gamma: [0, L] \rightarrow \mathcal{X}$ is a finite-length geodesic and that $\gamma(t)$ is ϵ -thin for all $t \in [0, L]$. Then*

$$d_{\mathcal{X}}(\gamma(0), \gamma(L)) \geq \log\left(\frac{1}{\epsilon}\right) \frac{d_{\mathcal{F}}(\gamma(0), \gamma(L)) - 20}{20}.$$

Proof We may suppose $N = d_{\mathcal{F}}(\gamma(0), \gamma(L)) > 11$, for otherwise there is nothing to prove. Set $a_0 = 0$. Supposing by induction that $a_i \in [0, L)$ has been defined for some $i \geq 0$, we then set

$$a_{i+1} = \sup\{t \in [a_i, L] : d_{\mathcal{F}}(\gamma(a_i), \gamma(t)) \leq 15\}.$$

In this way, we obtain a sequence of times $0 = a_0 < \dots < a_n = L$. Notice that provided $a_{i+1} < L$, we necessarily have $d_{\mathcal{F}}(\gamma(a_i), \gamma(a_{i+1} + \delta)) \geq 16$ for all $\delta > 0$. Furthermore, for all sufficiently small δ , the graphs $\gamma(a_{i+1})$ and $\gamma(a_{i+1} + \delta)$ necessarily have embedded loops representing the same conjugacy class, and so the projections $\pi(\gamma(a_{i+1}))$ and $\pi(\gamma(a_{i+1} + \delta))$ must overlap. Therefore the union of $\pi(\gamma(a_i))$ and $\pi(\gamma(a_{i+1}))$ has diameter at least 12. By Lemma 2.10, this implies that there is no nontrivial conjugacy class with length less than 1 in both graphs $\gamma(a_i)$ and $\gamma(a_{i+1})$. Since by assumption $\ell(\beta|\gamma(a_i)) < \epsilon$ for some nontrivial $\beta \in \mathbb{F}$, it follows that $\ell(\beta|\gamma(a_{i+1})) \geq 1$ and thus that

$$d_X(\gamma(a_i), \gamma(a_{i+1})) \geq \log \frac{\ell(\beta|\gamma(a_{i+1}))}{\ell(\beta|\gamma(a_i))} > \log \frac{1}{\epsilon}.$$

Therefore, since γ is a geodesic, we find that

$$(3) \quad d_X(\gamma(a_0), \gamma(a_n)) = \sum_{i=0}^{n-1} d_X(\gamma(a_i), \gamma(a_{i+1})) \geq (n-1) \log \frac{1}{\epsilon}.$$

On the other hand, for each $i > 0$ we can find arbitrarily small numbers $\delta > 0$ so that $d_{\mathcal{F}}(\gamma(a_{i-1}), \gamma(a_i - \delta)) \leq 15$. Since δ here can be taken arbitrarily small, it follows that $\gamma(a_i - \delta)$ and $\gamma(a_i)$ necessarily share an embedded loop. Consequently $\pi(\gamma(a_i))$ and $\pi(\gamma(a_i - \delta))$ overlap, and so we conclude

$$d_{\mathcal{F}}(\gamma(a_{i-1}), \gamma(a_i)) \leq 20.$$

By the triangle inequality, it follows that

$$(4) \quad d_{\mathcal{F}}(\gamma(0), \gamma(L)) = d_{\mathcal{F}}(\gamma(a_0), \gamma(a_n)) \leq 20n.$$

Combining equations (3) and (4) gives the claimed result. □

Lemma 4.6 *Let $\gamma: I \rightarrow \mathcal{X}$ be a K -quasigeodesic such that $\pi \circ \gamma: I \rightarrow \mathcal{F}$ is a K -quasigeodesic and $\gamma(I) \subset \mathcal{X}_{\epsilon}$. Then there exists $\epsilon' > 0$, depending only on ϵ and K , such that any standard geodesic $\rho: J \rightarrow \mathcal{X}$ with the same endpoints as γ is ϵ' -thin, ie $\rho(J) \subset \mathcal{X}_{\epsilon'}$.*

Proof Let $E \geq 1$ be the maximum of K and the constant D_0 that is provided by Proposition 4.4, and choose $\epsilon_1 \leq \epsilon$ sufficiently small, so that $\log(1/\epsilon_1) \geq 40E^2$. Notice that ϵ_1 depends only on K and ϵ . The facts that $\gamma(I) \subset \mathcal{X}_{\epsilon}$ and that ρ and γ have

finite Hausdorff distance (since they share the same endpoints) imply that there is some ϵ_0 such that $\rho(\mathbf{J}) \subset \mathcal{X}_{\epsilon_0}$. Choosing $\epsilon_0 < \epsilon$, we then have $\rho(\mathbf{J}), \gamma(\mathbf{I}) \subset \mathcal{X}_{\epsilon_0}$.

Let us write $G_t = \rho(t)$ for $t \in \mathbf{J}$. Suppose now that $(a', b') \subset \mathbf{J}$ is a subinterval such that $G_t \notin \mathcal{X}_{\epsilon_1}$ for all $t \in (a', b')$ (ie G_t has an immersed loop of length less than ϵ_1). Since $\rho|_{[a', b']}$ is a geodesic, Lemma 4.5 implies that

$$d_X(G_{a'}, G_{b'}) \geq \log\left(\frac{1}{\epsilon_1}\right) \frac{d_{\mathcal{F}}(G_{a'}, G_{b'}) - 20}{20}.$$

By Proposition 4.4, we can find points $a, b \in \mathbf{I}$ so that $d_X(\gamma(a), G_{a'}) \leq E$ and $d_X(\gamma(b), G_{b'}) \leq E$. Together with the fact that $\pi: \mathcal{X} \rightarrow \mathcal{F}$ is coarsely 80-Lipschitz, this implies

$$\begin{aligned} d_X(G_{a'}, G_{b'}) &\geq \log\left(\frac{1}{\epsilon_1}\right) \frac{d_{\mathcal{F}}(\gamma(a), \gamma(b)) - 160E - 20}{20} \\ &\geq \log\left(\frac{1}{\epsilon_1}\right) \frac{|b - a|/E - 161E - 20}{20}. \end{aligned}$$

On the other hand, since $\gamma(a)$ and $G_{a'}$ are ϵ_0 -thick, $d_X(G_{a'}, \gamma(a)) \leq E \cdot M_{\epsilon_0}$, for M_{ϵ_0} as in Lemma 2.3. So by the triangle inequality,

$$\begin{aligned} d_X(G_{a'}, G_{b'}) &\leq d_X(G_{a'}, \gamma(a)) + d_X(\gamma(a), \gamma(b)) + d_X(\gamma(b), G_{b'}) \\ &\leq E \cdot M_{\epsilon_0} + E|b - a| + 2E. \end{aligned}$$

Combining these inequalities, and using $\log(1/\epsilon_1) \geq 40E^2$, we find that

$$|b - a| \leq M_{\epsilon_0} + 2 + 322E^2 + 40E.$$

By the triangle inequality it follows that

$$\begin{aligned} |b' - a'| &\leq d_X(G_{a'}, \gamma(a)) + d_X(\gamma(a), \gamma(b)) + d_X(\gamma(b), G_{b'}) \\ &\leq E \cdot M_{\epsilon_0} + E|b - a| + E + E \\ &\leq 2E \cdot M_{\epsilon_0} + 4E + 322E^3 + 40E^2. \end{aligned}$$

In particular, this shows that \mathbf{J} cannot contain an infinite-length subinterval on which ρ is ϵ_1 -thin. Thus $\mathbf{J}' := \{t \in \mathbf{J} : G_t \notin \mathcal{X}_{\epsilon_1}\}$ is a disjoint union of finite subintervals of \mathbf{J} . Each component of \mathbf{J}' thus has the form $(c', d') \subset \mathbf{I}'$ where $G_{c'}, G_{d'} \in \mathcal{X}_{\epsilon_1}$ but $G_t \notin \mathcal{X}_{\epsilon_1}$ for all $t \in (c', d')$. (Note that if $\mathbf{I}_{\pm} \neq \pm\infty$, then $\gamma(\mathbf{I}_{\pm}) \in \mathcal{X}_{\epsilon_1}$ by choice of $\epsilon_1 \leq \epsilon$.) Since $G_{c'}, G_{d'} \in \mathcal{X}_{\epsilon_1}$, a repetition of the above argument now implies

$$|d' - c'| \leq L,$$

where $L := 2E \cdot M_{\epsilon_1} + 4E + 322E^3 + 40E^2$ depends only on E and ϵ_1 (and hence only on K and ϵ). Consequently, since ρ is a geodesic, for any $t \in [c', d']$ and $\alpha \in \mathbb{F}$ we have

$$\epsilon_1 \leq \ell(\alpha|G_{d'}) \leq e^L \ell(\alpha|G_t),$$

which implies that $G_t \in X_{\epsilon'}$ for $\epsilon' := \epsilon_1 e^{-L}$. Since this estimate holds for every point $t \in J'$ and ϵ' depends only on K and ϵ , the result follows. \square

Before proving Theorem 4.1 in its full generality, we focus on the case where the geodesic ρ is a standard geodesic.

Proposition 4.7 *The conclusions of Theorem 4.1 hold under the additional assumption that $\rho: J \rightarrow X$ is a standard geodesic.*

Proof Let $\gamma: I \rightarrow X$ be a K -quasigeodesic whose projection $\pi \circ \gamma: I \rightarrow \mathcal{F}$ is a K -quasigeodesic, and let $\rho: J \rightarrow X$ be any standard geodesic with the same endpoints as γ . By Lemma 4.3, γ is ϵ -thick for some $\epsilon \geq 0$ depending only on K (and on the injectivity radius of $\gamma(I_+)$ when $I_+ < \infty$). Lemma 4.6 therefore provides an $\epsilon' \geq 0$, depending only on K and ϵ , such that $\rho(t) \in X_{\epsilon'}$ for all $t \in J$. Thus conclusion (i) holds.

By applying Proposition 4.4 in conjunction with the symmetrization estimate from Lemma 2.3, we see that for each $i \in I$ there is a $t_i \in J$ with $d_x^{\text{sym}}(\gamma(i), \rho(t_i)) \leq M_{\epsilon'} D_0$. Similarly for every $j \in J$ there is some $s_j \in I$ such that $d_x^{\text{sym}}(\gamma(s_j), \rho(j)) \leq M_{\epsilon'} D_0$. Thus conclusion (ii) holds with $A = M_{\epsilon'} D_0$ since we have shown that

$$d_{\text{Haus}}(\gamma(I), \rho(J)) \leq M_{\epsilon'} D_0.$$

It is now easy to see that $\pi \circ \rho: J \rightarrow \mathcal{F}$ is a parametrized quasigeodesic: Consider any times $a, b \in J$ with $a < b$. Since π is coarsely 80-Lipschitz, we automatically have

$$d_{\mathcal{F}}(\rho(a), \rho(b)) \leq 80 \cdot d_x(\rho(a), \rho(b)) + 80 = 80|b - a| + 80.$$

On the other hand, by the above there exist times $s, t \in I$ such that $d_x^{\text{sym}}(\gamma(s), \rho(a))$ and $d_x^{\text{sym}}(\gamma(t), \rho(b))$ are both bounded by $M_{\epsilon'} D_0$. By the triangle inequality, it follows that

$$d_x(\gamma(s), \gamma(t)) \geq d_x(\rho(a), \rho(b)) - 2M_{\epsilon'} D_0 = |b - a| - 2M_{\epsilon'} D_0.$$

Since γ is a directed K -quasigeodesic by assumption, this implies

$$(t - s) \geq \frac{1}{K} d_x(\gamma(s), \gamma(t)) - K \geq \frac{1}{K} |b - a| - \frac{2M_{\epsilon'} D_0}{K} - K.$$

Since $\pi \circ \gamma: J \rightarrow \mathcal{F}$ is also a K -quasigeodesic, we may extend this to conclude

$$|b - a| \leq K(t - s) + 2M_{\epsilon'} D_0 + K^2 \leq K(Kd_{\mathcal{F}}(\rho(a), \rho(b)) + K) + 2M_{\epsilon'} D_0 + K^2.$$

Therefore, $\pi \circ \rho$ is a K' -quasigeodesic for $K' = \max\{80, 2K^2 + 2M_{\epsilon'} D_0\}$. This proves conclusion (iii). \square

4.1 More on Bestvina–Feighn projections

Proposition 4.7 above suffices to prove our main result on hyperbolic extensions of free groups (Theorem 1.1). However for completeness, and to strengthen the quasiconvexity

results in Section 5, it is desirable to prove the more general result Theorem 4.1 which applies to arbitrary geodesics. This subsection is devoted to that purpose.

Heuristically, Theorem 4.1 follows easily from Proposition 4.7 and some ideas in Bestvina and Feighn [5]. Specifically, as remarked in [5, Corollary 7.3], Bestvina and Feighn’s Proposition 7.2 (Proposition 2.11 here) essentially says that folding paths that make definite progress in the factor complex are strongly contracting in Outer space, which generalizes Algom-Kfir’s result [1]. One should then apply this notion of strong contracting to conclude that such folding paths are stable (using standard arguments). However, to make this precise, we first require a more detailed discussion of the projection $\text{Pr}_\gamma: \mathcal{X} \rightarrow \gamma(I)$.

Following [5], given a free factor $A \in \mathcal{F}^0$ and a point $G \in \mathcal{X}$, we write $A|G$ for the core subgraph of the cover of G corresponding to the conjugacy class of A in $\mathbb{F} \cong \pi_1(G)$. We say that $A|G$ is *the core of the A -cover*. Restricting the covering map thus gives a canonical immersion $A|G \rightarrow G$ that identifies $\pi_1(A|G)$ with $A \leq \pi_1(G)$. The graph $A|G$ is equipped with a metric structure by pulling back the edge lengths from G . Similarly, whenever G is given an illegal turn structure (eg if G lies on a folding path), we may pull back this structure via $A|G \rightarrow G$, equipping $A|G$ with an illegal turn structure as well. When A is a cyclic free factor generated by a primitive element $\alpha \in \mathbb{F}$, we note that $A|G$ agrees with our already defined $\alpha|G$.

Setting $I = (18\check{m}(3r - 3) + 6)(2r - 1)$, where $r = \text{rk}(\mathbb{F})$ and \check{m} denotes the maximum number of illegal turns in any train track structure on any $G \in \mathcal{X}$, Bestvina and Feighn then define the following projections from \mathcal{F} to folding paths in \mathcal{X} :

Definition 4.8 Let $\gamma: I \rightarrow \mathcal{X}$ be a folding path, and let $A \in \mathcal{F}^0$ be a proper free factor. The *left* and *right* projections of A to γ are respectively given by

$$\begin{aligned} \text{left}_\gamma(A) &:= \inf\{t \in I : A|G_t \text{ has an immersed legal segment of length } 3\} \in I, \\ \text{right}_\gamma(A) &:= \sup\{t \in I : A|G_t \text{ has an immersed illegal segment of length } I\} \in I, \end{aligned}$$

where here an *illegal segment* means a segment that does not contain a legal segment of length 3.

Using this, the Bestvina–Feighn projection Pr_γ is defined as follows:

Definition 4.9 (Bestvina–Feighn projection) Let $\gamma: I \rightarrow \mathcal{X}$ be a folding path. For $H \in \mathcal{X}$, the left and right projections of H are defined to be

$$\text{left}_\gamma(H) := \inf_{A \in \pi(H)} \text{left}_\gamma(A) \quad \text{and} \quad \text{right}_\gamma(H) := \sup_{A \in \pi(H)} \text{right}_\gamma(A).$$

The *projection* of H to $\gamma(I)$ is then given by $\text{Pr}_\gamma(H) := \gamma(\text{left}_\gamma(H))$.

Note that every candidate conjugacy class $\alpha \in \mathcal{C}_H$ at $H \in \mathcal{X}$ is primitive and thus generates a cyclic free factor of \mathbb{F} ; thus we may view α as a point in \mathcal{F}^0 . Since the immersion $\alpha|H \rightarrow H$ lands in a proper subgraph of H , we additionally have $\alpha \leq A$ for some $A \in \pi(H)$. Therefore, Bestvina and Feighn’s Proposition 6.4 and Corollary 6.11 immediately give the following estimates regarding the above projections.

Proposition 4.10 (Bestvina and Feighn) *Let $\gamma: I \rightarrow \mathcal{X}$ be a folding path and let $H \in \mathcal{X}$ be any point. Then for every candidate $\alpha \in \mathcal{C}_H$ of H , we have*

$$[\text{left}_\gamma(\alpha), \text{right}_\gamma(\alpha)] \subset [\text{left}_\gamma(H), \text{right}_\gamma(H)] \subset I.$$

Moreover, the set

$$\pi(\gamma([\text{left}_\gamma(H), \text{right}_\gamma(H)])) \subset \mathcal{F}$$

has uniformly bounded diameter depending only on $\text{rk}(\mathbb{F})$.

As a consequence, we may deduce that $\text{Pr}_\gamma(H)$ coarsely agrees with the closest-point projection of H to $\gamma(I)$ in the case that γ makes definite progress in \mathcal{F} .

Lemma 4.11 *Let $\gamma: I \rightarrow \mathcal{X}$ be a folding path whose projection $\pi \circ \gamma: I \rightarrow \mathcal{F}$ is a K -quasigeodesic. Then there exists $D \geq 0$, depending only on K and $\text{rk}(\mathbb{F})$ (and the injectivity radius of $\gamma(I_+)$ when $I_+ < \infty$) satisfying the following: If $H \in \mathcal{X}$ and $t_0 \in I$ are such that*

$$d_{\mathcal{X}}(H, \gamma(t_0)) = \inf\{d_{\mathcal{X}}(H, \gamma(t)) : t \in I\},$$

then $d_{\mathcal{X}}^{\text{sym}}(\gamma(t_0), \text{Pr}_\gamma(H)) \leq D$.

Proof We write $G_t = \gamma(t)$ for $t \in I$. Let us define

$$L = \inf\{\text{left}_\gamma(\alpha) : \alpha \in \mathcal{C}_H\} \quad \text{and} \quad R = \sup\{\text{right}_\gamma(\alpha) : \alpha \in \mathcal{C}_H\}.$$

Note that each candidate $\alpha \in \mathcal{C}_H$ is a simple class and that, by definition of $\text{left}_\gamma(\alpha)$, the loop $\alpha|G_s$ cannot contain a legal segment of length 3 for any $s < L$. Therefore, Lemma 5.8 of [5] and the fact that $\pi \circ \gamma$ is a K -quasigeodesic together imply that there exists $T \geq 0$ depending only on K and $\text{rk}(\mathbb{F})$ such that for all $t \geq T$ we have

$$\ell(\alpha|G_{L-t}) > 2\ell(\alpha|G_L).$$

Since this estimate holds for each candidate, Proposition 2.5 implies that $2d_{\mathcal{X}}(H, G_L) < d_{\mathcal{X}}(H, G_{L-t})$ for all $t \geq T$. Similarly, for all $s > R$ the loop $\alpha|G_s$ contains immersed legal segments contributing to a definite fraction of $\ell(\alpha|G_s)$. Thus, by Corollary 4.8 of [5], the length $\ell(\alpha|G_s)$ grows exponentially beyond R and so after increasing T if necessary we have

$$\ell(\alpha|G_{R+t}) > 2\ell(\alpha|G_R)$$

and consequently $2 d_{\mathcal{X}}(H, G_R) < d_{\mathcal{X}}(H, G_{R+t})$ for all $t \geq T$. Given any time $t_0 \in I$ satisfying

$$d_{\mathcal{X}}(H, G_{t_0}) = \inf\{d_{\mathcal{X}}(H, \gamma(t)) : t \in I\},$$

it follows that t_0 necessarily lies in $[L - T, R + T]$.

By Proposition 4.10, we know that $\pi(\gamma([L, R]))$ has bounded diameter and bounded \mathcal{F} -distance from $\pi(\text{Pr}_{\gamma}(H))$. Therefore, since $\pi \circ \gamma$ is a K -quasigeodesic, there exists D' , depending only on K and $\text{rk}(\mathbb{F})$, such that $|s_0 - t_0| \leq D'$, where $s_0 \in I$ is the time for which $G_{s_0} = \text{Pr}_{\gamma}(H)$. By Lemma 4.3, we additionally know $\gamma(I) \subset \mathcal{X}_{\epsilon}$ for some $\epsilon > 0$ depending on K (and the injectivity radius of $\gamma(I_+)$ when $I_+ < \infty$). Therefore, since γ is a directed geodesic, we may conclude $d_{\mathcal{X}}^{\text{sym}}(G_{t_0}, \text{Pr}_{\gamma}(H)) \leq M_{\epsilon} D'$, as desired. \square

Lemma 4.11 shows that whenever $\gamma: I \rightarrow \mathcal{X}$ is a standard geodesic for which $\pi \circ \gamma$ is a K -quasigeodesic, then the closest-point projection $\mathcal{X} \rightarrow \gamma(I)$ coarsely agrees with $\text{Pr}_{\gamma}: \mathcal{X} \rightarrow \gamma(I)$. Thus, since γ makes definite progress in \mathcal{F} , Proposition 2.11 implies that γ is strongly contracting. That is, there exists D , depending only on $\text{rk}(\mathbb{F})$ and K (and the injectivity radius of $\gamma(I_+)$ if $I_+ < \infty$), such that if $d_{\mathcal{X}}(H, H') \leq d_{\mathcal{X}}(H, \gamma(I))$, then any closest-point projections of H and H' to $\gamma(I)$ are at most $d_{\mathcal{X}}^{\text{sym}}$ -distance D apart. We are therefore in the situation of the standard Morse lemma (see eg Section 5.4 of [1]), which gives the following stability result.

Lemma 4.12 (Morse lemma for \mathcal{F} -progressing folding paths) *Suppose $\gamma: I \rightarrow \mathcal{X}$ is a standard geodesic for which $\pi \circ \gamma: I \rightarrow \mathcal{F}$ is a K -quasigeodesic. Then for any $K' \geq 1$ there exists B depending only on $\text{rk}(\mathbb{F})$, K and K' (and the injectivity radius of $\gamma(I_+)$ when $I_+ < \infty$) such that $d_{\text{Haus}}(\gamma(I), \rho(J)) \leq B$ for every K' -quasigeodesic $\rho: J \rightarrow \mathcal{X}$ with the same endpoints as γ .*

Using this, we may finally give the proof of Theorem 4.1:

Proof of Theorem 4.1 Let $\gamma: I \rightarrow \mathcal{X}$ be a K -quasigeodesic such that $\pi \circ \gamma$ is also a K -quasigeodesic, and let $\epsilon, A > 0$ and $K' \geq 1$ be the corresponding constants provided by Proposition 4.7. Choose a standard geodesic $\rho': J' \rightarrow \mathcal{X}$ with the same endpoints as γ . Then by Proposition 4.7 we know that $\rho'(J') \subset \mathcal{X}_{\epsilon}$ and that $\pi \circ \rho'$ is a K' -quasigeodesic. Now consider an arbitrary geodesic $\rho: J \rightarrow \mathcal{X}$ with the same endpoints as γ , and thus also ρ' . Applying Lemma 4.12 to ρ and the folding path ρ' , we find that

$$d_{\text{Haus}}(\rho(J), \rho'(J')) \leq B$$

for some B depending only on ϵ and K' . Consequently $\rho(J) \subset \mathcal{X}_{\epsilon'}$, where $\epsilon' = e^{-B}\epsilon$. Since $\rho'(J')$ and $\gamma(I)$ have Hausdorff distance at most A by Proposition 4.7, it also

follows that $d_{\text{Haus}}(\rho(\mathbf{J}), \gamma(\mathbf{I})) \leq B + A$. Finally, as in the proof of Proposition 4.7 above, these two facts easily show that $\pi \circ \rho$ is a K'' -quasigeodesic for some K'' depending only on ϵ' and $A + B$. \square

5 Quasi-isometric into \mathcal{F} implies quasiconvex in \mathcal{X}

Consider a finitely generated subgroup $\Gamma \leq \text{Out}(\mathbb{F})$. For any finite generating set $S \subset \Gamma$, we then consider the word metric $d_\Gamma = d_{\Gamma,S}$ on Γ defined by $d_\Gamma(g, h) = |g^{-1}h|_S$, where $|\cdot|_S$ denotes word length with respect to S . This is just the restriction of the path metric on the Cayley graph $\text{Cay}(\Gamma, S)$ to $\Gamma = (\text{Cay}(\Gamma, S))^0$. In this section we explain various ways in which the geometry of Γ relates to that of \mathcal{X} or \mathcal{F} .

For any free factor $A \in \mathcal{F}^0$, we may consider the orbit map $(\Gamma, d_\Gamma) \rightarrow (\mathcal{F}, d_{\mathcal{F}})$ given by $g \mapsto g \cdot A$. We say that this map is a qi-embedding if it is a K -quasi-isometric embedding for some $K \geq 1$. We remark that if some orbit map into \mathcal{F} is a quasi-isometric embedding, then so is any orbit map into \mathcal{F} .

Definition 5.1 We say $\Gamma \leq \text{Out}(\mathbb{F})$ *qi-embeds into \mathcal{F}* if Γ is finitely generated and any orbit map into \mathcal{F} is a qi-embedding.

Given a point $H \in \mathcal{X}$, we say that the orbit $\Gamma \cdot H$ is *quasiconvex* if it is A -quasiconvex for some $A \geq 0$, meaning that every (directed) geodesic between points of $\Gamma \cdot H$ lies in the (symmetric) A -neighborhood $\mathcal{N}_A(\Gamma \cdot H)$ (see Section 2.6). We record the following straightforward consequence of quasiconvexity.

Lemma 5.2 *Let $\Gamma \leq \text{Out}(\mathbb{F})$ be finitely generated with corresponding word metric d_Γ , and suppose $H \in \mathcal{X}$ is such that $\Gamma \cdot H \subset \mathcal{X}$ is quasiconvex. Then the orbit map $g \mapsto g \cdot H$ defines a quasi-isometric embedding $(\Gamma, d_\Gamma) \rightarrow (\mathcal{X}, d_{\mathcal{X}})$.*

Proof Let $S \subset \Gamma$ be the generating set inducing the word metric d_Γ . By assumption, there exists $A \geq 0$ such that $\Gamma \cdot H$ is A -quasiconvex. Choose $\epsilon = \epsilon(H, A) > 0$ so that $\mathcal{N}_A(\Gamma \cdot H) \subset \mathcal{X}_\epsilon$. Since $\text{Out}(\mathbb{F})$ acts properly discontinuously on \mathcal{X} , the set

$$D = \{g \in \Gamma : d_{\mathcal{X}}^{\text{sym}}(H, g \cdot H) \leq 2A + M_\epsilon\}$$

is finite, and we may set $K = \max_{g \in D} d_\Gamma(1, g)$.

Letting $\gamma: [0, L] \rightarrow \mathcal{X}$ be a (directed) geodesic from $g \cdot H$ to $g' \cdot H$, our hypothesis implies $\gamma \subset \mathcal{N}_A(\Gamma \cdot H)$ and thus that $\gamma(t) \in \mathcal{X}_\epsilon$ for all $t \in [0, L]$. Setting $N = \lfloor L \rfloor$, we may find $h_0, \dots, h_{N+1} \in \Gamma$ so that $h_0 = g$, $h_{N+1} = g'$ and $d_{\mathcal{X}}^{\text{sym}}(\gamma(i), h_i \cdot H) < A$ for all $i = 0, \dots, N$. In particular, we see that for each $i = 0, \dots, N$ the element $h_i^{-1}h_{i+1}$

translates H by at most $d_{\mathcal{X}}^{\text{sym}}$ -distance $2A + M_{\epsilon}$ and hence has $d_{\Gamma}(1, h_i^{-1}h_{i+1}) \leq K$. Thus

$$\begin{aligned} d_{\Gamma}(g, g') &\leq d_{\Gamma}(h_0, h_1) + \dots + d_{\Gamma}(h_N, h_{N+1}) \\ &= \sum_{i=0}^N d_{\Gamma}(1, h_i^{-1}h_{i+1}) \\ &\leq K(N + 1) \leq K(L + 1) \\ &= K d_{\mathcal{X}}(g \cdot H, g' \cdot H) + K. \end{aligned}$$

On the other hand, if $K' = \max\{d_{\mathcal{X}}(H, s \cdot H) : s \in S\}$, then $d_{\mathcal{X}}(g \cdot H, g' \cdot H) \leq K' d_{\Gamma}(g, g')$. Therefore $g \mapsto g \cdot H$ is a $\max\{K', K\}$ -quasi-isometric embedding. \square

Definition 5.3 A subgroup $\Gamma \leq \text{Out}(\mathbb{F})$ is said to be *quasiconvex in \mathcal{X}* if the orbit $\Gamma \cdot H$ is quasiconvex for every $H \in \mathcal{X}$.

We remark that knowing a single orbit $\Gamma \cdot H$ is quasiconvex in \mathcal{X} does not necessarily seem to imply that Γ is quasiconvex: it is conceivable that some other orbit $\Gamma \cdot H'$ could fail to be quasiconvex.

We now employ the results of Section 4 to show that every subgroup that qi-embeds into the factor complex is quasiconvex in Outer space:

Theorem 5.4 *Let $\Gamma \leq \text{Out}(\mathbb{F})$ be finitely generated. If Γ qi-embeds into \mathcal{F} , then Γ is quasiconvex in \mathcal{X} .*

Proof Let $H \in \mathcal{X}$ be arbitrary and let $A \in \pi(H) \subset \mathcal{F}$. Since $\pi: \mathcal{X} \rightarrow \mathcal{F}$ is coarsely Lipschitz and $g \mapsto g \cdot A$ gives a quasi-isometric embedding $\Gamma \rightarrow \mathcal{F}$, the orbit map $\mathcal{O}: \Gamma \rightarrow \mathcal{X}$ defined by $\mathcal{O}(g) = g \cdot H$ is also a quasi-isometric embedding. Let $g_1, g_2 \in \Gamma$ be given. For any (discrete) geodesic path $\rho: \{1, \dots, N\} \rightarrow \Gamma$ from g_1 to g_2 , the image $\mathcal{O} \circ \rho$ is thus a quasigeodesic path in \mathcal{X} joining $g_1 \cdot H$ and $g_2 \cdot H$ such that $\pi \circ \mathcal{O} \circ \rho$ is also a quasigeodesic in \mathcal{F} . Theorem 4.1 then implies that any geodesic $\gamma: I \rightarrow \mathcal{X}$ from $g_1 \cdot H$ to $g_2 \cdot H$ stays uniformly close to the image of $\mathcal{O} \circ \rho$, which is contained in $\Gamma \cdot H$. Hence, Γ is quasiconvex in \mathcal{X} . \square

6 Quasiconvex orbit implies conjugacy flaring

Consider a subgroup $\Gamma \leq \text{Out}(\mathbb{F})$ with finite generating set $S \subset \Gamma$ and corresponding word length $|\cdot|_S$. Fix also a basis X of \mathbb{F} . We say that Γ has (λ, M) -conjugacy flaring

for the given $\lambda > 1$ and positive integer $M \in \mathbb{N}$ if the following condition is satisfied:

For all $\alpha \in \mathbb{F}$ and $g_1, g_2 \in \Gamma$ with $|g_i|_S \geq M$ and $|g_1 g_2|_S = |g_1|_S + |g_2|_S$, we have

$$\lambda \|\alpha\|_X \leq \max\{\|g_1(\alpha)\|_X, \|g_2^{-1}(\alpha)\|_X\},$$

where $\|\cdot\|_X$ denotes conjugacy length (ie the shortest word length with respect to X of any element in the given conjugacy class).

In this section we show that any purely hyperbolic subgroup $\Gamma \leq \text{Out}(\mathbb{F})$ that qi-embeds into \mathcal{F} has conjugacy flaring. In fact, our argument only relies on the following weaker hypothesis. Before making the definition, we first recall that a (finite) geodesic in Γ may be encoded by a sequence of group elements (g_0, \dots, g_N) such that $d_\Gamma(g_i, g_j) = |i - j|$ for all $i, j = 0, \dots, N$. For $R \in \mathcal{X}$, the image of this geodesic in the orbit $\Gamma \cdot R$ is simply the set of points $g_0 \cdot R, \dots, g_N \cdot R$.

Definition 6.1 (QCX condition) Consider a subgroup $\Gamma \leq \text{Out}(\mathbb{F})$ and point $R \in \mathcal{X}$. We say that the orbit $\Gamma \cdot R$ is A -QCX if for any geodesic (g_0, \dots, g_N) in Γ there exists a folding path $\rho: \mathbf{J} \rightarrow \mathcal{X}$ with Hausdorff distance at most A from the image of (g_0, \dots, g_N) , that is

$$d_{\text{Haus}}(\rho(\mathbf{J}), \{g_0 \cdot R, \dots, g_N \cdot R\}) \leq A,$$

such that $d_{\mathcal{X}}^{\text{sym}}(\rho(\mathbf{J}_-), g_0 \cdot R) \leq A$ and $d_{\mathcal{X}}^{\text{sym}}(\rho(\mathbf{J}_+), g_N \cdot R) \leq A$.

We summarize this property by saying the image of the geodesic (g_0, \dots, g_N) in $\Gamma \cdot R \subset \mathcal{X}$ has Hausdorff distance at most A from a folding path in \mathcal{X} with the correct orientation. Note that for an arbitrary subgroup Γ , there is no direct correspondence between quasiconvexity and this QCX condition. However, we have the following relationship when Γ is hyperbolic.

Lemma 6.2 Suppose that $\Gamma \leq \text{Out}(\mathbb{F})$ is finitely generated, δ -hyperbolic, and that $\Gamma \cdot R \subset \mathcal{X}$ is A -quasiconvex. Then $\Gamma \cdot R$ is A' -QCX for some A' .

Proof Let d_Γ be a word metric on Γ such that (Γ, d_Γ) is δ -hyperbolic. By Lemma 5.2 the orbit map $g \mapsto g \cdot R$ defines a K -quasi-isometric embedding $(\Gamma, d_\Gamma) \rightarrow (\mathcal{X}, d_{\mathcal{X}})$ for some K . Let (g_0, \dots, g_N) be any geodesic in Γ and let $\gamma_0: \mathbf{I}_0 \rightarrow \mathcal{X}$ be a standard geodesic from $g_0 \cdot R$ to $g_N \cdot R$. Then by quasiconvexity we have $\gamma_0(\mathbf{I}_0) \subset \mathcal{N}_A(\Gamma \cdot R)$. Note that $\mathcal{N}_A(\Gamma \cdot R) \subset \mathcal{X}_\epsilon$ for some $\epsilon > 0$ (since R has positive injectivity radius). The scaling image $\text{Im}^{\text{sc}}(\gamma_0)$ of γ_0 thus lives in \mathcal{X}_ϵ and hence has length at most $\log(2/\epsilon)$ by Lemma 2.6. Setting $A_0 = A + M_\epsilon(\log(2/\epsilon) + 1)$, it follows that if $\gamma: [0, L] \rightarrow \mathcal{X}$ is the

folding portion of γ_0 and $m = \lfloor L \rfloor$, then we may find group elements h_0, \dots, h_m with $h_0 = g_0$ and $h_m = g_N$ such that $d_{\mathcal{X}}^{\text{sym}}(\gamma(i), h_i \cdot R) \leq A_0$ for all $i = 0, \dots, m \in [0, L]$. It follows that for $i < j$,

$$j - i - 2A_0 \leq d_{\mathcal{X}}(h_i \cdot R, h_j \cdot R) \leq j - i + 2A_0.$$

Thus the map $i \mapsto h_i \cdot R$ is a discrete $2A_0$ -quasigeodesic in $(\mathcal{X}, d_{\mathcal{X}})$; consequently, the sequence $g_0 = h_0, \dots, h_m = g_N$ is a K' -quasigeodesic in Γ for some $K' = K'(K, A_0)$. Since Γ is δ -hyperbolic, Proposition 2.1 implies that for each $j \in \{0, \dots, m\}$ there exists $i \in \{0, \dots, N\}$ with $d_{\Gamma}(h_j, g_i) \leq R_0 = R_0(K', \delta)$. Noting that

$$d_{\mathcal{X}}^{\text{sym}}(h_j \cdot R, g_i \cdot R) \leq M_{\epsilon}(Kd_{\Gamma}(h_j, g_i) + K),$$

it follows that

$$\{h_0 \cdot R, \dots, h_m \cdot R\} \subset \mathcal{N}_{M_{\epsilon}(KR_0 + K)}(\{g_0 \cdot R, \dots, g_N \cdot R\}).$$

Because we also have $\gamma([0, L]) \subset \mathcal{N}_{M_{\epsilon} + A_0}(\{h_0 \cdot R, \dots, h_m \cdot R\})$ by the selection of h_0, \dots, h_m , the claim follows with $A' = M_{\epsilon}(KR_0 + K) + M_{\epsilon} + A_0$. \square

Corollary 6.3 *If $\Gamma \leq \text{Out}(\mathbb{F})$ qi-embeds into \mathcal{F} , then for every $R \in \mathcal{X}$ there exists $A \geq 0$ such that the orbit $\Gamma \cdot R$ is A -QCX.*

Proof By Theorem 5.4 we know that every orbit $\Gamma \cdot R$ is quasiconvex in \mathcal{X} . Since \mathcal{F} is hyperbolic, the hypothesis that Γ qi-embeds into \mathcal{F} also implies that Γ is finitely generated and δ -hyperbolic for some $\delta \geq 0$. Lemma 6.2 thus implies the claim. \square

We also have the following simple consequence of being A -QCX:

Lemma 6.4 *Suppose $\Gamma \leq \text{Out}(\mathbb{F})$ is finitely generated and that the orbit $\Gamma \cdot R \subset \mathcal{X}$ is A -QCX. Then $g \mapsto g \cdot R$ gives a quasi-isometric embedding $(\Gamma, d_{\Gamma}) \rightarrow (\mathcal{X}, d_{\mathcal{X}})$.*

Proof Let g_1, \dots, g_N be a geodesic in Γ from $g = g_1$ to $g' = g_N$. By using a folding path $\gamma: I \rightarrow \mathcal{X}$ with Hausdorff distance at most A from the image of (g_1, \dots, g_N) , an argument exactly as in Lemma 5.2 shows that $d_{\Gamma}(g, g')$ and $d_{\mathcal{X}}(g \cdot R, g' \cdot R)$ agree up to bounded additive and multiplicative error depending only on R and A . \square

Having established this terminology, we now turn to the main result of this section:

Theorem 6.5 *Suppose that $\Gamma \leq \text{Out}(\mathbb{F})$ is finitely generated and purely hyperbolic and that for some $R \in \mathcal{X}$ the orbit $\Gamma \cdot R$ is A -QCX. Then Γ has $(2, M)$ -conjugacy flaring for some $M \in \mathbb{N}$ depending only on A and R .*

The proof of Theorem 6.5 will take several steps. We first show in Proposition 6.11 that, provided Γ is purely hyperbolic, a corresponding flaring property holds for the length of any conjugacy class along any folding path that remains within the symmetric A -neighborhood of the orbit $\Gamma \cdot R \subset \mathcal{X}$. When the orbit $\Gamma \cdot R$ is A -QCX we use this flaring on folding paths to deduce a similar flaring in the orbit $\Gamma \cdot R$. Measuring this flaring from R , where $\ell(\cdot|R)$ coarsely agrees with the conjugacy length $\|\cdot\|_X$, then yields Theorem 6.5. We first require the following lemma, which is central to this section. It implies that there is a uniform bound on how long a conjugacy class can stay short along our folding paths.

Lemma 6.6 *Fix $\Gamma \leq \text{Out}(\mathbb{F})$ and $R \in \mathcal{X}$. For any $L_0 \geq 0$ and $A_0 \geq 0$, there is a $D_0 \geq 0$ satisfying the following: if $\alpha \in \mathbb{F}$ is nontrivial and $\gamma: I \rightarrow \mathcal{X}$ is a folding path with $G_t = \gamma(t) \in \mathcal{N}_{A_0}(\Gamma \cdot R)$ for all $t \in I$, then either*

$$\text{diam}\{t \in I : \ell(\alpha|G_t) \leq L_0\} \leq D_0$$

or there is an **infinite-order** element $\phi \in \Gamma$ with $\phi([\alpha]) = [\alpha]$.

Proof Let Γ_α be the subgroup of elements of Γ that fix the conjugacy class of α . If Γ_α is a torsion group, then $|\Gamma_\alpha| \leq e_r$ by Lemma 2.14.

Let a and b be the infimum and supremum of the set $\{t \in I : \ell(\alpha|G_t) \leq L_0\}$. Then, by Proposition 3.2, for all $t \in [a, b]$ we have $\ell(\alpha|G_t) \leq ML_0$, where $M = 6 \text{rk}(\mathbb{F})$. It follows that if $d_0 \geq 3A_0$, then for all $t, t + d_0 \in [a, b]$ the points G_t and G_{t+d_0} cannot both be A_0 -close (in symmetric distance) to the same orbit point of $\Gamma \cdot R$ (since $d_X^{\text{sym}}(G_t, G_{t+d_0}) \geq d_X(G_t, G_{t+d_0}) = d_0 > 2A_0$).

Set $N = \lfloor (b - a)/d_0 \rfloor$ and for each $0 \leq n \leq N$ select $\phi_n \in \Gamma$ such that

$$d_X^{\text{sym}}(\phi_n \cdot R, G_{a+d_0n}) \leq A_0.$$

By our choice of d_0 , we have, for $0 \leq i, j \leq N$, that $\phi_i = \phi_j$ if and only if $i = j$. By assumption, $\alpha \in \mathbb{F}$ has length at most ML_0 in G_{a+d_0n} ; therefore we have that $\ell(\phi_n^{-1}(\alpha)|R) \leq e^{A_0}ML_0$ for all $0 \leq n \leq N$. Let C denote the number of immersed loops in R of length at most $e^{A_0}ML_0$; we note that C depends only on R, A_0 and L_0 . It follows that if $N > C(e_r + 1)$ then we may find distinct $0 \leq k_0 < \dots < k_{e_r} \leq N$ such that

$$\phi_{k_0}^{-1}(\alpha) = \phi_{k_1}^{-1}(\alpha) = \dots = \phi_{k_{e_r}}^{-1}(\alpha).$$

Since the ϕ_{k_i} are all distinct, this implies that Γ_α contains at least $e_r + 1$ elements, and hence, an infinite-order element. Otherwise $N \leq C(e_r + 1)$ and thus we conclude

$$b - a \leq d_0(N + 1) \leq d_0(C(e_r + 1) + 1).$$

Setting $D_0 = d_0(C(e_r + 1) + 1)$ completes the proof. □

We next examine how the length of a loop varies over a folding path G_t that is near the orbit of Γ . Our arguments are inspired by Section 5 of [5], however, the use of Lemma 6.6 greatly simplifies our analysis.

For a folding path G_t and a conjugacy class α , recall that $\alpha|G_t$ is the core of the α -cover of G_t . We think of $\alpha|G_t$ as having edge lengths and illegal turn structure induced from G_t . As such, $\alpha|G_t$ is composed of legal segments separated by illegal turns. We say that a collection of consecutive illegal turns in $\alpha|G_t$ survive to $\alpha|G_{t'}$ for $t \leq t'$ if no illegal turn in the collection becomes legal in the process of folding from G_t to $G_{t'}$ nor do two illegal turns of the collection collide. In other words, a collection of consecutive illegal turns of $\alpha|G_t$ survive to $\alpha|G_{t'}$ if and only if there is a collection of consecutive illegal turns of $\alpha|G_{t'}$ and a bijection between the illegal turns in both collections induced by the process of unfolding an illegal turn of $\alpha|G_{t'}$ to an illegal turn of $\alpha|G_t$ (see Lemma 2.7 and the surrounding discussion). Set \check{m} equal to the maximum number of illegal turns in any train track structure on any $G \in \mathcal{X}$. Note that $\check{m} \geq 2 \operatorname{rk}(\mathbb{F}) - 2$.

Lemma 6.7 (illegal turns don't survive) *Suppose that $\Gamma \leq \operatorname{Out}(\mathbb{F})$ is purely hyperbolic and that $R \in \mathcal{X}$. For each $l \geq 0$ and $A_0 \geq 0$ there exists $D_l \geq 0$ satisfying the following property. If G_t is a folding path with $G_t \in \mathcal{N}_{A_0}(\Gamma \cdot R)$ for all $t \in [a, b]$ and α is a conjugacy class such that $\alpha|G_a$ has a segment containing $\check{m} + 1$ consecutive illegal turns that survive to $\alpha|G_b$ and the length of each legal segment between these illegal turns in $\alpha|G_b$ is no greater than l , then $b - a \leq D_l$.*

Proof Let s_t be the segment spanning the consecutive surviving illegal turns in $\alpha|G_t$ for $a \leq t \leq b$. Since the number of illegal turns in s_b is greater than the total number of illegal turns in G_b , there are a pair of illegal turns of s_b that project to the same illegal turn of G_b under the immersion $\alpha|G_b \rightarrow G_b$. Let s'_b be the subsegment between two such turns and let σ_b denote the loop obtained by projecting s'_b to G_b and identifying its endpoints.

By the unfolding principle of [5] (Lemma 2.7), there is a subsegment s'_t of s_t that maps to the segment s'_b after folding and tightening and such that the illegal turn endpoints of s'_t map to the same illegal turn in G_t (just as in G_b). Hence, we may form the loop σ_t by identifying these endpoints in G_t . We note σ_t is immersed for each $a \leq t \leq b$ except possibly at the illegal turn corresponding to the endpoints of s'_t and that the conjugacy class of σ_t maps to the conjugacy class of σ_b under the folding map $G_t \rightarrow G_b$, again by the unfolding principle. Let σ denote this conjugacy class in \mathbb{F} .

By construction, the length of σ_b is bounded by $l \cdot (\check{m} + 1)$ and the number of illegal turns of σ_a is no more than $\check{m} + 1$, since these illegal turns all survive in G_b by assumption.

By Lemma 3.1, $\ell(\sigma_t) \leq 2l \cdot (\check{m} + 1)$ for all $a \leq t \leq b$. Then, by Lemma 6.6 either $\phi(\sigma) = \sigma$ for some infinite-order $\phi \in \Gamma$ or we have $b - a \leq D_l$ for some D_l depending only on A_0 , l and R . Since Γ is purely hyperbolic, the claim follows. \square

Recall the notation from Section 3: if G_t is a folding path and α is a conjugacy class, then $k_t = k(\alpha|G_t)$ denotes the number of illegal turns of $\alpha|G_t$ and m_t denotes the illegality of G_t . The following lemma is similar to Lemma 5.4 of [11]. Again, we use that our folding path is near the orbit of Γ as a replacement for having a single train track map, as was the case in [11]. Let $r = \text{rk}(\mathbb{F})$.

Lemma 6.8 *Let G_t be a folding path with $G_t \in \mathcal{N}_{A_0}(\Gamma \cdot R)$ for $t \in [a, b]$ and let p_b be an immersed path in G_b whose endpoints are illegal turns such that $k(p_b) \geq 2(2r - 2)$ and p_b contains no legal segment of length $L \geq 3$. Let p_t be the corresponding path in G_t whose endpoints are illegal turns which is obtained from p_b by unfolding. Then*

$$\frac{\epsilon_0 \cdot k(p_t)}{2(2r - 2)} \leq \ell(p_t) \leq L \cdot k(p_t),$$

where ϵ_0 is the minimal injectivity radius of any graph in $\mathcal{N}_{A_0}(\Gamma \cdot R)$.

Proof Any path in G_t with at least $2r - 2$ illegal turns contains a loop in G_t which has length at least ϵ_0 . The lemma now easily follows. \square

We find the following terminology helpful. Suppose that G_t , $t \in [a, b]$, is a folding path and that α is a nontrivial conjugacy class in \mathbb{F} . As mentioned earlier, the immersed loop $\alpha|G_t \rightarrow G_t$ consists of legal segments separated by illegal turns. We let α_t^{leg} denote the subset of $\alpha|G_t$ consisting of maximal legal segments of length at least 3, and we write $\text{leg}(\alpha|G_t)$ for the length of α_t^{leg} . This is the *legal length* of $\alpha|G_t$. The complement $\alpha|G_t - \alpha_t^{\text{leg}}$ consists of finitely many disconnected segments, and we write $\text{ilg}(\alpha|G_t)$ for the sum of the lengths of the components of $\alpha|G_t - \alpha_t^{\text{leg}}$ that contain at least $\check{m} + 1$ illegal turns (counting the endpoints). This is the *illegal length* of $\alpha|G_t$. Finally we write $\text{ntr}(\alpha|G_t)$ for the sum of the lengths of the remaining components of $\alpha|G_t - \alpha_t^{\text{leg}}$, that is, those components with fewer than $\check{m} + 1$ illegal turns. This is the *neutral length* of $\alpha|G_t$. By construction we thus have

$$\ell(\alpha|G_t) = \text{leg}(\alpha|G_t) + \text{ilg}(\alpha|G_t) + \text{ntr}(\alpha|G_t).$$

Notice that, since every component of α_t^{leg} has length at least 3, there are at most $(\text{leg}(\alpha|G_t)/3) + 1$ components of $\alpha|G_t - \alpha_t^{\text{leg}}$. On the other hand, each component contributing to $\text{ntr}(\alpha|G_t)$ has length at most $3\check{m}$ by definition, and so we find that

$$\text{ntr}(\alpha|G_t) \leq \check{m}(\text{leg}(\alpha|G_t) + 3).$$

The previous two lemmas allow us to show that the illegal length of $\alpha|G_t$ decreases exponentially fast along a folding path that remains close to the orbit of Γ .

Lemma 6.9 (illegal turn mortality rate) *Suppose that Γ is purely hyperbolic and that $\gamma: [a, b] \rightarrow \mathcal{X}$ is a folding path with $G_t = \gamma(t) \in \mathcal{N}_{A_0}(\Gamma \cdot R)$ for all t . Then for every nontrivial conjugacy class α we have*

$$\text{ilg}(\alpha|\gamma(a)) \geq \frac{\epsilon_0 \check{m}}{3(2r-2)(2\check{m}+1)} \left(\frac{2\check{m}+1}{2\check{m}} \right)^{(b-a)/D_3} \cdot \text{ilg}(\alpha|\gamma(b)),$$

where ϵ_0 is the minimal injectivity radius of any point in $\mathcal{N}_{A_0}(\Gamma \cdot R)$, $r = \text{rk}(\mathbb{F})$, and D_3 is the constant from Lemma 6.7.

Proof Let p_b be a component of $\alpha|G_b - \alpha_b^{\text{leg}}$ contributing to $\text{ilg}(\alpha|G_b)$, and write p_t for the corresponding path in $\alpha|G_t$ (ie $p_{t'}$ unfolds to p_t for $t \leq t'$). First note that for $t \in [a, b]$, the hypotheses on p_b imply that every legal subsegment of p_t has length less than 3 (since legal segments of length at least 3 grow under folding) and the number of illegal turns in p_t is at least $\check{m} + 1$ (since $k(p_t)$ is nonincreasing in t).

Suppose that $t \in [a, b]$ is such that $t - D_3 \in [a, b]$. Partition p_t into $s + 1$ subpaths

$$p_t = p_t^1 \cdots p_t^s \cdot q_t,$$

where each p_t^i has $\check{m} + 1$ illegal turns and q_t has fewer than $\check{m} + 1$ illegal turns (counting endpoints). Thus the number of illegal turns in p_t is $k(p_t) = s\check{m} + k(q_t)$, where in the case that q_t is degenerate we view it as a segment with one illegal turn so that $k(q_t) = 1$. By our condition on q_t and the assumption that $k(p_t) \geq (\check{m} + 1)$, it follows that

$$\frac{k(p_t)}{\check{m}} \leq s + 1 \quad \text{and} \quad 1 \leq \frac{k(p_t) - 1}{\check{m}} \leq \frac{k(p_t)}{2\check{m}}.$$

Unfolding these p_t^i to subsegments of p_{t-D_3} and applying Lemma 6.7, we conclude that the number of illegal turns in each subsegment increases by at least 1. Thus

$$\begin{aligned} k(p_{t-D_3}) &\geq s(\check{m} + 1) + k(q_t) = k(p_t) + s \\ &\geq k(p_t) + \frac{k(p_t)}{\check{m}} - 1 \\ &\geq k(p_t) + \frac{k(p_t)}{2\check{m}} \\ &= \left(\frac{2\check{m}+1}{2\check{m}} \right) k(p_t). \end{aligned}$$

So long as $a \leq t - nD_3 \leq b$, we may inductively apply this argument to conclude that

$$k(p_a) \geq k(p_{t-nD_3}) \geq \left(\frac{2\check{m}+1}{2\check{m}}\right)^n k(p_t).$$

Using Lemma 6.8 to compare lengths with number of illegal turns, we conclude that

$$\begin{aligned} \ell(p_a) &\geq \left(\frac{\epsilon_0}{4r-4}\right)k(p_a) \\ &\geq \left(\frac{\epsilon_0}{4r-4}\right)\left(\frac{2\check{m}+1}{2\check{m}}\right)^{\lfloor (b-a)/D_3 \rfloor} k(p_t) \\ &\geq \left(\frac{\epsilon_0\check{m}}{3(2r-2)(2\check{m}+1)}\right)\left(\frac{2\check{m}+1}{2\check{m}}\right)^{(b-a)/D_3} \ell(p_b). \end{aligned}$$

Summing these estimates over each component of $\alpha|G_b$ contributing to $\text{ilg}(\alpha|G_b)$ gives the desired result. \square

There is a similar estimate for the growth of legal length in the forward direction.

Lemma 6.10 *For any folding path G_t , $t \in [a, b]$, every nontrivial conjugacy class $\alpha \in \mathbb{F}$ satisfies*

$$\text{leg}(\alpha|G_b) \geq \text{leg}(\alpha|G_a)\left(\frac{1}{3}\right)e^{b-a}.$$

Proof Let p_a be a component of α_a^{leg} and let p_b be the corresponding segment in $\alpha|G_b$ (so that p_b unfolds to p_a). Then $\ell(p_a) \geq 3$ by assumption, so Corollary 4.8 of [5] gives

$$\ell(p_b) \geq 2 + (\ell(p_a) - 2)e^{b-a} \geq \ell(p_a)\left(1 - \frac{2}{\ell(p_a)}\right)e^{b-a} \geq \frac{1}{3}\ell(p_a)e^{b-a}.$$

Summing over the segments contributing to $\text{leg}(\alpha|G_a)$ now proves the claim. \square

Combining these estimates easily leads to uniform flaring along folding paths that stay close to the orbit $\Gamma \cdot R$:

Proposition 6.11 (flaring in folding paths) *Suppose $\Gamma \leq \text{Out}(\mathbb{F})$ is purely hyperbolic and that $R \in \mathcal{X}$ is such that $\Gamma \cdot R$ is A_0 -QCX. Then for all $\lambda_1 \geq 1$ there exists $D_1 \geq 1$ such that the following holds: for any nontrivial conjugacy class α , any folding path $\gamma: I \rightarrow \mathcal{X}$ with $G_s = \gamma(s) \in \mathcal{N}_{A_0}(\Gamma \cdot R)$, and any parameters $t \in \mathbb{R}$ and $d \geq D_1$ satisfying $[t - d, t + d] \subset I$, we have*

$$\lambda_1 \cdot \ell(\alpha|G_t) \leq \max\{\ell(\alpha|G_{t-d}), \ell(\alpha|G_{t+d})\}.$$

Proof Fix $t \in I$.

Case 1 Suppose $\text{ilg}(\alpha|G_t) \geq \ell(\alpha|G_t)/2$. Then Lemma 6.9 provides a constant D' such that for all $d \geq D'$ with $[t - d, t] \subset \mathbf{I}$ we have

$$\ell(\alpha|G_{t-d}) \geq \text{ilg}(\alpha|G_{t-d}) \geq 2\lambda_1 \cdot \text{ilg}(\alpha|G_t) \geq \lambda_1 \cdot \ell(\alpha|G_t).$$

Case 2 Suppose $\text{ilg}(\alpha|G_t) < \ell(\alpha|G_t)/2$ and $\text{leg}(\alpha|G_t) \neq 0$. In this case we have

$$\begin{aligned} \ell(\alpha|G_t) &= \text{ilg}(\alpha|G_t) + \text{leg}(\alpha|G_t) + \text{ntr}(\alpha|G_t) \\ &\leq \frac{1}{2}\ell(\alpha|G_t) + \text{leg}(\alpha|G_t) + \check{m}(\text{leg}(\alpha|G_t) + 3), \end{aligned}$$

which gives $\ell(\alpha|G_t) < 2(1 + \check{m}) \text{leg}(\alpha|G_t) + 6$. Note that $3 \leq \text{leg}(\alpha|G_t)$ by definition of (nonzero) legal length. Lemma 6.10 now provides a constant D'' such that for all $d \geq D''$ with $[t, t + d] \subset \mathbf{I}$ we similarly have

$$\ell(\alpha|G_{t+d}) \geq \text{leg}(\alpha|G_{t+d}) \geq \lambda_1 4(1 + \check{m}) \text{leg}(\alpha|G_t) \geq \lambda_1 \cdot \ell(\alpha|G_t).$$

Case 3 Suppose $\text{ilg}(\alpha|G_t) < \ell(\alpha|G_t)/2$ and $\text{leg}(\alpha|G_t) = 0$. Then the above shows $\ell(\alpha|G_t) \leq 6$. Thus by Lemma 6.6, applied with $L_0 = 6\lambda_1$ shows that there exists a constant D_0 such that for all $d \geq D_0$ we have $\ell(\alpha|G_{t+d}) > L_0 \geq \lambda_1 \ell(\alpha|G_t)$. \square

We are now prepared to prove the main result of this section:

Proof of Theorem 6.5 Fix a finite generating set $S \subset \Gamma$ and a free basis X of \mathbb{F} . We must produce $M \in \mathbb{N}$ such that for every nontrivial $\alpha \in \mathbb{F}$ and all $g_1, g_2 \in \Gamma$ with $|g_i|_S \geq M$ and $|g_1 g_2| = |g_1|_S + |g_2|_S$ we have

$$2\|\alpha\|_X \leq \max\{\|g_1(\alpha)\|_X, \|g_2^{-1}(\alpha)\|_X\}.$$

Recall first that, by Proposition 2.5, there exists a constant $K = K(X, R)$ such that $(1/K)\|\alpha\|_X \leq \ell(\alpha|R) \leq K\|\alpha\|_X$ for every conjugacy class α in \mathbb{F} . We apply Proposition 6.11 with $\lambda_1 = 2K^2 e^{2A}$ and obtain a corresponding constant D_1 . By Lemma 6.4, we know that $g \mapsto g \cdot R$ defines a quasi-isometric embedding of (Γ, d_Γ) into $(\mathcal{X}, d_\mathcal{X})$. Thus we may choose $M \in \mathbb{N}$ sufficiently large that every $g \in \Gamma$ with $|g|_S \geq M$ satisfies $d_\mathcal{X}(R, g \cdot R) \geq D_1 + 2A$. We claim Γ has $(2, M)$ -conjugacy flaring.

Let $g_1, g_2 \in \Gamma$ be any elements with $|g_i|_S \geq M$ and $|g_1 g_2|_S = |g_1|_S + |g_2|_S$. It follows that there exists a geodesic $(h_{-k}, \dots, h_{-1}, e, h_1, \dots, h_j)$ in Γ with $h_{-k} = g_1^{-1}$ and $h_j = g_2$. In particular, $k = |g_1|_S$ and $j = |g_2|_S$. Since $\Gamma \cdot R$ is A -QCX by hypothesis, there exists a folding path $\gamma: \mathbf{I} \rightarrow \mathcal{X}$ that has Hausdorff distance at most A from the image of (h_{-k}, \dots, h_j) . Writing $G_t = \gamma(t)$, we may thus choose times $a < s < b$ in \mathbf{I} so that

$$(5) \quad d_{\mathcal{X}}^{\text{sym}}(G_a, h_{-k} \cdot R) \leq A, \quad d_{\mathcal{X}}^{\text{sym}}(G_s, R) \leq A, \quad d_{\mathcal{X}}^{\text{sym}}(G_b, h_j \cdot R) \leq A.$$

Since $|g_1^{-1}|_S, |g_2|_S \geq M$, the above remarks imply that

$$d_X(h_{-k} \cdot R, R) = d_X(R, g_1 \cdot R) \quad \text{and} \quad d_X(R, h_j \cdot R) = d_X(R, g_2 \cdot R)$$

are both bounded below by $D_1 + 2A$. Thus by the triangle inequality we have

$$d_X(G_a, G_s) \geq D_1 \quad \text{and} \quad d_X(G_s, G_b) \geq D_1,$$

which is equivalent to $s - a \geq D_1$ and $b - s \geq D_1$. Since the folding path $\gamma(I)$ lies in $\mathcal{N}_A(\Gamma \cdot R)$ and the orbit $\Gamma \cdot R$ is A -QCX by assumption, Proposition 6.11 now ensures that

$$2K^2 e^{2A} \cdot \ell(\alpha|_{G_s}) \leq \max\{\ell(\alpha|_{G_a}), \ell(\alpha|_{G_b})\}$$

for every nontrivial $\alpha \in \mathbb{F}$. Finally, since $d_X \leq d_X^{\text{sym}}$, (5) implies that

$$\ell(\alpha|_{G_a}) \leq e^A \ell(\alpha|_{g_1^{-1} \cdot R}), \quad \ell(\alpha|_{G_b}) \leq e^A \ell(\alpha|_{g_2 \cdot R}), \quad \ell(\alpha|_R) \leq e^A \ell(\alpha|_{G_s}).$$

Combining the above estimates and using the rule $\ell(\alpha|_{g^{-1} \cdot R}) = \ell(g(\alpha)|_R)$ yields

$$\begin{aligned} 2\|\alpha\|_X &\leq 2K \ell(\alpha|_R) \\ &\leq \frac{1}{K} e^{-A} \max\{\ell(\alpha|_{G_a}), \ell(\alpha|_{G_b})\} \\ &\leq \frac{1}{K} \max\{\ell(g_1(\alpha)|_R), \ell(g_2^{-1}(\alpha)|_R)\} \\ &\leq \max\{\|g_1(\alpha)\|_X, \|g_2^{-1}(\alpha)\|_X\}. \end{aligned}$$

Since this holds for every nontrivial $\alpha \in \mathbb{F}$, we have proved the claim. □

7 The Cayley graph bundle of a free group extension

Fix $\Gamma \leq \text{Out}(\mathbb{F})$ with finite generating set $S = \{s_1, \dots, s_n\}$, and fix a free basis $X = \{x_1, \dots, x_r\}$ for \mathbb{F} . Recalling that the extension E_Γ is naturally a subgroup of $\text{Aut}(\mathbb{F})$, choose lifts $t_i \in \text{Aut}(\mathbb{F})$ of s_i for each $1 \leq i \leq n$ so that E_Γ is generated as a subgroup of $\text{Aut}(\mathbb{F})$ by $W = \{i_{x_1}, \dots, i_{x_r}, t_1, \dots, t_n\}$. That is,

$$E_\Gamma = \langle i_{x_1}, \dots, i_{x_r}, t_1, \dots, t_n \rangle \leq \text{Aut}(\mathbb{F}).$$

Here, i_x is the inner automorphism given by conjugation by $x \in \mathbb{F}$, ie $i_x(\alpha) = x\alpha x^{-1}$ for $\alpha \in \mathbb{F}$. Note that by construction,

$$t_i x t_i^{-1} = i_{t(x)} \in \text{Aut}(\mathbb{F})$$

for each $x \in \mathbb{F}$ and each $t \in \text{Aut}(\mathbb{F})$. For convenience, set $\hat{X} = \{i_{x_1}, \dots, i_{x_r}\}$ and $\hat{\mathbb{F}} = \langle \hat{X} \rangle$, so that $\hat{\mathbb{F}}$ is the image of \mathbb{F} in $\text{Aut}(\mathbb{F})$. Note that $\hat{\mathbb{F}}$ is also the kernel

of the homomorphism $E_\Gamma \rightarrow \Gamma$. In general, for $g \in \Gamma$ we denote a lift of g to an automorphism in the extension E_Γ by \tilde{g} .

Let $T = \text{Cay}(\mathbb{F}, X)$, $\mathcal{E} = \text{Cay}(E_\Gamma, W)$ and $\mathcal{B} = \text{Cay}(\Gamma, S)$, where $\text{Cay}(\cdot, \cdot)$ denotes the Cayley graph with the specified generating set equipped with the path metric in which each edge has length one. Set \mathcal{R} to be the standard rose on the generating set X , so that $\mathcal{R} = T/\mathbb{F}$. There is an obvious equivariant simplicial map

$$p: \mathcal{E} \rightarrow \mathcal{B}$$

obtained from the surjective homomorphism $E_\Gamma \rightarrow \Gamma$. In detail, $p: \mathcal{E} \rightarrow \mathcal{B}$ is defined to be the homomorphism $E_\Gamma \rightarrow \Gamma$ on the vertices of \mathcal{E} and maps edges of \mathcal{E} to either vertices or edges of \mathcal{B} , depending on whether the given edge corresponds to a generator in X or S , respectively. Note that for each $b \in \Gamma$, the preimage $T_b = p^{-1}(b)$ is the simplicial tree (isomorphic to T) with vertices labeled by the coset $\tilde{b}\hat{\mathbb{F}}$ (\tilde{b} any lift of b) and edges labeled by \hat{X} . We write d_b for the induced path metric on the fiber T_b over $b \in \Gamma$.

In Example 1.8 of [28], it is verified that $p: \mathcal{E} \rightarrow \mathcal{B}$ is a metric graph bundle. We provide the details here for completeness. We first make the following observation.

Lemma 7.1 *Let g_1, g_2 be vertices of $p^{-1}(b) = T_b$. Then $g_1^{-1}g_2 \in \hat{\mathbb{F}} \cong \mathbb{F}$ and*

$$d_b(g_1, g_2) = |g_1^{-1}g_2|_{\hat{X}}.$$

Proof Since T_b is a graph (it is a tree), $d_b(g_1, g_2)$ counts the minimal number of edges traversed by any path from g_1 to g_2 that remains in T_b . Such a path consists of edges labeled by generators in W coming from \hat{X} . As any such path represents $g_1^{-1}g_2$, we have $g_1^{-1}g_2 \in \hat{\mathbb{F}}$ and $d_b(g_1, g_2) \geq |g_1^{-1}g_2|_{\hat{X}}$. Conversely, writing $g_1^{-1}g_2$ in terms of $\{i_{x_1}^\pm, \dots, i_{x_r}^\pm\}$ produces a path in T_b from g_1 to g_2 . Thus $d_b(g_1, g_2) \leq |g_1^{-1}g_2|_{\hat{X}}$. \square

Lemma 7.2 *The equivariant map of Cayley graphs $p: \mathcal{E} \rightarrow \mathcal{B}$ is a metric graph bundle.*

Proof For each $n \in \mathbb{N}$, the n -ball $\{g \in E_\Gamma : |g|_W \leq n\}$ is finite. Thus we may define the properness function $f: \mathbb{N} \rightarrow \mathbb{N}$ by setting $f(n) = \max\{|i_\alpha|_{\hat{X}} : i_\alpha \in \hat{\mathbb{F}} \text{ and } |i_\alpha|_W \leq n\}$. Then for any $b \in \Gamma$ and any g_1, g_2 in $T_b = p^{-1}(b)$, Lemma 7.1 implies that

$$d_b(g_1, g_2) = |g_1^{-1}g_2|_{\hat{X}} \leq f(|g_1^{-1}g_2|_W) = f(d_{\mathcal{E}}(g_1, g_2)),$$

as required. Lastly, suppose $b_1, b_2 \in \mathcal{B}$ are adjacent vertices and that $g_1 \in T_{b_1}$ is any vertex over b_1 . Then $b_2 = b_1s$ for some $s \in S$. If $t \in W$ is the chosen lift of s , then g_1t is adjacent to g_1 in \mathcal{E} and satisfies $p(g_1t) = b_1s = b_2$, as desired. This completes the proof that $p: \mathcal{E} \rightarrow \mathcal{B}$ is a metric graph bundle. \square

Using our choice of generators in W , we may define canonical lifts of paths in \mathcal{B} through any particular point in a fiber. For $N \in \mathbb{N}$, let $\gamma: [-N, N] \rightarrow \mathcal{B}$ be any edge path in \mathcal{B} (by which we mean a path that maps each integer j to a vertex and each intervening interval $[j, j + 1]$ isometrically onto an edge) and let $\tilde{\gamma}(0)$ be any vertex in the fiber $T_{\gamma(0)}$. For each integer $-N \leq j < N$, the product $s_j = \gamma(j)^{-1}\gamma(j + 1)$ then lies in the generating set S , and we let t_j be the chosen lift of s_j to W . Thus for $j > 0$ we have $\gamma(j) = \gamma(0)s_0 \cdots s_{j-1}$ and $\gamma(-j) = \gamma(0)s_{-1}^{-1} \cdots s_{-j}^{-1}$. Accordingly, the canonical lift of γ through $\tilde{\gamma}(0) \in T_{\gamma(0)}$ is defined to be the edge path $\tilde{\gamma}: [-N, N] \rightarrow \mathcal{E}$ given by

$$\tilde{\gamma}(j) = \tilde{\gamma}(0)t_0 \cdots t_{j-1} \quad \text{and} \quad \tilde{\gamma}(-j) = \tilde{\gamma}(0)t_{-1}^{-1} \cdots t_{-j}^{-1}$$

for each integer $0 \leq j \leq N$. Observe that $p(\tilde{\gamma}(j)) = \gamma(j)$, so that $\tilde{\gamma}$ is in fact a lift of γ . Moreover, since $p: \mathcal{E} \rightarrow \mathcal{B}$ is 1-Lipschitz, when the original path $\gamma: [-N, N] \rightarrow \mathcal{B}$ is a geodesic, so is the canonical lift of γ through any point in $T_{\gamma(0)}$. These lifts will be instrumental in establishing the flaring property for the metric graph bundle $\mathcal{E} \rightarrow \mathcal{B}$, which we do in Proposition 8.1.

8 Conjugacy flaring implies hyperbolicity of E_Γ

In this section we complete the proof of our main theorem and show that the \mathbb{F} -extension group E_Γ is hyperbolic when $\Gamma \leq \text{Out}(\mathbb{F})$ is purely hyperbolic and qi-embeds into the factor complex \mathcal{F} . We first show that conjugacy flaring for the group Γ implies that the metric bundle $\mathcal{E} \rightarrow \mathcal{B}$ defined in Section 7 has the flaring property. Combining with Theorem 2.2, this will show that \mathcal{E} , and consequently E_Γ , is hyperbolic.

Proposition 8.1 (conjugacy flaring implies the flaring property) *Suppose a finitely generated subgroup $\Gamma \leq \text{Out}(\mathbb{F})$ satisfies (λ, N) -conjugacy flaring for some $\lambda > 1$ and $N \in \mathbb{N}$. Then the corresponding metric graph bundle $p: \mathcal{E} \rightarrow \mathcal{B}$ satisfies the flaring condition.*

Proof By hypothesis, there is a finite generating set $S = \{s_1, \dots, s_n\}$ of Γ and a free basis $X = \{x_1, \dots, x_r\}$ of \mathbb{F} with respect to which Γ has (λ, N) -conjugacy flaring (see Section 6). As in Section 7, we then consider the generating set $W = \{i_{x_1}, \dots, i_{x_r}, t_1, \dots, t_n\}$ of E_Γ , where t_i denotes a chosen lift of s_i , and the natural simplicial surjection $p: \mathcal{E} \rightarrow \mathcal{B}$, where $\mathcal{E} = \text{Cay}(E_\Gamma, W)$ and $\mathcal{B} = \text{Cay}(\Gamma, S)$. As before, set \hat{X} equal to the subset of the generators of W coming from X and denote the isomorphic image of \mathbb{F} in E_Γ by $\hat{\mathbb{F}} = \langle \hat{X} \rangle$.

To establish the flaring property, we must show that for every $k \geq 1$ there exists $\lambda_k > 1$ and $n_k, M_k \in \mathbb{N}$ such that for any geodesic $\gamma: [-n_k, n_k] \rightarrow \mathcal{B}$ and any two k -qi lifts

$\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ satisfying $d_{\gamma(0)}(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) \geq M_k$ we have

$$\lambda_k \cdot d_{\gamma(0)}(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) \leq \max\{d_{\gamma(n_k)}(\tilde{\gamma}_1(n_k), \tilde{\gamma}_2(n_k)), d_{\gamma(-n_k)}(\tilde{\gamma}_1(-n_k), \tilde{\gamma}_2(-n_k))\}.$$

In fact, we show that in terms of the given conjugacy flaring constants (λ, N) we may take $\lambda_k = (\lambda + 1)/2$ and $n_k = N$ (each independent of k) so that given any $k \geq 1$, if

$$M_k = 2(\lambda + 2e_k)/(\lambda - 1),$$

then the flaring condition holds with these constants. Here $e_k = f(N + 1 + kN + k)$, where $f(\cdot)$ is the properness function for the bundle $\mathcal{E} \rightarrow \mathcal{B}$.

Let $\gamma: [-N, N] \rightarrow \mathcal{B}$ be a geodesic and set $b = \gamma(0)$. Suppose that two k -qi lifts $\tilde{\gamma}_1, \tilde{\gamma}_2: [-N, N] \rightarrow \mathcal{E}$ are given. (Hence, $p(\tilde{\gamma}_i(j)) = \gamma(j)$ for $i = 1, 2$ and each integer j .) Recall from Section 7 that $T_{\gamma(j)} = p^{-1}(\gamma(j))$ is a simplicial tree with edges labeled by the free basis \hat{X} of $\hat{\mathbb{F}}$. With respect to this basis, the element $\tilde{\gamma}_1(0)^{-1}\tilde{\gamma}_2(0) \in \hat{\mathbb{F}}$ may not be cyclically reduced. However, there is some $x \in \hat{X}$ such that $i_\alpha = \tilde{\gamma}_1(0)^{-1}\tilde{\gamma}_2(0)x \in \hat{\mathbb{F}}$ is cyclically reduced. Then i_α has the property that $\|i_\alpha\|_{\hat{X}} = |i_\alpha|_{\hat{X}}$ and that $|i_\alpha|_{\hat{X}}$ differs from $d_b(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) = |\tilde{\gamma}_1(0)^{-1}\tilde{\gamma}_2(0)|_{\hat{X}}$ by at most 1. Set $z_1 = \tilde{\gamma}_1(0)$ and $z_2 = \tilde{\gamma}_2(0)x \in T_b$, so that by construction,

$$z_1 i_\alpha = z_2.$$

For each integer $-N \leq j < N$, let us set $s_j = \gamma(j)^{-1}\gamma(j + 1) \in S$. Since γ is a geodesic, the products $g = s_{-N} \cdots s_{-1} \in \Gamma$ and $h = s_0 \cdots s_{N-1} \in \Gamma$ satisfy $|g|_S = |h|_S = N$ and $|gh|_S = |g|_S + |h|_S$. Thus (λ, N) -conjugacy flaring implies

$$\begin{aligned} \max\{\|g(\alpha)\|_X, \|h^{-1}(\alpha)\|_X\} &\geq \lambda \cdot \|\alpha\|_X \\ &= \lambda \cdot |\alpha|_X \\ &= \lambda \cdot |i_\alpha|_{\hat{X}} \\ &\geq \lambda \cdot (d_b(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) - 1). \end{aligned}$$

Let $\tilde{\gamma}_{z_1}, \tilde{\gamma}_{z_2}: [-N, N] \rightarrow \mathcal{E}$ be the canonical (geodesic) lifts of $\gamma: [-N, N] \rightarrow \mathcal{B}$ through the points z_1 and z_2 , respectively. Let us also write $\tilde{g} = t_{-N} \cdots t_{-1}$ and $\tilde{h} = t_0 \cdots t_{N-1}$, where t_i is the chosen lift of $s_i \in S$ in the generating set W of E_Γ . By construction, \tilde{g} and \tilde{h} are also lifts of $g, h \in \Gamma \leq \text{Out}(\mathbb{F})$ to $E_\Gamma \leq \text{Aut}(\mathbb{F})$. Recall that the canonical lifts $\tilde{\gamma}_{z_j}$ are defined so that

$$\tilde{\gamma}_{z_j}(-N) = z_j t_{-1}^{-1} \cdots t_{-N}^{-1} = z_j \tilde{g}^{-1} \quad \text{and} \quad \tilde{\gamma}_{z_j}(N) = z_j t_0 \cdots t_{N-1} = z_j \tilde{h}$$

for $j = 1, 2$. Therefore

$$\begin{aligned} \tilde{\gamma}_{z_1}(-N)^{-1}\tilde{\gamma}_{z_2}(-N) &= (\tilde{g}z_1^{-1})(z_2\tilde{g}^{-1}) = \tilde{g}i_\alpha\tilde{g}^{-1} = i_{\tilde{g}(\alpha)}, \\ \tilde{\gamma}_{z_1}(N)^{-1}\tilde{\gamma}_{z_2}(N) &= (\tilde{h}^{-1}z_1^{-1})(z_2\tilde{h}) = \tilde{h}^{-1}i_\alpha\tilde{h} = i_{\tilde{h}^{-1}(\alpha)}. \end{aligned}$$

Hence, the endpoints of our canonical lifts of γ satisfy

$$d_{\gamma(-N)}(\tilde{\gamma}_{z_1}(-N), \tilde{\gamma}_{z_2}(-N)) = |\tilde{\gamma}_{z_1}(-N)^{-1} \tilde{\gamma}_{z_2}(-N)|_{\hat{X}} = |i_{\tilde{g}(\alpha)}|_{\hat{X}} = |\tilde{g}(\alpha)|_X \geq \|g(\alpha)\|_X$$

and

$$d_{\gamma(N)}(\tilde{\gamma}_{z_1}(N), \tilde{\gamma}_{z_2}(N)) = |\tilde{\gamma}_{z_1}(N)^{-1} \tilde{\gamma}_{z_2}(N)|_{\hat{X}} = |i_{\tilde{h}^{-1}(\alpha)}|_{\hat{X}} = |\tilde{h}^{-1}(\alpha)|_X \geq \|h^{-1}(\alpha)\|_X.$$

In light of conjugacy flaring, it follows that we have

$$\max\{d_{\gamma(-N)}(\tilde{\gamma}_{z_1}(-N), \tilde{\gamma}_{z_2}(-N)), d_{\gamma(N)}(\tilde{\gamma}_{z_1}(N), \tilde{\gamma}_{z_2}(N))\} \geq \lambda \cdot (d_b(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) - 1).$$

Let us now estimate the distances between our canonical lifts $\tilde{\gamma}_{z_j}$ and the given lifts $\tilde{\gamma}_j$ of γ . By metric properness, for $j = 1, 2$ we have

$$\begin{aligned} d_{\gamma(N)}(\tilde{\gamma}_{z_j}(N), \tilde{\gamma}_j(N)) &\leq f(d_{\mathcal{E}}(\tilde{\gamma}_{z_j}(N), \tilde{\gamma}_j(N))) \\ &\leq f(d_{\mathcal{E}}(\tilde{\gamma}_{z_j}(N), \tilde{\gamma}_{z_j}(0)) + d_{\mathcal{E}}(\tilde{\gamma}_{z_j}(0), \tilde{\gamma}_j(0)) + d_{\mathcal{E}}(\tilde{\gamma}_j(0), \tilde{\gamma}_j(N))) \\ &\leq f(N + 1 + kN + k) = e_k. \end{aligned}$$

We similarly have $d_{\gamma(-N)}(\tilde{\gamma}_{z_j}(-N), \tilde{\gamma}_j(-N)) \leq e_k$ for $j = 1, 2$. The triangle inequality thus gives

$$d_{\gamma(*)}(\tilde{\gamma}_1(*), \tilde{\gamma}_2(*)) \geq d_*(\tilde{\gamma}_{z_1}(*), \tilde{\gamma}_{z_2}(*)) - 2e_k \quad \text{for } * = \pm N.$$

Combining with our above estimate, it follows that the given lifts $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ satisfy

$$\max\{d_{\gamma(-N)}(\tilde{\gamma}_1(-N), \tilde{\gamma}_2(-N)), d_{\gamma(N)}(\tilde{\gamma}_1(N), \tilde{\gamma}_2(N))\} \geq \lambda \cdot d_b(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) - \lambda - 2e_k.$$

Therefore whenever $d_b(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) \geq M_k = 2(\lambda + 2e_k)/(\lambda - 1)$, so that

$$\begin{aligned} \lambda \cdot d_b(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) - \lambda - 2e_k &\geq \lambda \cdot d_b(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) - \frac{1}{2}(\lambda - 1)d_b(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)) \\ &= \frac{1}{2}(\lambda + 1)d_b(\tilde{\gamma}_1(0), \tilde{\gamma}_2(0)), \end{aligned}$$

we obtain the inequality required by the flaring property. This completes the proof. \square

Theorem 8.2 (hyperbolic extensions) *Suppose that $\Gamma \leq \text{Out}(\mathbb{F})$ is purely hyperbolic and that there exists $R \in \mathcal{X}$ such that $\Gamma \cdot R$ is A-QCX. Then the corresponding extension group E_Γ is hyperbolic.*

Proof Since \mathcal{E} is the Cayley graph of E_Γ , it suffices to show that \mathcal{E} is hyperbolic. We show that the metric graph bundle $\mathcal{E} \rightarrow \mathcal{B}$ satisfies the three conditions for hyperbolicity appearing in Theorem 2.2 (the Mj–Sardar theorem). Conditions (1) and (2) are obvious since each fiber is isomorphic to the universal cover of a $\text{rk}(\mathbb{F})$ –petal rose. Since the hypotheses imply that Γ has conjugacy flaring (Theorem 6.5), condition (3) follows from Proposition 8.1. Hence, E_Γ is hyperbolic. \square

Corollary 8.3 *Suppose $\Gamma \leq \text{Out}(\mathbb{F})$ is purely hyperbolic and qi-embeds into \mathcal{F} . Then E_Γ is hyperbolic.*

Proof This follows immediately from Corollary 6.3 and Theorem 8.2. □

9 Applications

In this section, we produce examples of hyperbolic extensions of the free group \mathbb{F} using the main result of this paper. We begin by defining a version of the intersection graph \mathcal{I} for \mathbb{F} , which is an $\text{Out}(\mathbb{F})$ -graph introduced by Kapovich and Lustig in [22]. First, let \mathcal{I}' be the graph whose vertices are conjugacy classes of \mathbb{F} and where two vertices are joined by an edge if there is a very small simplicial tree $\mathbb{F} \curvearrowright T$ in which each conjugacy class fixes a point. (Recall that a simplicial tree is very small if edge stabilizers are maximal cyclic and tripod stabilizers are trivial.) Define \mathcal{I} to be the connected component of \mathcal{I}' that contains the primitive conjugacy classes. We note that there is a coarsely Lipschitz surjective map $\Theta: \mathcal{F} \rightarrow \mathcal{I}$ given by mapping the free factor A to the set of primitive conjugacy classes that are contained in A . Note that $\Theta: \mathcal{F} \rightarrow \mathcal{I}$ is $\text{Out}(\mathbb{F})$ -equivariant.

Second, recall that the action of a group G on a hyperbolic metric space X is *WPD* if G is not virtually cyclic and for every $g \in G$ with positive translation length on X , the following property holds: for every $R \geq 0$ and every $x \in X$ there is an $N \geq 1$ such that the set

$$\{\phi \in G : d_X(x, \phi(x)) \leq R \text{ and } d_X(g^N(x), \phi(g^N(x))) \leq R\}$$

is finite. It is further required that the group G contains an element that acts with positive translation length on X . This property was first defined by Bestvina and Fujiwara in [8], where it was shown that the action of the mapping class group on the curve complex is WPD. The following theorem was communicated to us by Patrick Reynolds. For complete proofs see Mann [26] and [15, Theorem 4.2, Proposition 4.4].

Theorem 9.1 (Mann and Reynolds) *The graph \mathcal{I} is hyperbolic and $f \in \text{Out}(\mathbb{F})$ acts with positive translation length on \mathcal{I} if and only if f is atoroidal and fully irreducible. Moreover the action $\text{Out}(\mathbb{F}) \curvearrowright \mathcal{I}$ is WPD.*

Following Bestvina and Fujiwara, we say that loxodromic elements $f_1, f_2 \in G$ are *independent* if their quasigeodesic axes in X do not contain rays that have finite Hausdorff distance from one another. Said differently, f_1 and f_2 are independent if they determine four distinct points on the Gromov boundary of X . The WPD

condition can be used to understand how distinct loxodromic elements can fail to be independent. In particular, Proposition 6 of [8], implies that f_1 and f_2 are independent if and only if they do not have a common power. Since Theorem 9.1 states that the action $\text{Out}(\mathbb{F}) \curvearrowright \mathcal{I}$ is WPD, two hyperbolic, fully irreducible automorphisms $f_1, f_2 \in \text{Out}(\mathbb{F})$ are independent if and only if they have no common power. Thus the notion of independence of two fully irreducibles (with respect to the action $\text{Out}(\mathbb{F}) \curvearrowright \mathcal{I}$) is intrinsic to the algebra of $\text{Out}(\mathbb{F})$.

Using Theorem 9.1, we have an (a priori weaker) version of our main theorem:

Theorem 9.2 *Let $\Gamma \leq \text{Out}(\mathbb{F})$ be a finitely generated subgroup such that some (any) orbit map into \mathcal{I} is a quasi-isometric embedding. Then the corresponding extension E_Γ is hyperbolic.*

Proof Fix $A \in \mathcal{F}$ and let $\mathcal{O}: \Gamma \rightarrow \mathcal{F}$ be the corresponding orbit map into the free factor complex. By assumption $\Theta \circ \mathcal{O}: \Gamma \rightarrow \mathcal{I}$ is a quasi-isometric embedding. Since Θ is coarsely Lipschitz, \mathcal{O} must also be a quasi-isometric embedding. Moreover, since all outer automorphisms with positive translation length of \mathcal{I} are hyperbolic, Γ must be purely hyperbolic, ie each infinite-order element is atoroidal. Now apply Corollary 8.3 to conclude that E_Γ is hyperbolic. \square

We remark that our subsequent work [15] implies that Theorem 9.2 is equivalent to our main theorem Theorem 1.1.

Our first application is a new proof of the following theorem of Bestvina, Feighn and Handel [7], where we allow for any number of hyperbolic, fully irreducible automorphisms.

Theorem 9.3 *Let $f_1, \dots, f_k \in \text{Out}(\mathbb{F})$ be a collection of pairwise independent, hyperbolic, fully irreducible outer automorphisms. Then for sufficiently large $N \geq 1$, every nonidentity element of*

$$\Gamma = \langle f_1^N, \dots, f_k^N \rangle$$

is hyperbolic and fully irreducible. Moreover, Γ is isomorphic to the free group of rank k and the extension E_Γ is hyperbolic.

Proof The proof that the subgroup quasi-isometrically embeds into \mathcal{I} follows from a standard geometric ping-pong argument for groups acting on hyperbolic spaces, exactly as in the proof of Theorem 1.4 (abundance of Schottky groups) in Kent and Leininger [25]. One can also deduce the result from [34, Lemma 3.2]. The point is that we are dealing with a collection of independent loxodromic automorphisms of a hyperbolic graph. To conclude that E_Γ is hyperbolic, apply Theorem 9.2. \square

Our next application produces the first examples, to the authors' knowledge, of hyperbolic \mathbb{F} -extensions E_Γ where Γ has torsion and is not virtually cyclic. First, for a finite group $H \leq \text{Out}(\mathbb{F})$ say that a hyperbolic, fully irreducible $f \in \text{Out}(\mathbb{F})$ is *independent for H* if f and hfh^{-1} are independent for each $h \in H$. Hence, f is independent for H if and only if $H \cap \text{comm}(f) = \emptyset$, where $\text{comm}(f)$ is the commensurator of f in $\text{Out}(\mathbb{F})$.

Theorem 9.4 *Let H be a finite subgroup of $\text{Out}(\mathbb{F})$ and let $f \in \text{Out}(\mathbb{F})$ be a hyperbolic, fully irreducible outer automorphism that is independent for H . Then for all sufficiently large $N \geq 1$, the subgroup*

$$\Gamma = \langle H, f^N \rangle$$

*is isomorphic to $H * \mathbb{Z}$ and quasi-isometrically embeds into \mathcal{I} . Therefore, the \mathbb{F} -by- $(H * \mathbb{Z})$ extension E_Γ is hyperbolic.*

Proof Fix $x \in \mathcal{I}$ and for each $h \in H$ set $f_h = hfh^{-1}$. Let $D = \max_{h \in H} d(x, hx)$. Consider the Cayley graph \mathcal{C}_h of $\langle f_h \rangle$ and the equivariant orbit map $\mathcal{C}_h \rightarrow \mathcal{I}$ obtained by mapping f_h^i to $f_h^i(hx)$ and edges to geodesic segments. Since f has positive translation length on \mathcal{I} by Theorem 9.1, the maps $\mathcal{C}_h \rightarrow \mathcal{I}$ are all K_0 -quasi-isometric embeddings (for some $K_0 \geq 1$). Let us write $\rho_h^\pm: [0, \infty) \rightarrow \mathcal{I}$ for the positive and negative subrays of $\mathcal{C}_h \rightarrow \mathcal{I}$ based at hx . Since the f_h for $h \in H$ are all pairwise independent, no distinct pair of rays in the set $\{\rho_h^+, \rho_h^-\}_{h \in H}$ have finite Hausdorff distance.

Similar to [25], we now consider the following set of paths in \mathcal{I} . For any $h_1, h_2 \in H$ and $\epsilon_1, \epsilon_2 \in \{+, -\}$ with $\rho_{h_1}^{\epsilon_1} \neq \rho_{h_2}^{\epsilon_2}$, we may build a bi-infinite path in \mathcal{I} by traversing $\rho_{h_1}^{\epsilon_1}$ with the reverse parametrization, then following a geodesic from h_1x to h_2x (which has length at most D), and lastly traversing the ray $\rho_{h_2}^{\epsilon_2}$ with the usual parametrization. As there are finitely many such paths and the chosen rays $\rho_{h_1}^{\epsilon_1}$ and $\rho_{h_2}^{\epsilon_2}$ have infinite Hausdorff distance, there exists a uniform constant $K_1 \geq 1$ such that each of these paths is a K_1 -quasigeodesic in \mathcal{I} . Subpaths of these K_1 -quasigeodesics, and their images under the isometric action of $\text{Out}(\mathbb{F})$ on \mathcal{I} , we call *f -pieces*.

Since \mathcal{I} is hyperbolic, there exist $L, K_2 \geq 1$ such that any L -local, K_1 -quasigeodesic is a K_2 -quasigeodesic [10]. In particular, if $\gamma: I \rightarrow \mathcal{I}$ is any path that agrees with some f -piece on every length- L subinterval of I , then γ is a K_2 -quasigeodesic.

Now take N to be an integer larger than L and D , and let $\theta: H * \mathbb{Z} \rightarrow \text{Out}(\mathbb{F})$ be the homomorphism that restricts to the identity on H and maps the generator t of \mathbb{Z} to f^N . Let $\Gamma = \langle H, f^N \rangle$ be the image of this homomorphism and let \mathcal{C} be the Cayley graph of $H * \mathbb{Z}$ for the generating set $\{t, h : h \in H\}$, metrized so that each

edge labeled $h \in H$ has length D and each edge labeled t has length N . We define a θ -equivariant map $\mathcal{O}: \mathcal{C} \rightarrow \mathcal{I}$ as follows: For each vertex $w \in H * \mathbb{Z}$ of \mathcal{C} , we set $\mathcal{O}(w) = \theta(w)x$. For $h \in H$, the edge in \mathcal{C} from 1 to h is mapped by \mathcal{O} to any geodesic from x to hx , and the edge in \mathcal{C} from 1 to t is mapped by \mathcal{O} to the f -piece from x to $f^N x$ using the parametrization coming from the quasigeodesic $\mathcal{C}_1 \rightarrow \mathcal{I}$. Now extend \mathcal{O} by equivariance. Observe that \mathcal{O} maps the length- $((a+b)N + D)$ path in \mathcal{C} from 1 to $t^a h t^b$ to an f -piece in \mathcal{I} , for all $a, b \in \mathbb{Z}$ and $h \in H$. Thus by construction, \mathcal{O} maps any geodesic in \mathcal{C} to a path that agrees with f -pieces on all subintervals of length at most L . Using the constant K_2 obtained above, it follows that \mathcal{O} sends every geodesic path in \mathcal{C} to a K_2 -quasigeodesic in \mathcal{I} and thus that $\mathcal{O}: \mathcal{C} \rightarrow \mathcal{I}$ is a K_2 -quasi-isometric embedding. Since the metric on \mathcal{C} differs from the word metric on $H * \mathbb{Z}$ (with our chosen generators) by a multiplicative factor of no more than N , we conclude that $\theta: H * \mathbb{Z} \rightarrow \mathcal{I}$ is an NK_2 -quasi-isometric embedding.

Finally, to see that θ is an isomorphism, note that θ itself is a quasi-isometric embedding into $\text{Out}(\mathbb{F})$. This is a simple consequence of the fact that any orbit map from $\text{Out}(\mathbb{F})$ to \mathcal{I} is coarsely Lipschitz. Hence, θ must have finite kernel. Since each finite-order $g \in H * \mathbb{Z}$ is conjugate into H , and H injects into $\Gamma \leq \text{Out}(\mathbb{F})$, we must have that $\theta: H * \mathbb{Z} \rightarrow \Gamma$ is an isomorphism. Since Theorem 9.2 implies that E_Γ is hyperbolic, this completes the proof. □

Remark Note that for $\Gamma = \langle H, f^N \rangle \cong H * \mathbb{Z}$ as in Theorem 9.4, the subgroup

$$\Gamma_0 = \langle H, f^N H f^{-N} \rangle$$

is undistorted and isomorphic to $H * H$. Hence, the \mathbb{F} -by- $(H * H)$ extension E_{Γ_0} is also hyperbolic. In the situation of surface group extensions, Honglin Min has constructed convex cocompact subgroups of the mapping class group that are isomorphic to the free product of two finite groups [27].

Finally, we show how to construct examples of hyperbolic, fully irreducible $f \in \text{Out}(\mathbb{F})$ that are independent for a given finite group $H \leq \text{Out}(\mathbb{F})$. First, say that the finite group $H \leq \text{Out}(\mathbb{F})$ is *projectively good* if its image under the surjective homomorphism $\text{Out}(\mathbb{F}) \rightarrow \text{GL}_r(\mathbb{Z})$ does not contain $-I$ (where $r = \text{rk}(\mathbb{F})$). Note that any finite group H embeds into the outer automorphism group $\text{Out}(\mathbb{F}(H))$ with projectively good image, where $\mathbb{F}(H)$ is the free group on H . This may be achieved by using the left action of H on itself to embed H into $\text{Aut}(\mathbb{F}(H))$ as permutation automorphisms whose images in $\text{GL}_{|H|}(\mathbb{Z})$ are permutation matrices.

Example 9.5 Let H be any projectively good, finite subgroup of $\text{Out}(\mathbb{F})$ satisfying $\text{rk}(\mathbb{F}) \geq 3$. We show that there is a hyperbolic, fully irreducible $f \in \text{Out}(\mathbb{F})$ that is

independent for H . By Theorem 9.4, this shows that there is a hyperbolic group G fitting into the exact sequence $1 \rightarrow \mathbb{F} \rightarrow G \rightarrow H * \mathbb{Z} \rightarrow 1$. As any finite group embeds into the outer automorphism group of some free group with projectively good image, this shows that there exist extensions of the above form for any finite group H .

Suppose that $H \leq \text{Out}(\mathbb{F})$ is a finite, projectively good subgroup. Write $r = \text{rk}(\mathbb{F})$. As in Lemma 2.14, the restriction of the homomorphism $\text{Out}(\mathbb{F}) \rightarrow \text{GL}_r(\mathbb{Z})$ to H is injective and we identify H with its image in $\text{GL}_r(\mathbb{Z})$.

Claim 9.6 *There is a matrix $A \in \text{GL}_r(\mathbb{Z})$ such that for any $h \in H \setminus 1$, the matrices hAh^{-1} and A have no common power.*

We complete the argument before proving the claim. Let A be a matrix as in the claim. Now an application of the main result of Clay and Pettet [13] implies that there is a hyperbolic, fully irreducible outer automorphism f whose image in $\text{GL}_r(\mathbb{Z})$ is A . We then have that f is independent for the finite group H . Otherwise, there is an $h \in H \setminus 1$ and integers r, s such that $hf^r h^{-1} = f^s$. Applying the homomorphism $\text{Out}(\mathbb{F}) \rightarrow \text{GL}_r(\mathbb{Z})$ we see that this equation contradicts our choice of A . Hence, f is independent for H . To complete the example, it now suffices to prove the claim.

Proof of Claim 9.6 By assumption, the finite subgroup $H \leq \text{GL}_r(\mathbb{Z})$ does not contain $-I$. Hence, the action $H \curvearrowright \mathbb{RP}^{r-1}$ is effective and if we denote the fixed subspace of $h \in H$ by V_h , we have that $V_H = \bigcup_{h \in H \setminus 1} V_h$ is a union of positive-codimension projective hyperplanes. Hence, $\mathbb{RP}^{r-1} \setminus V_H$ is open.

Now let $B \in \text{GL}_r(\mathbb{Z})$ be the block diagonal matrix consisting of $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$ in the upper left 2×2 corner and the identity matrix in the lower right corner. The eigenvalues for B are λ , 1 and λ^{-1} , where λ is the golden ratio. Moreover, the λ -eigenspace is one-dimensional and so defines a point $[v] \in \mathbb{RP}^{r-1}$. Since $\mathbb{RP}^{r-1} \setminus V_H$ is open and every orbit of $\text{GL}_r(\mathbb{Z}) \curvearrowright \mathbb{RP}^{r-1}$ is dense [30, Lemma 8.5], there is a $C \in \text{GL}_r(\mathbb{Z})$ such that $C[v] \notin V_H$. Setting $A = CBC^{-1}$, we see that the λ -eigenspace of A is one-dimensional and is not projectively fixed by any $h \in H \setminus 1$. Hence, no power of A can equal any power of hAh^{-1} for $h \in H \setminus 1$. This completes the proof. \square

References

- [1] **Y Algom-Kfir**, *Strongly contracting geodesics in outer space*, *Geom. Topol.* 15 (2011) 2181–2233 MR
- [2] **Y Algom-Kfir**, **M Bestvina**, *Asymmetry of outer space*, *Geom. Dedicata* 156 (2012) 81–92 MR

- [3] **M Bestvina**, *A Bers-like proof of the existence of train tracks for free group automorphisms*, *Fund. Math.* 214 (2011) 1–12 MR
- [4] **M Bestvina**, **M Feighn**, *A combination theorem for negatively curved groups*, *J. Differential Geom.* 35 (1992) 85–101 MR
- [5] **M Bestvina**, **M Feighn**, *Hyperbolicity of the complex of free factors*, *Adv. Math.* 256 (2014) 104–155 MR
- [6] **M Bestvina**, **M Feighn**, *Subfactor projections*, *J. Topol.* 7 (2014) 771–804 MR
- [7] **M Bestvina**, **M Feighn**, **M Handel**, *Laminations, trees, and irreducible automorphisms of free groups*, *Geom. Funct. Anal.* 7 (1997) 215–244 MR
- [8] **M Bestvina**, **K Fujiwara**, *Bounded cohomology of subgroups of mapping class groups*, *Geom. Topol.* 6 (2002) 69–89 MR
- [9] **M Bestvina**, **M Handel**, *Train tracks and automorphisms of free groups*, *Ann. of Math.* 135 (1992) 1–51 MR
- [10] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, *Grundlehr. Math. Wissen.* 319, Springer (1999) MR
- [11] **P Brinkmann**, *Hyperbolic automorphisms of free groups*, *Geom. Funct. Anal.* 10 (2000) 1071–1089 MR
- [12] **S Buyalo**, **V Schroeder**, *Elements of asymptotic geometry*, *Eur. Math. Soc.*, Zürich (2007) MR
- [13] **M Clay**, **A Pettet**, *Current twisting and nonsingular matrices*, *Comment. Math. Helv.* 87 (2012) 385–407 MR
- [14] **M Culler**, **K Vogtmann**, *Moduli of graphs and automorphisms of free groups*, *Invent. Math.* 84 (1986) 91–119 MR
- [15] **S Dowdall**, **S J Taylor**, *The co-surface graph and the geometry of hyperbolic free group extensions*, *J. Topol.* 10 (2017) 447–482 MR
- [16] **B Farb**, **L Mosher**, *Convex cocompact subgroups of mapping class groups*, *Geom. Topol.* 6 (2002) 91–152 MR
- [17] **S Francaviglia**, **A Martino**, *Metric properties of outer space*, *Publ. Mat.* 55 (2011) 433–473 MR
- [18] **U Hamenstädt**, *Word hyperbolic extensions of surface groups*, preprint (2005) arXiv
- [19] **U Hamenstädt**, *Stability of quasi-geodesics in Teichmüller space*, *Geom. Dedicata* 146 (2010) 101–116 MR
- [20] **M Handel**, **L Mosher**, *The expansion factors of an outer automorphism and its inverse*, *Trans. Amer. Math. Soc.* 359 (2007) 3185–3208 MR
- [21] **A Hatcher**, **K Vogtmann**, *The complex of free factors of a free group*, *Quart. J. Math. Oxford Ser.* 49 (1998) 459–468 MR

- [22] **I Kapovich, M Lustig**, *Geometric intersection number and analogues of the curve complex for free groups*, *Geom. Topol.* 13 (2009) 1805–1833 MR
- [23] **I Kapovich, M Lustig**, *Ping-pong and outer space*, *J. Topol. Anal.* 2 (2010) 173–201 MR
- [24] **R P Kent, IV, C J Leininger**, *Subgroups of mapping class groups from the geometrical viewpoint*, from “In the tradition of Ahlfors–Bers, IV” (D Canary, J Gilman, J Heinonen, H Masur, editors), *Contemp. Math.* 432, Amer. Math. Soc., Providence, RI (2007) 119–141 MR
- [25] **R P Kent, IV, C J Leininger**, *Shadows of mapping class groups: capturing convex cocompactness*, *Geom. Funct. Anal.* 18 (2008) 1270–1325 MR
- [26] **B Mann**, *Some hyperbolic $\text{Out}(F_N)$ -graphs and nonunique ergodicity of very small F_N -trees*, PhD thesis, The University of Utah (2014) MR Available at <http://search.proquest.com/docview/1614969728>
- [27] **H Min**, *Hyperbolic graphs of surface groups*, *Algebr. Geom. Topol.* 11 (2011) 449–476 MR
- [28] **M Mj, P Sardar**, *A combination theorem for metric bundles*, *Geom. Funct. Anal.* 22 (2012) 1636–1707 MR
- [29] **L Mosher**, *A hyperbolic-by-hyperbolic hyperbolic group*, *Proc. Amer. Math. Soc.* 125 (1997) 3447–3455 MR
- [30] **G D Mostow**, *Strong rigidity of locally symmetric spaces*, *Annals of Mathematics Studies* 78, Princeton Univ. Press (1973) MR
- [31] **F Paulin**, *The Gromov topology on \mathbb{R} -trees*, *Topology Appl.* 32 (1989) 197–221 MR
- [32] **J R Stallings**, *Topology of finite graphs*, *Invent. Math.* 71 (1983) 551–565 MR
- [33] **S J Taylor**, *Right-angled Artin groups and $\text{Out}(\mathbb{F}_n)$, I: Quasi-isometric embeddings*, *Groups Geom. Dyn.* 9 (2015) 275–316 MR
- [34] **S J Taylor, G Tiozzo**, *Random extensions of free groups and surface groups are hyperbolic*, *Int. Math. Res. Not.* 2016 (2016) 294–310 MR
- [35] **J Väisälä**, *Gromov hyperbolic spaces*, *Expo. Math.* 23 (2005) 187–231 MR

Department of Mathematics, Vanderbilt University
Nashville, TN, United States

Department of Mathematics, Temple University
Philadelphia, PA, United States

spencer.dowdall@vanderbilt.edu, samuel.taylor@temple.edu

Proposed: Walter Neumann
Seconded: Benson Farb, Bruce Kleiner

Received: 9 March 2016
Revised: 20 October 2016