

The free splitting complex of a free group, I

Hyperbolicity

MICHAEL HANDEL

LEE MOSHER

We prove that the free splitting complex of a finite rank free group, also known as Hatcher’s sphere complex, is hyperbolic.

20F65; 57M07

Given a free group F_n of finite rank $n \geq 2$, a *free splitting* over F_n is a minimal, simplicial action of the group F_n on a simplicial tree T such that the stabilizer of each edge of T is the trivial subgroup of F_n . A free splitting is denoted $F_n \curvearrowright T$, or just T , when the group and its action are understood. Although the tree T is allowed to have vertices of valence 2, there is a unique *natural cell structure* on T the vertices of which are the points of valence at least 3. We say that T is a *k-edge free splitting* if k is the number of natural edge orbits, a number which can take on any value from 1 to $3n - 3$. The equivalence relation amongst free splittings is *conjugacy*, where two free splittings of F_n are conjugate if there exists an F_n -equivariant homeomorphism between them. See the beginning of Section 1 for the details of these definitions.

The *free splitting complex* of F_n , denoted $\mathcal{FS}(F_n)$, is a simplicial complex of dimension $3n - 4$ having a simplex $\langle T \rangle$ of dimension k for each conjugacy class of $k + 1$ -edge free splittings $F_n \curvearrowright T$. Given another free splitting $F_n \curvearrowright S$, the simplex $\langle S \rangle$ is a face of $\langle T \rangle$ if and only if there is a *collapse map* $T \mapsto S$, which collapses to a point each edge in some F -invariant set of edges of T . We write $T \succ S$ for the relation “ T collapses to S ”, and $S \prec T$ for the inverse relation “ S expands to T ”. There is a natural left action of the outer automorphism group $\text{Out}(F_n)$ on $\mathcal{FS}(F_n)$, where $\phi \in \text{Out}(F_n)$ acts on the conjugacy class of a free splitting $F_n \curvearrowright T$ by precomposing the action by an automorphism of F_n representing ϕ . The free splitting complex was introduced by Hatcher in [9] in its role as the sphere complex of a connected sum of n copies of the 3-manifold $S^2 \times S^1$. A careful construction of an isomorphism between the 1-skeletons of $\mathcal{FS}(F_n)$ and Hatcher’s sphere complex can be found in Aramayona and Souto [1], and that proof extends with little trouble to the entire complexes. In Section 1.3 we shall give a rigorous construction of the free splitting complex given purely in tree language.

The complex $\mathcal{FS}(F_n)$ is regarded as one of several $\text{Out}(F_n)$ analogues of the curve complex of a surface; another competing analogue is the free factor complex of F_n introduced by Hatcher and Vogtmann in [10]. The analogies are imperfect in each case: Hatcher and Vogtmann showed that the free factor complex, like the curve complex, has the homotopy type of a wedge of spheres of constant dimension [10]; by contrast, Hatcher showed that $\mathcal{FS}(F_n)$ is contractible [9]. On the other hand we showed in [8] that simplex stabilizers of $\mathcal{FS}(F_n)$ are all undistorted subgroups of $\text{Out}(F_n)$, just as simplex stabilizers of the curve complex of a surface are undistorted subgroups of its mapping class group; by contrast, we also showed that the simplex stabilizers of the free factor complex of F are, most of them, distorted in $\text{Out}(F_n)$.

Here is our main result, an analogue to the theorem of Masur and Minsky [15] on the hyperbolicity of the curve complex:

Main Theorem *The free splitting complex $\mathcal{FS}(F_n)$, with its geodesic simplicial metric, is Gromov hyperbolic.*

By comparison Bestvina and Feighn have proved that the free factor complex is Gromov hyperbolic [2].

In rank $n = 2$, the Main Theorem is well known, because the simplicial complex $\mathcal{FS}(F_2)$ contains the Farey graph as a coarsely dense subcomplex, and the Farey graph is quasi-isometric to an \mathbb{R} -tree and is therefore Gromov hyperbolic (see eg Manning [14, Example 5.2]).

One should contrast the Main Theorem with the result of Sabalka and Savchuk [17], which says that the “edge splitting graph” of F_n is not hyperbolic; this is the 1-dimensional subcomplex of $\mathcal{FS}(F_n)$ spanned by the 0-simplices corresponding to those 1-edge free splittings $F \curvearrowright T$ that have 2 vertex orbits. Their result has an analogue in a theorem of Schleimer [18] that on a closed, oriented surface of genus at least 3, the subcomplex of the curve complex spanned by separating curves is not hyperbolic.

The Main Theorem has spurred some recent developments. It has been applied in work of Kapovich and Rafi [12] who have used it, together with a novel projection technique, to obtain a new proof of the theorem of [2] that the free factor complex is hyperbolic. Still more recent work of Mann [13] has used the Main Theorem coupled with the projection technique of [12] to prove hyperbolicity of the cyclic splitting complex. Also recently, Hilion and Horbez [11] have given another proof of the Main Theorem from the point of view of the sphere complex, in which concepts of Stallings fold paths as used in this paper are re-interpreted in terms of sphere surgery paths.

In Part II of this work we shall determine the dynamics of the action of elements of $\text{Out}(F_n)$ on $\mathcal{FS}(F_n)$, showing in particular that $\phi \in \text{Out}(F_n)$ acts loxodromically on $\mathcal{FS}(F_n)$ if and only if, in the terminology and notation of Bestvina, Feighn and Handel [5], there exists an element Λ of the set $\mathcal{L}(\phi)$ of attracting laminations such that the free factor support of Λ is the whole group F_n .

Outline of the proof

Outside of applying the hyperbolicity axioms of Masur and Minsky, our methods of proof, although intricate, are mostly self-contained, depending on basic tools from the theory of group actions on trees including Bass–Serre theory and Stallings folds. Beyond the methods, there are important motivations coming from the proof of Masur and Minsky, in particular the definition of the projection maps that play a role in verifying the Masur–Minsky axioms.

Section 1 We give the basic concepts underlying the construction of the free splitting complex $\mathcal{FS}(F_n)$, including definitions of collapse maps. We give Gilbert Levitt’s short proof of Lemma 1.3, which contains the technical results about free splittings that are needed to verify that $\mathcal{FS}(F_n)$ is, indeed, a simplicial complex. Collapse maps are also needed to understand the first barycentric subdivision $\mathcal{FS}'(F_n)$, which is what we actually use in our proof of hyperbolicity. In brief, $\mathcal{FS}'(F_n)$ has a *vertex* for each conjugacy class of free splitting $F \curvearrowright T$, and an oriented edge for each collapse relation $T > S$. Since the composition of two collapse maps is a collapse map, the collapse relation is transitive, from which it follows that each geodesic in the 1–skeleton of $\mathcal{FS}'(F_n)$ is a “zig-zag path” that alternates between collapses and expansions.

Sections 2 and 3 Following Stallings’ method [20] as extended by Bestvina and Feighn [3], we define a system of paths in $\mathcal{FS}'(F)$ called *fold paths*. We also review the criterion for hyperbolicity due to Masur and Minsky [15], which is concerned with families of paths and projection maps to those paths that satisfy certain axioms, which we refer to as the *coarse retraction*, *coarse Lipschitz*, and *Strong Projection* axioms.

The first step of progress on the Main Theorem is the statement of Proposition 3.3, which asserts the existence of a system of projection maps, one such map from the ambient space $\mathcal{FS}'(F_n)$ to each fold path, that satisfy the Masur–Minsky axioms.

Section 4 We introduce the concept of combing of fold paths. The combing process has as input a fold path $S_0 \mapsto \cdots \mapsto S_K$ plus a single edge in $\mathcal{FS}'(F_n)$ with one endpoint S_K and opposite endpoint denoted S'_K , which can be either a collapse $S_K > S'_K$ or an expand $S_K < S'_K$. The output is a fold path (roughly speaking) from some S'_0 to S'_K that stays a uniformly bounded distance from the input path, and

which has the following rather strong asynchronous fellow traveller property: every free splitting along the input fold path from S_0 to S_K is connected by a single edge to some free splitting along the output path from S'_0 to S'_K . The result of the combing process is a *combing rectangle*, the general form of which is depicted in Figure 1. These rectangles are certain commutative diagrams of fold maps and collapse maps that can be viewed as living in the 1–skeleton of $\mathcal{FS}'(F_n)$. We use many such diagrams throughout the paper, both as formal tools and as visualization aids.

Section 4.1 contains basic definitions and properties regarding combing rectangles. In this section we also take the next step of progress in the proof of the Main Theorem, by using combing to define the system of projection maps to fold paths, and we state Proposition 4.2, which asserts that these particular projection maps satisfy the Masur–Minsky axioms. Section 4.2 contains the statements and proofs of various useful constructions of combing rectangles.

Section 5 We introduce *free splitting units* as a way of subdividing a fold path into subpaths each of which has uniformly bounded diameter in $\mathcal{FS}'(F_n)$ (see Lemma 5.11) but which nevertheless measure progress through $\mathcal{FS}'(F_n)$ (as stated later in Proposition 6.2). Section 5.1 contains important diameter bounds for subsegments of fold paths. Section 5.2 uses these diameter bounds to formulate the definition of free splitting units. Once they are defined, we are able to use the diameter bounds to quickly verify the *Coarse Retraction* axiom; see Proposition 5.9.

Section 6 We verify the *Coarse Lipschitz* and *Strong Projection* axioms, completing the proof of the Main Theorem. In this section we also verify that when a fold path is parametrized by free splitting units it becomes a quasigeodesic in $\mathcal{FS}'(F_n)$; see Proposition 6.2. See the beginning of Section 6 for a sketch of the proof of the Main Theorem.

1 The free splitting complex

We begin with some basic notations used throughout the paper.

For the rest of the paper we shall fix a free group F of finite rank at least 2.

A *graph* is a 1–dimensional simplicial complex equipped with the CW topology. A *tree* T is a contractible graph. *Simplicial maps* between graphs and trees are maps taking each vertex to a vertex, and taking each edge to a vertex or to another edge–preserving barycentric coordinates. We use $G \curvearrowright T$ to denote an action of a group G on T , which by definition is a homomorphism $G \mapsto \text{Aut}(T)$ from G to the group of simplicial automorphisms of T . The action associates to each $\gamma \in G$ a simplicial automorphism

of T denoted $x \mapsto \gamma \cdot x$, a notation that extends to subsets of T by $\gamma \cdot A = \{\gamma \cdot x \mid x \in A\}$. The *stabilizer* of a subset $A \subset T$ is the subgroup $\text{Stab}_T(A) = \{\gamma \in G \mid \gamma \cdot A = A\}$. Given two actions $G \curvearrowright S, T$, a function $f: S \rightarrow T$ is said to be *equivariant* if $f(\gamma \cdot x) = \gamma \cdot f(x)$ for all $x \in S$, $\gamma \in G$.

Given a set A and a subset $B \subset A$, we denote the set theoretic complement as $A - B$. Given a graph X and a subgraph $Y \subset X$, we denote the graph theoretic complement as $X \setminus Y$, whose topological description is the closure of $X - Y$.

1.1 Free splittings, maps, natural vertices and edges, edgelets

Recall from the introduction that a *free splitting of F* is an action $F \curvearrowright T$, where T is a tree that is not a point, the action is *minimal* meaning that there is no proper F -invariant subtree, and for every edge $e \subset T$ the subgroup $\text{Stab}_T(e)$ is trivial. We use without comment the basic fact that every homeomorphism of a tree T either fixes a point or translates along a properly embedded copy of \mathbb{R} called its *axis*, and that minimality of an action $F \curvearrowright T$ is equivalent to the statement that T is the union of the axes of the elements of F that have no fixed point in T . We also use without comment the fact that every free splitting is cocompact, that is, there is a finite number of orbits of vertices and of edges; this follows from Bass–Serre theory (Scott and Wall [19]) combined with the fact that the rank of F is finite.

Given a free splitting $F \curvearrowright T$, from Bass–Serre theory [19] it follows that the set of conjugacy classes in F of nontrivial vertex stabilizers of T forms a free factor system in the sense of [5], which means that by appropriate choice of representatives $H_1 = \text{Stab}_T(v_1), \dots, H_k = \text{Stab}_T(v_k)$ of each conjugacy class — where v_1, \dots, v_k are the corresponding vertex orbit representatives — there exists a free factorization of the form $F = H_1 * \dots * H_k * B$, with B possibly trivial. We refer to this free factor system as the *vertex group system* of $F \curvearrowright T$, and denote it $\mathcal{F}(T)$.

Given a free splitting $F \curvearrowright T$, proper discontinuity of the action is equivalent to the property that all vertex stabilizers are trivial, that is, $\mathcal{F}(T) = \emptyset$, and it is also equivalent to every vertex of T having finite valence. We will use these equivalences without comment.

Definition 1.1 (Maps between free splittings) Given free splittings $F \curvearrowright S, T$, a *map* from S to T is defined to be an F -equivariant simplicial map $f: S \rightarrow T$.

We will encounter several different kinds of maps, most commonly “collapse maps” defined in Section 1.2, “foldable maps” defined in Section 2.1, and “folds” defined in Section 2.3. The category of maps will usually suffice for much of this paper, but

we will occasionally have to consider more general equivariant continuous functions between free splittings, for example conjugacies.

We will sometimes emphasize the role of the action of F by referring to a “free splitting over F ” or a “map over F ”, and we shall use similar terminology for more complicated objects introduced later on that are built out of free splittings and maps over F .

Recall from the introduction that a *conjugacy* between free splittings $F \curvearrowright S, T$ is an equivariant homeomorphism between S and T . A conjugacy *need not be* a map as just defined, i.e., it need not take vertices to vertices or edges to edges, and even if it does it need not preserve barycentric coordinates. Notice that if one is given a map $f: S \rightarrow T$ as just defined — an equivariant simplicial map — then f is a conjugacy if and only if it is locally injective: for if f is locally injective then it is evidently injective, and it is surjective by minimality of the action $F \curvearrowright T$, and so f is a simplicial isomorphism and hence a homeomorphism.

Given a free splitting $F \curvearrowright T$, recall also from the introduction the *natural cell structure* on T , a CW structure whose 0–skeleton is the set of *natural vertices* that are the vertices of valence at least 3. Implicit in the definition of the natural cell structure is the fact that each point of T that is not a natural vertex is contained in the interior of a unique *natural edge*, which is an arc of T each of whose endpoints is a natural vertex and none of whose interior points is a natural vertex. If this fact were not true then T would contain a valence 1 vertex, violating minimality, or T would contain arbitrarily long simplicial arcs with no natural vertices. In the latter case, by cocompactness it would follow that T is homeomorphic to a line: but then either the action would be properly discontinuous implying that F has rank 1 which is a contradiction; or the kernel of the action would be a free factor of corank 1, contradicting that edge stabilizers are trivial. We have also defined the notion of a k –edge free splitting $F \curvearrowright T$ meaning that T has k orbits of natural edges; this notion is invariant under conjugacy. In terms of Bass–Serre theory [19], the number of orbits of natural vertices of a free splitting $F \curvearrowright T$ equals the number of points in the quotient graph of groups T/F which either have a nontrivial group or have valence at least 3.

The word “natural” in this context refers to naturality in the category of free splittings and conjugacies: every conjugacy is an automorphism of the natural cell structure, and in particular preserves the numbers of orbits of natural vertices and edges. On this basis one might have wished to refer to a valence 1 vertex as “natural”, were it not for the fact that T has no vertices of valence 1, by virtue of minimality of the action $F \curvearrowright T$.

Remark on terminology Outside of discussions involving natural cell structures and nonsimplicial conjugacies, we work primarily in the simplicial category: a free splitting $F \curvearrowright T$ comes equipped with a simplicial structure on the tree T , which is invariant

under the action of F ; maps between free splittings are F -equivariant simplicial maps. This will be particularly convenient when we encounter subcomplexes of the simplicial structure which are not subcomplexes of the natural cell structure, for example in the results of Sections 6.2 and 6.3 where the heart of the proof of the Main Theorem resides.

For any free splitting $F \curvearrowright T$, in order to distinguish between the natural edges of T and the edges of the given simplicial structure on T we shall refer to the latter as the *edgelets* of T . This word is meant to evoke the phenomenon that, fairly often, there are many, many, many edgelets in a single natural edge, and we often visualize the edgelets as being very, very, very tiny.

1.2 Collapse maps

In order to define the free splitting complex of F rigorously we need some preliminaries regarding collapse maps.

Given two free splittings $F \curvearrowright S, T$, a map $f: S \rightarrow T$ is called a *collapse map* if f is injective over the interior of each edgelet of T . The *collapsed subgraph* $\sigma \subset S$ is the F -equivariant subgraph that is the union of those edgelets of F that are collapsed to a vertex by the map f . We put σ into the notation by writing

$$f: S \xrightarrow{[\sigma]} T,$$

the square brackets highlighting that σ is the name of the collapsed graph, whereas the notation

$$S \xrightarrow{f} T$$

tells us the name of the collapse map f itself. Note that $\sigma \subset S$ is a *proper* subgraph, meaning that $\sigma \neq S$.

Here are some basic facts about collapse maps. Items (1) and (2) will be used without mention throughout the paper. Item (3) will be needed for the proof of Proposition 4.4.

Lemma 1.2 *For any free splittings $F \curvearrowright S, T$, any collapse map*

$$f: S \xrightarrow{[\sigma]} T,$$

and any vertex $v \in T$, the following hold:

- (1) *The subgraph $f^{-1}(v)$ is connected.*
- (2) *$f^{-1}(v)$ does not degenerate to a point if and only if it is a component of σ .*
- (3) *$f^{-1}(v)$ is the convex hull of its frontier in S .*

Proof Denote $\sigma_v = f^{-1}(v)$. Given vertices $w_1 \neq w_2 \in \sigma_v$, if the segment $[w_1, w_2]$ does not map to v then $f[w_1, w_2]$ is a nondegenerate finite tree and there must exist two edgelets in $[w_1, w_2]$ with the same image in that tree, contradicting the definition of a collapse map; this proves that σ_v is connected. If σ_v is nondegenerate, i.e., if it contains an edgelet, then, each of its edgelets being in σ , it follows by connectivity that σ_v is a subset of σ . It is moreover a maximal connected subset of σ — a component of σ — because any edgelet of S incident to a vertex of σ_v but not in σ_v does not have constant image under f and so is not contained in σ . This proves (1) and (2).

To prove (3), let Fr be the frontier of σ_v in S and let $H \subset S$ be the convex hull of Fr . By connectivity we have $H \subset \sigma_v$. If the opposite conclusion did not hold then there would be an edgelet $e \subset \sigma_v \setminus H$. Only one of its two complementary components $S \setminus e = S_0 \sqcup S_1$ can contain a point of Fr , and so up to interchanging indices we have $H \subset S_0$. Since S_1 is disjoint from Fr but contains the point $x = e \cap S_1 \subset e \subset \sigma_v$, it follows that $S_1 \subset \sigma_v \subset \sigma$. The point x is the unique frontier point of S_1 . Choose $\gamma \in F$ having an axis L contained in S_1 . Let z be the point of L closest to x . For each $y \in S \setminus S_1$, z is also the point of L closest to y , and so $\gamma(z)$ is the point of L closest to $\gamma(y)$. But $\gamma(z) \neq z$ and so $\gamma(y) \in S_1 \subset \sigma$, implying that $y \in \sigma$ and contradicting properness of σ . \square

From Lemma 1.2(1), given a collapse map

$$f: S \xrightarrow{[\sigma]} T$$

it follows that σ determines T up to simplicial conjugacy, in that the map $S \mapsto T$ induces a simplicial isomorphism between T and the quotient tree obtained from S by collapsing each component of σ to a point, and furthermore this simplicial isomorphism is F -equivariant. In this situation we often say that T is obtained by collapsing σ .

Furthermore, any choice of collapsed subgraph may be used, in the sense that for any free splitting $F \curvearrowright S$ and any F -equivariant, proper subgraph $\sigma \subset S$ there exists a free splitting T and a collapse map

$$S \xrightarrow{[\sigma]} T.$$

The tree T is defined as the quotient of S obtained by collapsing to a point each component of σ . Since σ is proper, T is not a point. Since σ is equivariant, the action $F \curvearrowright S$ descends to an action $F \curvearrowright T$. This action is minimal because T is a union of axes of elements of F : for each edge $e \subset T$ there exists a unique pre-image edge $e' \subset S$ such that e' maps to e , and there exists $\gamma \in F$ whose axis in S contains e' , so the axis of γ in T contains e . The stabilizer of an edge $e \subset T$ equals the stabilizer of

the pre-image edge and so is trivial. This shows that $F \curvearrowright T$ is a free splitting, and by construction the quotient map

$$S \xrightarrow{[\sigma]} T$$

is a collapse map.

The (nonsimplicial) conjugacy type of the collapsed tree actually depends only on the “natural core” of the collapsed subgraph. To be precise, given a free splitting $F \curvearrowright S$ and a proper, F -equivariant subgraph $\sigma \subset S$, define the *natural core* of σ to be the largest natural subcomplex of S contained in σ whose components are all nondegenerate. For any collapse maps

$$S \xrightarrow{[\sigma]} T, \quad S \xrightarrow{[\sigma']} T',$$

if σ, σ' have the same natural core then there exists a conjugacy $T \rightarrow T'$, although this conjugacy need not be a simplicial map with respect to the given simplicial structures of T, T' .

Given free splittings $F \curvearrowright S, T$, we say that S *collapses* to T or that T *expands* to S , denoted $S \succ T$ or $T \prec S$, if there exists a function $S \mapsto T$ that is a collapse map with respect to some simplicial subdivisions of the natural cell structures on S and T . These relations are well-defined on the conjugacy classes of S, T , indeed $S \succ T$ if and only if there exists a function $S \mapsto T$ that is a collapse map with respect to the natural cell structures themselves. Even when it is known that $S \succ T$, notice that there might not exist a collapse map $S \mapsto T$ without first changing the simplicial structures on S and/or T , for example if T is subdivided so finely that it has more edgelet orbits than S . The collapse and expand relations are transitive, eg if $S \succ S' \succ S''$ then $S \succ S''$, for if $S \mapsto S' \mapsto S''$ are collapse maps of natural cell structures then the composition $S \mapsto S''$ is a collapse map of natural cell structures.

In several places throughout the paper we use without comment the fact that every free splitting $F \curvearrowright T$ has a *properly discontinuous expansion* $T \prec S$, meaning that the free splitting $F \curvearrowright S$ is properly discontinuous; see [8, Section 3.2] for a proof, under the heading “How to construct trees in \mathcal{K}_n^T ”, Steps 1 and 2. When a properly discontinuous expansion $T \prec S$ is chosen, with collapse map

$$S \xrightarrow{[\sigma]} T,$$

the free factor system $\mathcal{F}(T)$ is represented in S as the conjugacy classes of the stabilizers of the infinite components of σ .

1.3 The free splitting complex in terms of collapse maps

The following result contains the technical facts needed to justify the construction of the simplicial complex $\mathcal{FS}(F)$. For any free splitting $F \curvearrowright T$ and any proper F -invariant natural subgraph $\sigma \subset T$ let

$$T \xrightarrow{[\sigma]} T_\sigma$$

be the corresponding collapse map, the quotient map obtained by collapsing to a point each component of σ . If T is a $(K+1)$ -edge free splitting then for each $k = 0, \dots, K$ let $\mathcal{F}_k(T)$ be the set of conjugacy classes of $(k+1)$ -edge free splittings of the form T_σ , indexed by those natural subgraphs $\sigma \subset T$ that contain exactly $K-k$ natural edge orbits of T . For fixed k there are exactly

$$\binom{K+1}{k+1} = \frac{(K+1)!}{(k+1)!(K-k)!}$$

choices of such σ , and as $k = 0, \dots, K$ varies there are $2^{K+1} - 1$ choices of σ .

Lemma 1.3 *For any free splittings $F \curvearrowright T, T'$ the following hold:*

- (1) *For any two F -equivariant natural subgraphs $\sigma_1, \sigma_2 \subset T$ we have $\sigma_1 = \sigma_2$ if and only if $T_{\sigma_1}, T_{\sigma_2}$ are conjugate.*
- (2) *$\mathcal{F}_0(T) = \mathcal{F}_0(T')$ if and only if T, T' are conjugate.*

We thank Gilbert Levitt for describing the following proof to us based on his paper with Guirardel [7]. This proof replaces a much longer argument in the previous version of this paper.

Proof Let $\mathbb{C}(F)$ be the set of conjugacy classes of nontrivial elements of F . In this proof every free splitting $F \curvearrowright T$ is equipped with its natural cell structure and with a geodesic metric assigning length 1 to each natural edge. Let $\ell_T \in \mathbb{R}^{\mathbb{C}(F)}$ be the translation length function of T : for nontrivial $\gamma \in F$ with conjugacy class $[\gamma]$, the quantity $\ell_T[\gamma]$ equals the translation length of the action of $\gamma \in F$ on T . By Culler and Morgan [6], the function ℓ_T determines and is determined by the conjugacy class of T .

Fix a free splitting T , let E_0, \dots, E_K be an enumeration of representatives of the natural edge orbits of T , and for $k = 0, \dots, K$ let T_k be the one-edge free splitting obtained from T by collapsing $T \setminus (F \cdot E_k)$. By construction the conjugacy class of T determines the conjugacy classes of T_0, \dots, T_K . By [7, Lemma 3.18], the trees T_0, \dots, T_K are pairwise nonconjugate; incidentally this proves (1) in the special case

that σ, σ' each contain all but one of E_0, \dots, E_K . By [6] the functions $\ell_{T_0}, \dots, \ell_{T_K} \in \mathbb{R}^{\mathbb{C}(F)}$ are pairwise distinct, and as pointed out in [7, Definition 3.17], the equation holds $\ell_T = \sum_{i=0}^K \ell_{T_k}$ in the vector space $\mathbb{R}^{\mathbb{C}(F)}$. It follows that the subset $\{\ell_{T_0}, \dots, \ell_{T_K}\} \subset \mathbb{R}^{\mathbb{C}(F)}$ determines ℓ_T . Putting it all together, the conjugacy classes of T_0, \dots, T_K determine the conjugacy class of T , proving (2).

Now we prove (1) in general. Consider a proper F -equivariant natural subgraph $\sigma \subset T$. By construction the subgraph σ determines and is determined by the subset $\mathcal{E}(\sigma) \subset \{E_0, \dots, E_K\}$ consisting of those edges not contained in σ . Under the collapse map

$$T \xrightarrow{[\sigma]} T_\sigma,$$

the set $\mathcal{E}(\sigma)$ corresponds bijectively with representatives of the natural edge orbits of T_σ . Clearly the conjugacy classes of one-edge collapses of T_σ — which by (2) determine and are determined by the conjugacy class of T_σ — are precisely the conjugacy classes of the free splittings T_k for $E_k \in \mathcal{E}(\sigma)$. It follows that the conjugacy class of T_σ determines and is determined by σ , proving (1). \square

By applying item (1) of this lemma we may define a collapse $T \succ U$ to be *proper* if it satisfies either of the following equivalent conditions: U, T are not conjugate, or *for any map*

$$T \xrightarrow{[\sigma]} U$$

that is a collapse map with respect to some subdivision of the natural cell structures, the natural core of σ is nonempty. We also refer to the collapse maps of the latter type as *proper collapse maps*. Notice that properness of a collapse relation $T \succ U$ is also equivalent to the statement that *there exists* a map

$$T \xrightarrow{[\sigma]} U$$

that is a collapse map with respect to some subdivision of the natural structures, such that the natural core of σ is nonempty. A collapse relation $T \succ U$ that is not proper is *improper*.

We now apply Lemma 1.3 to the construction of $\mathcal{FS}(F)$. From item (1) it follows that we can associate an abstract K -simplex denoted $\langle T \rangle$ to the conjugacy class of each $(K + 1)$ -edge free splitting $F \curvearrowright T$, where the k -dimensional faces of $\langle T \rangle$ are labelled by the conjugacy classes of those free splittings of the form T_σ such that σ contains exactly $K - k$ natural edge orbits of T , and where T_σ is a face of $T_{\sigma'}$ if and only if $\sigma' \subset \sigma$. We can then glue these simplices together, where for each collapse relation $T \succ U$ the simplex $\langle U \rangle$ is glued to the unique face of the simplex $\langle T \rangle$ that

is labelled by the conjugacy class of U and where the gluing preserves the labelling of subfaces. From item (1) it follows that the result of these gluings is a simplicial complex. We have proved:

Corollary 1.4 *There exists a simplicial complex $\mathcal{FS}(F)$ whose K -simplices $\langle T \rangle$ are in one-to-one correspondence with the conjugacy classes of $K + 1$ -edge free splittings $F \curvearrowright T$, such that for any pair of simplices $\langle T \rangle, \langle U \rangle$ we have $\langle U \rangle \subset \langle T \rangle$ if and only if $U \prec T$.*

The alternate and more well-known approach to this corollary is to appeal to Hatcher's construction of the sphere complex [9]; see for example Aramayona and Souto [1], which constructs the 1-skeleton of $\mathcal{FS}(F)$ in this manner.

The dimension of $\mathcal{FS}(F)$ equals $3 \cdot \text{rank}(F) - 4$, the number $3 \cdot \text{rank}(F) - 3$ being the maximum number of natural edge orbits of a free splitting $F \curvearrowright T$, the maximum occurring if and only if every natural vertex of T has valence 3 (which implies that the free splitting $F \curvearrowright T$ is properly discontinuous).

We usually work with the first barycentric subdivision of $\mathcal{FS}(F)$, denoted $\mathcal{FS}'(F)$. Gromov hyperbolicity of $\mathcal{FS}(F)$ and $\mathcal{FS}'(F)$ are equivalent because, as with any connected simplicial complex, the identity map is a quasi-isometry between their geodesic simplicial metrics (connectivity follows from Hatcher's proof of contractibility [9], or from the construction of Stallings fold paths reviewed in Section 2). The simplicial complex $\mathcal{FS}'(F)$ has one 0-simplex associated to each conjugacy class of free splittings, and it has a k -simplex associated to each sequence of conjugacy classes of free splittings obtained from any chain of k proper expansions $T_0 \prec T_1 \prec \cdots \prec T_k$. In particular, an edge in $\mathcal{FS}'(F)$ oriented from S to T can be written uniquely as either an expand $S \prec T$ or a collapse $S \succ T$; uniqueness follows from asymmetry of the collapse relation, which is a consequence of Lemma 1.3(1).

As mentioned earlier, the relations of collapse and expand are transitive. It follows that every geodesic in the 1-skeleton of $\mathcal{FS}'(F)$ can be written as an alternating sequence of expands and collapses, for example starting with an expand $T_0 \prec T_1 \succ T_2 \prec T_3 \succ T_4 \prec T_5 \succ \cdots$ or starting with a collapse $T_0 \succ T_1 \prec T_2 \succ T_3 \prec T_4 \succ T_5 \prec \cdots$. Any edge path in $\mathcal{FS}'(F)$ that alternates between expands and collapses is called a *zig-zag path* in $\mathcal{FS}'(F)$.

Throughout the paper, given free splittings $F \curvearrowright S, T$, we use the notation $d(S, T)$ to denote the length of the shortest edge path in the simplicial complex $\mathcal{FS}'(F)$ between the vertices represented by S and T . We must prove that this metric is Gromov hyperbolic.

2 Fold paths

We define the class of fold paths between vertices of $\mathcal{FS}'(F)$, using a method pioneered by Stallings [20] for factoring maps of graphs into products of folds. This method was extended to the category of group actions on trees by Bestvina and Feighn [3]. We refer to the latter paper for some details, although these details are considerably simplified in the category of free splittings.

2.1 Directions, gates and foldable maps

First we set up some of the basic definitions which are used throughout the paper. We will also prove a tree-theoretic version of the first derivative test, Lemma 2.3.

Given any graph X and a vertex $v \in X$, the set of *directions of X at v* , denoted $D_v X$, is defined to be the set of germs of oriented arcs in X with initial vertex v . Each direction at v is uniquely represented by an oriented edgelet with initial vertex v . The union of the sets $D_v X$ over all vertices $v \in X$ is denoted DX . Given a subgraph $Y \subset X$, the subset of DX represented by oriented edgelets $e \subset X \setminus Y$ having initial vertex in Y is denoted $D_Y X$.

Given two free splittings $F \curvearrowright S, T$ and a map $f: S \rightarrow T$, the *derivative* of f is a partially defined map $df: DS \rightarrow DT$ whose domain is the set of directions of oriented edgelets e on which f is nonconstant, and whose value on the direction of e is the direction of the oriented edgelet $f(e)$. Given a subgraph $W \subset S$, if f is nonconstant on each edgelet representing a direction in the set $D_W S$ then we obtain by restriction a map $d_W f: D_W S \rightarrow DT$; as a special case, when $W = \{v\}$ is a vertex we obtain a map $d_v f: D_v S \rightarrow D_{f(v)} T$.

Suppose now that the map $f: S \rightarrow T$ is nonconstant on all edgelets of S , so $df: DS \rightarrow DT$ has full domain of definition. For each vertex $v \in S$ the set $D_v S$ partitions into *gates*, which are the nonempty subsets of the form $(d_v f)^{-1}(\delta)$ for $\delta \in D_{f(v)} T$. Every gate is a finite set; indeed, we have:

Lemma 2.1 *For any free splittings $F \curvearrowright S, T$, for any map $f: S \rightarrow T$ that is nonconstant on each edgelet of S , and for any vertex $v \in S$, the cardinality of each gate of $D_v S$ is at most $2 \operatorname{rank}(F)$.*

Proof Let $e_1, \dots, e_M \subset S$ be oriented edgelets with initial vertex v representing a gate of $D_v S$. These oriented edgelets are all in distinct orbits under the action of F , for otherwise their common image in T would have a nontrivial stabilizer. It follows that in the quotient graph of groups S/F , the quotients of e_1, \dots, e_M represent M

distinct directions at the quotient of v . It therefore suffices to bound the valence of each vertex in the quotient graph of groups of a free splitting. Without decreasing the valence at the quotient of v , one can blow up all other vertex orbits so that the only vertex orbit with nontrivial stabilizers is the orbit of v . Then, still without decreasing quotient valence, one can inductively collapse natural edges whose endpoints are in different vertex orbits. When this process stops, the quotient graph of groups is a rose with one natural vertex (possibly having nontrivial vertex group) and with $\leq \text{rank}(F)$ edges, whose natural vertex has valence $\leq 2 \text{rank}(F)$. \square

Definition 2.2 (Foldable maps and edgelets) A map $f: S \rightarrow T$ is *foldable* if it satisfies either of the following two equivalent statements:

Natural edge definition of foldable f is injective on each natural edge of S and f has at least 3 gates at each natural vertex of S .

Edgelet definition of foldable f is injective on every edgelet, f has at least 2 gates at every vertex, and f has at least 3 gates at every natural vertex.

We will without warning switch between these two definitions whenever it is convenient. Notice that the restrictions on the number of gates are significant only at vertices of finite valence, because every gate is a finite set; for example, if every natural vertex of S has nontrivial stabilizer then every map defined on S that is injective on natural edges is foldable. Notice also that foldability of f depends only on the natural cell structures on S and T , not on the given simplicial structures; to put it more formally, foldability is an invariant of f in the category of equivariant continuous functions between free splittings of F .

Given free splittings $F \curvearrowright S, T$, a foldable map $f: S \rightarrow T$ and an edgelet $e \subset T$, an *e -edgelet of f* is an edgelet of S that is mapped to e by f .

In Lemma 2.4 below we shall prove the existence of foldable maps in the appropriate context.

Remark In other treatments of Stallings folds we have not seen any analogue of our gate at least 3 condition on natural vertices. This condition is crucial to the diameter bound obtained in Lemma 5.5, as well as in the heart of the proof of the Main Theorem, particularly in the proof of Proposition 6.5, Step 3.

The first derivative test The first derivative test of calculus implies that if the derivative of a function has no zeroes then local extreme values occur only at endpoints of the domain.

Lemma 2.3 (The first derivative test) *Suppose that $f: S \rightarrow T$ is a foldable map of free splittings. Given a connected subgraph $W \subset S$ and a vertex $v \in W$, if $f(v)$ has valence 1 in the subgraph $f(W) \subset T$ then v is a frontier point of W .*

Proof If v is an interior point of W then $D_v W = D_v S$, and since f has at least 2 gates at v it follows that $d_v f(D_v W)$ has cardinality at least 2, implying that $f(v)$ has valence at least 2 in $f(W)$. \square

2.2 Construction of foldable maps

Given free splittings $F \curvearrowright S, T$, a fold path from S to T will be defined by factoring a foldable map $S \mapsto T$. Although a foldable map does not always exist, one will exist after moving S a distance at most 2 in $\mathcal{FS}'(F)$.

Lemma 2.4 *For any free splittings $F \curvearrowright S, T$ there exist free splittings S', S'' and a foldable map $S'' \mapsto T$ such that $S \prec S' \succ S''$.*

Proof Fix the free splitting $F \curvearrowright T$. Given a free splitting $F \curvearrowright R$, let $\mathcal{M}(R, T)$ denote the set of all equivariant continuous functions $f: R \rightarrow T$ that take each natural vertex of R to a vertex of T and whose restriction to each natural edge of R is either injective or constant. It follows that f is a map with respect to the pullback simplicial structure on R whose vertex set consists of all points that map to vertices of T and that are not in the interior of a natural edge of R that is collapsed by f . The edges of this simplicial structure on R will be referred to as *pullback edgelets of f* .

Choose any properly discontinuous expansion $S \prec S'$, which implies that the set $\mathcal{M}(S', T)$ is nonempty. Amongst all elements of $\mathcal{M}(S', T)$ choose $f: S' \rightarrow T$ to maximize the number of orbits of natural edges of S' on which f is constant. By collapsing each such natural edge we define a collapse map $S' \mapsto S''$ and an induced function, which is an element of the set $\mathcal{M}(S'', T)$. By maximality of f it follows that any element of $\mathcal{M}(S'', T)$ is injective on each natural edge of S'' , for otherwise by composing the collapse map $S' \mapsto S''$ with an element of $\mathcal{M}(S'', T)$ that collapses some natural edge of S'' we obtain an element of $\mathcal{M}(S', T)$ that collapses a larger number of natural edge orbits than f does, a contradiction.

We find a foldable element of $\mathcal{M}(S'', T)$ by solving optimization problems. First we prove that if $g \in \mathcal{M}(S'', T)$ minimizes the number of orbits of pullback edgelets then g has at least 2 gates at each vertex of S'' . Suppose there is a vertex $v \in S''$ at which g has only 1 gate. Let K be the valence of v ; note that $K \geq 3$ because g is injective on natural edges. Let η_1, \dots, η_K be the oriented natural edges of S''

with initial vertex v . Let e_1, \dots, e_K be the initial pullback edgelets of η_1, \dots, η_K , and let w_1, \dots, w_K be the terminal endpoints of e_1, \dots, e_K , respectively. We have $g(e_1) = \dots = g(e_K) = e$ for some oriented edge $e \subset T$ with initial vertex $g(v)$ and opposite vertex $w = g(w_1) = \dots = g(w_K)$. Consider first the case that $e_i \neq \eta_i$ for each i , and so we can isotope each restricted map $g|_{\eta_i}$ by pushing $g(v)$ across e to w by an isotopy supported in a neighborhood of e_i , and we can extend these isotopies to an equivariant homotopy of g , to produce an element of $\mathcal{M}(S'', T)$ that has K fewer orbits of pullback edgelets than g has, a contradiction. Consider next the case that $e_i = \eta_i$ for certain values of $i = 1, \dots, K$. If v, w_i are in distinct F -orbits for each such i then we can equivariantly homotope g , pushing $g(v)$ across e to w , so as to collapse each e_i for which $e_i = \eta_i$, to produce an element of $\mathcal{M}(S'', T)$ that collapses each of the natural edges η_i such that $e_i = \eta_i$, a contradiction. In the remaining case there exists some $i = 1, \dots, K$ such that $e_i = \eta_i$ and $w_i = \gamma \cdot v$ for some $\gamma \in F$, and it follows that $w = \gamma \cdot g(v)$. The edges $e_i \subset S''$ and $e \subset T$ are therefore fundamental domains for the actions of γ on its axes in S'', T , respectively. It follows that the direction of $\gamma^{-1} \cdot e_i$ at v maps to the direction of $\gamma^{-1} \cdot e$ at $g(v)$, which is *not equal* to the direction of e at $g(v)$, contradicting that g has a single gate at v .

Next we prove that among all $g \in \mathcal{M}(S'', T)$ that minimize the number of orbits of pullback edges, there is at least one which is foldable, having at least 3 gates at each natural vertex. This is achieved, mostly, by solving another optimization problem. Define the *edgelet vector* of g to be the vector of positive integers L_g indexed by the natural edge orbits of S'' , whose entry $L_g(e)$ corresponding to the orbit of a natural edge $e \subset S''$ is the number of pullback edgelets in e . Define $\text{Length}(L_g)$ to be the sum of its entries, which equals the number of pullback edgelet orbits of g , a number which has already been minimized so as to guarantee at least 2 gates at each vertex. Define $\text{Energy}(L_g)$ to be the sum of the squares of its entries. We have the inequality $\text{Energy}(L_g) \leq (\text{Length}(L_g))^2$. Amongst all $g \in \mathcal{M}(S'', T)$ with minimal value of $\text{Length}(L_g)$, choose g so as to maximize $\text{Energy}(L_g)$.

We claim that with energy maximized as above, one of the following holds:

- (1) g has at least 3 gates at each natural vertex, and so g is foldable.
- (2) S'' has exactly one natural vertex orbit, g has two gates at every natural vertex and each natural edge of S'' has its two directions lying in distinct gate orbits.

To prove this dichotomy, suppose that g has exactly two gates at some natural vertex v . The gates must have the same cardinality: otherwise, by doing a valence 2 homotopy, pushing $g(v)$ across one edge of T in the image direction of the larger of the two gates at v , one reduces the total number of pullback edgelets. Now consider $g_1, g_2 \in \mathcal{M}(S'', T)$ defined by the two possible valence 2 homotopies at v , pushing $g(v)$ across

the two edges of T in the two image directions of the two gates at v . Note that the average of the two vectors L_{g_1}, L_{g_2} is the vector L_g . It follows that $L_g = L_{g_1} = L_{g_2}$, for otherwise, by convexity of energy, one of $\text{Energy}(L_{g_1})$ or $\text{Energy}(L_{g_2})$ would be larger than $\text{Energy}(L_g)$. It also follows that S'' has exactly one natural vertex orbit, for otherwise v would be connected across a natural edge e to some natural vertex in a different orbit, implying that one of $L_{g_1}(e), L_{g_2}(e)$ equals $L_g(e) + 1$ and the other equals $L_g(e) - 1$. It also follows that each natural edge e has one end in the orbit of one gate at v and opposite end in the orbit of the other gate at v , for otherwise one of $L_{g_1}(e), L_{g_2}(e)$ would equal $L_g(e) + 2$ and the other equals $L_g(e) - 2$. This shows that g satisfies item (2).

To finish up we show that if g satisfies (2) then there exists $g' \in \mathcal{M}(S'', T)$ that satisfies (1). Item (2) implies that there is an orientation of the natural edges of S'' such that at each natural vertex $v \in S''$, the directions with initial vertex v form one gate of g' at v denoted D_v^+ , and the directions with terminal vertex v form the other gate denoted D_v^- .

Pick a natural vertex $v \in S''$. Let τ be the subtree of S'' consisting of the union of all oriented rays in S'' with initial vertex v . The restriction of g to each such ray is injective and proper, and their initial directions all map to the same direction in T , so it follows that the subtree $g(\tau) \subset T$ has a valence 1 vertex at $g(v)$ and no other valence 1 vertex. Also, if we orient each edge of $g(\tau)$ to point away from the vertex $g(v)$ then the map $g: \tau \rightarrow g(\tau)$ preserves orientation. Furthermore $g(\tau)$ is not itself just a ray, for if it were then T would be just a line, an impossibility for a free splitting of a free group of rank at least 2. Let $w \in g(\tau)$ be the vertex of $g(\tau)$ of valence at least 3 that is closest to $g(v)$. Define $g': S'' \rightarrow T$ by mapping v to w , extending equivariantly to the orbit of v , and extending equivariantly to an embedding on each edge of S'' .

We claim that g' has one gate at v corresponding to each direction of $g(\tau)$ at w , which implies that g' is foldable. To see why, first note that the set D_v^- is mapped by $d_v g'$ to the unique direction of the segment $[w, g(v)]$ at w . Next note that each direction in the set D_v^+ is mapped by $d_v g'$ to one of the directions of T at w distinct from the direction of $[w, g(v)]$; furthermore, each such direction is in the image of $d_v g'$ because g' maps τ onto $g(\tau) \setminus [w, g(v)]$ by an orientation-preserving map.

This completes the proof of Lemma 2.4. □

2.3 Folds

Given free splittings $F \curvearrowright S, T$ and a foldable map $f: S \rightarrow T$, we say that f is a *fold* if there exist oriented natural edges $\eta, \eta' \subset S$ with the same initial vertex v , and

there exist nondegenerate initial segments $e \subset \eta$, $e' \subset \eta'$ that are subcomplexes of S with the same positive number of edgelets, such that if we let $\phi: e \rightarrow e'$ denote the unique orientation-preserving simplicial isomorphism, then for all $x \neq x' \in S$ we have $f(x) = f(x')$ if and only if there exists $\gamma \in F$ such that (up to interchanging x, x') $\gamma \cdot x \in e$ and $\phi(\gamma \cdot x) = \gamma \cdot x' \in e'$. We also say that the map f folds the segments e and e' .

The pair of segments e, e' determines the free splitting $F \curvearrowright T$ up to simplicial conjugacy, namely $F \curvearrowright T$ is conjugate to the equivariant quotient complex of S obtained by equivariantly identifying e and e' via $\phi: e \rightarrow e'$. In this context we shall say that the free splitting T is determined by *folding the segments* e, e' . Letting $d, d' \in D_v S$ denote the initial directions of e, e' respectively, we also say that f folds the directions d, d' , although d, d' do not determine the segments e, e' and they need not determine T up to conjugacy. Notice that d, d' are in different orbits under the action $\text{Stab}_S(v) \curvearrowright D_v S$ (equivalently under the action $F \curvearrowright DS$), for otherwise the segment $f(e) = f(e') \subset T$ would have nontrivial stabilizer. Folds are classified according to the properness of the inclusions $e \subset \eta$, $e' \subset \eta'$, as follows. If e, e' are both proper initial segments of η, η' then we say that f is a *partial fold*; otherwise f is a *full fold*. If f is a full fold and exactly one of e, e' is proper then we say that f is a *proper full fold*; otherwise, when $e = \eta$ and $e' = \eta'$, we say that f is an *improper full fold*. For later purposes we note that if f is a full fold then every natural vertex of T is the image of a natural vertex of S ; and even when f is a partial fold, every natural vertex of T that is not in the orbit of the image of the terminal endpoints of the folded edges e, e' is the image of a natural vertex of S .

In the terminology of [3], folds between free splittings can also be classified into two types as follows. If the opposite vertices w, w' of e, e' are in different F -orbits one gets a type IA fold; in this case the stabilizer of the vertex $W = f(w) = f(w')$ is the subgroup generated by the stabilizers of w, w' , which (if nontrivial) is a free factor whose rank is the sum of the ranks of the stabilizers of w and w' . If w, w' are in the same F -orbit then one gets a type IIIA fold, and the stabilizer of the vertex W is the subgroup generated by the stabilizer of w and an element $\gamma \in F$ such that $\gamma(w) = w'$, which is a free factor whose rank is 1 plus the rank of the stabilizer of w . Notice that a type IIIA fold is only possible if f is a partial fold or an improper full fold, because a natural and an unnatural vertex can never be in the same orbit. We refer to [3] for an understanding of the map on quotient graphs of groups $S/F \rightarrow T/F$ that is induced by a fold $f: S \rightarrow T$.

The following lemma and its proof are well known in the narrower context of the first barycentric subdivision of the spine of outer space.

Lemma 2.5 For any fold $f: S \rightarrow T$, the distance in $\mathcal{FS}'(F)$ from S to T equals 1 or 2.

Proof Let f fold oriented segments e, e' with common initial endpoint v and opposite endpoints w, w' . After possibly subdividing S and T so that e, e' each contain at least 2 edgelets, the map f can be factored into two maps as

$$S \xrightarrow{g} U \xrightarrow{h} T,$$

where g folds the initial edgelets $e_0 \subset e, e'_0 \subset e'$, and h folds the g -images of the terminal segments $e_1 = e \setminus e_0, e'_1 = e' \setminus e'_0$. Letting $\hat{e} = g(e_0) = g(e'_0) \subset U$ and $\sigma_0 = F \cdot \hat{e} \subset U$, resubdividing S there is an expansion $S \prec U$ defined by a collapse map

$$U \xrightarrow{[\sigma_0]} S.$$

Also, letting $\sigma_1 = F \cdot (g(e_1) \cup g(e'_1)) \subset U$, after resubdividing T there is a collapse $U \succ T$ defined by a collapse map

$$U \xrightarrow{[\sigma_1]} T.$$

It follows that $d(S, T) \leq 2$ in $\mathcal{FS}'(F)$.

It remains to show that $d(S, T) \neq 0$, that is, S, T are not conjugate free splittings. Since each fold map is foldable, the natural vertex v has at least 3 gates with respect to f . It therefore has at least 3 gates with respect to g , and so $g(v) \in U$ is natural. It follows that \hat{e} is a natural edge of U , having one endpoint at $g(v)$ and opposite endpoint of valence 3 in U . The subgraph $\sigma_0 \subset U$ is therefore natural, and it follows from Lemma 1.3 that S is not conjugate to U . The free splittings U, T may or may not be conjugate, depending on whether at least one of $g(e_1), g(e_2) \subset U$ is a natural edge. If neither of $g(e_1), g(e_2)$ is natural then T is conjugate to U , and so T is not conjugate to S . If one or both of $g(e_1), g(e_2)$ is natural then (after resubdividing T) the collapse $U \succ T$ may also be defined by collapsing the natural subgraph $\hat{\sigma}_1 \subset U$ that is the union of the F orbits of whichever of $g(e_1), g(e_2)$ is natural, but $\sigma_0 \neq \hat{\sigma}_1$ and so by Lemma 1.3 we conclude that S, T are not conjugate. \square

2.4 Fold sequences and fold paths

Consider free splittings $F \curvearrowright S, T, U$ and a sequence of maps of the form

$$S \xrightarrow{h} U \xrightarrow{g} T.$$

Letting $f = g \circ h: S \rightarrow T$, we say that h is a *maximal fold factor* of f if the following hold: h is a fold map that folds oriented initial segments $e, e' \subset S$ of oriented natural edges $\eta, \eta' \subset S$, respectively, and e, e' are the maximal initial subsegments of η, η' such that in T we have $f(e) = f(e')$. Recall from the definition of a fold that e, e' are edgelet paths with the same number of edgelets.

Fold sequences Consider a sequence of free splittings and maps of the form

$$S_0 \xrightarrow{f_1} S_1 \xrightarrow{f_2} \cdots \xrightarrow{f_K} S_K,$$

$K \geq 0$. In this context we will always denote

$$f_j^i = f_j \circ \cdots \circ f_{i+1}: S_i \rightarrow S_j, \quad \text{for } 0 \leq i < j \leq K,$$

so that expressions such as $f_j^i: S_i \rightarrow S_j$ obey the Einstein indexing convention. Note also that the convention $f_i^{i-1} = f_i: S_{i-1} \rightarrow S_i$ amounts to suppression of the superscript.

We say that this is a *fold sequence* if the following holds:

- (1) $f_K^0: S_0 \rightarrow S_K$ is a foldable map.
- (2) Each map $f_{i+1}: S_i \rightarrow S_{i+1}$ is a maximal fold factor of the map $f_K^i: S_i \rightarrow S_K$, for $0 \leq i < K$.

It follows from (1) and (2) that

- (3) $f_j^i: S_i \rightarrow S_j$ is a foldable map for each $0 \leq i < j \leq K$.

To prove (3), starting from the base assumption (1), and assuming by induction that $f_K^{i-1} = f_K^i \circ f_i$ is foldable, we prove that f_K^i is foldable. By (2) the map f_i is a maximal fold factor of f_K^{i-1} . The map f_K^{i-1} is injective on each edgelet of S_{i-1} , and each edgelet of S_i is the f_i image of some edgelet of S_{i-1} , so f_K^i is injective on each edgelet. Consider a vertex $v \in S_i$ and a vertex $u \in S_{i-1}$ for which $f_i(u) = v$. The number of f_K^i -gates at v is greater than or equal to the number of f_K^{i-1} -gates at u , which is at least 2, and furthermore if u is natural then this number is at least 3. This covers all cases except for when v is natural and each such u has valence 2. Since f_i is a maximal fold factor of f_K^{i-1} , this is only possible if f_i is a partial fold that folds segments $e, e' \subset S_{i-1}$ such that if w, w' denote the terminal endpoints of e, e' then $v = f_i(w) = f_i(w')$. If f_i is a type IA fold, that is, if w, w' are in different orbits, then v has valence 3, and by maximality of the fold f_i it follows that the three directions at v are all in different gates with respect to f_K^i . If f_i is a type IIIA fold, that is, if w, w' are in the same orbit, say $\gamma \cdot w = w'$ for a nontrivial $\gamma \in F$, then $\text{Stab}_{S_i}(v)$ contains γ and so is nontrivial, and hence v has infinitely many gates with respect to f_K^i . This

proves by induction that each f_K^i is foldable. Next, to prove that f_j^i is foldable, given a vertex $v \in S_i$ we simply note that the decomposition of $D_v S_i$ into f_j^i -gates is a refinement of the decomposition into f_K^i -gates, of which there are at least 2, and at least 3 if v is natural. This completes the proof that (1) and (2) imply (3).

In this proof we have shown the following fact, which will be useful in Lemma 2.7 below when we construct fold sequences:

Lemma 2.6 For any foldable map $f: S \rightarrow T$ and any factorization of f into two maps of the form

$$S \xrightarrow{k} U \xrightarrow{g} T,$$

if k is a maximal fold factor of f then the map $g: U \rightarrow T$ is also foldable. □

The implication of this lemma is *false* if one allows k to be a partial fold that is not a maximal fold factor of f , for in that case the map $g: U \rightarrow T$ will have only 2 gates at the valence 3 vertex that is the k -image of the terminal endpoints of oriented segments e, e' that are folded by k .

Fold paths A *fold path* in $\mathcal{FS}'(F)$ is any sequence of vertices represented by free splittings $F \curvearrowright S_0, S_1, \dots, S_K$ for which there exists a fold sequence

$$S_0 \mapsto S_1 \mapsto \dots \mapsto S_K;$$

we also say that this fold path has K -steps.

Strictly speaking a fold path need not be the sequence of vertices along an actual edge path in the simplicial complex $\mathcal{FS}'(F)$, because the size of the step from S_{i-1} to S_i is either 1 or 2; see Lemma 2.5. If one so desires, one can easily interpolate the gap between S_{i-1} and S_i by an edge path of length 1 or 2, to get an actual edge path from S_0 to S_K .

We define two fold sequences to be *equivalent* if they have the same length and there is a commutative diagram of the form

$$\begin{array}{ccccccc} S_0 & \longrightarrow & S_1 & \longrightarrow & \cdots & \longrightarrow & S_{K-1} & \longrightarrow & S_K \\ \downarrow & & \downarrow & & & & \downarrow & & \downarrow \\ S'_0 & \longrightarrow & S'_1 & \longrightarrow & \cdots & \longrightarrow & S'_{K'-1} & \longrightarrow & S'_K \end{array}$$

where the top and bottom rows are the two given fold sequences and each vertical arrow is a conjugacy. Note that the vertical arrows are *not* required to be “maps” as we have defined them, in that they need not be simplicial. For example, if the bottom row

is obtained by taking the 400th barycentric subdivision of each 1–simplex in the top row then the two fold sequences are equivalent.

Equivalent fold sequences determine the same fold path, but the converse is false. A counterexample consisting of a 1–step fold path is given at the end of this section.

Construction of fold factorizations Having constructed many foldable maps in Lemma 2.4, to construct many fold paths it suffices to factor each foldable map as a fold sequence.

Given free splittings $F \curvearrowright S, T$ and a foldable map $f: S \rightarrow T$, a *fold factorization* of f is any fold sequence $S_0 \mapsto S_1 \mapsto \dots \mapsto S_K$, which factors f as shown in the following commutative diagram:

$$\begin{array}{ccccccc}
 & & & & f & & \\
 & & & & \curvearrowright & & \\
 S & \xlongequal{\quad} & S_0 & \xrightarrow{f_1} & S_1 & \xrightarrow{f_2} & \dots & \xrightarrow{f_K} & S_K & \xlongequal{\quad} & T
 \end{array}$$

A fold factorization of any foldable map can be constructed by an inductive process described in [3], with considerable simplification arising from the fact that all edgelet stabilizers are trivial in T . We give this simplified argument here.

Lemma 2.7 *For any free splittings $F \curvearrowright S, T$, every foldable map $f: S \mapsto T$ has a fold factorization.*

Proof If f is a simplicial isomorphism then we are done, with a fold factorization of length $K = 0$. Otherwise, we use the following obvious but key fact:

Local to global principle Any simplicial map between trees which is locally injective is globally injective. If furthermore it is surjective then it is a simplicial isomorphism.

For the inductive step we show that every foldable map $f: S \rightarrow T$ that is not a homeomorphism factors into maps as

$$S \xrightarrow{k} U \xrightarrow{g} T,$$

where k is a maximal fold factor of f . By the *local to global principle*, plus the fact that $F \curvearrowright T$ is minimal, it follows that f is surjective and so f is not locally injective. We may therefore find a vertex $v \in S$ and two directions $d, d' \in D_v S$ such that $d_v f(d) = d_v f(d')$. Let η, η' be the oriented natural edges with initial vertex v and initial directions d, d' . Let $e \subset \eta, e' \subset \eta'$ be the maximal initial segments such that $f(e) = f(e')$. Noting that e, e' are subcomplexes with the same number of edgelets,

let $h: e \rightarrow e'$ be the unique orientation-preserving simplicial homeomorphism. Define $k: S \rightarrow U$ to be the quotient map obtained by equivariantly identifying e and e' , and let $g: U \rightarrow T$ be the induced map. As indicated in [3], U is a tree and the induced action $F \curvearrowright U$ is minimal. The map k is simplicial by construction, from which it follows that g is simplicial as well. The stabilizer of each edgelet of U is trivial because it is contained in the stabilizer of its image in T under g , which is trivial, and so $F \curvearrowright U$ is a free splitting. By construction the map $k: S \rightarrow U$ is a maximal fold factor of the foldable map f .

To support the inductive step we must prove that U has fewer edgelet orbits than S , which follows from the fact that the initial edgelets of e and e' are in different orbits of the action $F \curvearrowright S$, because they have the same image edgelet in T and its stabilizer is trivial.

The fold factorization of $f = f_T^0: S = S_0 \rightarrow T$ may now be constructed as follows. Assuming f_T^0 is not locally injective, factor f_T^0 into maps as

$$S_0 \xrightarrow{f_1} S_1 \xrightarrow{f_T^1} T$$

where f_1 is a maximal fold factor of f_T^0 . The induced map f_T^1 is foldable by Lemma 2.6, and the number of edgelet orbits of S_1 is smaller than the number of edgelet orbits of S_0 . The process therefore continues by induction on the number of edgelet orbits, stopping at

$$S = S_0 \xrightarrow{f_1} S_1 \xrightarrow{f_2} \dots \xrightarrow{f_K} S_K \xrightarrow{f_T^K} T$$

when f_T^K is locally injective and therefore a simplicial conjugacy, and we identify $S_K = T$. □

Remark The *local to global principle* may be used to construct fold factorizations with various special properties. In particular, if $\beta \subset S$ is a subtree on which f is not locally injective then we may choose the folded edges η, η' to lie in β . This is used in the proof of Lemma 5.2.

Counterexample: inequivalent folds We describe two inequivalent folds

$$\tilde{f}, \tilde{f}': S_0 \rightarrow S_1$$

that determine the same 1 step fold path S_0, S_1 in $\mathcal{FS}'(F)$. Both of the free splittings $F \curvearrowright S_0, S_1$ are properly discontinuous. We first describe the quotient marked graphs $G_0 = S_0/F, G_1 = S_1/F$ and the induced homotopy equivalences $f, f': G_0 \rightarrow G_1$. The marked graph G_0 has a valence 4 vertex v with the following incident directions:

directed natural edges a, b with initial vertex v , and a directed natural edge c with initial and terminal vertex v ; subject to this description, G_0 is then filled out to be a marked graph in an arbitrary manner. The marked graph G_1 is defined to have the same underlying unmarked graph as G_0 . The homotopy equivalences $f, f': G_0 \rightarrow G_1$ are defined so that $f(a) = ca$, $f'(b) = c^{-1}b$ and f, f' are the identity elsewhere. Clearly f, f' are homotopic, by a homotopy that spins the c loop once around itself and is stationary on $G_0 \setminus (a \cup b \cup c)$. The marking on G_1 is defined by pushing forward the marking on G_0 via either of f, f' , and so each of f, f' preserves marking. Consider the universal covering maps $S_i \mapsto G_i$, $i = 0, 1$. We may choose F -equivariant lifts $\tilde{f}, \tilde{f}': S_0 \rightarrow S_1$, which are the two fold maps at issue. If they were equivalent then, since any self-conjugacy of S_0 or of S_1 fixes each vertex and each oriented natural edge (see the *Remark* at the end of Section 1), each direction in DS_0 would have the same image in DS_1 under $d\tilde{f}$ and $d\tilde{f}'$. However, fixing a lift \tilde{v} and lifts $\tilde{a}, \tilde{b}, \tilde{c}$ of a, b, c with initial vertex \tilde{v} and a lift \tilde{c}' of c with terminal vertex \tilde{v} , we have $d\tilde{f}(\tilde{a}) = d\tilde{f}(\tilde{c})$ but $d\tilde{f}'(\tilde{a}) \neq d\tilde{f}'(\tilde{c})$.

3 The Masur–Minsky axioms

Our proof that $\mathcal{FS}(F)$ is hyperbolic uses the axioms introduced by Masur and Minsky [15] for their proof that the curve complex of a finite type surface is hyperbolic. The axioms require existence of a family of paths that satisfy a strong projection property. For this purpose we shall use fold paths: Proposition 3.3 stated at the end of this section says, roughly speaking, that fold paths in $\mathcal{FS}'(F)$ satisfy the Masur–Minsky axioms.

First we give an intuitive explanation of the content of Proposition 3.3 by giving an outline of the Masur–Minsky axioms as they would apply to fold paths. The axioms require that a map be defined that is a kind of projection from $\mathcal{FS}'(F)$ to each fold path S_0, S_1, \dots, S_K . To make things work the range of the projection is taken to be the parameter interval $[0, K]$ of the fold path, giving the projection map the form $\pi: \mathcal{FS}'(F) \rightarrow [0, K]$. When one projects two vertices of $\mathcal{FS}'(F)$ to two parameters $l \leq k \in [0, K]$, one is interested in the “diameter (of the subpath) between these two parameters”, which means the diameter of the set $\{S_l, S_{l+1}, \dots, S_k\}$ in $\mathcal{FS}'(F)$. There are three axioms. The *coarse retraction* axiom bounds the diameter between each $k \in [0, K]$ and its projection $\pi(S_k) \in [0, K]$. The *coarse Lipschitz* axiom bounds the diameter between the projections $\pi(T), \pi(T') \in [0, K]$ of two nearby vertices $T, T' \in \mathcal{FS}'(F)$. The *strong contraction* axiom says roughly that, for each metric ball in $\mathcal{FS}'(F)$ that stays a bounded distance away from the fold path, if one takes the sub-ball having a certain proportion of the total radius, the diameter between the

projections of any two vertices in the sub-ball is bounded. All the bounds occurring in this discussion must be uniform, depending only on the rank of F .

In fact rather than using fold paths themselves, we use fold sequences. As we have seen in the counterexample at the end of Section 2, the same fold path S_0, \dots, S_K can be represented by inequivalent fold sequences, and the projection maps $\mathcal{FS}'(F) \rightarrow [0, K]$ of these two fold sequences may be different. This kind of situation is handled formally by expressing the Masur–Minsky axioms in terms of “families” of paths that allow a path to occur repeatedly in the family.

Given integers i, j we adopt interval notation $[i, j]$ for the set of all integers between i and j inclusive, regardless of the order of i, j .

Consider a connected simplicial complex X with the simplicial metric. A *path* in X is just a finite sequence of 0–simplices $\gamma(0), \gamma(1), \dots, \gamma(K)$, which we write in function notation as $\gamma: [0, K] \rightarrow X$. A *family of paths* in X is an indexed collection $\{\gamma_i\}_{i \in I}$ of paths in X ; we allow repetition in the family. A family of paths in X is said to be *almost transitive* if there exists a constant A such that for any 0–simplices v, w there is a path $\gamma: [0, K] \rightarrow X$ in the family such that all of the distances $d(v, \gamma(0)), d(\gamma(0), \gamma(1)), \dots, d(\gamma(K-1), \gamma(K)), d(\gamma(K), w)$ are at most A .

Given a path $\gamma: [0, K] \rightarrow X$ and a function $\pi: X \rightarrow [0, K]$, called the “projection map” to the path γ , and given constants $a, b, c > 0$, consider the following three axioms:

Coarse retraction For all $k \in [0, K]$ the set $\gamma[k, \pi(\gamma(k))]$ has diameter at most c .

Coarse Lipschitz For all vertices $v, w \in X$, if $d(v, w) \leq 1$ then the set $\gamma[\pi(v), \pi(w)]$ has diameter at most c .

Strong contraction For all vertices $v, w \in X$, if $d(v, \gamma[0, K]) \geq a$, and if $d(w, v) \leq b \cdot d(v, \gamma[0, K])$, then the set $\gamma[\pi(v), \pi(w)]$ has diameter at most c .

Theorem 3.1 [15, Theorem 2.3] *Given a connected simplicial complex X , if there exists an almost transitive family of paths $\{\gamma_i\}_{i \in I}$ in X and for each $i \in I$ a projection map $\pi_i: X \rightarrow [0, K]$ to the path $\gamma_i: [0, K] \rightarrow X$ such that the coarse retraction, the coarse Lipschitz and the strong contraction axioms all hold with uniform constants $a, b, c > 0$ for all $i \in I$, then X is hyperbolic.*

Remarks Our notion of “almost transitivity” is not quite the same as “coarse transitivity” used in [15], which requires that the paths in the set be continuous and that there is a constant D such that any two points at distance at least D are connected by a path in the set. However, the proof of equivalence of the two forms of the axioms, one with “almost transitive” and the other with “coarse transitive”, is very simple, and is left to

the reader. The set of fold paths in $\mathcal{FS}'(F)$ is almost transitive with constant $A = 2$: for any free splittings S, T , by moving S a distance at most 2 one obtains a naturally foldable morphism to T (Lemma 2.4), which has a fold factorization (Section 2.3), and consecutive free splittings in such a factorization have distance at most 2 (Lemma 2.5).

The concept of a “family of paths” is left undefined in [15] but the proof of the above theorem and the application to curve complexes given in [15] clearly indicate that an indexed family with repetition is allowed. On top of that, given any indexed family satisfying the hypothesis of the theorem, if we removed repetition by kicking out all but one copy of each path then the resulting family would still satisfy the hypotheses of the theorem. In our situation, although we use fold paths in our application of the above theorem, we shall index them by (equivalence classes of) fold sequences; thus, we allow for the possibility that two inequivalent fold sequences representing the same fold path might have somewhat different projection maps.

Notice that the *strong contraction* axiom, unlike the *coarse Lipschitz* axiom, is not symmetric in the variables v, w . For our proof we shall need to extend the applicability of the *strong contraction* axiom by further desymmetrizing it:

Desymmetrized strong contraction For all vertices $v, w \in X$, if $\pi(w) \leq \pi(v)$ in the interval $[0, K]$, if $d(v, \gamma[0, K]) \geq a$, and if $d(w, v) \leq b \cdot d(v, \gamma[0, K])$, then the set $\gamma[\pi(v), \pi(w)]$ has diameter at most c .

Lemma 3.2 For all constants $a, b, c > 0$ there exist constants $A, B > 0$ such that the desymmetrized strong contraction axiom with constants a, b and c implies the strong contraction axiom with constants A, B and $C = c$.

Proof Set $A = 4a$ and $B = \min\{1/4, 3b/4\}$. We need only show that if $\pi(w) > \pi(v)$ in $[0, K]$, if $d(v, \gamma[0, K]) \geq A$ and if $d(w, v) \leq B \cdot d(v, \gamma[0, K])$, then $d(w, \gamma[0, K]) \geq a$ and $d(v, w) \leq b \cdot d(w, \gamma[0, K])$. We have

$$\begin{aligned} d(w, \gamma[0, K]) &\geq d(v, \gamma[0, K]) - d(w, v) \\ &\geq d(v, \gamma[0, K]) - \frac{1}{4} \cdot d(v, \gamma[0, K]) \\ &\geq \frac{3}{4} \cdot d(v, \gamma[0, K]) \geq 3a \geq a \end{aligned}$$

and

$$\begin{aligned} d(v, w) &\leq \frac{3}{4} \cdot b \cdot d(v, \gamma[0, K]) \\ &\leq \frac{3}{4} \cdot b \cdot \frac{4}{3} d(w, \gamma[0, K]) = b \cdot d(w, \gamma[0, K]). \quad \square \end{aligned}$$

We now define the path family $\{\gamma_i\}_{i \in \mathcal{I}}$ in $\mathcal{FS}'(F)$ that we use to prove the Main Theorem. Having associated to each fold sequence a fold path, which clearly depends only

on the equivalence class of that fold sequence, the index set is defined to be the set of equivalence classes of fold sequences.

To prove the Main Theorem, by combining the Masur–Minsky Theorem, almost transitivity of fold paths, and Lemma 3.2, it therefore suffices to prove:

Proposition 3.3 *Associated to each fold sequence $S_0 \mapsto \cdots \mapsto S_K$ in $\mathcal{FS}'(F)$ there is a projection map $\pi: \mathcal{FS}'(F) \rightarrow [0, K]$, depending only on the equivalence class of the fold sequence, such that the coarse retraction, the coarse Lipschitz and the desymmetrized strong contraction axioms all hold, with constants a, b, c depending only on $\text{rank}(F)$.*

The next step in the proof of the Main Theorem will be taken with the formulation of Proposition 4.2, where the projection maps are defined.

Remark [15, Theorem 2.3] contains an additional conclusion, which in our context says that fold paths may be reparametrized to become uniform quasigeodesics in $\mathcal{FS}'(F_n)$, although the reparametrization does not fall out explicitly from their proof. Our method of proof will actually yield an explicit quasigeodesic reparametrization of fold paths, in terms of the “free splitting units” introduced in Section 5. See Proposition 6.2 for the statement and proof regarding this reparametrization.

4 Combing

In this section we define a combing method for fold sequences. Roughly speaking, given a fold sequence $S_0 \mapsto \cdots \mapsto S_K$ and a free splitting T' that differs from S_K by a single edge in $\mathcal{FS}'(F)$, we want a construction that combs backwards to produce a fold sequence $T_0 \mapsto \cdots \mapsto T_K = T'$ in which each T_k differs from the corresponding S_k by at most a single edge in $\mathcal{FS}'(K)$. We would like to give this construction in two cases, depending on whether the oriented edge from S_K to T' is a collapse $S_K > T'$ or an expansion $S_K < T'$. In the case of a collapse $S_K > T'$ there is indeed a process of “combing by collapse”, which produces a fold sequence as stated; see Proposition 4.3. In the case of an expansion $S_K < T'$, although there is a process of “combing by expansion”, the sequence $T_0 \mapsto \cdots \mapsto T_K = T'$ produced need not be a fold sequence; instead, it belongs to a broader class of map sequences that we refer to as “foldable sequences”; see Proposition 4.4. It is an important part of our theory that both combing processes are closed on the collection of foldable sequences; combing by collapse is closed as well on the smaller collection of fold sequences.

In Section 4.1 we define the collection of foldable sequences on which combing will be defined, and we define *combing rectangles*, which are the commutative diagrams of

foldable sequences and collapse maps that are used to describe combing; see Figure 1. We also prove Lemma 4.1, which says that combing by collapse is closed on foldable sequences.

The two main combing processes — combing by collapse, and combing by expansion — are described in Section 4.2. In Section 4.3 we will also give some methods for constructing new combing rectangles by composing or decomposing old ones.

Also in Section 4.1, combing rectangles will be used to define the projection map from $\mathcal{FS}'(F)$ to each fold path $S_0 \mapsto \cdots \mapsto S_K$, and we will state Proposition 4.2, which says that these projection maps satisfy the requirements of the Masur–Minsky axioms.

Combing rectangles will be important structures for the rest of the paper. Much of the geometric intuition behind our methods involves visualizing combing rectangles and other, more complicated diagrams of free splittings and maps as objects sitting in the complex $\mathcal{FS}'(F)$, and visualizing various methods for geometrically manipulating these objects. The technical details of the proof of the Main Theorem will involve a calculus of combing rectangles, which is based on the constructions of combing rectangles given in Sections 4.2 and 4.3.

4.1 Combing rectangles and the projection onto fold paths

Foldable sequences Consider a sequence of free splittings and maps over F of the form

$$S_0 \xrightarrow{f_1} S_1 \xrightarrow{f_2} \cdots \xrightarrow{f_K} S_K,$$

and recall the notation $f_j^i = f_j \circ \cdots \circ f_{i+1}: S_i \rightarrow S_j$ for each $0 \leq i < j \leq K$. This sequence is said to be a *foldable sequence* over F if for each $i = 0, \dots, K$ the map $f_K^i: S_i \rightarrow S_K$ is a foldable map. It follows that each of the maps $f_j^i: S_i \rightarrow S_j$ is a foldable map, $0 \leq i < j \leq K$, because for each vertex $v \in S_i$, the f_j^i -gate decomposition of $D_v S_i$ is a refinement of the f_K^i -gate decomposition.

Combing rectangles A *combing rectangle* over F is a commutative diagram of maps over F having the form depicted in Figure 1, such that:

- (1) The top horizontal row is a foldable sequence.
- (2) Each vertical arrow $\pi_i: S_i \rightarrow T_i$ is a collapse map with collapsed subgraph $\sigma_i \subset S_i$ indicated in the notation.
- (3) For all $i = 1, \dots, K$ we have $\sigma_{i-1} = f_i^{-1}(\sigma_i)$. Equivalently, for all $0 \leq i < j \leq K$, we have $\sigma_i = (f_j^i)^{-1}(\sigma_j)$.




Figure 1: A combing rectangle. Horizontal sequences are foldable: the top by definition and the bottom by Lemma 4.1. Vertical arrows are collapses and $\sigma_{i-1} = f_i^{-1}(\sigma_i)$.

As mentioned earlier, combing is not closed on the set of fold sequences. We will eventually prove that combing is closed on the set of all foldable sequences; the following proves this in part, by showing closure under “combing by collapse”.

Lemma 4.1 *For any combing rectangle notated as in Figure 1, the bottom row is a foldable sequence.*

We put off the proof of Lemma 4.1 until after the definition of the projection map.

Projecting onto fold paths Given a free splitting $F \curvearrowright T$, a fold sequence $S_0 \mapsto \cdots \mapsto S_K$, and an integer $k \in [0, K]$, a *projection diagram from T to $S_0 \mapsto \cdots \mapsto S_K$ of depth k* is a commutative diagram of free splittings and maps over F of the form depicted in Figure 2, such that each horizontal row is a foldable sequence, and the two rectangles shown are combing rectangles.




Figure 2: A projection diagram of depth k from T to $S_0 \mapsto \cdots \mapsto S_K$

The *projection* $\pi(T) \in [0, \dots, K]$ of T to $S_0 \mapsto \cdots \mapsto S_K$ is defined to be the maximum depth of any projection diagram from a free splitting conjugate to T to a fold sequence equivalent to $S_0 \mapsto \cdots \mapsto S_K$, if such a diagram exists, and $\pi(T) = 0$

otherwise. It is clear that this gives a well-defined function $\pi: \mathcal{FS}'(F) \rightarrow [0, \dots, K]$ that depends only on the equivalence class of the fold sequence $S_0 \mapsto \dots \mapsto S_K$.

One way to understand this definition is to think of $\mathcal{FS}'(F)$ as being Gromov hyperbolic and to think of fold paths as being quasigeodesic, all of which are true a posteriori assuming that Proposition 3.2 is true. That being so, given a fold path $S_0 \mapsto \dots \mapsto S_K$ and T projecting to $\pi(T) \in [0, \dots, K]$, by moving to some point S'_0 nearby S_0 we should obtain a fold path from S'_0 to T having an initial segment that fellow travels the given fold path from S_0 to $S_{\pi(T)}$ and no farther. In defining the projection map as above, the intuition is that combing rectangles provide an adequate definition of fellow traveling. The technical details of the definition were crafted to what would work in our proofs, but also received some original motivation from results of [15], which amount to a proof that for any finite type oriented surface S , splitting sequences of train tracks on S define quasigeodesics in the curve complex of S . In particular, [15, Lemma 4.4] — which can be regarded as a verification of the *coarse Lipschitz* axiom — contains the statement “ $\beta \in PE(\sigma)$ ”, and if one works out the train track diagram for that statement one obtains a rather strong analogue of our projection diagram above.

The rest of the paper is devoted to the proof of the following, which immediately implies Proposition 3.3 and therefore implies the Main Theorem:

Proposition 4.2 *There exist $a, b, c > 0$ depending only on $\text{rank}(F)$ such that for any fold sequence $S_0 \mapsto \dots \mapsto S_K$ in $\mathcal{FS}'(F)$, the projection map $\pi: \mathcal{FS}'(F) \rightarrow [0, \dots, K]$ defined above satisfies the coarse retraction, coarse Lipschitz and desymmetrized strong contraction axioms with constants a, b, c .*

The *coarse retraction* axiom is proved in Proposition 5.9 and the other two axioms are proved in Section 6.

We now turn to:

Proof of Lemma 4.1 Following the notation of Figure 1, we must show that each map $g_K^i: T_i \rightarrow T_K$ is foldable. First note that g_K^i is injective on each edgelet $e \subset T_i$, because $e = \pi_i(\tilde{e})$ for some edgelet $\tilde{e} \subset S_i \setminus \sigma_i$, so $f_K^i(\tilde{e}) \subset S_K \setminus \sigma_K$, so $\pi_K(f_K^i(\tilde{e})) = g_K^i(\pi_i(\tilde{e})) = g_K^i(e)$ is an edgelet of T_K .

Given a vertex $w \in T_i$, we must show that g_K^i has at least 2 gates at w , and that if w is natural then g_K^i has at least 3 gates. Let $w' = g_K^i(w) \in T_K$. We have a subgraph $W' = \pi_K^{-1}(w') \subset S_K$, and a subgraph $W = \pi_i^{-1}(w) \subset S_i$ such that $f_K^i(W) \subset W'$. Note that each direction in $D_W S_i$ is based at a frontier vertex of W and is represented by an edgelet of $S_i \setminus \sigma_i$, and similarly for $D_{W'} S_K$, and so these direction sets are in

the domain of definition of the derivative maps $d\pi_i, d\pi_K$, respectively. We have a commutative diagram of derivatives

$$\begin{array}{ccc} D_W S_i & \xrightarrow{df_K^i} & D_{W'} S_K \\ d\pi_i \downarrow & & \downarrow d\pi_K \\ D_w T_i & \xrightarrow{d_w g_K^i} & D_{w'} T_K \end{array}$$

in which the vertical maps are bijections and so $d\pi_i$ induces a bijection between gates of $d_w g_K^i$ and point pre-images of the map in the top row. The valence of w therefore equals the cardinality of the set $D_W S_i$, and the number of gates of g_K^i at w equals the cardinality of the image of the map in the top row. If w has valence at least 2 (resp. at least 3) then we must prove that the image of the map in the top row has cardinality at least 2 (resp. at least 3).

Suppose that w is a valence 2 vertex contained in the interior of a natural edge $\eta \subset T_i$. The subgraph W is either a point or a segment contained in the interior of a natural edge $\tilde{\eta} \subset S_i$ such that $\pi_i(\tilde{\eta}) = \eta$. Let $e_1, e_2 \subset \eta$ be the two oriented edgelets incident to w , representing the two directions of the set $D_w T_i$. Let $\tilde{e}_1, \tilde{e}_2 \subset \tilde{\eta} \setminus W$ be the two oriented edgelets incident to the endpoints of W representing the two elements of the set $D_W S_i$, indexed so that $\pi_i(\tilde{e}_j) = e_j, j = 1, 2$. Since f_K^i is injective on $\tilde{\eta}$ it follows that $f_K^i(\tilde{e}_1), f_K^i(\tilde{e}_2)$ are distinct edgelets of S_K . Noting that

$$g_K^i(e_j) = g_K^i(\pi_i(\tilde{e}_j)) = \pi_K(f_K^i(\tilde{e}_j))$$

for $j = 1, 2$, it follows that these are two distinct edgelets of T_K , and so g_K^i has 2 gates at w .

Suppose now that w is a vertex of valence at least 3, so the set $D_W S_i$ has cardinality at least 3. If W is a point then it has valence at least 3 and, since f_K^i is foldable, there are at least 3 gates of f_K^i in $D_W S_i$; it follows that there are at least 3 gates of g_K^i in $D_w S_i$. If W has infinite diameter then $D_W S_i$ is infinite and so w has infinite valence, implying that g_K^i has infinitely many gates at w . If W does not contain a natural vertex of S_i then it is a segment in the interior of a natural edge of S_i implying that w has valence 2, a contradiction.

We have reduced to the case that the graph W has finite diameter, is not a point, and contains a natural vertex of S_i . The graph $f_K^i(W)$ also has finite diameter and is not a point, and so has $P \geq 2$ vertices of valence 1 (the cardinality P may be countably infinite). Let $X \subset W$ be a set consisting of one vertex of W in the preimage of each valence 1 vertex of $f_K^i(W)$. By the first derivative test, each $x \in X$ is a frontier vertex

of W . Choosing a direction $\delta_x \in D_W S_i$ based at each $x \in X$, it follows that the directions $df_K^i(\delta_x)$ are based at P distinct points of S_K and are therefore P distinct directions in the set $D_{W'} S_K$. If $P \geq 3$ then we are done.

We have reduced further to the subcase that $P = 2$, and so $f_K^i(W)$ is a segment with endpoints u_1, u_2 . We have $X = \{x_1, x_2\}$ with $f_K^i(x_j) = u_j$. Consider a natural vertex $v \in S_i$ such that $v \in W$, and its image $v' = f_K^i(v) \in f_K^i(W)$. Since f_K^i is foldable, there are at least 3 gates at v with respect to f_K^i . If $v' = u_j$ then at least one of the gates at v maps to a direction at u_j that is distinct from the direction $df_K^i(\delta_{x_j})$ and from the unique direction of the segment $f_K^i(W)$ at u_j , and so we have found a third direction in the set $D_{W'} S_K$. If v' is an interior point of the segment $f_K^i(W)$ then at least one of the gates at v maps to a direction at v' distinct from the two directions of the segment $f_K^i(W)$ at v' and again we have found a third direction in $D_{W'} S_K$. \square

4.2 Combing by collapse and combing by expansion

In approaching the proof of Proposition 4.2, one immediately confronts the need for constructions of combing rectangles, in order to construct the projection diagrams needed to compute projection maps. This section and the next contain the constructions of combing rectangles that we use for this purpose.

Our first construction of combing rectangles shows how to comb a foldable sequence followed by a collapse map.

Proposition 4.3 (Combing by collapse) *For each foldable sequence*

$$S_0 \xrightarrow{f_1} S_1 \xrightarrow{f_2} \dots \xrightarrow{f_K} S_K,$$

and for each collapse

$$S_K \xrightarrow{[\sigma_K]} T'$$

there exists a combing rectangle of the form shown in Figure 1 such that $T_K = T$.

Proof Define an equivariant subgraph $\sigma_i \subset S_i$ using the definition of a combing rectangle: starting with $\sigma_K \subset S_K$, by induction define $\sigma_i = f_{i+1}^{-1}(\sigma_{i+1})$. Since $\sigma_K \subset S_K$ is a proper equivariant subgraph it follows by induction that each $\sigma_i \subset S_i$ is a proper equivariant subgraph, and so free splittings $F \curvearrowright T_i$ with collapse maps

$$S_i \xrightarrow{[\sigma_i]} T_i$$

and induced maps $g_i: T_{i-1} \rightarrow T_i$ are all defined, and the squares are all evidently commutative. \square

We remark that the cheapness of the above proof is slightly offset by the modest expense of proving that the T_i sequence is foldable, which was done back in Lemma 4.1.

Next we explain how to comb a foldable sequence followed by an expansion. In sharp contrast to the case of combing by collapse, both the construction of the combing rectangle and the proof that the resulting map sequence is foldable are very intricate in the case of combing by expansion.

Proposition 4.4 (Combing by expansion) *For each foldable sequence*

$$S_0 \xrightarrow{f_1} S_1 \xrightarrow{f_2} \dots \xrightarrow{f_K} S_K,$$

each expansion $S_K < T'$, and each collapse map $\pi': T' \rightarrow S_K$, there exists a combing rectangle of the form:

$$\begin{array}{ccccccc} S_0 & \xrightarrow{f_1} & \dots & \xrightarrow{f_{i-1}} & S_{i-1} & \xrightarrow{f_i} & S_i & \xrightarrow{f_{i+1}} & \dots & \xrightarrow{f_K} & S_K \\ \uparrow [\sigma_0] \pi_0 & & & & \uparrow [\sigma_{i-1}] \pi_{i-1} & & \uparrow [\sigma_i] \pi_i & & & & \uparrow [\sigma_K] \pi_K = \pi' \\ T_0 & \xrightarrow{g_1} & \dots & \xrightarrow{g_{i-1}} & T_{i-1} & \xrightarrow{g_i} & T_i & \xrightarrow{g_{i+1}} & \dots & \xrightarrow{g_K} & T_K = T' \end{array}$$

Remark Implicit in the conclusion via the definition of combing rectangle is that the sequence

$$T_0 \xrightarrow{g_1} \dots \xrightarrow{g_K} T_K$$

is foldable.

Proof We will construct this combing rectangle in two steps. In Step 1 we produce a commutative diagram of free splittings and maps of the form

$$\begin{array}{ccccccc} S_0 & \xrightarrow{f_1} & \dots & \xrightarrow{f_{i-1}} & S_{i-1} & \xrightarrow{f_i} & S_i & \xrightarrow{f_{i+1}} & \dots & \xrightarrow{f_K} & S_K \\ \uparrow [\sigma'_0] \pi'_0 & & & & \uparrow [\sigma'_{i-1}] \pi'_{i-1} & & \uparrow [\sigma'_i] \pi'_i & & & & \uparrow [\sigma'_K] \pi'_K = \sigma' \\ U_0 & \xrightarrow{h_1} & \dots & \xrightarrow{h_{i-1}} & U_{i-1} & \xrightarrow{h_i} & U_i & \xrightarrow{h_{i+1}} & \dots & \xrightarrow{h_K} & U_K = T' \end{array}$$

in which each π'_i is a collapse and $h_i^{-1}(\sigma'_i) = \sigma'_{i-1}$, but the U row slightly fails to be foldable in that certain explicitly described natural vertices of U_i are “bad” by fault of having only 2 gates with respect to $h_K^i: U_i \rightarrow U_K$. One of these gates will always be a singleton, and so each “bad natural vertex” will be incident to a “bad natural edge”. In Step 2 we will repair this problem by splitting each bad natural edge, to produce a

commutative diagram of the form:

$$\begin{array}{ccccccccccc}
 U_0 & \xrightarrow{h_1} & \cdots & \xrightarrow{h_{i-1}} & U_{i-1} & \xrightarrow{h_i} & U_i & \xrightarrow{h_{i+1}} & \cdots & \xrightarrow{h_K} & U_K & \equiv & T' \\
 \mu_0 \uparrow & & & & \mu_{i-1} \uparrow & & \mu_i \uparrow & & & & \mu_K \parallel & & \\
 T_0 & \xrightarrow{g_1} & \cdots & \xrightarrow{g_{i-1}} & T_{i-1} & \xrightarrow{g_i} & T_i & \xrightarrow{g_{i+1}} & \cdots & \xrightarrow{g_K} & T_K & \equiv & T
 \end{array}$$

The T row will be a foldable sequence. The μ_i maps are not collapses but instead are “multifolds” that invert the splitting process. The desired combing rectangle will be obtained by concatenating these two rectangles: the composition

$$\pi_i: T_i \xrightarrow{\mu_i} U_i \xrightarrow{\pi'_i} S_i$$

will indeed be a collapse map, which collapses the subgraph $\sigma_i = \mu_i^{-1}(\sigma'_i) \subset T_i$.

Step 1 The free splitting $F \curvearrowright U_i$ is defined to be the minimal subtree of the pushout of S_i and T' . Here are more details. As a set, the pushout of S_i and T' is

$$\wedge(S_i, T') = \{(x, y) \in S_i \times T' \mid f_K^i(x) = \pi'(y)\}.$$

The action $F \curvearrowright \wedge(S_i, T')$ is obtained by restricting the diagonal action $F \curvearrowright S_i \times T'$. The restrictions of the two projection maps are denoted

$$\pi'_i: \wedge(S_i, T') \rightarrow S_i \quad \text{and} \quad h_{T'}^i: \wedge(S_i, T') \rightarrow T'.$$

Both are clearly F -equivariant and we have $f_K^i \circ \pi'_i = \pi' \circ h_{T'}^i: \wedge(S_i, T') \rightarrow S_K$. As a graph, the vertices and edgelets of the pushout are as follows. A vertex is a pair $(v, w) \in \wedge(S_i, T')$ such that v is a vertex of S_i and w is a vertex of T' . Edgelets are of two types. First, a *collapsed edgelet* is one of the form $v \times e'$ where $v \in S_i$ is a vertex and $e' \subset \sigma' \subset T'$ is an edgelet such that $\pi'(e') = f_K^i(v)$; the barycentric coordinates on e' induce those on $v \times e'$ via the projection $h_{T'}^i$. Second, to each edgelet $e \subset S_i$ there corresponds a unique edgelet $e' \subset T'$ with the property that $f_K^i(e) = \pi'(e')$ (uniqueness follows since π' is a collapse map), and there corresponds in turn an *uncollapsed edgelet* $\tilde{e} = \wedge(e, e') = \{(x, y) \in \wedge(S_i, T') \mid x \in e, y \in e'\}$ of $\wedge(S_i, T')$ with barycentric coordinates induced via the map $f_K^i \circ \pi'_i = \pi' \circ h_{T'}^i$, which takes \tilde{e} bijectively to the edgelet $f_K^i(e) = \pi'(e')$ of S_K . The action of F on $\wedge(S_i, T')$ and the projection maps $\pi'_i, h_{T'}^i$ are each simplicial. The simplicial complex $\wedge(S_i, T')$ is 1-dimensional by construction. It is furthermore a tree, in that removal of the interior of each edgelet separates, because the simplicial map $\pi'_i: \wedge(S_i, T') \rightarrow S_i$ is injective over the interior of each edgelet of S_i , and for each vertex $x \in S_i$ the subcomplex $(\pi'_i)^{-1}(x)$ is a tree (mapped by a simplicial isomorphism to the tree $(\pi')^{-1}(f_K^i(x)) \subset T'$). The action $F \curvearrowright S_i$ has no point fixed by each element of F , and so neither does the

action $F \curvearrowright \wedge(S_i, T')$; it follows that the F -tree $\wedge(S_i, T')$ contains a unique minimal F -invariant subtree, which, by definition, is U_i . For each edgelet $e \subset \wedge(S_i, T')$, its stabilizer is contained in $\text{Stab}_{S_i}(\pi'_i(e))$ if e is uncollapsed and in $\text{Stab}_{T'}(h^i_{T'}(e))$ if e is collapsed, and in either case is trivial. This proves that $F \curvearrowright U_i$ is a free splitting.

Here are some structural facts about the tree U_i . For each edgelet $e \subset S_i$, the edgelet $\tilde{e} \subset \wedge(S_i, T')$ is the unique one mapped to e via π'_i , and since $F \curvearrowright S_i$ is minimal the map $\pi'_i: U_i \rightarrow S_i$ is surjective, which forces \tilde{e} to be contained in U_i . This also shows that π'_i is a collapse map. The union of the collapsed edgelets of the pushout $\wedge(S_i, T')$ forms a subgraph $\Sigma_i \subset \wedge(S_i, T')$ with one component $\Sigma_{i,v} = (\pi'_i)^{-1}(v)$ for each vertex $v \in S_i$ such that $(\pi')^{-1}(f^i_K(v))$ is a component of σ' ; the map $h^i_{T'}$ restricts to a simplicial isomorphism between these components. The subgraph of $\sigma'_i \subset U_i$ that is collapsed by $\pi'_i: U_i \rightarrow S_i$ is the union of those components of $\Sigma_i \cap U_i$ that contain at least one edge. Each of these components has the form $\sigma'_{i,v} = \Sigma_{i,v} \cap U_i$ when this intersection contains at least one edge; by convention we set $\sigma'_{i,v} = \emptyset$ otherwise. See below for a more detailed description of various features of $\sigma'_{i,v}$.

There is an induced map $h_i: \wedge(S_{i-1}, T') \rightarrow \wedge(S_i, T')$ that is defined by the formula $h_i(x, y) = (f_i(x), y)$, which makes sense because for each $(x, y) \in \wedge(S_{i-1}, T')$ we have $f^i_K(f_i(x)) = f^{i-1}_K(x) = \pi'(y)$. The commutativity equation $\pi'_i \circ h_i = f_i \circ \pi'_{i-1}$ is immediate. Since U_i is the minimal subtree of $\wedge(S_i, T')$ it follows that $h_i(U_{i-1}) \supset U_i$, but we are not yet in a position to prove the opposite inclusion, not until we have established that the map $h^i_{T'}: U_i \rightarrow T'$ has at least 2 gates at each vertex.

Preparation for Step 2 Here are some structural facts about the components of σ'_i . Consider a vertex $v \in S_i$ for which $\sigma'_{i,v} \neq \emptyset$ and so is a component of σ'_i . Given an oriented edgelet $e \subset S_i$ we abuse notation by writing $e \in D_v S_i$ to mean that v is the initial vertex of e . There is a function $\xi_{i,v}: D_v S_i \rightarrow U_i$ where for each $e \in D_v S_i$ we define $\xi_{i,v}(e) \in \sigma'_{i,v}$ to be the initial vertex of the corresponding oriented edgelet $\tilde{e} \subset U_i$. Note that the set $\text{image}(\xi_{i,v})$ is the topological frontier of the subtree $\sigma'_{i,v}$ in the tree U_i . By Lemma 1.2(3) it follows that $\sigma'_{i,v}$ is the convex hull of the set $\text{image}(\xi_{i,v})$ in U_i . Notice also that the function $\xi_{i,v}$ is constant on each gate of $D_v S_i$ with respect to the map f^i_K , for if $e_1, e_2 \in D_v S_i$ are in the same gate then $f^i_K(e_1) = f^i_K(e_2)$ is a single edgelet in S_K , which lifts to a unique edgelet $e' \subset T'$ and we have

$$h^i_K(\tilde{e}_1) = \wedge(f^i_K(e_1), e') = \wedge(f^i_K(e_2), e') = h^i_K(\tilde{e}_2)$$

and so the initial endpoints of \tilde{e}_1 and \tilde{e}_2 have the same image under h^i_K . But these initial endpoints are in the graph $\sigma'_{i,v}$ on which h^i_K is injective, so these initial endpoints are equal. Letting $\Gamma_v S_i$ denote the set of gates of f^i_K in $D_v S_i$, the map $\xi_{i,v}$ induces a map, which we also denote $\xi_{i,v}: \Gamma_v S_i \rightarrow \sigma'_{i,v}$, whose image is also the frontier of $\sigma'_{i,v}$.

We now study the extent to which the maps $h_K^i: U_i \rightarrow U_K$ are foldable. Note first that we may identify T' with the pushout $\wedge(S_K, T')$ and so we may identify $U_K = T'$ and $\sigma'_K = \sigma'$ up to simplicial conjugacy and we may identify $h_K^i = h_{T'}^i$; in particular, the gates of h_K^i and of $h_{T'}^i$ are therefore identical. We will show that $h_{T'}^i$ has at least 2 gates at each vertex of U_i , so a vertex is either *good*, meaning it has valence at least 3 and at least 3 gates or valence 2 and 2 gates, or *bad*, meaning it has valence at least 3 but only 2 gates. We shall do this through a case analysis, going through various cases of good vertices and narrowing down to one remaining case, which is bad. This will yield an explicit description of the bad vertices, which will be used in describing the free splitting $F \curvearrowright T_i$.

Fix a vertex $u = (v, w) \in U_i$, so if $\sigma'_{i,v} \neq \emptyset$ then $u \in \sigma'_{i,v}$. Denote $x = f_K^i(v) = \pi'(w)$.

Consider first the case that $\sigma'_{i,v} = \emptyset$; we shall show that u is good. We have a commutative diagram of derivative maps

$$\begin{array}{ccc}
 D_v S_i & \xrightarrow{d_v f_K^i} & D_x S_K \\
 d_u \pi'_k \uparrow & & \uparrow d_w \pi' \\
 D_u U_i & \xrightarrow{d_u h_{T'}^i} & D_w T'
 \end{array}$$

where the left arrow is a bijection, i.e., the valences of u and v are equal. Also, the set $\text{image}(d_u h_{T'}^i)$ is in the domain of definition of the right arrow and the right arrow is an injection on its domain of definition. The number of gates at u, v are therefore equal. Since f_K^i is foldable it follows that u is good.

Consider now the case that $\sigma'_{i,v} \neq \emptyset$. To simplify notation we denote $W = \sigma'_i$ and $W_v = \sigma'_{i,v}$. Each gate of $h_{T'}^i$ in DU_i is contained either in DW or its complement $D(U_i \setminus W) = DU_i - DW$, because $W = \sigma'_i = (h_{T'}^i)^{-1}(\sigma')$ implying that $h_{T'}^i$ never maps a direction of W and a direction of $U_i \setminus W$ to the same direction of T' . Since $h_{T'}^i$ is injective on W_v , each direction in the set $D_u W_v$ constitutes an entire gate of $D_u U_i$. Gates at u in the complement $D_u U_i - D_u W_v$ exist if and only if u is a frontier vertex of W_v , if and only if u is in the image of $\xi_{i,v}: D_v S_i \rightarrow W_v$.

Consider the subcase that v has valence 2 in S_i . The graph W_v is then a segment contained in the interior of a natural edge of U_i . The vertex u therefore has valence 2 in U_i , with either 2 directions in W_v or one each in W_v and in $U_i \setminus W_v$, and in either case these 2 directions are mapped by $h_{T'}^i$ to two different directions in T' and so u is good.

Consider the subcase that v has valence at least 3 in S_i . If the valence of u in W_v plus the number of gates at u in the complement of W_v is at least 3 then $h_{T'}^i$ has at

least 3 gates at u , so u is good. If u is an interior vertex of W_v then u has valence at least 2 in W_v by minimality of $F \curvearrowright U_i$; furthermore, the valences of u in W_v and in U_i are equal and the number of gates of h_T^i , at u equals the valence, so u is good. If u is a frontier vertex of valence at least 2 in W_v then u has at least 1 gate in the complement of W_v and we considered this case already and showed that u is good. If u is a frontier vertex of valence 1 in W_v and if u has at least 2 gates in the complement of W_v then we have also considered this case already and showed that u is good. If u is a frontier vertex of valence 1 in W_v and u has exactly 1 direction in the complement of W_v then u has valence 2 in U_i and 2 gates, so u is good.

The only case that remains, and the case that characterizes when u is bad, is when v has valence at least 3 in S_i , u is a frontier vertex of W_v , u has valence 1 in W_v , u has exactly one gate in the complement of W_v and that gate has cardinality $\xi_u \geq 2$ called the *external valence* of u . When in this case, let ζ_u be the unique natural edge of U_i with endpoint u and with direction at u equal to the unique direction of W_v at u . We call ζ_u the *bad natural edge* incident to u . Let z_u be the natural endpoint of ζ_u opposite u .

We claim that for each bad natural vertex $u \in U_i$ we have $\zeta_u \subset W_v$; the only way this could fail is if W_v is an edgelet path whose vertices apart from u all have valence 2 in U_i , implying that f_K^i has 2 gates at the natural vertex v , contradicting that f_K^i is foldable. We claim also that z_u is good; otherwise it would follow that $W_v = \zeta_u = \zeta_{z_u}$, which again would imply the contradiction that f_K^i has 2 gates at v .

The union of the bad natural edges of U_i forms an equivariant natural subgraph denoted $Z_i = \bigcup \zeta_u \subset U_i$. The natural edges of its complement $U_i \setminus Z_i$ are the *good natural edges* of U_i , some of which may be contained in W , some in $U_i \setminus W$, and some in neither. The endpoints of a good natural edge need not be good. From the description of bad natural edges it follows that each component of Z_i contains a unique good vertex z and is the union of some number $m \geq 1$ of bad natural edges with endpoint z , forming a star graph with m valence 1 vertices apart from z .

Step 2 Ignoring the simplicial structure for the moment, define the free splitting $F \curvearrowright T_i$ to be the one obtained from $F \curvearrowright U_i$ by collapsing the bad subgraph $Z_i \subset U_i$. Let

$$\rho_i: U_i \xrightarrow{[Z_i]} T_i$$

be the collapse map. Note that ρ_i restricts to an equivariant bijection from the good natural vertices of U_i to the natural vertices of T_i , because Z_i is a natural subgraph, each of whose components contains exactly one good natural vertex. Also, ρ_i induces a bijection from the good natural edges of U_i — those in $U_i \setminus Z_i$ — to the natural edges

of T_i : denote this correspondence by $\tilde{\eta} \leftrightarrow \eta$ for each good natural edge $\eta \subset T_i$, and note that ρ_i maps $\tilde{\eta}$ homeomorphically to η .

Define the map $\mu_i: T_i \rightarrow U_i$ as follows. The restriction of μ_i to the natural vertices of T_i is the equivariant bijection onto the good natural vertices of U_i that is obtained by inverting ρ_i . The endpoints of each natural edge of T_i map to distinct points of U_i , and so μ_i may be extended equivariantly and continuously to be an injection on each natural edge of T_i .

Define the simplicial structure on T_i to be the unique one with respect to which μ_i is a simplicial map: its vertices are the inverse image under μ_i of the vertices of U_i ; each of its edgelets maps via μ_i by simplicial isomorphism to an edgelet of U_i .

Define the subgraph $\sigma_i \subset T_i$ to be $\mu_i^{-1}(\sigma'_i)$; we will see below that $\pi'_i \circ \mu_i: T_i \rightarrow S_i$ is a collapse map that collapses the subgraph σ_i .

Knowing that μ_i is injective on each natural edge of T_i , we describe the image of each natural edge as follows. The notation $u \mapsto z_u$, which so far defines an equivariant function from the bad natural vertices of U_i to the good natural vertices of U_i , extends to all natural vertices of U_i by defining $z_u = u$ when u is good. For each natural vertex $u \in U_i$ we have $\mu_i(\rho_i(u)) = z_u$: if $u = z_u$ is good this is because μ_i and ρ_i are inverse bijections between good natural vertices of U_i and all natural vertices of T_i ; if u is bad then u and z_u are contained in the same component of Z_i so $\rho_i(u) = \rho_i(z_u)$ and hence $\mu_i(\rho_i(u)) = \mu_i(\rho_i(z_u)) = z_u$. Given a natural edge $\eta \subset T_i$ with corresponding good natural edge $\tilde{\eta} \subset U_i$, letting $u_1, u_2 \in U_i$ be the endpoints of $\tilde{\eta}$ and letting $z_1 = z_{u_1} \in U_i$, it follows that $\mu_i(\eta) = \mu_i(\rho_i(\tilde{\eta}))$ is the arc in U_i connecting z_1 to z_2 , which is just the union of $\tilde{\eta}$ together with the bad natural edges incident to whichever of u_1, u_2 are bad.

From this description of μ_i we derive a few more properties of μ_i , giving details about its behavior over good and bad natural edges of U_i , and its behavior on natural edges and natural vertices of T_i .

- (a) **μ_i over good natural edges of U_i** The map μ_i is injective over the interior of each good natural edge $\tilde{\eta} \subset U_i$, the closure of $\mu_i^{-1}(\text{int}(\eta))$ is an edgelet path contained in η , and the restriction of μ_i to this edgelet path is a simplicial isomorphism onto $\tilde{\eta}$.
- (b) **μ_i over bad natural edges of U_i** For each bad natural edge $\zeta_u \subset U_i$ oriented to have terminal point u and initial point z_u , letting χ_u be the external valence of u , letting $\tilde{\eta}_\ell \subset U_i$ ($\ell = 1, \dots, \chi_u$) be the oriented good natural edges with common initial point u , and letting $\eta_\ell = \rho_i(\tilde{\eta}_\ell) \subset T_i$ be the corresponding oriented natural edges with common initial point $w = \rho_i(u)$, there exist initial

segments $[w, w_\ell] \subset \eta_\ell$, $\ell = 1, \dots, \chi_u$, such that μ_i maps each $[w, w_\ell]$ to ζ_u by a simplicial isomorphism and such that $\mu_i^{-1}(\zeta_u) = \bigcup_{\ell=1}^{\chi_u} [w, w_\ell] \subset \sigma_i$. Furthermore each w_ℓ is a valence 1 vertex of σ_i .

Intuitively (a) and (b) together say that μ_i is a “partial multifold”, which for each of its gates identifies proper initial segments of the oriented natural edges representing that gate. Perhaps the only nonobvious part of (a) and (b) is the last sentence of (b). For each bad natural vertex $u \in U_i$, from (a) and the previous sentences of (b) it follows that $\mu_i^{-1}(u) = \{w_1, \dots, w_{\chi_u}\}$, and that for each $\ell = 1, \dots, \chi_u$ the vertex w_ℓ is contained in the interior of the natural edge η_ℓ , one direction being in the segment $[w_u, w_\ell] \subset \sigma_i$ and the other direction being in the closure of $\mu_i^{-1}(\text{int}(\eta_\ell))$, which is in $T_i \setminus \sigma_i$, and so w_ℓ has valence 1 in σ_i .

- (c) **μ_i on natural edges of T_i** The restriction of μ_i to each good natural edge of T_i is injective. Furthermore, an embedded edgelet path $\alpha \subset U_i$ is the μ_i -image of some good natural edge of T_i if and only if the endpoints of α are good natural vertices of U_i , no interior point of α is a good natural vertex, and $h_K^i \upharpoonright \alpha$ is injective.

Only the “if” part of (c) is not obvious. Let $\alpha \subset U_i$ be an embedded edgelet path whose only good natural vertices are its endpoints, and suppose that $h_K^i \upharpoonright \alpha$ is injective. If α contains no bad natural vertex then $\alpha = \tilde{\eta}$ is a good natural edge with associated natural edge $\eta \subset T_i$ and $\alpha = \mu_i(\eta)$. If $u \in \alpha$ is a bad natural vertex then $u \in \text{int}(\alpha)$, and since $h_K^i \upharpoonright \alpha$ is injective it follows that one direction of α at u is the direction of the bad natural arc ζ_u , whose opposite good natural endpoint z_u must be an endpoint of α ; the edgelet path α is therefore the concatenation of some natural edge $\tilde{\eta} \subset U_i \setminus Z_i$ with any bad natural edges incident to the endpoints of $\tilde{\eta}$, and it follows that $\alpha = \mu_i(\eta)$.

- (d) **$d\mu_i$ at natural vertices of T_i** For each natural vertex $v \in T_i$, the map $d_v\mu_i: D_vT_i \rightarrow D_{\mu_i(v)}U_i$ is surjective.

To justify (d), the vertex $\mu_i(v)$ is a good natural vertex of U_i . Consider a direction $d \in D_{\mu_i(v)}U_i$. If d is the initial direction of some oriented good natural edge $\tilde{\eta} \subset U_i$ corresponding to an oriented natural edge $\eta \subset T_i$, it follows that the initial vertex of η equals v and the initial direction of η maps to d . If d is the initial direction of some bad oriented natural edge $\zeta_u \in U_i$ with opposite bad natural vertex u , let $\tilde{\eta}$ be any of the good natural edges incident to u oriented with initial vertex u , and let $\eta \subset T_i$ be the corresponding oriented natural edge, and it follows that the initial vertex of η again equals v and that the initial direction of η maps to d .

We now prove that we have a collapse map

$$\pi_i = \pi'_i \circ \mu_i: T_i \xrightarrow{\sigma_i = (\mu_i)^{-1}(\sigma'_i)} S_i.$$

Clearly an edgelet of T_i is in σ_i if and only if its image under μ_i is in σ'_i if and only if its image under $\pi_i = \pi'_i \circ \mu_i$ is a point. Given an edgelet $e \subset S_i$, the collapse map π'_i is injective over the interior of e , so there is a unique edgelet $e' \subset U_i$ mapped to e by π'_i , and $e' \not\subset \sigma'_i$; it follows that $e' \not\subset Z_i$ and so by item (a) above the map μ_i is injective over the interior of e' ; therefore π_i is injective over the interior of e .

Putting off for the moment the issue of defining the maps $g_i: T_{i-1} \rightarrow T_i$, we define the maps $g_K^i: T_i \rightarrow T_K$ as follows. First note that the map $\mu_K: T_K \rightarrow U_K$ is evidently a simplicial isomorphism, and so we may identify T_K with U_K and with T' . We now define g_K^i to be the composition

$$T_i \xrightarrow{\mu_i} U_i \xrightarrow{h_K^i} U_K \xrightarrow{(\mu_K)^{-1}} T_K.$$

The map g_K^i is foldable, equivalently $h_K^i \circ \mu_i: T_i \rightarrow U_K$ is foldable, for the following reasons: by (c) the map g_K^i is injective on natural edges of T_i ; for each natural vertex $v \in T_i$, its image $\mu_i(v) \in U_i$ is a good natural vertex and so has at least 3 gates with respect to h_K^i , and by (d) the derivative map $d_v \mu_i: D_v T_i \rightarrow D_{\mu_i(v)} U_i$ is surjective, which implies that $h_K^i \circ \mu_i$ has at least 3 gates at v .

All that remains is to define a map $g_i: T_{i-1} \rightarrow T_i$ so that the commutativity equation $h_i \circ \mu_{i-1} = \mu_i \circ g_i$ holds, for by combining this with the equation $h_K^{i-1} = h_K \circ \dots \circ h_i$ it immediately follows that $g_K^{i-1} = g_K \circ \dots \circ g_i$ and so the map sequence

$$T_0 \xrightarrow{g_1} \dots \xrightarrow{g_K} T_K$$

is defined and is foldable.

Consider a natural vertex $v \in T_{i-1}$. Its image $\mu_{i-1}(v) \in U_{i-1}$ is a good natural vertex and so has at least 3 gates with respect to h_K^{i-1} , implying that $h_i(\mu_{i-1}(v)) \in U_i$ has at least 3 gates with respect to h_K^i and so is a good natural vertex, and hence there is a unique natural vertex in T_i that maps to $h_i(\mu_{i-1}(v))$, which we take to be $g_i(v)$. We have thus defined g_i so as to satisfy the commutativity equation on each natural vertex $v \in T_{i-1}$.

Consider a natural edge $\eta \subset T_{i-1}$ with natural endpoints $v_0 \neq v_1$. Its image $\mu_{i-1}(\eta) \subset U_{i-1}$ is the arc with good natural endpoints $\mu_{i-1}(v_0) \neq \mu_{i-1}(v_1)$. By (c) above the map $h_K^{i-1} = h_K^i \circ h_i$ is injective on the arc $\mu_{i-1}(\eta)$, implying that h_i is injective on $\mu_{i-1}(\eta)$ and that h_K^i is injective on the arc $h_i(\mu_{i-1}(\eta)) \subset U_i$, the latter of which has good

natural endpoints $h_i(\mu_{i-1}(v_0)) \neq h_i(\mu_{i-1}(v_1))$. Subdividing the arc $h_i(\mu_{i-1}(\eta))$ at all interior good natural vertices of U_i , we write it as a concatenation:

$$h_i(\mu_{i-1}(\eta)) = \alpha_1 * \dots * \alpha_M.$$

Each of the arcs $\alpha_m, m = 1, \dots, M$ has good natural endpoints, no good natural interior points, and the map h_K^i is injective on α_m , and so by (c) there is a unique natural edge $\hat{\alpha}_m \subset T_i$ mapped by μ_i to α_m by a simplicial isomorphism. Since every good natural vertex in U_i has a unique natural pre-image in T_i , it follows that we may concatenate to obtain an arc $\hat{\alpha}_1 * \dots * \hat{\alpha}_m \subset T_i$, and furthermore the restriction $\mu_i | \hat{\alpha}_1 * \dots * \hat{\alpha}_m$ is a simplicial isomorphism onto $h_i(\mu_{i-1}(\eta))$. Inverting this restriction, we may then define

$$g_i | \eta = (\mu_i | \hat{\alpha}_1 * \dots * \hat{\alpha}_m)^{-1} \circ (h_i \circ \mu_{i-1}) | \eta,$$

which is a simplicial isomorphism with image $\hat{\alpha}_1 * \dots * \hat{\alpha}_m$. We have thus defined g_i so as to satisfy the commutativity equation on each natural edge $\eta \subset T_{i-1}$.

This completes the proof of Proposition 4.4. □

4.3 Composition and decomposition of combing rectangles

Lemma 4.5 (Composition of combing rectangles) *Given two combing rectangles of the form*

$$\begin{array}{ccccccc}
 S_0 & \xrightarrow{f_1} & \dots & \xrightarrow{f_i} & S_i & \xrightarrow{f_{i+1}} & \dots & \xrightarrow{f_K} & S_K \\
 \downarrow \pi_0 & & & & \downarrow \pi_i & & & & \downarrow \pi_K \\
 T_0 & \xrightarrow{g_1} & \dots & \xrightarrow{g_i} & T_i & \xrightarrow{g_{i+1}} & \dots & \xrightarrow{g_K} & T_K \\
 \downarrow \rho_0 & & & & \downarrow \rho_i & & & & \downarrow \rho_K \\
 U_0 & \xrightarrow{h_1} & \dots & \xrightarrow{h_i} & U_i & \xrightarrow{h_{i+1}} & \dots & \xrightarrow{h_K} & U_K
 \end{array}$$

their composition, which is the commutative diagram

$$\begin{array}{ccccccc}
 S_0 & \xrightarrow{f_1} & \dots & \xrightarrow{f_i} & S_i & \xrightarrow{f_{i+1}} & \dots & \xrightarrow{f_K} & S_K \\
 \downarrow \rho_0 \circ \pi_0 & & & & \downarrow \rho_i \circ \pi_i & & & & \downarrow \rho_K \circ \pi_K \\
 U_0 & \xrightarrow{h_1} & \dots & \xrightarrow{h_i} & U_i & \xrightarrow{h_{i+1}} & \dots & \xrightarrow{h_K} & U_K
 \end{array}$$

is a combing rectangle. The collapsed subgraph of $\rho_i \circ \pi_i$ is the union of the collapsed subgraph of π_i with the inverse image under π_i of the collapsed subgraph of ρ_i .

Proof For each edgelet $e \subset U_i$, the map ρ_i is injective over the interior of e , and so there is a unique edgelet $e' \subset T_i$ such that $\rho_i(e') = e$. The map π_i is injective over the interior of e' , and it follows that $\rho_i \circ \pi_i$ is injective over the interior of e . This proves that $\rho_i \circ \pi_i$ is a collapse map and that the second diagram in the statement above is a combing rectangle.

Given an edgelet of S_i , clearly its image under $\rho_i \circ \pi_i$ is a vertex of U_i if and only if its image under π_i is a vertex of T_i or an edgelet of T_i whose image under ρ_i is a vertex of U_i . □

Lemma 4.6 (Decomposition of combing rectangles) *Given a combing rectangle of the form*

$$\begin{array}{ccccccc}
 S_0 & \xrightarrow{f_1} & \cdots & \xrightarrow{f_i} & S_i & \xrightarrow{f_{i+1}} & \cdots & \xrightarrow{f_K} & S_K \\
 [\sigma_0] \downarrow v_0 & & & & [\sigma_i] \downarrow v_i & & & & [\sigma_K] \downarrow v_K \\
 U_0 & \xrightarrow{h_1} & \cdots & \xrightarrow{h_i} & U_i & \xrightarrow{h_{i+1}} & \cdots & \xrightarrow{h_K} & U_K
 \end{array}$$

and given equivariant subgraphs $\sigma'_i \subset \sigma_i$ ($i = 0, \dots, K$) having the property that $f_i^{-1}(\sigma'_i) = \sigma'_{i-1}$ for each $i = 1, \dots, K$, there exist two combing rectangles of the form

$$\begin{array}{ccccccc}
 S_0 & \xrightarrow{f_1} & \cdots & \xrightarrow{f_i} & S_i & \xrightarrow{f_{i+1}} & \cdots & \xrightarrow{f_K} & S_K \\
 [\sigma'_0] \downarrow \pi_0 & & & & [\sigma'_i] \downarrow \pi_i & & & & [\sigma'_K] \downarrow \pi_K \\
 T_0 & \xrightarrow{g_1} & \cdots & \xrightarrow{g_i} & T_i & \xrightarrow{g_{i+1}} & \cdots & \xrightarrow{g_K} & T_K \\
 \downarrow \rho_0 & & & & \downarrow \rho_i & & & & \downarrow \rho_K \\
 U_0 & \xrightarrow{h_1} & \cdots & \xrightarrow{h_i} & U_i & \xrightarrow{h_{i+1}} & \cdots & \xrightarrow{h_K} & U_K
 \end{array}$$

whose composition (as in Lemma 4.5) is the given combing rectangle.

Proof Define the collapse map

$$\pi_i: S_i \xrightarrow{[\sigma'_i]} T_i$$

to be the quotient map obtained by collapsing each component of σ'_i to a point. Since $f_i^{-1}(\sigma'_i) = \sigma'_{i-1}$, there exists a map $g_i: T_{i-1} \rightarrow T_i$ induced from $f_i: S_{i-1} \rightarrow S_i$ under the quotient, which makes the top rectangle with the S row and the T row commutative, and this rectangle is therefore a combing rectangle. By Lemma 4.1, the T sequence is foldable. Define a subgraph $\tau_i = \pi_i(\sigma_i) \subset T_i$. We have $g_i^{-1}(\tau_i) = g_i^{-1}(\pi_i(\sigma_i)) = \pi_{i-1}(f_i^{-1}(\sigma_i)) = \pi_{i-1}(\sigma_{i-1}) = \tau_{i-1}$, where the second equation is

verified by a diagram chase using that the map π_{i-1} is surjective, and that π_i is injective over the interior of each edgelet of T_i . Clearly the collapse map

$$v_i: S_i \xrightarrow{[\sigma_i]} U_i$$

factors as the composition of

$$\pi_i: S_i \xrightarrow{[\sigma'_i]} T_i \quad \text{and a collapse map } \rho_i: T_i \xrightarrow{[\tau_i]} U_i,$$

making the bottom diagram with the T row and the U row commutative, and this row is therefore a combing rectangle. □

5 Free splitting units

In this section we study how to break a fold sequence into natural units called *free splitting units*. Our story of free splitting units grew in the telling. The original concept was motivated by units along train track splitting paths that are implicit in the “nested train track argument” of [15] and refinements of that argument in Masur, Mosher and Schleimer [16]. The details of the definition were tailored to fit the proofs of our two major results: our Main Theorem on hyperbolicity of the free splitting complex, via the arguments of Sections 6.2, and Proposition 6.2, which says that free splitting units give a uniformly quasigeodesic parametrization of fold paths in $\mathcal{FS}'(F)$.

The main results of this section are Proposition 5.9, which verifies the *coarse retraction* axiom of Masur and Minsky, and Lemma 5.11, which says that free splitting units give a uniformly coarse Lipschitz parametrization of fold paths in $\mathcal{FS}'(F)$. Underlying Lemma 5.11 are Lemmas 5.2 and 5.5, which give two methods of finding diameter bounds along foldable sequences.

The diameter bounds, which are stated and proved in Section 5.1, arise from finding “invariant natural structures” along the foldable sequence. The first diameter bound, Lemma 5.2, occurs when each free splitting along the fold path decomposes equivariantly into a pair of natural subgraphs in a manner that is “invariant” with respect to the foldable maps (see Definition 5.1). The second diameter bound, Lemma 5.5, occurs when each free splitting has a particular orbit of natural edges that is “almost invariant” with respect to the foldable maps (see Definition 5.4).

The combinatorial structures underlying the two diameter bounds are used to formulate the definition of free splitting units along a fold sequence (see Definitions 5.6 and 5.10). The diameter bounds are not applied directly to the fold sequence itself, but instead to foldable sequences obtained by transforming the given fold sequence via an application

of “combing by collapse” followed by an application of “combing by expansion”. One can already see this kind of transformation in the “nested train track argument” of [15].

5.1 Diameter bounds along foldable sequences

In this section we describe a pair of techniques for finding upper bounds on the diameter of foldable sequences.

Diameter bounds from natural red-blue decompositions Consider a free splitting $F \curvearrowright T$ and a nonempty, proper, F -invariant subgraph $\beta \subset T$ having no degenerate components. The conjugacy classes of nontrivial stabilizers of connected components of β form a free factor system $\mathcal{F}(\beta)$, as one can see by forming the collapse map

$$T \xrightarrow{[\beta]} U$$

and noting that $\mathcal{F}(\beta)$ is a subset of $\mathcal{F}(U)$. Passing further to the quotient graph of groups $X = U/F_n$, the image of β under the composition $T \mapsto U \mapsto X$ is a subset V_β of the vertex set of X . Let $C_1(\beta)$ be the number of F -orbits of components of β , equivalently, the cardinality of V_β . Let $C_2(\beta)$ be the sum of the ranks of the components of $\mathcal{F}(\beta)$, equivalently, the sum of the ranks of the subgroups labelling the vertices V_β in the graph of groups X , and so we have $0 \leq C_2(\beta) \leq \text{rank}(F)$. Defining the *complexity* of β to be $C(\beta) \equiv C_1(\beta) + (\text{rank}(F) - C_2(\beta))$, we have $C(\beta) \geq 1$. If furthermore β is a natural subgraph of S then $C_1(\beta) \leq 3 \text{rank}(F) - 3$, because the number of component orbits of β is at most the number of natural edge orbits in β , and $3 \text{rank}(F) - 3$ is an upper bound for the number of natural edge orbits of any free splitting of F . Altogether this shows that the complexity of any nonempty, proper, natural, F -invariant subgraph $\beta \subset T$ satisfies

$$1 \leq C(\beta) \leq 4 \text{rank}(F) - 3.$$

Definition 5.1 (Invariant blue–red decompositions) An *invariant blue–red decomposition* for a foldable sequence

$$T_0 \xrightarrow{g_1} T_1 \xrightarrow{g_2} \dots \xrightarrow{g_k} T_K,$$

also called an *invariant decomposition* for short, is a decomposition $\beta_k \cup \rho_k = T_k$ for each $k = 0, \dots, K$ such that for $0 \leq i \leq j \leq K$ we have $(g_j^i)^{-1}(\beta_j) = \beta_i$ and $(g_j^i)^{-1}(\rho_j) = \rho_i$ (where in expressions like $(g_j^i)^{-1}(\beta_j)$ we abuse notation by deleting degenerate components). Notice that any choice of final decomposition $\beta_K \cup \rho_K = T_K$ determines a unique invariant decomposition by the equations $\beta_i = (g_K^i)^{-1}(\beta_K)$ and $\rho_i = (g_K^i)^{-1}(\rho_K)$.

An invariant decomposition is *natural* if either of the following two equivalent properties holds: β_0, ρ_0 are natural subgraphs of T_0 , which holds if and only if β_k, ρ_k are natural subgraphs of T_k for all $k = 0, \dots, K$. The “only if” direction follows by observing that the image of each natural vertex under a foldable map is a natural vertex, and so the image of a natural subgraph is a natural subgraph.

Because an invariant decomposition is determined by the final decomposition, a general invariant decomposition carries little information about the foldable sequence. The typical behavior is that the edgelets within a natural edge $e \subset T_i$ will alternate many times between red and blue, that is, the number of components of $e \cap \beta_i$ and $e \cap \rho_i$ will be very large. Exploiting the difference between this typical behavior and the contrasting special behavior of a natural invariant decomposition is at the heart of the proof of the Main Theorem, specifically in the proof of Proposition 6.5 Step 2.

Here is our first diameter bound:

Lemma 5.2 *Given a foldable sequence*

$$T_0 \xrightarrow{g_1} T_1 \xrightarrow{g_2} \dots \xrightarrow{g_k} T_K$$

with an invariant natural decomposition $\beta_k \cup \rho_k = T_k$, the following hold:

- (1) *The complexity $C(\beta_k)$ is nonincreasing as a function of $k = 0, \dots, K$.*
- (2) *The interval $0 \leq k \leq K$ subdivides into at most $4 \operatorname{rank}(F) - 3$ subintervals on each of which $C(\beta_k)$ is constant.*
- (3) *If $C(\beta_k)$ is constant on the subinterval $a \leq k \leq b$, where $0 \leq a \leq b \leq K$, then*

$$\operatorname{diam}\{T_a, \dots, T_b\} \leq 4.$$

Remark When

$$T_0 \xrightarrow{g_1} T_1 \xrightarrow{g_2} \dots \xrightarrow{g_k} T_K$$

is a *fold* sequence, one obtains a diameter bound for the entire sequence as follows. Subdivide the interval $0, \dots, K$ into at most $4 \operatorname{rank}(F) - 3$ subintervals on which $C(\beta_k)$ is constant. On each subinterval one has a diameter bound of 4. At each of the at most $4 \operatorname{rank}(F) - 4$ fold maps where one subinterval transitions to another, one has an additional distance bound of 2 coming from Lemma 2.5. Putting these together,

$$\operatorname{diam}\{T_0, \dots, T_K\} \leq 4(4 \operatorname{rank}(F) - 3) + 2(4 \operatorname{rank}(F) - 4) = 24 \operatorname{rank}(F) - 20.$$

However, the manner in which we actually apply Lemma 5.2 to fold sequences is via concepts of free splitting units in the next section; see Lemma 5.11.

Before turning to the proof proper of Lemma 5.2, we first state a sublemma about the behavior of complexity of invariant subforests under fold maps.

Sublemma 5.3 *If $f: S \rightarrow T$ is a fold map of free splittings, if $\beta_T \subset T$ is a nonempty, proper, F -invariant subgraph, and if $\beta_S = f^{-1}(\beta_T)$ (as usual ignoring degenerate components), then $C_1(\beta_S) \geq C_1(\beta_T)$, and $C_2(\beta_S) \leq C_2(\beta_T)$, and so $C(\beta_S) \geq C(\beta_T)$. Furthermore, equality holds if and only if f restricts to a bijection of component sets of β_S and β_T .*

We delay the proof of this sublemma and meanwhile turn to:

Proof of Lemma 5.2 Item (1) follows from Sublemma 5.3 by factoring each foldable map $g_k: T_{k-1} \rightarrow T_k$ into folds. Item (2) follows from (1) and the fact that $1 \leq C(\beta_K) \leq C(\beta_0) \leq 4 \operatorname{rank}(F) - 3$.

To prove (3), fixing i, j with $a \leq i < j \leq b$, it suffices to prove that $d(T_i, T_j) \leq 4$. By the hypothesis of (3), $C(\beta_k)$ is constant for $i \leq k \leq j$. For each $i < k \leq j$, factoring $g_k: T_{k-1} \rightarrow T_k$ into folds, applying (1) to get constant complexity on the fold factorization, and applying Sublemma 5.3 to each of those folds, it follows that g_k induces a bijection between the component sets of β_{k-1} and β_k . By composing, it follows that $g_j^i = g_j \circ \cdots \circ g_{i+1}$ induces a bijection between the component sets of β_i and β_j .

Now we may factor g_j^i into a fold sequence of the form

$$T_i = U_0 \xrightarrow{h_1} \cdots \xrightarrow{h_P} U_P \xrightarrow{h_{P+1}} \cdots \xrightarrow{h_Q} U_Q = T_j$$

by prioritizing folds of blue edge pairs over folds of red edge pairs up until U_P when there are no more blue edge pairs to fold, with the result that if $0 < p \leq P$ then an edge pair of U_{p-1} folded by f_p is blue, whereas if $P < q \leq Q$ then an edge pair of S_{q-1} folded by h_q is red. To see that prioritizing in this manner is possible, if g_j^i does not already restrict to a simplicial isomorphism from β_i to β_j then, using that g_j^i induces a bijection of components of β_i and β_j , together with the *local to global principle* (see the proof of Lemma 2.7 and the following remark), it follows that some pair of oriented natural edges $\eta_1, \eta_2 \subset \beta_i$ with the same initial vertex have images in β_j with the same initial direction. We may therefore define the first fold h_1 to be a maximal fold factor of g_j^i obtained by folding η_1, η_2 , producing a factorization

$$T_i = U_0 \xrightarrow{h_1} U_1 \mapsto T_j.$$

Pushing the natural blue–red decomposition of U_0 forward (or equivalently pulling that of T_j back), we obtain a natural blue–red decomposition of U_1 , and the map $U_1 \mapsto T_j$

still induces a bijection of component sets of blue graphs. We may then continue by induction, stopping when the map $U_P \mapsto T_j$ restricts to a simplicial isomorphism of blue graphs. If the map $U_P \mapsto T_j$ is not already a simplicial isomorphism then one continues the fold factorization arbitrarily, with the result that all folds from U_P to T_j are red.

For $0 \leq p \leq P$, by collapsing all blue edges of U_p , we obtain a free splitting X_p with a collapse map $U_p \mapsto X_p$. Also, for $P \leq q \leq Q$, by collapsing red edges of U_q we obtain a free splitting Y_q with a collapse map $U_q \rightarrow Y_q$.

We claim that, up to equivalence, X_p is independent of $p = 0, \dots, P$ and Y_q is independent of $q = P, \dots, Q$. From this claim it follows that T_i, T_j are connected in $\mathcal{FS}'(F_n)$ by a path of length at most 4 as follows:

$$[T_i] = [U_0] \succ [X_0] = [X_P] \prec [U_P] \succ [Y_P] = [Y_Q] \prec [U_Q] = [T_j]$$

This completes the proof.

We prove for each $p = 1, \dots, P$ that X_{p-1}, X_p are equivalent, and for $q = P+1, \dots, Q$ that Y_{q-1}, Y_q are equivalent; the two cases are similar and we do just the first. Let e_1, e_2 be maximal oriented segments with the same initial vertex that are identified by the fold $U_{p-1} \mapsto U_p$. Recall that the fold map $U_{p-1} \mapsto U_p$ can be factored as

$$U_{p-1} \xrightarrow{q'} U' \xrightarrow{q''} U_p$$

where q' identifies proper initial segments of e_1, e_2 and q'' folds the remaining unidentified segments. Since e_1, e_2 are blue, by pushing forward the blue–red decomposition of U_{p-1} , or pulling back that of U_p , we obtain a blue–red decomposition of U' . Furthermore, there is a collapse map $U' \mapsto U_{p-1}$ that collapses the blue segment resulting from partially identifying e_1, e_2 , and a collapse map $U' \mapsto U_p$ that collapses the remaining unidentified segments, also blue. By composition we obtain collapse maps $U' \mapsto U_{p-1} \rightarrow X_{p-1}$ and $U' \mapsto U_p \mapsto X_p$, each of which collapses the entire blue subgraph of U' . It follows that X_{p-1} and X_p are equivalent. \square

Proof of Sublemma 5.3 Let $e_1, e_2 \subset S$ be oriented natural edges with the same initial vertex that are folded by the map f . Let $\eta_1 \subset e_1, \eta_2 \subset e_2$ be maximal initial subsegments that are identified by f . Let $v_1 \in \eta_1, v_2 \in \eta_2$ be the terminal endpoints. Note that either $\eta_1 \cup \eta_2 \subset \beta_S$ or $\eta_1 \cup \eta_2 \subset S \setminus \beta_S$. If $\eta_1, \eta_2 \subset \beta_S$, or if $\eta_1, \eta_2 \subset S \setminus \beta_S$ and either $v_1 \notin \beta_S$ or $v_2 \notin \beta_S$, then all inequalities are equations and f is a bijection of component sets.

We are reduced to the case that $\eta_1 \cup \eta_2 \subset S \setminus \beta_S$ and $v_1, v_2 \in \beta_S$, and so f is not a bijection of component sets because the two components $\beta_{S,1}, \beta_{S,2}$ of β_S containing

v_1, v_2 are mapped to the one component $\beta_{T,0}$ of β_T that contains $f(v_1) = f(v_2)$. We must prove that the inequalities $C_1(\beta_S) \geq C_1(\beta_T)$ and $C_2(\beta_S) \leq C_2(\beta_T)$ both hold and that at least one of them is strict.

Let the fold map $f: S \rightarrow T$ be factored as $S \mapsto U \mapsto T$, where $S \mapsto U$ folds short initial segments of η_1, η_2 , and $U \mapsto T$ folds the remaining segments, as in the proof of Lemma 2.5. Let $u_1, u_2 \in U$ be the images of v_1, v_2 . In order to compare the complexities of $\beta_S \subset S$ and $\beta_T \subset T$ we shall move them both into U where we can make the comparison directly.

Letting $\beta_U \subset U$ be the image of β_S , equivalently the preimage of β_T , the fold map $S \mapsto U$ clearly induces an equivariant bijection from the component set of β_S to that of β_U , and so the values of C_1, C_2 and C on β_S, β_U are all equal. Letting $\beta_U^+ = \beta_U \cup F \cdot [u_1, u_2]$, the fold map $U \mapsto T$ induces an equivariant bijection from the component set of β_U^+ to that of β_T , and so the values of C_1, C_2 and C on β_U^+, β_T are equal. So now we must prove the inequalities $C_1(\beta_U) \geq C_1(\beta_U^+)$ and $C_2(\beta_U) \leq C_2(\beta_U^+)$ and that at least one of them is strict.

Let $\beta_{U,1}, \beta_{U,2}$ be the images of $\beta_{S,1}, \beta_{S,2}$, respectively, under the fold map $S \mapsto U$. In the quotient graph of groups U/F , notice that β_U^+/F is the union of β_U/F with the segment obtained by projecting $[u_1, u_2]$, that segment is disjoint from β_U/F except at its endpoints, it has one endpoint on $\beta_{U,1}/F$ and the other endpoint at $\beta_{U,2}/F$, and the stabilizer of the interior vertex of that segment is trivial. It follows that if $C_1(\beta_U) > C_1(\beta_U^+)$, that is, if $\beta_{U,1}, \beta_{U,2}$ are in different component orbits, then $C_1(\beta_U) = C_1(\beta_U^+) + 1$ and $C_2(\beta_U) = C_2(\beta_U^+)$. On the other hand if $C_1(\beta_U) = C_1(\beta_U^+)$, that is if $\beta_{U,1}$ and $\beta_{U,2}$ are in the same component orbit, then $C_1(\beta_U) = C_1(\beta_U^+)$ and $C_2(\beta_U) + 1 = C_2(\beta_U^+)$. □

Diameter bounds from almost invariant edges Consider a foldable map $f: S \rightarrow T$ and a natural edge $e_T \subset T$. By ignoring unnatural vertices in e_T and their pre-images in S we may speak about e_T -edgelets in S ; these are the closures of the components of $f^{-1}(\text{int}(e_T))$, each of which is a subsegment of a natural edge of S . If S contains a unique e_T -edgelet and if $e_S \subset S$ is the natural edge containing that edgelet then we say that the pair e_S, e_T is an *almost invariant edge* of the foldable map f .

Definition 5.4 (Almost invariant edge) An *almost invariant edge* for a foldable sequence

$$T_0 \xrightarrow{f_1} T_1 \xrightarrow{f_2} \dots \xrightarrow{f_k} T_K$$

is a sequence of natural edges $e_k \subset T_k, k = 0, \dots, K$, such that for $0 \leq i < j \leq K$ the edges $e_i \subset T_i$ and $e_j \subset T_j$ are an almost invariant edge for the foldable map

$f_j^i: T_i \rightarrow T_j$. Note that an almost invariant edge exists for the whole foldable sequence if and only if one exists for the map $f_K^0: T_0 \rightarrow T_K$. To see why, observe that for any natural edge $e_K \subset T_K$, letting m_k be the number of e_K edgelets in T_k , the sequence m_k is nonincreasing as a function of $k \in 0, \dots, K$. If there is a natural edge $e_0 \subset T_0$ so that e_0, e_K is an almost invariant edge for the map f_K^0 , then $m_0 = 1$, and so m_k has constant value equal to 1. Letting $e_k \subset T_k$ be the unique natural edge containing an e_K edgelet in T_k , it follows that $(e_k)_{0 \leq k \leq K}$ is an almost invariant edge for the whole foldable sequence. This argument also shows that each almost invariant edge for a foldable sequence $T_0 \mapsto \dots \mapsto T_K$ is determined by its last term $e_K \subset T_K$.

Here is our second diameter bound:

Lemma 5.5 *Given a foldable sequence $T_0 \mapsto \dots \mapsto T_K$, the following are equivalent:*

- (1) *The foldable map $T_0 \mapsto T_K$ has an almost invariant edge.*
- (2) *The foldable sequence $T_0 \mapsto \dots \mapsto T_K$ has an almost invariant edge.*
- (3) *There exists a one-edge free splitting R such that $d(T_k, R) \leq 1$ for all k between 0 and K .*
- (4) *There is a one-edge free splitting R such that $d(T_0, R) \leq 1$, and $d(T_K, R) \leq 1$.*

Furthermore if these hold then $\text{diam}\{T_0, \dots, T_K\} \leq 2$.

Proof The bound in the last sentence clearly follows from (3). We have seen that (1) \implies (2), and clearly (3) \implies (4).

We next prove (2) \implies (3). Let $(e_k)_{k=0, \dots, K}$ be an almost invariant edge. Let $\sigma_k \subset T_k$ be the complement of the orbit of the natural edge e_k . Define a collapse map

$$T_k \xrightarrow{[\sigma_k]} R_k,$$

so R_k is a one-edge free splitting. It suffices to prove for each $k = 1, \dots, K$ that $[R_{k-1}] = [R_k]$. Letting $e'_{k-1} \subset e_{k-1}$ be the unique e_k -edgelet in T_{k-1} , letting $\sigma'_{k-1} \subset T_{k-1}$ be the complement of the orbit of e'_{k-1} , and defining a collapse map

$$T_{k-1} \xrightarrow{[\sigma'_{k-1}]} R'_{k-1},$$

clearly the map $T_{k-1} \mapsto T_k$ induces an equivariant homeomorphism $R'_{k-1} \rightarrow R_k$, and so $[R'_{k-1}] = [R_k]$. Also, since σ_{k-1} is the maximal natural subgraph of σ'_{k-1} , the identity map on T_{k-1} induces a collapse map $R_{k-1} \rightarrow R'_{k-1}$, which is a bijection on natural vertices and which, on each natural edge of R_{k-1} , collapses an initial and/or

terminal segment and is otherwise injective. It follows that the collapse map $R_{k-1} \mapsto R'_{k-1}$ is equivariantly homotopic to a conjugacy, and so $[R_{k-1}] = [R'_{k-1}] = [R_k]$.

It remains to prove (4) \implies (1). After rewording, this says that if $f: S \rightarrow T$ is a foldable map of free splittings, and if there exists a one-edge free splitting R such that $d(R, S), d(R, T) \leq 1$, then $f: S \rightarrow T$ has an almost invariant edge. Fix an oriented natural edge $e_R \subset R$ with initial and terminal vertices r_{\pm} , and oriented natural edges $e_S \subset S, e_T \subset T$ with initial and terminal vertices s_{\pm}, t_{\pm} respectively, so that there are collapse maps $S, T \mapsto R$ that collapse the complement of the orbits of e_S, e_T and which take e_S, e_T homeomorphically to e_R . We shall prove that e_S, e_T is an almost invariant edge for $f: S \rightarrow T$.

There is a component decomposition $R \setminus e_R = R_- \sqcup R_+$, where R_{\pm} contains the vertex r_{\pm} and there are corresponding component decompositions $S \setminus e_S = S_- \sqcup S_+, T \setminus e_T = T_- \sqcup T_+$ so that S_{\pm}, T_{\pm} are the inverse images of R_{\pm} , respectively, under the collapse maps $S, T \mapsto R$ (in general the “ \pm ” notation means “+ or –, respectively”; for instance “ S_{\pm} is the inverse image of R_{\pm} ” means “ S_+, S_- is the inverse image of R_+, R_- , respectively”). Note that $R_{\pm}, S_{\pm}, T_{\pm}$ are natural subgraphs of R, S, T , respectively. Also, r_{\pm} is the unique point on the topological frontier of R_{\pm} in R , and similarly for S_{\pm}, T_{\pm} . Also, each vertex in each of these subgraphs has valence at least 2 within the subgraph: in, say, R_- , this is obvious for all interior vertices, and the frontier vertex r_- is a natural vertex in R having only one R -direction not in R_- , namely the direction of e_R .

It suffices to prove that $f(S_{\pm}) \subset T_{\pm}$, which immediately implies that e_S, e_T is an almost invariant edge for $f: S \rightarrow T$. Assuming that either $f(S_-) \not\subset T_-$ or $f(S_+) \not\subset T_+$, we shall produce a contradiction. The arguments are similar in either case, so we shall assume that $f(S_-) \not\subset T_-$.

Given a free splitting $F \curvearrowright U$ and a nontrivial $\gamma \in F$ let $\alpha_U(\gamma)$ denote either the axis of γ in U or the unique vertex of γ fixed by U . Let F_{\pm} denote the set of nontrivial $\gamma \in F$ such that $\alpha_R(\gamma) \subset R_{\pm}$. Note that for each natural edge $e \subset R_{\pm}$ there exists $\gamma \in F_{\pm}$ whose axis under the action $F \curvearrowright R$ contains e . It follows that

$$R_{\pm} = \bigcup_{\gamma \in F_{\pm}} \alpha_R(\gamma).$$

Note also that

$$(1) \quad S_{\pm} = \bigcup_{\gamma \in F_{\pm}} \alpha_S(\gamma) \quad \text{and} \quad T_{\pm} = \bigcup_{\gamma \in F_{\pm}} \alpha_T(\gamma).$$

To prove this for S_- , say, note first that the collapse map $S \mapsto R$ takes S_{\pm} to R_{\pm} and its restriction to $\alpha_S(\gamma)$ has image $\alpha_R(\gamma)$ for each $\gamma \in F$. If $\alpha_S(\gamma) \subset S_-$ then

$\alpha_R(\gamma) \subset R_-$ and hence $\gamma \in F_-$, and since the axes contained in S_- cover S_- we get one inclusion $S_- \subset \cup_{\gamma \in F_-} \alpha_S(\gamma)$. For the other inclusion, if $\alpha_S(\gamma) \not\subset S_-$ then either $\alpha_S(\gamma)$ crosses e_S and so $\alpha_R(\gamma)$ crosses e_r , or $\alpha_S(\gamma) \subset S_+$ and so $\alpha_R(\gamma) \subset R_+$, and in either case $\gamma \notin F_-$.

Next we show:

- (2) There exists a finite number $A \geq 0$ such that $T_- \subset f(S_-) \subset N_A(T_-)$.

Applying the inclusion $f(\alpha_S(\gamma)) \supset \alpha_T(\gamma)$ to all $\gamma \in F_-$ and using (1) we obtain one inclusion $T_- \subset f(S_-)$. The opposite inclusion follows by applying the *bounded cancellation lemma* to the map $f: S \rightarrow T$. The version of the lemma that we need comes from Bestvina, Feighn and Handel [4, Lemma 3.1], and although the hypothesis there requires that the free splitting $F \curvearrowright S$ be properly discontinuous (called there a “free simplicial tree”), the first paragraph of that proof works exactly as stated for a map like f that factors as a fold sequence. The conclusion of that first paragraph is that there exists A , a *bounded cancellation constant* for f , such that for any vertices $x, y \in S$, in the tree T the set $f[x, y]$ is contained in the A neighborhood of the segment $[f(x), f(y)]$. Applying this to our situation, we conclude that for any $\gamma \in F$ we have $f(\alpha_S(\gamma)) \subset N_A(\alpha_T(\gamma))$. Applying this to all $\gamma \in F_-$ and using (1), it follows that $f(S_-) \subset N_A(T_-)$, completing the proof of (2).

We show that the only way for $f(S_-)$ to cross e_T is to do so rather sharply:

- (3) If $f(S_-) \not\subset T_-$ then $f(S_-) = T_- \cup [t_-, f(s_-)]$. Recalling that t_- is the unique frontier point of T_- , it follows that $T_- \cap [t_-, f(s_-)] = \{t_-\}$.

To see why, by (2) the tree $f(S_-) \setminus T_-$ has finite diameter, by assumption of (3) it is nondegenerate, and so it has at least two vertices of valence 1, at least one being distinct from t_- . The graph $f(S_-)$ therefore has at least one vertex of valence 1. But s_- is the unique frontier vertex of S_- so by the First Derivative Test the point $f(s_-)$ is the unique valence 1 vertex of $f(S_-)$. Combining this with $T_- \subset f(S_-)$, (3) follows immediately.

But from (3) we deduce that $f: S \rightarrow T$ has at most 2 gates at the natural vertex s_- , because all of the directions at s_- distinct from the direction of e_S are mapped by f to a single direction at $f(s_-)$, namely, the direction of the segment $[f(s_-), t_-]$. This contradicts that a foldable map has at least 3 gates at every natural vertex. □

5.2 Definitions and properties of free splitting units

Given a fold sequence

$$S_0 \xrightarrow{f_1} S_1 \xrightarrow{f_2} \dots \xrightarrow{f_K} S_K,$$

we shall first define what it means for S_i, S_j to “differ by less than 1 free splitting unit” for $i, j \in 0, \dots, K$, and we prove an appropriate stability result for this definition. With this in hand, for any $i, j \in 0, \dots, K$ we then define the number of free splitting units between S_i and S_j . Lemma 5.8 proves that the free splitting parametrization along the fold sequence is a Lipschitz parametrization with respect to distance in $\mathcal{FS}'(F)$.

Definition 5.6 (Less than 1 free splitting unit) Given a fold sequence

$$S_0 \xrightarrow{f_1} \dots \xrightarrow{f_K} S_K$$

and $0 \leq i < j \leq K$, we say that S_i, S_j differ by less than 1 free splitting unit if there exists a commutative diagram of the form

$$\begin{array}{ccccccc}
 T_i & \longrightarrow & T_{i+1} & \longrightarrow & \dots & \longrightarrow & T_{j-1} & \longrightarrow & T_j \\
 [\tau_i] \downarrow & & [\tau_{i+1}] \downarrow & & & & \downarrow [\tau_{j-1}] & & \downarrow [\tau_j] \\
 S'_i & \longrightarrow & S'_{i+1} & \longrightarrow & \dots & \longrightarrow & S'_{j-1} & \longrightarrow & S'_j \\
 [\sigma_i] \uparrow & & [\sigma_{i+1}] \uparrow & & & & \uparrow [\sigma_{j-1}] & & \uparrow [\sigma_j] \\
 S_i & \xrightarrow{f_{i+1}} & S_{i+1} & \xrightarrow{f_{i+2}} & \dots & \xrightarrow{f_{j-1}} & S_{j-1} & \xrightarrow{f_j} & S_j
 \end{array}$$

whose top and bottom rectangles are combed rectangles, so that foldable sequence $T_i \mapsto \dots \mapsto T_j$ on the top row has either an invariant natural blue–red decomposition of constant complexity or an almost invariant edge (by combining Lemmas 5.2 and 5.5, this holds if and only if the foldable map $T_i \mapsto T_j$ has either an invariant natural blue–red decomposition of constant complexity or an almost invariant edge). To complete the definition, we symmetrize the concept by requiring that S_j, S_i differ by less than 1 free splitting unit if and only if S_i, S_j differ by less than 1 free splitting unit.

The following is an immediate consequence of the definition, by restricting to the appropriate subdiagram of the above commutative diagram:

Lemma 5.7 (Stability of free splitting units) Given a fold sequence $S_0 \mapsto \dots \mapsto S_K$ and $0 \leq i \leq i' \leq j' \leq j \leq K$, if S_i, S_j differ by less than 1 free splitting unit then $S_{i'}, S_{j'}$ differ by less than 1 free splitting unit. □

Using these concepts we get a diameter bound as follows:

Lemma 5.8 *Given a fold sequence $S_0 \mapsto \dots \mapsto S_K$ and $0 \leq i \leq j \leq K$, if S_i, S_j differ by less than 1 free splitting unit then $\text{diam}\{S_i, \dots, S_j\} \leq 8$.*

Proof Consider the commutative diagram in the definition of less than 1 free splitting unit. Combining Lemmas 5.2 and 5.5, it follows that $\text{diam}\{T_i, \dots, T_j\} \leq 4$. Since $d(S_k, T_k) \leq 2$ for each k , we have $\text{diam}\{S_i, \dots, S_j\} \leq 8$. \square

The coarse retract axiom As an application of the concepts of free splitting units, particularly Lemma 5.5, we now prove that our definition for projecting $\mathcal{FS}'(S)$ onto fold paths satisfies the first of the three Masur–Minsky axioms:

Proposition 5.9 *For any fold sequence $S_0 \mapsto \dots \mapsto S_K$, the associated projection map $\pi: \mathcal{FS}'(F) \rightarrow [0, \dots, K]$ satisfies the coarse retraction axiom with the constant $c = 6$: for any $i = 0, \dots, K$ we have $i \leq \pi(S_i)$ and the diameter of the set $\{S_i, \dots, S_{\pi(S_i)}\}$ is at most 6. Furthermore, there is less than 1 free splitting unit between S_i and $S_{\pi(S_i)}$.*

Proof We start by noticing that a projection diagram from S_i to $S_0 \mapsto \dots \mapsto S_K$ of depth i certainly exists, where all vertical arrows are conjugacies and all collapse graphs are empty; see Figure 3.




Figure 3: A projection diagram from S_i to $S_0 \mapsto \dots \mapsto S_K$ of depth i

By definition, $\pi(S_i)$ is the largest integer in the set $[0, \dots, K]$ such that (after rechoosing the free splitting $F \curvearrowright S_i$ in its conjugacy class, and after rechoosing the fold sequence $S_0 \mapsto \dots \mapsto S_K$ in its conjugacy class) a projection diagram from S_i to $S_0 \mapsto \dots \mapsto S_K$ of depth $\pi(S_i)$ exists. This largest integer therefore satisfies $i \leq \pi(S_i)$ and yields a projection diagram as in Figure 4. Let $e' \subset S'_i$ be any natural edge, and let R be the one-edge free splitting obtained from S'_i by collapsing the complement of the orbit of e' . Then we have collapse maps $T_i \mapsto S'_i \mapsto R$ and $S_i \mapsto S'_i \mapsto R$, proving that $d(T_i, R) \leq 1$ and $d(S_i, R) \leq 1$. Applying Lemma 5.5,

the foldable sequence on the top row from T_i to S_i has an almost invariant edge, and by restriction there is an almost invariant edge from T_i to $T_{\pi(S_i)}$. Also by Lemma 5.5, the set $\{T_i, \dots, T_{\pi(S_i)}\}$ has diameter at most 2, and since $d(S_k, T_k) \leq 2$ for each k it follows that $\text{diam}\{S_i, \dots, S_{\pi(S_i)}\} \leq 6$. And by Definition 5.6, it follows that there is less than 1 free splitting unit between S_i and $S_{\pi(S_i)}$. \square




Figure 4: A maximal depth projection diagram from S_i to $S_0 \mapsto \cdots \mapsto S_K$

Definition 5.10 (General count of free splitting units) Given a fold sequence $S_0 \mapsto \cdots \mapsto S_K$, for $0 \leq i, j \leq K$ we say that S_i, S_j differ by at least 1 free splitting unit if they do not differ by less than 1 free splitting unit. Then, for $0 \leq I \leq J \leq K$, the number of free splitting units between S_I and S_J is defined to be the maximum integer $\Upsilon \geq 0$ for which there exists a sequence of integers $I \leq i(0) < \cdots < i(\Upsilon) \leq J$ of length $\Upsilon + 1$, parametrized by integers $0 \leq u \leq \Upsilon$, such that if $1 \leq u \leq \Upsilon$ then $S_{i(u-1)}, S_{i(u)}$ differ by at least 1 free splitting unit. Notice that our definitions are consistent in that $\Upsilon = 0$ if and only if, following the earlier definition, there is less than 1 free splitting unit between S_I and S_J . Also, we symmetrize the definition by saying that the number of free splitting units between S_J and S_I equals the number between S_I and S_J .

Remark In counting the number of free splitting units between S_i and S_j , although this number depends on the fold sequence $S_i \mapsto \cdots \mapsto S_j$ that connects S_i to S_j , that fold sequence will always be clear by context and we suppress this dependence in our terminology. Notice that this number does *not* depend on any other details of an ambient fold sequence of which $S_i \mapsto \cdots \mapsto S_j$ might be a subinterval. In particular, the number of free splitting units between S_i and S_j is unaffected if the ambient fold sequence is truncated by deleting an initial segment before S_i and/or a terminal segment after S_j .

Notice that with the notation as above, if $0 \leq u \leq v \leq \Upsilon$ then the number of free splitting units between $S_{i(u)}$ and $S_{i(v)}$ equals $v - u$. To see why, first note that this

number is at least $v - u$ by construction. If it were at least $v - u + 1$ then one could alter the sequence $i(0) < \dots < i(\Upsilon)$ by removing the entries $i(u), \dots, i(v)$ and inserting an increasing sequence of at least $v - u + 2$ entries in the interval $[i(u), i(v)]$, which amongst themselves have at least 1 free splitting unit between any consecutive two. By *stability of free splitting units* the new entries would have at least 1 free splitting units with the remaining entries outside of the interval $[i(u), i(v)]$. The new sequence would therefore still have at least 1 free splitting units between consecutive terms, but would have length at least $\Upsilon + 2$, contradicting the maximality of Υ .

One can count free splitting units between S_I and S_J in several ways. For example, define the *front greedy subsequence* from I to J to be the sequence $I = j(0) < j(1) < \dots < j(\Upsilon') \leq J$ obtained by induction as follows: assuming $j(u)$ is defined, and assuming $S_{j(u)}$ and S_J differ by at least 1 free splitting unit, let $j(u + 1)$ be the least integer greater than $j(u)$ such that $S_{j(u)}$ and $S_{j(u+1)}$ differ by at least 1 free splitting unit; the sequence stops when $S_{j(\Upsilon')}, S_J$ differ by less than 1 free splitting unit. We claim that Υ' , the length of the front greedy subsequence, is equal to the number of free splitting units between S_I and S_J . When S_I, S_J differ by less than 1 free splitting unit the claim is immediate. In the case where S_I, S_J differ by at least 1 free splitting unit, clearly $\Upsilon' \geq 1$; then, noting by stability that $S_{j(u)}, S_{j(v)}$ differ by at least 1 free splitting unit for $1 \leq u < v \leq \Upsilon'$, and using maximality of Υ , it follows that $\Upsilon \geq \Upsilon'$. For the opposite inequality we argue by contradiction assuming that $\Upsilon \geq \Upsilon' + 1$. Consider any subsequence $I \leq i(0) < i(1) < \dots < i(\Upsilon) \leq J$ such that $S_{i(u-1)}, S_{i(u)}$ differ by at least 1 free splitting unit for each $u = 1, \dots, \Upsilon$. By maximality of Υ it follows that between each of the pairs $S_I, S_{i(0)}$ and $S_{i(\Upsilon)}, S_J$ there is less than 1 free splitting unit. By stability it follows that between S_I and $S_{i(1)}$ there is at least 1 free splitting unit. By definition of $j(1)$ we have $j(1) \leq i(1)$. By stability it follows that $S_{j(1)}$ and $S_{i(2)}$ differ by at least 1 free splitting unit from which it follows that $j(2) \leq i(2)$. Continuing by induction we see that $j(u) \leq i(u)$ for $u = 1, \dots, \Upsilon'$. But since $j(\Upsilon') \leq i(\Upsilon') < i(\Upsilon' + 1) \leq i(\Upsilon) \leq J$ and since $S_{i(\Upsilon')}, S_{i(\Upsilon'+1)}$ differ by at least 1 free splitting unit, it follows by stability that $S_{j(\Upsilon')}, S_J$ differ by at least 1 free splitting unit, which contradicts the definition of Υ' .

In a similar fashion one proves that the number of free splitting units is equal to the length of the *back greedy subsequence* $I \leq \ell(\Upsilon'') < \ell(\Upsilon'' - 1) < \dots < \ell(1) < \ell(0) = J$, defined as follows: assuming by induction that $\ell(u)$ is defined and that S_I and $S_{\ell(u)}$ differ by at least 1 free splitting unit, $\ell(u + 1)$ is the greatest integer less than $\ell(u)$ such that $S_{\ell(u+1)}$ and $S_{\ell(u)}$ differ by at least 1 free splitting unit; the sequence stops when $S_I, S_{\ell(\Upsilon'')}$ differ by less than 1 free splitting unit.

The following result says that a fold path which is parametrized by free splitting units is a coarse Lipschitz path in $\mathcal{FS}(F)$:

Lemma 5.11 For any fold path $S_0 \mapsto \cdots \mapsto S_K$ and any $0 \leq I \leq J \leq K$, if the number of free splitting units between S_I and S_J equals Υ then the diameter in $\mathcal{FS}'(F)$ of the set $\{S_I, \dots, S_J\}$ is at most $10\Upsilon + 8$.

Proof If $\Upsilon = 0$, that is, if S_I, S_J differ by less than 1 free splitting unit, then by Lemma 5.8 we have $\text{diam}\{S_I, \dots, S_J\} \leq 8$.

If $\Upsilon \geq 1$, from S_I to S_J let $I = i(0) < \cdots < i(\Upsilon) \leq J$ be the front greedy sequence. For $u = 1, \dots, \Upsilon$, the free splittings $S_{i(u-1)}$ and $S_{i(u)}$ differ by less than 1 free splitting unit, and so $\text{diam}\{S_{i(u-1)}, \dots, S_{i(u)}\} \leq 8$. By Lemma 2.5 we have $d(S_{i(u-1)}, S_{i(u)}) \leq 2$ and so $\text{diam}\{S_{i(u-1)}, \dots, S_{i(u)}\} \leq 10$. It follows in turn that $\text{diam}\{S_I = S_{i(0)}, \dots, S_{i(\Upsilon)}\} \leq 10\Upsilon$. Since $S_{i(\Upsilon)}, S_J$ differ by less than 1 free splitting unit we have $\text{diam}\{S_{i(\Upsilon)}, \dots, S_J\} \leq 8$, and putting it all together, $\text{diam}\{S_I, \dots, S_J\} \leq 10\Upsilon + 8$. \square

We also need the following lemma, which gives a coarse triangle inequality for free splitting units within a fold path:

Lemma 5.12 Given a fold path $S_0 \mapsto \cdots \mapsto S_K$ and $i, j, k \in \{0, \dots, K\}$, if Υ_1 is the number of free splitting units between S_i and S_j and Υ_2 is the number between S_j and S_k then the number Υ between S_i and S_k satisfies $\Upsilon \leq \Upsilon_1 + \Upsilon_2 + 1$.

Proof In the case where j is between i and k , using *symmetry of free splitting units* we may assume that $i \leq j \leq k$. Let $i = i(0) < \cdots < i(\Upsilon) \leq k$ be the front greedy sequence from S_i to S_k . Clearly the front greedy sequence from S_i to S_j is an initial segment, implying that $i(\Upsilon_1) \leq j$ and $i(\Upsilon_1 + 1) > j$, and so we have a subsequence $S_{i(\Upsilon_1+1)}, \dots, S_{i(\Upsilon)}$ of S_j, \dots, S_k with the property that between any two adjacent elements of this subsequence there is at least 1 free splitting unit. By Definition 5.10 and the hypothesis on Υ_2 , the length of this subsequence is therefore at most $\Upsilon_2 + 1$, giving us $\Upsilon - \Upsilon_1 \leq \Upsilon_2 + 1$.

In the case where $j > \max\{i, k\}$, again using *symmetry* we may assume $i \leq k < j$. Let $i = i(0) < \cdots < i(\Upsilon_1) \leq j$ be the front greedy subsequence between S_i and S_j . Again the front greedy subsequence between S_i and S_k is an initial subsegment and so $\Upsilon \leq \Upsilon_1 \leq \Upsilon_1 + \Upsilon_2 + 1$.

In the case where $j < \min\{i, k\}$, using *symmetry* we assume $j < k \leq i$, and we proceed similarly using the *back greedy subsequence* between S_j and S_i . \square

6 Proof of the Main Theorem

We begin with a quick sketch of the proof.

Consider a free splitting T , a fold sequence $S_0 \mapsto \dots \mapsto S_K$, and a maximal depth projection diagram that defines the projection $k_T \in \{0, \dots, K\}$ from T to this fold sequence. The form of this projection diagram can be viewed in Section 4.1, Figure 2, the top row of which is a foldable sequence $T_0 \mapsto \dots \mapsto T_{k_T} \mapsto T$. We then apply Lemma 2.7 to factor the final foldable map $T_{k_T} \mapsto T$ as a fold sequence of the form $T_{k_T} \mapsto \dots \mapsto T_L = T$, which we then paste into the foldable sequence on the top row of the projection diagram to get an “augmented” projection diagram. Figure 5 shows the original, unaugmented projection diagram and the augmented version in the same picture. Note that the top row of the augmented projection diagram is the foldable sequence $T_0 \mapsto \dots \mapsto T_{k_T} \mapsto \dots \mapsto T_L = T$. See Section 6.1 for more details on augmented projection diagrams.

Consider also a geodesic in the 1–skeleton of $\mathcal{FS}'(F)$ starting with T and ending with some free splitting R . This geodesic is a zig-zag path; suppose for concreteness that it starts with a collapse and ends with an expand, $T = T_L^0 > T_L^1 < T_L^2 > T_L^3 < T_L^4 > \dots < T_L^D = R$, and so $D = d(T, R) = d(T_L^0, T_L^D)$ is even. By combing the foldable sequence $T_0 \mapsto \dots \mapsto T_{k_T} \mapsto \dots \mapsto T_L = T$ across each collapse and expansion in this zig-zag path one at a time, we obtain “the Big diagram, Step 0” depicted in Section 6.3, Figure 15, which is built out of the projection diagram and an $L \times D$ rectangle composed of D combing rectangles. Note that the interior even terms along the zig-zag path, the free splittings $T_L^2, T_L^4, \dots, T_L^{D-2}$, are “peaks” of the zig-zag. The big $L \times D$ rectangle has the form of a corrugated aluminum roof in which the interior even horizontal rows are peaks of the corrugations.

Our technique can be described as “pushing down the peaks”. In brief, we prove that if one backs up from T_L to some earlier term in the fold path $T_{k_T} \mapsto \dots \mapsto T_L$, moving back a certain fixed number of free splitting units, then the Big diagram can be simplified by pushing the first corrugation peak down, reducing the number of corrugation peaks by 1, as shown in “the Big diagram, Step 1”. These “back up–pushdown” arguments are found in Section 6.2. Therefore, if the number of free splitting units between T_{k_T} and T_L is greater than a certain multiple of the number of peaks in the zig-zag path from T_L to T_L^D then the number of corrugation peaks in the Big diagram can be reduced to zero. With one final “back up–push down” step that uses up some of the original projection diagram for T_L , one obtains a projection diagram from R to $S_0 \mapsto \dots \mapsto S_K$, from which one concludes that the projection of R to $S_0 \mapsto \dots \mapsto S_K$ is not much further back (measured in free splitting units) than S_{k_T} , which is the projection of T .

The exact statement proved by these arguments is contained in Proposition 6.1, which can be regarded as a reformulation of the *coarse Lipschitz* and *desymmetrized strong contraction* axioms in terms of free splitting units, and which quickly implies those axioms and the Main Theorem as shown in Section 6.1. The proof of Proposition 6.1 itself is carried out in Sections 6.2 and 6.3.

6.1 Desymmetrized strong contraction reformulated and applied

In Proposition 6.1 we reformulate the *coarse Lipschitz* and *desymmetrized strong contraction* axioms as a joint statement expressed in terms of free splitting units. The proposition will be proved in later subsections of Section 6.

After stating the proposition, we use it to finish off the proof of the main theorem. We also use it to prove Proposition 6.2, which describes precisely how to reparametrize fold paths in terms of free splitting units so as to obtain uniform quasigeodesics in $\mathcal{FS}'(F)$.

To set up Proposition 6.1, consider any fold path $S_0 \mapsto \cdots \mapsto S_K$, any free splitting $F \curvearrowright T$ and any projection diagram of maximal depth $\pi(T) = k_T \in [0, \dots, K]$ as depicted in Figure 5. Applying Lemma 2.7, we may factor the foldable map $f: T_{k_T} \rightarrow T$ as a fold sequence, and then replace f with this factorization in the top line of the projection diagram, to obtain a sequence of maps:

$$T_0 \xrightarrow{f_1} \cdots \xrightarrow{f_{k_T}} T_{k_T} \xrightarrow{f_{k_T+1}} \cdots \xrightarrow{f_L} T_L = T.$$

This sequence of maps is still foldable: if $0 \leq k \leq k_T$ then f_L^k is foldable by virtue of being a map in the original foldable sequence on the top line of the unaugmented projection diagram; and if $k_T < k \leq L$ then f_L^k is foldable by virtue of being a map in the newly inserted fold sequence (note that if one replaces any but the last map in a foldable sequence with a fold factorization, this trick does not work: the resulting sequence need not be foldable). We therefore obtain an *augmented projection diagram* from T to $S_0 \mapsto \cdots \mapsto S_K$ of maximal depth, as depicted also in the Figure 5.

Proposition 6.1 (Strong contraction in terms of free splitting units) *Letting $b_1 = 4 \operatorname{rank}(F) - 3$, the following holds. Consider a fold path $S_0 \mapsto \cdots \mapsto S_K$, a free splitting $F \curvearrowright T$ with projection $\pi(T) = k_T \in [0, \dots, K]$, and an augmented projection diagram of maximal depth k_T as notated in Figure 5. Let Υ be the number of free splitting units between T_{k_T} and $T_L = T$. If $F \curvearrowright R$ is a free splitting such that $d(T, R) \leq \max\{2\lfloor \Upsilon/b_1 \rfloor, 1\}$, and if the number of free splitting units between S_0 and S_{k_T} is at least b_1 , then there exists $l \in [0, \pi(R)]$ such that the number of free splitting units between S_l and S_{k_T} is at most b_1 .*




Figure 5: An *augmented* projection diagram from T to $S_0 \mapsto \dots \mapsto S_K$ of maximal depth k_T (with the straight arrows from T_{k_T} to T) is obtained from a maximal depth projection diagram (with the curved arrow from T_{k_T} to T labelled f) by inserting a fold sequence factorization of the foldable map $f: T_{k_T} \rightarrow T$. After this insertion the whole sequence $T_0 \mapsto \dots \mapsto T_{k_T} \mapsto \dots \mapsto T_L = T$ in the top row is still a foldable sequence.

Remark To put it more plainly, Proposition 6.1 says that the projection of R to the fold path $S_0 \mapsto \dots \mapsto S_K$ is no farther to the left of the projection of T than a bounded number of free splitting units, as long as $d(T, R)$ is at most some bounded proportion of the number Υ . One can think of the number Υ as being a stand-in for the distance from T to the fold path $S_0 \mapsto \dots \mapsto S_K$ (a posterior one sees that Υ is indeed quasicomparable to that distance). Notice that the proposition does not apply if no projection diagram exists for T , nor if the number of free splitting units between S_0 and S_{k_T} is too small; in either of these cases the projection of T is close to S_0 in $\mathcal{FS}'(F)$. These special situations are handled in Case 1 of the proof of the Main Theorem.

Note that Proposition 6.1 is trivially true when $\pi(R) \geq k_T$, by taking $l = k_T$. The real meat of the proposition is when $\pi(R) < k_T$.

Proposition 6.1 is proved in Sections 6.2 and 6.3. For the rest of Section 6.1 we shall apply Proposition 6.1 to prove first the Main Theorem and then Proposition 6.2 regarding quasigeodesics in $\mathcal{FS}'(F)$.

Proof of the Main Theorem As we showed earlier, Proposition 4.2 implies Proposition 3.3, which implies the Main Theorem. To prove Proposition 4.2 we must prove that the projections to fold paths in $\mathcal{FS}'(F)$ satisfy the *coarse retraction*, *coarse Lipschitz*, and *desymmetrized strong contraction* axioms given in Section 3, with uniform constants

depending only on $\text{rank}(F)$. In Proposition 5.9 we already did this for the *coarse retraction* axiom. We turn to the other two axioms.

Fix the fold path $S_0 \mapsto \cdots \mapsto S_K$ and free splittings $F \curvearrowright T, R$ with projections $\pi(T), \pi(R) \in [0, \dots, K]$. For verifying both the *coarse Lipschitz* and *desymmetrized strong contraction* axioms we may assume that $\pi(R) \leq \pi(T)$. We seek to bound the diameter in $\mathcal{FS}'(F)$ of the set $\{S_{\pi(R)}, \dots, S_{\pi(T)}\}$. If $\pi(T) = 0$ then $\pi(R) = 0$ and we are done. Otherwise, after rechoosing T in its conjugacy class and rechoosing $S_0 \mapsto \cdots \mapsto S_K$ in its equivalence class, we may choose an augmented maximal depth projection diagram for T and $S_0 \mapsto \cdots \mapsto S_K$ as notated in Figure 5. Let Υ be the number of free splitting units between T_{k_L} and $T_L = T$.

Throughout the proof we denote the constants from Lemma 5.11 as

$$L = 10, \quad C = 8.$$

It follows that along any fold path, for any two terms of that path between which the number of free splitting units is at most

$$b_1 = 4 \text{rank}(F) - 3,$$

the diameter in $\mathcal{FS}'(F)$ of the segment between those two terms is at most

$$c = Lb_1 + C = 40 \text{rank}(F) - 22.$$

This is the value of c that will be used in verifying the two axioms.

Case 1 Suppose that the number of free splittings between S_0 and $S_{\pi(T)}$ is less than b_1 . Applying the inequality $0 \leq \pi(R) \leq \pi(T)$ together with *Stability of Free Splitting Units*, it follows that the number of free splitting units between $S_{\pi(R)}$ and $S_{\pi(T)}$ is $< b_1$. By Lemma 5.11 the diameter of the set $\{S_{\pi(R)}, \dots, S_{\pi(T)}\}$ is at most c , which is the common conclusion of the *coarse Lipschitz* and *desymmetrized strong contraction* axioms. In this case, those axioms are verified using any values of a, b .

Case 2 Suppose that the number of free splitting units between S_0 and $S_{\pi(T)}$ is at least $b_1 > 0$.

We claim that the following statement holds:

- (*) If $d(T, R) \leq \max\{2\lfloor \Upsilon/b_1 \rfloor, 1\}$ then the number of free splitting units between $S_{\pi(R)}$ and $S_{\pi(T)}$ is at most b_1 , and so the diameter in $\mathcal{FS}'(F)$ of the set $\{S_{\pi(R)}, \dots, S_{\pi(T)}\}$ is at most c .

To prove (*), assume that $d(T, R) \leq \max\{2\lfloor \Upsilon/b_1 \rfloor, 1\}$. Using the hypothesis of Case 2 we may apply Proposition 6.1, concluding that for some $l \in [0, \pi(R)]$ the

number of free splitting units between S_l and S_{k_T} is at most b_1 . Using *stability of free splitting units* it follows that the number of free splitting units between $S_{\pi(R)}$ and S_{k_T} is $\leq b_1$. Applying Lemma 5.11 we have $\text{diam}\{S_{\pi(R)}, \dots, S_{\pi(T)}\} \leq c$.

Since (*) applies whenever $d(T, R) \leq 1$, the *coarse Lipschitz* axiom follows immediately.

To prove *desymmetrized strong contraction* we shall produce constants $a, b > 0$ so that if $a \leq d(T, \{S_0, \dots, S_K\})$ and $d(T, R) \leq b \cdot d(T, \{S_0, \dots, S_K\})$ then $d(T, R) \leq 2\lfloor \Upsilon/b_1 \rfloor$, for then (*) applies and so $\text{diam}\{S_{\pi(R)}, \dots, S_{\pi(T)}\} \leq c$.

Consider first the case that $\Upsilon < 2b_1$. By Lemma 5.11 we have $d(T_{k_T}, T) < 2b_1L + C$ and so $d(T, S_0 \mapsto \dots \mapsto S_K) < 2b_1L + C + 2$. By taking $a = 2b_1L + C + 2 = 80 \text{rank}(F) - 52$ we may dispense with this case.

Consider next the case that $\Upsilon \geq 2b_1$. It follows that $\Upsilon \geq 1$. We have $\Upsilon/b_1 \leq 2(\Upsilon/b_1 - 1)$ from which it follows that

$$\Upsilon/b_1 \leq 2\lfloor \Upsilon/b_1 \rfloor.$$

The number of free splitting units between T_{k_T} and $T_L = T$ equals Υ and so by Lemma 5.11 we have $d(T, T_{k_T}) \leq L\Upsilon + C$. It follows that $d(T, S_{k_T}) \leq L\Upsilon + C + 2$, which implies that $d(T, S_0 \mapsto \dots \mapsto S_K) \leq L\Upsilon + C + 2$. Let

$$b = \frac{1}{80 \text{rank}(F) - 60} = \frac{1}{b_1(L + C + 2)} \leq \frac{1}{b_1(L + \frac{C+2}{\Upsilon})} = \frac{\Upsilon}{b_1(L\Upsilon + C + 2)},$$

where the inequality follows from $\Upsilon \geq 1$. We then have

$$b(L\Upsilon + C + 2) \leq \Upsilon/b_1.$$

It follows that if $d(T, R) \leq b \cdot d(T, S_0 \mapsto \dots \mapsto S_K)$ then $d(T, R) \leq \Upsilon/b_1 \leq 2\lfloor \Upsilon/b_1 \rfloor$ and we are done, subject to proving Proposition 6.1. □

Quasigeodesic reparametrization of fold paths We can also use these arguments to show how fold paths can be reparametrized, using free splitting units, to give a system of uniform quasigeodesics in $\mathcal{FS}'(F)$. Recall that each fold sequence $S_0 \mapsto \dots \mapsto S_M$ can be interpolated by a continuous edge path in $\mathcal{FS}'(F)$: for each fold $S_{m-1} \mapsto S_m$, the vertices S_{m-1}, S_m are connected in $\mathcal{FS}'(F)$ by an edge path of length 2, 1 or 0, by Lemma 2.5. Let Υ be the number of free splitting units from S_0 to S_M . Choose any sequence $0 \leq m_0 < m_1 < \dots < m_\Upsilon \leq M$ such that for $u = 1, \dots, \Upsilon$ there is at least 1 free splitting unit between $S_{m_{u-1}}$ and S_{m_u} . Notice that by *stability of free splitting units*, the number of free splitting units between S_0 and S_{m_1} , and between $S_{m_{\Upsilon-1}}$ and S_M is at least 1, and so we may rechoose the first and last terms of the

sequence so that $0 = m_0 < m_1 < \dots < m_\Upsilon = M$. Choose a continuous parametrization of the interpolating edge path of the form $\gamma: [0, \Upsilon] \rightarrow \mathcal{FS}'(F)$ such that $S_{m_u} = \gamma(u)$. We call this a *free splitting parametrization* of the fold sequence $S_0 \mapsto \dots \mapsto S_M$.

We use Proposition 6.1, in particular some details of the preceding proof, in order to prove the following result:

Proposition 6.2 *There exist constants k, c depending only on $\text{rank}(F)$ such that any free splitting parametrization $\gamma: [0, \Upsilon] \rightarrow \mathcal{FS}'(F)$ of any fold path $S_0 \mapsto \dots \mapsto S_M$ is a k, c quasigeodesic in $\mathcal{FS}'(F)$, that is,*

$$\frac{1}{k}|s-t| - c \leq d(\gamma(s), \gamma(t)) \leq k|s-t| + c \quad \text{for all } s, t \in [0, \Upsilon].$$

Proof We continue with the constants $L = 10$, $C = 8$, $b_1 = 4 \text{rank}(F) - 3$ from the previous proof.

As shown back in the definition of free splitting units, for each integer $u = 1, \dots, \Upsilon$ there is exactly 1 free splitting unit between $S_{m_{u-1}}$ and S_{m_u} . Applying Lemma 5.11, it follows that for each $u = 1, \dots, \Upsilon$ the set $\{S_{m_{u-1}}, \dots, S_{m_u}\}$ has diameter at most $L + C$. Combining this with the fact that the edge path interpolating each fold has length at most 2 it follows that:

$$(**) \quad \text{diam}(\gamma[u-1, u]) \leq L + C + 1 \quad \text{for each } u = 1, \dots, \Upsilon.$$

Given $s, t \in [0, \Upsilon]$, if there is no integer in the interval $[s, t]$ then $d(\gamma(s), \gamma(t)) \leq L + C + 1$. Otherwise we take $u, v \in [s, t]$ to be the smallest integer at least s and the largest integer at most t , respectively, and we have

$$\begin{aligned} d(\gamma(s), \gamma(t)) &\leq d(\gamma(u), \gamma(v)) + d(\gamma(s), \gamma(u)) + d(\gamma(t), \gamma(v)) \\ &\leq (L + C + 1)|v - u| + 2(L + C + 1) \\ &\leq k|s - t| + c \end{aligned}$$

using any $k \geq L + C + 1 = 19$ and any $c \geq 2(L + C + 1) = 38$ (and we note that this inequality also holds in the previous case where there is no integer in $[s, t]$). This proves the second inequality of the proposition.

To prove the first inequality, we first prove it for integer values $u \leq v \in [0, \dots, \Upsilon]$. Fix a geodesic edge path ρ of length $D = d(\gamma(u), \gamma(v))$ connecting $\gamma(u)$ to $\gamma(v)$ in $\mathcal{FS}'(F)$. Project ρ to the fold path $S_0 \mapsto \dots \mapsto S_M$. By the statement (*) above, within this fold path there are at most b_1 free splitting units between the projections of any two consecutive vertices of ρ . By applying Lemma 5.12, the coarse triangle inequality for free splitting units, it follows that there are at most

$D(b_1 + 1)$ free splitting units between $S_{\pi(\gamma(u))}$ and $S_{\pi(\gamma(v))}$, the projections of $\gamma(u)$ and $\gamma(v)$, respectively. By Proposition 5.9, where the *coarse retract* axiom was proved, the number of free splitting units between $S_{m_u} = \gamma(u)$ and $S_{\gamma(u)}$, and between $S_{m_v} = \gamma(v)$ and $S_{\gamma(v)}$, are both less than 1. By applying Lemma 5.12 again, the number of free splitting units between S_{m_u} and S_{m_v} is at most $D(b_1 + 1) + 2$, that is, $|u - v| \leq D(b_1 + 1) + 2 = (b_1 + 1)d(\gamma(u), \gamma(v)) + 2$.

For arbitrary $s < t \in [0, \dots, \Upsilon]$, letting $u \in [0, \Upsilon]$ be the largest integer at most s and $v \in [0, \Upsilon]$ be the smallest integer at least t , we have $\gamma(s) \in \gamma[u, u + 1]$ and $\gamma(t) \in \gamma[v - 1, v]$. By (**) we therefore have $d(\gamma(s), \gamma(u)), d(\gamma(t), \gamma(v)) \leq L + C + 1 = 19$. It follows that:

$$\begin{aligned} |s - t| &\leq |u - v| \\ &\leq (b_1 + 1)d(\gamma(u), \gamma(v)) + 2, \\ \frac{1}{b_1 + 1}|s - t| - \frac{2}{b_1 + 1} &\leq d(\gamma(u), \gamma(v)) \\ &\leq d(\gamma(u), \gamma(v)) + (19 - d(\gamma(u), \gamma(s))) \\ &\quad + (19 - d(\gamma(v), \gamma(t))), \\ \frac{1}{b_1 + 1}|s - t| - \left(\frac{2}{b_1 + 1} + 38\right) &\leq d(\gamma(s), \gamma(t)). \end{aligned}$$

This proves that the first inequality is true for any $k \geq b_1 + 1 = 4 \text{rank}(F) - 2$ and any

$$c \geq \frac{2}{b_1 + 1} + 38 = \frac{1}{2 \text{rank}(F) - 1} + 38.$$

Proposition 6.2 is therefore proved for $k = \max\{19, 4 \text{rank}(F) - 2\}$ and $c = 39$. \square

6.2 Pushing down peaks

Recall that every geodesic in $\mathcal{FS}'(F)$ is a zig-zag edge path. On a zig-zag subpath of the form $T^{i-1} \prec T^i \succ T^{i+1}$, where T^i is the domain of two incident collapse maps $T^i \mapsto T^{i-1}$ and $T^i \mapsto T^{i+1}$, we say that T^i is a *peak*. If on the other hand $T^{i-1} \succ T^i \prec T^{i+1}$ then T^i is a *valley*.

We start with a simplistic technique that can be used to shortcut a zig-zag path, and we work up to a technique, described in Proposition 6.5, that will be central to the proof of the Main Theorem. In each case the intuition is to “push down the peak”, thereby reducing length.

The peak of a W diagram A *W diagram* or a *W zig-zag* is a length 4 zig-zag path with a peak in the middle, sometimes depicted as in Figure 6.




Figure 6: A W diagram

We think of β, ρ as the “blue” and “red” subgraphs of T^2 . In this generality, an edgelet of T^2 may be in either, or both, or neither of β, ρ . The subgraphs β, ρ therefore do *not* necessarily form a blue–red decomposition of T^2 as in Definition 5.1, which requires that β, ρ have no edgelets in common and their union is all of T^2 ; furthermore, even if β, ρ did form a blue–red decomposition, they need not be a *natural* one, which requires in addition that they both be natural subgraphs of T^2 . Soon, though, we shall narrow down to a key special case where β, ρ is indeed a natural blue–red decomposition.

Pushing down the peak is easy when $\beta \cup \rho$ is a proper subgraph of T^2 , for in that case the given W diagram extends to a commutative diagram of collapse maps as shown in the diagram in Figure 7. In that diagram, collapse of $\beta \cup \rho \subset T^2$ produces T^h . The collapse map

$$T^2 \xrightarrow{[\rho]} T^1$$

takes the edgelets of the subgraph $\beta \setminus (\beta \cap \rho) \subset T^2$ bijectively to the edgelets of a subgraph of T^1 , which by convention is also denoted $\beta \setminus (\beta \cap \rho)$; collapse of this subgraph also produces T^h . Similarly, collapse of $\rho \setminus (\beta \cap \rho) \subset T^3$ produces T^h . Compositions of collapse maps being collapse maps, we obtain a length 2 zig-zag path $T^0 \rightarrow T^h \leftarrow T^4$ that cuts short the original length 4 zig-zag path; we have successfully “pushed down the peak”.




Figure 7: A simplistic pushdown works if $\beta \cup \rho \subset T^2$ is a proper subgraph.

The same argument works on a length 3 zig-zag path—which can be visualized by cutting off one of the terminal edges of a W zig-zag—with the result that if the union of the two collapse graphs at the peak of the zig-zag forms a proper subgraph then there is a length 2 path with the same endpoints. We summarize as follows:

Lemma 6.3 *Given a W zig-zag as notated in Figure 6 or a length 3 zig-zag obtained from Figure 6 by cutting off one of the terminal edges, if the path is geodesic then $T^2 = \beta \cup \rho$. \square*

Normalizing a W diagram We shall also need to push down the peak of certain W diagrams in the situation where $T^2 = \beta \cup \rho$. In this situation it is convenient to first alter the W diagram to ensure that $\beta \cap \rho$ contains no edgelet of T^2 , equivalently, β, ρ is a blue–red decomposition of T^2 as in Definition 5.1. If $\beta \cap \rho$ does contain an edgelet of T^2 then, since β, ρ are proper subgraphs, the given W diagram is contained in a commutative diagram of collapse maps as shown in the diagram in Figure 8, called a *normalization diagram*. In this diagram, subgraphs of T'^2 are labelled by the same convention as described above. Since $T^2 = \beta \cup \rho$ it follows that the two subgraphs $\beta \setminus (\beta \cap \rho)$ and $\rho \setminus (\beta \cap \rho)$ of T'^2 partition the edgelets of T'^2 .




Figure 8: A normalization diagram. The W zig-zag on the top of the diagram has the property that $T^2 = \beta \cup \rho$. The W zig-zag on the bottom of the diagram is normalized.

Motivated by this observation, we say that a zig-zag path in $\mathcal{FS}'(F)$ is *normalized* if at every free splitting $F \curvearrowright T$ along the path that forms a peak, the two subgraphs of T whose collapses define the vertices of the path incident to T form a blue–red decomposition of T . The argument we have just given shows that every geodesic zig-zag path in $\mathcal{FS}'(F)$ may be replaced by a normalized zig-zag path of the same length and with the same set of valleys.

Pushdown subgraphs and baseball diagrams We now turn to a more sophisticated technique for pushing down the peak of a W diagram. Consider a W diagram as notated in Figure 6 and suppose that $\beta \cup \rho = T^2$ is a blue–red decomposition. Consider also a subgraph $\kappa \subset T^2$ that satisfies the following:

κ is a pushdown subgraph κ is a proper, equivariant subgraph, and each natural edge of T^2 not contained in κ contains at least one red and one blue edgelet of T^2 that are not contained in κ .

No requirement is imposed that a pushdown subgraph be a natural subgraph; the proof of Proposition 6.5 produces pushdown subgraphs that are not natural. Note that a pushdown subgraph can *only* exist if $\beta \cup \rho = T^2$ is not a natural blue–red decomposition.





Figure 9: A baseball diagram

Given a normalized W diagram and a pushdown subgraph $\kappa \subset T^2$, we may extend the W diagram to a larger commutative diagram of collapse maps called a *baseball diagram*, as shown in Figure 9. Certain superscripts in this diagram represent various positions on a baseball diamond: T^1, T^2, T^3 represent 1st, 2nd and 3rd bases, T^p the pitcher’s mound, T^{h1} and T^{h3} the points halfway from home plate to 1st and 3rd bases. Collapsed subgraphs of the trees T^1, T^p, T^3 in this diagram are named following a convention similar to that used earlier. Because κ is a pushdown subgraph, neither of the two subgraphs $\rho \setminus (\kappa \cap \rho), \beta \setminus (\kappa \cap \beta) \subset T^p$ contains a natural edge of T^p . It follows that neither of the two collapse maps $\sigma: T^p \rightarrow T^{h1}, \gamma: T^p \rightarrow T^{h3}$ collapses an entire natural edge of T^p . Each of the maps σ, γ therefore induces by restriction a bijection of natural vertex sets, takes each natural edge onto a natural edge inducing a bijection of natural edge sets, and is homotopic to a conjugacy relative to natural vertex sets. By restricting to natural vertex sets we therefore obtain a well-defined bijection

$\gamma \circ \sigma^{-1}$ from the natural vertex set of T^{h1} to the natural vertex set of T^{h3} which extends to a conjugacy $\xi: T^{h1} \mapsto T^{h3}$. Since collapses are transitive, we have again successfully “pushed down the peak”, without even bothering to involve home plate as in the earlier scenario:



We record this as:

Lemma 6.4 (Pushing down peaks) *Given a normalized W diagram notated as in Figure 6, and given a pushdown subgraph $\kappa \subset T^2$, there exists a baseball diagram notated as in Figure 9, in which each map $\gamma: T^p \rightarrow T^{h3}$ and $\sigma: T^p \rightarrow T^{h1}$ induces by restriction a bijection of natural vertex sets and a bijection of natural edge sets, and is homotopic rel natural vertices to a conjugacy. By composition we therefore obtain a bijection $\gamma\sigma^{-1}$ from the natural vertex set of T^{h1} to the natural vertex set of T^{h3} that extends to a conjugacy $\xi: T^{h1} \rightarrow T^{h3}$. □*

We emphasize that the conjugacy in the conclusion of this lemma need not be a map, i.e., it need not be simplicial. Nonsimplicial conjugacies resulting from Lemma 6.4 will proliferate into the proof of Proposition 6.1 given in Section 6.3, and that proof will have a certain step dedicated to patching up this problem.

Pushing down corrugation peaks One key strategy occurring in the proof of Proposition 6.1 is to set up applications of Lemma 6.4 by finding pushdown subgraphs in peaks of normalized W diagrams. Of course this is impossible if the W diagram is geodesic. Nevertheless in Proposition 6.5 we will show that when combing a fold path across an arbitrary W diagram, even one which is geodesic, one can always locate enough pushdown subgraphs to carry out the pushdown process in a useful fashion, as long as the fold path is sufficiently long when measured in free splitting units.

Consider a fold sequence $T_0^0 \mapsto \dots \mapsto T_J^0$. Consider also a zig-zag path

$$T_J^0 \rightarrow T_J^1 \xleftarrow{[\rho_J]} T_J^2 \xrightarrow{[\beta_J]} T_J^3 \leftarrow T_J^4$$

in $\mathcal{FS}'(F)$, regarded as a W diagram, not assumed to be a geodesic, nor even to be normalized, but we do assume that $T_J^2 = \beta_J \cup \rho_J$. Consider finally a stack of four combing rectangles combined into one commutative diagram as shown in Figure 10, where the given fold sequence is the T^0 row along the bottom of the diagram, and

the W zig-zag is the T_J column along the right side (in such diagrams, in general we refer to rows by dropping subscripts, and to columns by dropping superscripts). Such a diagram can be constructed, for example, by starting with the bottom row and right side, and applying Propositions 4.3, then 4.4, then 4.3, then 4.4, in that order, to comb the given fold sequence along each of the four edges of the given zig-zag path. We will also encounter such diagrams constructed by other combing processes involving concatenation and deconcatenation of combing rectangles.




Figure 10: A diagram of four combing rectangles over F . The T^0 row along the bottom is assumed to be a fold sequence. In the T_J column we assume that $T_J^2 = \rho_J \cup \beta_J$.

We can visualize Figure 10 as a piece of corrugated metal. The T^2 row forms a peak of the corrugation that we wish to push down all at once, by parallel applications of Lemma 6.4. Of course this is impossible in general, for instance when the T_J column is a geodesic path in $\mathcal{FS}'(F)$. We describe a process that allows us to push down the corrugation peak along the T^2 row, at the expense of throwing away the portion of the diagram to the right of the T_I column that is depicted in Figure 10. The next lemma says that this is possible *as long as* the bottom row has enough free splitting units between T_I^0 and T_J^0 . As a consequence, the T_j columns for $0 \leq j \leq I$ are *not* geodesic paths in $\mathcal{FS}'(F)$ because $d(T_j^0, T_j^4) \leq 2$, even when the T_J on the far right is geodesic. We thus obtain a key indicator of “hyperbolic” behavior: local curve shortening.

The following proposition introduces the constant $4 \operatorname{rank}(F) - 3$, which is needed for the proof of Proposition 6.1.

Proposition 6.5 *For any diagram as in Figure 10, if the number of free splitting units between T_I^0 and T_J^0 is at least $4 \operatorname{rank}(F) - 3$ then there is a commutative diagram*

$$\begin{array}{ccccc}
 T_0^4 & \longrightarrow & \cdots & \longrightarrow & T_I^4 \\
 \downarrow & & & & \downarrow \\
 T_0^{h3} & \longrightarrow & \cdots & \longrightarrow & T_I^{h3} \\
 \parallel \xi_0 & & & & \parallel \xi_I \\
 T_0^{h1} & \longrightarrow & \cdots & \longrightarrow & T_I^{h1} \\
 \uparrow & & & & \uparrow \\
 T_0^0 & \longrightarrow & \cdots & \longrightarrow & T_I^0
 \end{array}$$

such that the top and bottom horizontal rows are the same foldable sequences as the top and bottom rows of Figure 10 between the T_0 and T_I columns, the T^{h1} and T^{h3} rows are foldable sequences, for each $j = 0, \dots, J$ the function ξ_j is a (nonsimplicial) conjugacy between T_j^{h1} and T_j^{h3} , and the top and bottom horizontal rectangles are combing rectangles obtained from the top and bottom combing rectangles of Figure 10 between the T_0 and T_I columns by application of the composition of combing rectangles Lemma 4.5.

Proof There are three steps to the proof: normalization, pullback and pushdown.

Step 1: Normalization Knowing that $T_J^2 = \beta_J \cup \rho_J$, and knowing for each $j = 0, \dots, J$ that β_j, ρ_j are the union of the edgelets mapped to β_J, ρ_J , respectively, under the foldable map $T_j^2 \mapsto T_J^2$, it follows that $T_j^2 = \beta_j \cup \rho_j$. If the T_J column is already normalized, that is, if $\beta_J \cup \rho_J = T_J$ is a blue–red decomposition, then the same is true of $\beta_j \cup \rho_j = T_j$, and so each T_j column is normalized and we pass directly to Step 2.

Otherwise, suppose that β_J, ρ_J have some edgelets in common, the union of which is a subgraph with nondegenerate components which, abusing notation, we denote $\beta_J \cap \rho_J \subset T_J^2$. It follows that for each $j = 0, \dots, J$ the graphs β_j, ρ_j have some edgelets in common, these edgelets mapping to $\beta_J \cap \rho_J$ by the foldable map $T_j^2 \mapsto T_J^2$; their union forms a subgraph $\beta_j \cap \rho_j \subset T_j^2$. For each $j = 0, \dots, J$ we may now construct normalization diagrams in parallel, the j^{th} such diagram shown in Figure 11.




Figure 11: Parallel normalization diagrams associated to the W zig-zags in Figure 10

We claim that for each of the seven arrows in Figure 11, as j varies from 0 to J we obtain a combing rectangle. One can visualize this statement as a 3–dimensional commutative diagram where the normalization diagrams are lined up in parallel vertical planes, connected up by six foldable sequences (one for each of the six positions in the normalization diagram) and seven combing rectangles (one for each of the seven arrows). The claim is true by hypothesis for the four arrows on the top of the diagram. To obtain the combing rectangle with vertical arrows from T_j^2 to $T_j'^2$, since $\beta_j \cap \rho_j$ is the inverse image of $\beta_J \cap \rho_J$ under the foldable map $T_j^2 \mapsto T_J^2$, by Proposition 4.3 the collapse maps

$$T_j^2 \xrightarrow{[\beta_j \cap \rho_j]} T_j'^2$$

fit together in a combing rectangle as follows:

$$\begin{array}{ccccccc} T_0^2 & \longrightarrow & \cdots & \longrightarrow & T_I^2 & \longrightarrow & \cdots & \longrightarrow & T_J^2 \\ \downarrow [\beta_0 \cap \rho_0] & & & & \downarrow [\beta_I \cap \rho_I] & & & & \downarrow [\beta_J \cap \rho_J] \\ T_0'^2 & \longrightarrow & \cdots & \longrightarrow & T_I'^2 & \longrightarrow & \cdots & \longrightarrow & T_J'^2 \end{array}$$

The two combing rectangles with vertical arrows from $T_j'^2$ to T_j^1 and from $T_j'^2$ to T_j^3 , respectively, are obtained by two applications of Lemma 4.6, *decomposition of combing rectangles*, the first application using the T_j^2 to T_j^1 and the T_j^2 to $T_j'^2$ combing rectangles, and the second using the T_j^2 to T_j^3 and the T_j^2 to $T_j'^2$ combing rectangles. This proves the claim.

The outcome of the claim is a commutative diagram of the form shown in Figure 12, in which the top and bottom rectangles are the same combing rectangles as in Figure 10.

By construction (see Figure 8), the zig-zag path on the right side of Figure 12 is normalized, completing Step 1.




Figure 12: The outcome of normalizing Figure 10, using the parallel normalization diagrams of Figure 11

Step 2: Pullback This is the central argument where the concepts of free splitting units are used to their maximal effect.

Having carried out Step 1, we may now go back to Figure 10 and assume that each T_j column is a normalized W zig-zag. In other words, for each j we have a blue–red decomposition $\beta_j^2 \cup \rho_j^2 = T_j^2$.

Let Υ be the number of free splitting units along the bottom row of the diagram between T_I^0 and T_J^0 , and choose a sequence $I \leq i(0) < \cdots < i(\Upsilon) \leq J$ so that for each $u = 1, \dots, \Upsilon$ there is at least 1 free splitting unit between $T_{i(u-1)}^0$ and $T_{i(u)}^0$. By hypothesis we have $\Upsilon \geq 4 \operatorname{rank}(F) - 3$.

We prove that the blue–red decomposition $\beta_I \cup \rho_I = T_I^2$ is not natural. Arguing by contradiction, suppose that $\beta_I \cup \rho_I = T_I^2$ is natural. By Definition 5.1, it follows that $\beta_i \cup \rho_i = T_i^2$ is natural for $I \leq i \leq J$. By Lemma 5.2, the interval $I \leq i \leq J$ breaks into no more than $4 \operatorname{rank}(F) - 3$ subintervals on each of which the complexity of β_i is constant. By Definition 5.6, on each of these subintervals there is less than 1 free splitting unit, and so each of these subintervals contains at most one entry from

the sequence $i(0) < \dots < i(\Upsilon)$. It follows that $\Upsilon \leq 4 \operatorname{rank}(F) - 4$, contradicting the hypothesis.

Remark The previous version of this paper contained an invalid argument, starting from the statement that the blue–red decomposition $\beta_i \cup \rho_i = T_i^2$ is natural for $I \leq i \leq J$. The erroneous statement, which incorrectly exploited β_i, ρ_i , said that if one expands T_i^2 by blowing up each vertex $v \in \beta_i \cap \rho_i$, pulling the blue and red edges at v apart to form two vertices connected by a gray edge, then the resulting tree with F –action is a free splitting. The error is that the inserted gray edges might have nontrivial stabilizers. Correcting this error led to a revamping of the theory of free splitting units presented in Section 5. In particular, the concept of an “invariant natural blue–red decomposition” in Definition 5.1, and the diameter bounds of Lemma 5.2, are new to this version of the paper and were concocted to correctly exploit the subgraphs $\beta_i, \rho_i \subset T_i^2$.

Step 3: Pushdown Having carried out Steps 1 and 2, we assume now that we have a commutative diagram as shown in Figure 13, in which each column is normalized and the blue–red decomposition $\beta_I \cup \rho_I = T_I^2$ is not natural. It follows that T_I^2 has a natural edge e that contains both red and blue edgelets. Using this, we shall produce the commutative diagram needed for the conclusion of Proposition 6.5. The argument will be a somewhat more intricate version of the parallel normalization process used in Step 1, using parallel baseball diagrams instead.

Define a proper F –equivariant natural subgraph $\kappa_I = T_I^2$ to be the complement of the orbit of e , and so every natural edge of T_I^2 not in κ_I contains both a red and a blue edgelet. By decreasing induction on $j \in \{0, \dots, I-1\}$ define an F –equivariant subgraph $\kappa_j \subset T_j^2$ to be the inverse image of κ_{j+1} under the foldable map $T_j^2 \mapsto T_{j+1}^2$ (ignoring degenerate components as usual); equivalently κ_j is the inverse image of κ_I under $T_j^2 \mapsto T_I^2$. It follows that the subgraphs $\kappa_j \subset T_j^2$ are proper for all $j = 0, \dots, I$.

We claim that for $j = 0, \dots, I$ the graph κ_j is a pushdown subgraph of T_j^2 . To prove this, given a natural edge $\eta_j \subset T_j^2$ such that $\eta_j \not\subset \kappa_j$, we must find a red and a blue edgelet in η_j , neither of which is in κ_j . Foldable maps take natural vertices to natural vertices and natural edges to nondegenerate natural edge paths, so the image of η_j under the foldable map $T_j^2 \mapsto T_I^2$ is a nondegenerate natural edge path denoted $\eta_I \subset T_I^2$. Since $\eta_j \not\subset \kappa_j$, it follows that $\eta_I \not\subset \kappa_I$, and so η_I contains a natural edge not in κ_I , which therefore has both a red and a blue edgelet. Since natural edges not in κ_I have interior disjoint from κ_I it follows that η_I contains a red and a blue edgelet neither of which is in κ_I . By pulling back under the foldable map $T_j^2 \mapsto T_I^2$ we obtain a red and a blue edgelet in η_j neither of which is in κ_j .




Figure 13: Each of the four horizontal rectangles is a combing rectangle. We assume that every column is a normalized W zig-zag and that the tree T_I^2 has an edge e containing both red and blue edgelets.

We now apply Lemma 6.4 in parallel to each column j of Figure 13 for $j = 0, \dots, I$. The resulting baseball diagrams, commutative diagrams of collapse maps, are shown in Figure 14 (compare Figure 9). Lemma 6.4 also produces conjugacies $T_j^p \mapsto T_j^{h3}$ and $T_j^p \mapsto T_j^{h1}$ and hence conjugacies $T_j^{h1} \rightarrow T_j^{h3}$. What we are still missing, however, are the conclusions of Proposition 6.5 concerned with combing rectangles and commutativity.

We claim that for each of the nine arrows in Figure 14, as j varies from 0 to I , we obtain a combing rectangle. As in Step 1, one visualizes this as a 3-dimensional commutative diagram by lining up the baseball diagrams in parallel vertical planes, connected up by eight foldable sequences (one for each of the eight positions in the baseball diagram) and nine combing rectangles (one for each of the nine arrows). The claim is true by hypothesis for the four arrows on the top of the diagram.

For the arrow from 2nd base to the pitcher’s mound, since κ_j is the inverse image of κ_J under the foldable map $T_j^2 \mapsto T_I^2$, by Proposition 4.3 the collapse maps

$$T_j^2 \xrightarrow{[\kappa_j]} T_j^p$$




Figure 14: The baseball diagram associated to the W-diagram from T_j^0 to T_j^4

fit together in a combing rectangle:

$$\begin{array}{ccc}
 T_0^2 & \longrightarrow \cdots \longrightarrow & T_I^2 \\
 [\kappa_0] \downarrow & & \downarrow [\kappa_I] \\
 T_0^p & \longrightarrow \cdots \longrightarrow & T_I^p
 \end{array}$$

Notice that for each $j = 0, \dots, I$, the subgraph $\kappa_j \cup \rho_j$ is proper, because any natural edge not in κ_j contains a blue edgelet not in κ_j , which is also not in $\kappa_j \cup \rho_j$. Similarly the subgraph $\kappa_j \cup \beta_j$ is proper. By Proposition 4.3, since $\kappa_j \cup \rho_j$ is the inverse image of $\kappa_{j+1} \cup \rho_{j+1}$, and since $\kappa_j \cup \beta_j$ is the inverse image of $\kappa_{j+1} \cup \beta_{j+1}$, we obtain combing rectangles

$$\begin{array}{ccc}
 T_0^2 & \longrightarrow \cdots \longrightarrow & T_I^2 \\
 [\kappa_0 \cup \beta_0] \downarrow & & \downarrow [\kappa_I \cup \beta_I] \\
 T_0^{h3} & \longrightarrow \cdots \longrightarrow & T_0^{h3}
 \end{array}
 \quad
 \begin{array}{ccc}
 T_0^2 & \longrightarrow \cdots \longrightarrow & T_I^2 \\
 [\kappa_0 \cup \rho_0] \downarrow & & \downarrow [\kappa_I \cup \rho_I] \\
 T_0^{h1} & \longrightarrow \cdots \longrightarrow & T_0^{h1}
 \end{array}$$

Rectangles II and III do not correspond to any of the nine arrows in the baseball diagram, but to invisible arrows going from 2nd base to the point halfway between 1st base and home plate and from 2nd base to the point halfway between 3rd base and home plate.

For the arrows going from the pitcher’s mound to the points halfway between 1st and home and halfway between 3rd and home, apply Lemma 4.6, *decomposition of combing rectangles*, first to combing rectangles II and I and then to combing rectangles III and I,

to obtain combing rectangles

$$\begin{array}{ccc}
 T_0^P & \longrightarrow \cdots \longrightarrow & T_I^P \\
 [\beta_0 \setminus (\kappa_0 \cap \beta_0)] \downarrow \gamma_0 & \text{IV} & \gamma_I \downarrow [\beta_I \setminus (\kappa_I \cap \beta_I)] \\
 T_0^{h3} & \longrightarrow \cdots \longrightarrow & T_0^{h3} \\
 \\
 T_0^P & \longrightarrow \cdots \longrightarrow & T_I^P \\
 [\rho_0 \setminus (\kappa_0 \cap \rho_0)] \downarrow \sigma_0 & \text{V} & \sigma_I \downarrow [\rho_I \setminus (\kappa_I \cap \rho_I)] \\
 T_0^{h1} & \longrightarrow \cdots \longrightarrow & T_0^{h1}
 \end{array}$$

where we follow the same notation convention for subgraphs of T_0^P as used in the original baseball diagram Figure 9.

For the arrows going from 1st base and 3rd base to the points halfway home, applying Lemma 4.6, *decomposition of combing rectangles* to combing rectangle II and the 2nd base to 3rd base combing rectangle, and then to combing rectangle III and the 2nd base to 1st base combing rectangle, we obtain combing rectangles:

$$\begin{array}{ccc}
 T_0^3 & \longrightarrow \cdots \longrightarrow & T_I^3 \\
 [\kappa_0 \setminus (\kappa_0 \cap \beta_0)] \downarrow \gamma_0 & \text{VI} & \downarrow [\kappa_I \setminus (\kappa_I \cap \beta_I)] \\
 T_0^{h3} & \longrightarrow \cdots \longrightarrow & T_0^{h3} \\
 \\
 T_0^1 & \longrightarrow \cdots \longrightarrow & T_I^1 \\
 [\kappa_0 \setminus (\kappa_0 \cap \rho_0)] \downarrow \sigma_0 & \text{VII} & \downarrow [\kappa_I \setminus (\kappa_I \cap \rho_I)] \\
 T_0^{h1} & \longrightarrow \cdots \longrightarrow & T_0^{h1}
 \end{array}$$

Applying Lemma 4.5, *composition of combing rectangles*, by composing the two combing rectangles corresponding to the arrows along the 1st base foul line in Figure 14 we obtain the combing rectangle from the T^0 row to the T^{h1} row needed for the conclusion of Proposition 6.5. Similarly, by composing the two combing rectangles corresponding to the arrows along the 3rd base foul line we obtain the combing rectangle from the T^4 row to the T^{h3} row.

To complete Step 3 and the proof of the proposition, it remains to construct the commutative diagram of conjugacy maps $\xi_j: T_j^{h1} \rightarrow T_j^{h3}$ in the conclusion of the lemma.

For this purpose it suffices to replace combing rectangles IV and V by commutative diagrams of conjugacies of the form

$$\begin{array}{ccc}
 T_0^P & \longrightarrow \cdots \longrightarrow & T_I^P \\
 \downarrow \bar{\gamma}_0 & \bar{\text{IV}} & \bar{\gamma}_I \downarrow \\
 T_0^{h3} & \longrightarrow \cdots \longrightarrow & T_I^{h3}
 \end{array}
 \qquad
 \begin{array}{ccc}
 T_0^P & \longrightarrow \cdots \longrightarrow & T_I^P \\
 \downarrow \bar{\sigma}_0 & \bar{\text{V}} & \bar{\sigma}_I \downarrow \\
 T_0^{h1} & \longrightarrow \cdots \longrightarrow & T_I^{h1}
 \end{array}$$

for then defining $\xi_j = \bar{\gamma}_j \circ \bar{\sigma}_j^{-1}: T_j^{h1} \rightarrow T_j^{h3}$, we will be done. While Lemma 6.4 produces conjugacies $T_j^{h1} \rightarrow T_j^{h3}$ for each $j = 0, \dots, J$, if that lemma is used crudely there is no guarantee that these conjugacies will form commutative diagrams as needed. With a little care in how Lemma 6.4 is applied we can get the needed guarantee. We construct diagram $\bar{\text{IV}}$ in detail, the construction of $\bar{\text{V}}$ being similar. The construction is by induction, starting from the T_I column on the far right and moving leftward.

First apply Lemma 6.4 to produce a conjugacy $\bar{\gamma}_I: T_I^P \rightarrow T_I^{h3}$ so that the restrictions of γ_I and $\bar{\gamma}_I$ to natural vertex sets are the same. Proceeding by decreasing induction on j , suppose that for some j we have produced all the conjugacies from column T_j to T_I in diagram $\bar{\text{IV}}$ making that portion of the diagram commute, and so that the restrictions to natural vertex sets of the conjugacies in diagrams IV and $\bar{\text{IV}}$ are the same from column T_j to column T_I . We must choose the conjugacy $\bar{\gamma}_{j-1}: T_{j-1}^P \rightarrow T_{j-1}^{h3}$ so as to fill in a commutative diagram of F -equivariant functions

$$\begin{array}{ccc}
 T_{j-1}^P & \xrightarrow{f_j} & T_j^P \\
 \bar{\gamma}_{j-1} \downarrow \cdots \downarrow & & \downarrow \bar{\gamma}_j \\
 T_{j-1}^{h3} & \xrightarrow{g_j} & T_j^{h3}
 \end{array}$$

where f_j, g_j are the foldable maps in rectangle IV, and where the restrictions of $\bar{\gamma}_{j-1}$ and γ_{j-1} to natural vertex sets are the same. This tells us how to define $\bar{\gamma}_{j-1}$ on natural vertex sets. Consider a natural edge $\eta \subset T_{j-1}^P$. By Lemma 6.4 its image $\gamma_{j-1}(\eta) \subset T_{j-1}^{h3}$ is a natural edge whose endpoints are the $\bar{\gamma}_{j-1}$ images of the endpoints of η . The foldable map $f_j: T_{j-1}^P \mapsto T_j^P$ is injective on η , the conjugacy $\bar{\gamma}_j$ is injective on $f_j(\eta)$, and we have the following equation of subsets:

$$g_j(\gamma_{j-1}(\eta)) = \gamma_j(f_j(\eta)) = \bar{\gamma}_j(f_j(\eta)).$$

The foldable map g_j is injective on the natural edge $\gamma_{j-1}(\eta)$, and therefore has a homeomorphic inverse $g_j^{-1}: \bar{\gamma}_j(f_j(\eta)) \rightarrow \gamma_{j-1}(\eta)$, and so we can define

$$\bar{\gamma}_{j-1} | \eta = (g_j^{-1} \circ \bar{\gamma}_j \circ f_j) | \eta.$$

This completes Step 3 and the proof of Proposition 6.5. □

6.3 Proof of Proposition 6.1

Prologue Consider a fold sequence $S_0 \mapsto \dots \mapsto S_K$ over F , a free splitting $F \curvearrowright T$, and an augmented projection diagram of maximal depth $k_T = \pi(T)$ as notated in Figure 5 of Section 6.1, whose top row has the fold sequence $T_{k_T} \mapsto \dots \mapsto T_L = T$ as a terminal segment. Let Υ be the number of free splitting units between T_{k_T} and $T_L = T$. Using the constant $b_1 = 4 \operatorname{rank}(F) - 3$ from Proposition 6.5, we list every b_1^{th} term of the back greedy subsequence of this fold sequence as

$$k_T \leq L_\Omega < L_{\Omega-1} < \dots < L_1 < L_0 = L,$$

where $\Omega = \lfloor \Upsilon/b_1 \rfloor$. Thus L_ω is the greatest integer less than $L_{\omega-1}$ such that there are exactly b_1 free splitting units between T_{L_ω} and $T_{L_{\omega-1}}$, for each $\omega = 1, \dots, \Omega$. Emphasizing only those T s with subscripts from the list L_Ω, \dots, L_0 , and assigning them a superscript 0 for later purposes, we may write the augmented projection diagram in the form

$$\begin{array}{ccccccccccc}
 T_0^0 & \longrightarrow & \cdots & \longrightarrow & T_{k_T}^0 & \longrightarrow & \cdots & \longrightarrow & T_{L_\Omega}^0 & \longrightarrow & T_{L_{\Omega-1}}^0 & \longrightarrow & \cdots & \longrightarrow & T_{L_1}^0 & \longrightarrow & T_{L_0}^0 = T \\
 \downarrow & & & & \downarrow & & & & & & & & & & & & & \\
 S_0' & \longrightarrow & \cdots & \longrightarrow & S_{k_T}' & & & & & & & & & & & & & \\
 \uparrow & & & & \uparrow & & & & & & & & & & & & & \\
 S_0 & \longrightarrow & \cdots & \longrightarrow & S_{k_T} & \longrightarrow & \cdots & \longrightarrow & S_K & & & & & & & & &
 \end{array}$$

where the foldable map $T_{k_T}^0 \rightarrow T_{L_\Omega}^0$ may just be the identity map.

Consider also a vertex $R \in \mathcal{FS}'(F)$ and a geodesic path from T to R in $\mathcal{FS}'(F)$. We shall assume here that $d(T, R) \geq 3$; the case that $d(T, R) \leq 2$ will be considered in the epilogue. If the path from T to R starts with an expansion of T , prefix the path with an improper collapse. The result is a zig-zag path of the form

$$T = T_{L_0}^0 \rightarrow T_{L_0}^1 \leftarrow T_{L_0}^2 \rightarrow T_{L_0}^3 \cdots T_{L_0}^D = R,$$

where $D = d(T, R)$ or $d(T, R) + 1$ and $D \geq 3$. The peaks along this zig-zag are the even terms strictly between 0 and D , the first such peak being $T_{L_0}^2$. For each peak along this path, applying Lemma 6.3 together with the assumption that $d(T, R) \geq 3$ it follows that the peak is the union of its two collapse graphs. The number of peaks along this zig-zag path equals $\lfloor (D-1)/2 \rfloor$, which equals $(D-2)/2$ if D is even and $(D-1)/2$ if D is odd.

By combing the foldable sequence $T_0^0 \mapsto \cdots \mapsto T_{L_0}^0$ across each collapse or expansion of the zig-zag path $T_{L_0}^0 \rightarrow T_{L_0}^1 \leftarrow \cdots \leftarrow T_{L_0}^D = R$, alternately applying *combing by collapse* Proposition 4.3 and *combing by expansion* Proposition 4.4, and by stacking the resulting combing rectangles atop the augmented projection diagram, we obtain the Big diagram, Step 0, shown in Figure 15.

Proposition 6.1 will be proved by explicitly transforming the Big diagram, Step 0 into a projection diagram from R onto $S_0 \mapsto \cdots \mapsto S_K$ of an appropriate depth l needed to verify the conclusions of the proposition. This transformation is primarily an induction that uses the pushdown tools of Section 6.2, followed by an epilogue which uses the pushdown tools one more time. As the proof progresses we will consider the truncated fold sequences $T_{k_T}^0 \mapsto \cdots \mapsto T_{L_\omega}^0$ for increasing values of ω , but such truncation will not affect measurements of free splitting units between T_i^0 and T_j^0 as long as $k_T \leq i \leq j \leq L_\omega$ (see the remark following Definition 5.10).

Induction We explain in detail how to carry out the first step of the induction. Under our assumption that $d(T, R) \geq 3$, the T_{L_0} column of the Big diagram, Step 0 has a peak at $T_{L_0}^2$. Assuming furthermore that $\Upsilon \geq b_1$, equivalently $\Omega \geq 1$, then L_1 is defined and there are at least $b_1 = 4 \operatorname{rank}(F) - 3$ free splitting units between $T_{L_1}^0$ and $T_{L_0}^0$. We may therefore apply Proposition 6.5 to the portion of the diagram between the T^0 and T^4 rows as follows: trim away all portions of the Big diagram, Step 0 that lie to the right of the T_{L_1} column and below the T^D row, and use the conclusion of Proposition 6.5 to replace the combing rectangles between the T^0 and T^4 rows, to get the Big diagram, Step 0.1, shown in Figure 16.

The rectangles of the Big diagram, Step 0.1 between the T^0 and T^{h_1} rows and between the T^{h_3} and T^4 rows are combing rectangles. Each $\xi_j: T_j^{h_1} \rightarrow T_j^{h_3}$ is a conjugacy, possibly nonsimplicial. Now we must pause to patch things up in order to make these conjugacies simplicial.

We claim that, by an operation of equivariant subdivision of simplicial structures and re-assignment of barycentric coordinates on edgelets, carried out over all free splittings in the Big diagram, Step 0.1, but *without* changing any of the functions, we may assume that the conjugacies ξ_i are indeed simplicial maps. Here are the details of this operation.

Consider first the conjugacy $\xi_{L_1}: T_{L_1}^{h_1} \rightarrow T_{L_1}^{h_3}$. We may subdivide $T_{L_1}^{h_1}$ at the pre-image of the vertex set of $T_{L_1}^{h_3}$, and simultaneously subdivide $T_{L_1}^{h_3}$ at the image of the vertex set of $T_{L_1}^{h_1}$, to obtain new equivariant vertex sets on which ξ_{L_1} is a bijection; it is also a bijection of edgelets, although it may not yet respect barycentric coordinates. We may then reassign the barycentric coordinates on one edgelet of $T_{L_1}^{h_1}$ in each F -orbit, and move these coordinates around by the F -action, to obtain a new simplicial structure on $T_{L_1}^{h_1}$. We may then push these coordinates forward under the map ξ_{L_1} to obtain new




Figure 15: Big diagram, Step 0. We emphasize the columns indexed by $L_\Omega, \dots, L_1, L_0$. Each horizontal row is a foldable sequence, and the rectangle between any two rows is a combing rectangle. The bottom row is a fold sequence, and the T^0 row from $T_{k_T}^0$ to $T_{L_0}^0$ is a fold sequence. Each peak of the T_{L_0} column is the union of its two collapse graphs. Rows in this and subsequent diagrams will be indicated by stripping off subscripts, for instance the “ T^0 row” refers to the foldable sequence $T_0^0 \mapsto \dots \mapsto T_{L_0}^0$; similarly, columns are indicated by stripping off superscripts. Since each peak of column T_{L_0} between rows T^0 and T^D is the union of its two collapse graphs, it follows that each peak of each column T_j between rows T^0 and T^D is the union of its two collapse graphs, because the two collapse graphs at a column j peak T_j^{2i} are the pullbacks under the foldable map $T_j^{2i} \mapsto T_{L_0}^{2i}$ of the two collapse graphs at the corresponding column L_0 peak $T_{L_0}^{2i}$.




Figure 16: Big diagram, Step 0.1

barycentric coordinates on the edgelets of $T_{L_1}^{h3}$. Having carried out these operations, the map ξ_{L_1} is now a simplicial conjugacy.

Now we move left one step: by a similar subdivision/re-assignment on $T_{L_{1-1}}^{h1}$, pulling back vertices and barycentric coordinates under the foldable map $T_{L_{1-1}}^{h1} \mapsto T_{L_1}^{h1}$, we may assume that this map is simplicial. Similarly, by a subdivision/re-assignment on $T_{L_{1-1}}^{h3}$, we may assume that the foldable map $T_{L_{1-1}}^{h3} \mapsto T_{L_1}^{h3}$ is simplicial. We have now verified that in the commutative diagram

$$\begin{array}{ccc}
 T_{L_{1-1}}^{h3} & \longrightarrow & T_{L_1}^{h3} \\
 \parallel \xi_{L_{1-1}} & & \parallel \xi_{L_1} \\
 T_{L_{1-1}}^{h1} & \longrightarrow & T_{L_1}^{h1}
 \end{array}$$

the top, bottom, and right sides are simplicial maps; by commutativity, the left side is therefore automatically simplicial.

Now we continue to move left: by similar subdivisions/re-assignments carried out one at a time on the trees in rows T^{h1} and T^{h3} , moving to the left one at a time from the last map in each row, we may assume that these rows are simplicial; having done this, by commutativity each of the maps $\xi_j: T_j^{h1} \rightarrow T_j^{h3}$ is automatically a simplicial conjugacy. Now we move up: by similar subdivisions/re-assignments carried out one at a time on the trees in rows T^4, \dots, T^D , starting with the collapse maps $T_j^4 \mapsto T_j^{h3}$ and moving upward, we may assume that each vertical arrow above row T^{h3} is simplicial; having done this, each of the horizontal arrows from row T^{h3} upward and between columns T_0 and T_{L_1} is automatically simplicial. Now, from $T_{L_1}^D$ we move to the right: by similar subdivisions/re-assignments we may assume that each of the maps $T_{L_1}^D \mapsto \dots \mapsto T_{L_0}^D = R$ is simplicial. Finally, in a similar manner moving down from row T^{h3} to row S , then moving right from S_{k_T} to S_K , we have proved the claim.

Knowing now that we have simplicial conjugacies $\xi_j: T_j^{h1} \rightarrow T_j^{h3}$, and using commutativity of the rectangle between rows T^{h1} and T^{h3} , we may identify T_j^{h1} and T_j^{h3} via the map ξ_j , replacing these two rows by a single row as shown in the *Big diagram, Step 1*.

In summary, when $d(T, R) \geq 3$ and $\Upsilon \geq b_1$, we have completed the first iteration of the induction argument: starting from the Big diagram, Step 0, by applying Proposition 6.5, trimming away everything to the right of column T_{L_1} and below row T^D , and replacing everything between rows T^0 and T^4 , we get the Big diagram, Step 0.1, and then by subdividing and re-assigning barycentric coordinates we may assume that the conjugacies between rows T^{h1} and T^{h3} are simplicial. Identifying rows T^{h1} and T^{h3} , we obtain the Big diagram, Step 1, shown in Figure 17. In the process we have decreased by 2 the lengths of all vertical zig-zag paths and the number of combing rectangles between the T^0 and T^D rows. Observe that the conjugacy class of the free splitting R , and the equivalence class of the fold sequence $S_0 \mapsto \dots \mapsto S_K$, have not been altered by these subdivision/re-assignment operations.

To complete the inductive step there is one last thing to do, namely to verify that along the zig-zag path in column T_{L_1} on the right side of the Big diagram, Step 1, each peak is the union of its two collapse graphs. This is evident for each peak from $T_{L_1}^6$ upward, since the collapse maps and collapse graphs are unchanged at those peaks from the Big diagram, Step 0. For the peak at $T_{L_1}^4$, one of the collapse graphs is unchanged from the Big diagram, Step 0, namely that of the map $T_{L_1}^4 \mapsto T_{L_1}^5$. For the collapse graph of the map $T_{L_1}^4 \mapsto T_{L_1}^h$, we use the part of the conclusion of Proposition 6.5, which tells us that the combing rectangle in the Big diagram, Step 1 between the T^4 and T^h




Figure 17: Big diagram, Step 1

rows is obtained by an application of *composition of combing rectangles*, Lemma 4.5, using the combing rectangle in the Big diagram, Step 0 between the T^4 and T^3 rows and between the T_0 and T_{L_1} columns. What Lemma 4.5 allows us to conclude is that the collapse graph of the Step 0 map $T_{L_1}^4 \mapsto T_{L_1}^3$ is contained in the collapse graph of Step 1 map $T_{L_1}^4 \mapsto T_{L_1}^h$. The union of the two collapse graphs of $T_{L_1}^4$ in the Big diagram, Step 1, is therefore still equal to $T_{L_1}^4$.

Remark The reader may wonder why the initial normalization step was necessary in the proof of Proposition 6.5: we could have started with a *normalized zig-zag* geodesic on the right side of the Big diagram, Step 0. This would imply that the T^4 column in that diagram is normalized at $T_{L_1}^4$. Nonetheless it is possible that the T^4 column in the Big diagram, Step 1 is not normalized at $T_{L_1}^4$, because the collapse graph for $T_{L_1}^4 \mapsto T_{L_1}^h$ may be strictly larger than the collapse graph for $T_{L_1}^4 \mapsto T_{L_1}^3$. If so then the normalization step of Proposition 6.5 is inescapable in the next step of the induction.

We now describe the induction step in general. From the hypothesis we have $d(T, R) \leq \max\{2\Omega, 1\}$. If $d(T, R) \leq 2$ then we refer to the epilogue below. Otherwise, under the

assumption $d(T, R) \geq 3$, we have $D \leq d(T, R) + 1 \leq 2\Omega + 1$, and so we may repeat the above argument inductively a total of $\lfloor (D - 1)/2 \rfloor$ times, removing the corrugation peaks one at a time. For each $\omega = 2, \dots, \lfloor (D - 1)/2 \rfloor$, the ω step of the induction will transform the Big diagram, Step $\omega - 1$ into the Big diagram, Step ω . The Big diagram Step $\omega - 1$, in comparison with the Big diagram, Step 1, will have $L_{\omega-1}$ in place of L_1 and $T^{2\omega}$ in place of T^4 , and will have a stack of $D - 2\omega + 2$ combing rectangles between the T^0 and T^D rows. To carry out the transformation, first we trim away the portion of the diagram to the right of column T_{L_ω} , on or above row T^0 , and below row T^D . Then we replace the four combing rectangles between rows T^0 and $T^{2\omega+2}$ by two combing rectangles and a commutative diagram of conjugacies. Then we carry out a subdivision/re-assignment operation, which allows us to assume that the conjugacies are simplicial. Then we collapse the commutative diagram of conjugacies, identifying its two rows into a single row. We have now produced the Big diagram, Step ω , with a stack of $D - 2\omega$ combing rectangles between the T^0 and T^D rows: we have decreased by 2 the lengths of all vertical zig-zag paths between the T^0 and T^D rows and decreased by 1 the number of corrugation peaks. Finally we verify that each peak along column T_{L_ω} is still the union of its two collapse graphs.

At each stage of the induction, we have not altered the conjugacy class of R nor the equivalence class of $S_0 \mapsto \dots \mapsto S_K$.

Epilogue If $d(T, R) \geq 3$, when the induction process stops we have backed up to column T_{L_ω} where $\omega = \lfloor (D - 1)/2 \rfloor$, and there are no remaining corrugation peaks above row T^0 . We obtain the Big diagram, Step $\lfloor (D - 1)/2 \rfloor$, a not-so-big diagram that comes in two cases. The Case 1 diagram occurs when D is even, and it has two combing rectangles between row T^0 and row T^D ; see Figure 18. The Case 2 diagram occurs when D is odd and has only one such combing rectangle; see Figure 19. In each of these diagrams, the conjugacy class of R and the equivalence class of the fold sequence $S_1 \mapsto \dots \mapsto S_K$ have not been changed from the initial setup in the prologue.

If $d(T, R) \leq 2$ then, starting from the augmented projection diagram depicted in the prologue, and depending on the nature of the geodesic from T to R , we proceed as follows. If $d(T, R) = 1$ and there is a collapse $T \succ R$, we comb the T^0 row along this collapse to obtain the Case 2 diagram with $\omega = 0$ and $T_{L_\omega}^D = T_{L_0}^D = R$. If $d(T, R) = 1$ and there is an expansion $T \prec R$ then we append an improper collapse $T \succ T$ to get a length 2 collapse–expand zig-zag $T \succ T \prec R$, and we comb the T^0 row along this collapse–expand to obtain the Case 1 diagram with similar notation changes. If $d(T, R) = 2$ and there is a collapse–expand from T to R then, combing the T^0 row along this collapse–expand, we produce the Case 1 diagram with similar notation changes. Finally, if $d(T, R) = 2$ and there is an expand–collapse from T




Figure 18: Case 1 a collapse–expand from T^0 to T^D




Figure 19: Case 2 a collapse from T^0 to T^D

to R , then combing the T^0 row along this expand–collapse, we obtain the Case 3 diagram in Figure 20.

We now finish off Case 1; afterwards we shall reduce Cases 2 and 3 to Case 1. In the Case 1 diagram, trim off everything to the right of column T_{k_T} , on or above row T^0 , and below row T^D , to obtain the diagram shown in Figure 21, which has a corrugation




Figure 20: Case 3 an expand-collapse from T^0 to T^2

peak along the T^0 row. We must consider two subcases, depending on whether the peak $T_{k_T}^0$ of the W zig-zag in column k_T is the union of its two collapse graphs b_{k_T}, r_{k_T} .

Suppose first that $T_{k_T}^0 \neq b_{k_T} \cup r_{k_T}$ in Figure 21. For each $j = 0, \dots, k_T$, in the tree T_j^0 , which is the peak of the W zig-zag in column j , the union of its two collapse graphs $b_j \cup r_j$ is a proper subgraph, that subgraph being the inverse image of $b_{k_T} \cup r_{k_T}$ under the foldable map $T_j^0 \mapsto T_{k_T}^0$. We may therefore carry out the simplistic pushdown depicted in Figure 7, in parallel as j varies from 0 to k_T , resulting in a diagram of the form depicted in Figure 22. In Figure 22, the T'' row is obtained by applying Proposition 4.3, *combing by collapse*, using the collapse graphs $b_j \cup r_j \subset T_j^0$, and the middle two combing rectangles are each obtained by an application of Lemma 4.6, *decomposition of combing rectangles*. By applications of Lemma 4.5, *composition of combing rectangles*, we may compose the lower two and the upper two combing rectangles of Figure 22 to produce a depth k_T projection diagram from R to $S_0 \mapsto \cdots \mapsto S_K$, and the proof of Proposition 6.1 is complete in this case.

Suppose next that $T_{k_T}^0 = b_{k_T} \cup r_{k_T}$ in Figure 21. From the hypothesis of Proposition 6.1, there are at least $b_1 = 4 \text{rank}(F) - 3$ free splitting units along the bottom row of the diagram between S_0 and S_{k_T} . Let $\ell \in \{0, \dots, k_T\}$ be the largest integer such that there are at least b_1 free splitting units between S_ℓ and S_{k_T} , from which it follows that there are exactly b_1 free splitting units between S_ℓ and S_{k_T} . We may now carry




Figure 21: The Case 1 diagram, trimmed down




Figure 22: The result of a parallel simplistic pushdown on Figure 21, in the case when $T_{k_T}^0 \neq \beta_{k_T} \cup \rho_{k_T}$. Concatenating the upper two combing rectangles into a single one, and the same for the lower two, we obtain a projection diagram.

out one last iteration of the induction. Applying Proposition 6.5, remove all portions of the diagram in Figure 21 to the right of column l , above the S row, and below the T^D row, and replace the four combing rectangles by two combing rectangles and a commutative diagram of conjugacies. After an operation of subdivision and re-assignment of barycentric coordinates, we may assume that the conjugacies are all simplicial. After collapsing the commutative diagram of conjugacies, identifying its two rows to a single row, we obtain the diagram depicted in Figure 23, in which the conjugacy class of the free splitting R and the equivalence class of the fold sequence $S_0 \mapsto \dots \mapsto S_K$ remain unchanged. This is the desired projection diagram from the free splitting R to the fold sequence $S_0 \mapsto \dots \mapsto S_K$, which completes the proof of Proposition 6.1 in Case 1.




Figure 23: The projection diagram resulting from one last iteration of the induction carried out on Figure 21, in the case when $T_{k_T}^0 = \beta_{k_T} \cup \rho_{k_T}$

Remark As was remarked earlier regarding the Big diagram, Step 1, depicted in Figure 17, in the context of Case 1 depicted in Figure 23, the initial normalization step in the proof of Proposition 6.5 cannot be avoided, because there is no guarantee that the S_{k_T} column is normalized at $T_{k_T}^0$.

We reduce Case 2 to Case 1 by producing a Case 1 diagram: just attach an improper combing rectangle to the top of the case 2 diagram, by defining the foldable sequence $T'_0 \mapsto \dots \mapsto T'_{L_\Omega}$ to equal the foldable sequence $T_0^D \mapsto \dots \mapsto T_{L_\Omega}^D$, and defining for each $j = 0, \dots, L_\Omega$ an improper collapse map $T_j^D \rightarrow T'_j$ that is just the identity map.

We also reduce Case 3 to Case 1. First trim away everything in the Case 3 diagram to the right of the k_T column, on or above the T^0 row, and below the T^2 row. Next, apply Lemma 4.5, *composition of combing rectangles*, to the two combing rectangles between the S' row and the T' row, concatenating them into a single combing rectangle. Finally, attach an improper combing rectangle to the top of the diagram as in Case 2. The result is a Case 1 diagram, completing the reduction.

References

- [1] **J Aramayona, J Souto**, *Automorphisms of the graph of free splittings*, Michigan Math. J. 60 (2011) 483–493 MR2861084
- [2] **M Bestvina, M Feighn**, *Hyperbolicity of the complex of free factors* arXiv:1107.3308
- [3] **M Bestvina, M Feighn**, *Bounding the complexity of simplicial group actions on trees*, Invent. Math. 103 (1991) 449–469 MR1091614
- [4] **M Bestvina, M Feighn, M Handel**, *Laminations, trees, and irreducible automorphisms of free groups*, Geom. Funct. Anal. 7 (1997) 215–244 MR1445386
- [5] **M Bestvina, M Feighn, M Handel**, *The Tits alternative for $\text{Out}(F_n)$, I: Dynamics of exponentially-growing automorphisms*, Ann. of Math. 151 (2000) 517–623 MR1765705
- [6] **M Culler, J W Morgan**, *Group actions on \mathbf{R} -trees*, Proc. London Math. Soc. 55 (1987) 571–604 MR907233
- [7] **V Guirardel, G Levitt**, *JSJ decompositions: Definitions, existence, uniqueness, II: Compatibility and acylindricity* arXiv:1002.4564
- [8] **M Handel, L Mosher**, *Lipschitz retraction and distortion for subgroups of $\text{Out}(F_n)$* , Geom. Topol. 17 (2013) 1535–1579
- [9] **A Hatcher**, *Homological stability for automorphism groups of free groups*, Comment. Math. Helv. 70 (1995) 39–62 MR1314940
- [10] **A Hatcher, K Vogtmann**, *The complex of free factors of a free group*, Quart. J. Math. Oxford Ser. 49 (1998) 459–468 MR1660045
- [11] **A Hillion, C Horbez**, *The hyperbolicity of the sphere complex via surgery paths* arXiv:1210.6183
- [12] **I Kapovich, K Rafi**, *On hyperbolicity of free splitting and free factor complexes* arXiv:1206.3626
- [13] **B Mann**, *Hyperbolicity of the cyclic splitting complex* arXiv:1212.2986
- [14] **J F Manning**, *Geometry of pseudocharacters*, Geom. Topol. 9 (2005) 1147–1185 MR2174263
- [15] **H Masur, Y Minsky**, *Geometry of the complex of curves, I: Hyperbolicity*, Invent. Math. 138 (1999) 103–149 MR1714338
- [16] **H Masur, L Mosher, S Schleimer**, *On train-track splitting sequences*, Duke Math. J. 161 (2012) 1613–1656 MR2942790
- [17] **L Sabalka, D Savchuk**, *On the geometry of the edge splitting complex*, to appear in Groups, Geometry, and Dynamics (2013)

- [18] **S Schleimer**, *Notes on the complex of curves* (2006) Available at <http://www.warwick.ac.uk/~masgar/Maths/notes.pdf>
- [19] **P Scott, T Wall**, *Topological methods in group theory*, from: “Homological group theory”, (T Wall, editor), London Math. Soc. Lecture Note Ser. 36, Cambridge Univ. Press (1979) 137–203 MR564422
- [20] **J R Stallings**, *Topology of finite graphs*, *Invent. Math.* 71 (1983) 551–565 MR695906

Glossary

Collapse and expansion: Inverse relations amongst free splittings, denoted $S \succ T$ and $T \prec S$ respectively, defined so that T is obtained from S by collapsing to a point each component of some proper, equivariant, natural subgraph of S .

Conjugacy: An equivariant homeomorphism between free splittings, which need not be a map.

Edgelet: A 1–cell of some given simplicial structure on a tree. The term is also used in a relative sense; given a foldable map $f: S \rightarrow T$ and an edgelet e of T , an e –edgelet of f is any edgelet of S mapped by f to e .

Foldable sequence: A sequence of maps of free splittings in which any composition of any subinterval of that sequence is a foldable map.

- A *fold sequence* is a special kind of foldable sequence in which each map is a fold.
- A *fold path* is the sequence of vertices in $\mathcal{FS}'(F)$ obtained from the conjugacy classes of the free splittings along a fold sequence.

Free splitting: A minimal action of a free group on a simplicial tree such that the stabilizer of each edge is trivial. The action is properly discontinuous if and only if the stabilizer of each vertex is trivial, if and only if each vertex has finite valence.

Map: An equivariant simplicial function between free splittings. Important types of maps include:

- A *collapse map* collapses to a point each edge in an equivariant subgraph.
- A *foldable map* is injective on each natural edge, and has at least 3 gates at each natural vertex.
- A *fold map* is a foldable map defined by identifying initial segments of some pair of natural edges with the same initial vertex.

Natural cell structure: Every tree with no isolated ends and no valence 1 vertices — in particular every free splitting of a free group of rank at least 2 — has a natural cell structure, whose vertices are the points that (locally) separate the tree into some number of components at least 3. A natural subgraph is a subcomplex of the natural cell structure. Any other cell structure on the graph is a refinement of the natural cell structure.

Zig-zag path: An edge path in $\mathcal{FS}'(F)$ which alternates between expansions and collapses. Examples include all geodesic edge paths in $\mathcal{FS}'(F)$.

Index

- back greedy subsequence, 1635
 baseball diagram, 1646
- collapse, 1587, 1589, *see also Glossary*
 map, *see also Glossary under Map*
 proper and improper, 1591
 combing rectangle, 1608
 conjugacy, 1581, 1586, *see also Glossary*
- derivative, 1593
 direction, 1593
- edgelet, 1587, *see also Glossary*
 of a foldable map, 1594
 expansion, 1589, *see also Glossary*
 properly discontinuous, 1589
- fold, 1597
 full, 1598
 improper, 1598
 proper, 1598
 type IA, 1598
 type IIIA, 1598
 fold factorization, 1602
 fold map, 1597, *see also Glossary under Map*
 fold path, 1601, *see also Glossary under Foldable sequence*
- fold sequence, 1600, *see also Glossary under Foldable sequence*
 equivalence, 1601
 foldable map, 1594, *see also Glossary under Map*
 foldable sequence, 1608, *see also Glossary*
 free splitting, 1581, 1585, *see also Glossary*
 free splitting unit, 1634
 front greedy subsequence, 1635
- gate, 1593
- map, 1585, *see also Glossary*
 collapse, 1581, 1587
 fold, 1597
 foldable, 1594
- natural cell structure, 1586, *see also Glossary*
 natural core, 1589
 normalization diagram, 1645
- projection diagram, 1609
 augmented, 1638
- W diagram, 1643
- zig-zag, 1592, *see also Glossary*

*Mathematics & Computer Science Department, Herbert H Lehman College (CUNY)
 Bronx, NY 10468-1589, USA*

*Department of Mathematics and Computer Science, Rutgers University, Newark
 Newark, NJ 07102, USA*

michael.handel@lehman.cuny.edu, mosher@rutgers.edu

Proposed: David Gabai
 Seconded: Danny Calegari, Walter Neumann

Received: 16 March 2012
 Revised: 16 January 2013