

Splitting Line Patterns in Free Groups

CHRISTOPHER H. CASHEN

We construct a boundary of a finite rank free group relative to a finite list of conjugacy classes of maximal cyclic subgroups. From the cut points and uncrossed cut pairs of this boundary we construct a simplicial tree on which the group acts cocompactly. We show that the quotient graph of groups is the JSJ decomposition of the group relative to the given collection of conjugacy classes.

This provides a characterization of virtually geometric multiwords: they are the multiwords that are built from geometric pieces. In particular, a multiword is virtually geometric if and only if the relative boundary is planar.

20F65; 20F67, 20E05, 57M05, 20E06

1 Introduction

Let $F = F_n$ be a free group of finite rank $n > 1$. Let $[w] = \{[\langle w_1 \rangle], \dots, [\langle w_k \rangle]\}$ be a *multiclass*, a non-empty collection of distinct conjugacy classes of maximal cyclic subgroups.

The goal of this paper is to find splittings of F relative to $[w]$ (rel $[w]$), that is, splittings of F as a free product or as an amalgam over cyclic subgroups in such a way that each $[\langle w_i \rangle] \in [w]$ is elliptic. We do this by analyzing the topology of a certain relative boundary \mathcal{D} of F , defined as follows:

The free group F has a well-defined Gromov boundary ∂F that is homeomorphic to a Cantor set. Left multiplication of F on itself extends continuously to an action of F on ∂F by homeomorphisms. For each non-trivial element $f \in F$, the f -action on ∂F has an attracting fixed point f^∞ and a repelling fixed point, which is the attracting fixed point of $\bar{f} = f^{-1}$, and is denoted \bar{f}^∞ .

Definition 1.1 Define the *boundary pattern associated to $[w]$* to be:

$$\partial[w] = \{\{w^\infty, \bar{w}^\infty\} \mid \langle w \rangle \text{ is a maximal cyclic subgroup with } [\langle w \rangle] \in [w]\}$$

Definition 1.2 The *decomposition space* $\mathcal{D} = \mathcal{D}_{[w]}$ of ∂F associated to $[w]$ is the quotient of ∂F obtained by identifying the two points ξ_0 and ξ_1 for each $\{\xi_0, \xi_1\} \in \partial[w]$. Let $\Delta: \partial F \rightarrow \mathcal{D}$ be the quotient map.

The action of F on ∂F preserves the boundary pattern $\partial[w]$, so it induces an action of F on \mathcal{D} by homeomorphisms.

It is not hard to see that if F has a free splitting rel $[w]$ then \mathcal{D} is not connected. Similarly, if F splits over $\langle f \rangle$ rel $[w]$, where $\langle f \rangle$ is a maximal cyclic subgroup, then $\mathcal{D} \setminus \Delta(\partial\langle f \rangle)$ is not connected. If $[\langle f \rangle] \in [w]$ then $\Delta(\partial\langle f \rangle)$ is a single point in \mathcal{D} , so there is a point whose removal disconnects \mathcal{D} . If f is non-trivial and $[\langle f \rangle] \notin [w]$ then $\Delta(f^\infty)$ and $\Delta(f^{-\infty})$ are distinct points in \mathcal{D} , so there is a pair of points whose removal disconnects \mathcal{D} .

We will focus on the case that F does not split freely rel $[w]$, in which case we will see that \mathcal{D} is connected. The previous paragraph then suggests that we analyze cut points and cut pairs of \mathcal{D} .

We show that the cut points and cut pairs of \mathcal{D} encode a simplicial tree on which F acts cocompactly. The quotient graph of groups gives us a canonical decomposition of F rel $[w]$. The main result, the Relative JSJ-Decomposition Theorem (Theorem 4.25), is that this canonical graph of groups decomposition of F obtained from \mathcal{D} is the *JSJ decomposition of F relative to $[w]$* (the rJSJ). That is, it is the decomposition that encodes all cyclic splittings of F relative to $[w]$, and satisfies certain universality and maximality properties, in the sense of Guirardel and Levitt [9] (see also [12]).

In Section 2 we introduce some preliminaries, including various versions of Whitehead graphs.

In Section 3 we use generalized Whitehead graphs to investigate topological features of \mathcal{D} . We regard this section as semi-preliminary, as it is a development of ideas that were present in [4]. However, some of the proofs are technical, so they are presented here in detail in the interests of rigor and of making this paper self-contained.

In Section 4 we construct a simplicial tree from the cut points and cut pairs of \mathcal{D} and show that the quotient graph is the rJSJ.

We apply these results in Section 5 to characterize virtually geometric multiclasses, which will be introduced in Section 1.2.

A benefit of our approach to relative splittings via the decomposition space and generalized Whitehead graphs is that the arguments end up being combinatorial. It follows that not only is the rJSJ algorithmically constructible, which was already known

by work of Kharlampovich and Miasnikov [12], but there is a combinatorial algorithm that is actually implementable. In subsequent work with Manning, we have extended the methods of this paper to get an implementable algorithm to construct the rJSJ, and we have written a computer program [5] (see also [6]) that will compute the rJSJ and decide whether or not a given multiclass is virtually geometric.

1.0.1 First Examples

We give two examples to give an idea of what decomposition spaces and relative JSJ decompositions can look like. These examples are of a very special type: the free group F can be viewed as the fundamental group of a compact, connected, orientable surface with boundary, and the multiclass includes the conjugacy class of each of the boundary curves of the surface. In this case there are no relative free splittings, and we can ‘see’ the relative cyclic splittings—they correspond to essential, non-peripheral simple closed curves¹ in the surface that can be homotoped to be disjoint from a multicurve representing the multiclass.

Much of the work in Section 3 is about how to ‘see’ relative cyclic splittings when no ambient surface topology is available.

Example 1.3 Consider $F = \langle a, b \rangle$ and $[w] = \{[\langle ab\bar{a}\bar{b} \rangle]\}$. The decomposition space is homeomorphic to the circle. To see this, view F as the fundamental group of a complete, finite volume hyperbolic punctured torus Σ , and represent $ab\bar{a}\bar{b}$ as a simple closed curve running around the puncture. The universal cover of Σ is the hyperbolic plane \mathbb{H}^2 , see Figure 1. The group F acts by deck transformations on \mathbb{H}^2 , and the




Figure 1: Horocycles in \mathbb{H}^2 demonstrating $\mathcal{D} = \partial\mathbb{H}^2 = S^1$.

¹A curve is *essential* if it not homotopic to a curve that bounds a disc and *non-peripheral* if it not homotopic to a curve that bounds a once-punctured disc or an annulus.

action extends to a continuous surjection $\partial F \rightarrow \partial\mathbb{H}^2$ that is 2 to 1 on the parabolic points and 1 to 1 off them. The curve representing $ab\bar{a}\bar{b}$ is freely homotopic to the quotient of a horocycle. The element $ab\bar{a}\bar{b}$ acts parabolically, fixing a point $\xi \in \partial\mathbb{H}^2$, and the preimage of ξ in ∂F is exactly the two points $(ab\bar{a}\bar{b})^\infty$ and $(ab\bar{a}\bar{b})^{-\infty}$ of ∂F . Since there is only one orbit of parabolic points, we conclude $\mathcal{D} = \partial\mathbb{H}^2 = \mathbb{S}^1$.

Alternatively, we could take Σ' to be a hyperbolic one-holed torus with geodesic boundary component representing $ab\bar{a}\bar{b}$. The universal cover sits inside of \mathbb{H}^2 as a thickened tree, and ∂F embeds into $\partial\mathbb{H}^2$. The points $(ab\bar{a}\bar{b})^\infty$ and $(ab\bar{a}\bar{b})^{-\infty}$ are sent to the endpoints of an interval in $\partial\mathbb{H}^2$ not containing any other points of ∂F , so the quotient map $\Delta: \partial F \rightarrow \mathcal{D} = \mathbb{S}^1$ can be viewed as a circular analogue of the map collapsing missing intervals of the ternary Cantor set to get the unit interval.

In this example the rJSJ is trivial — a single vertex stabilized by F . This is due to the universality requirement, see Theorem 4.25. The reason is that cyclic splittings of F rel $\{[\langle ab\bar{a}\bar{b} \rangle]\}$ correspond to essential, non-peripheral simple closed curves in Σ' , but for every such curve there is another intersecting it, so none of these splittings are universal. \diamond

Example 1.4 Consider the marked surface Σ in Figure 2. The labeled curves are generators of the fundamental group $F_5 = \langle a, b, c, d, e \rangle$.




Figure 2: A marked surface




Figure 3: A multi-curve representing $[w]$

Consider the multiclass:

$$[w] = \{[\langle a \rangle], [\langle c \rangle], [\langle d \rangle], [\langle e \rangle], [\langle \bar{a}c \rangle], [\langle d\bar{c}\bar{d} \rangle], [\langle d\bar{c}\bar{c}ab\bar{a}\bar{b}\bar{e} \rangle]\}$$

Figure 3 shows a *multicurve* representing $[w]$.

In this example, all boundary curves of the surface belong to the multicurve, so \mathcal{D} is connected as in Example 1.3.




Figure 4: Decomposition of a surface.




Figure 5: Corresponding graph of groups

Consider the decomposition of Σ into subsurfaces given in Figure 4. The corresponding graph of groups in Figure 5 is the rJSJ for this example. In Figure 5 the vertices are labelled with their stabilizer subgroups. The edge stabilizers are all infinite cyclic, and the label at each end of each edge indicates the image of a fixed generator of the edge stabilizer in the vertex group.

This is the rJSJ because every essential, non-peripheral simple closed curve in Σ that does not intersect a curve of the multicurve is either homotopic to an essential, non-peripheral simple closed curve in the subsurface Σ_2 or to the core curve of annulus A_3 or annulus A_4 .

Let us examine the four subsurfaces:

- (1) Subsurface Σ_1 is ‘filled’ by the multicurve, in the sense that every essential, non-peripheral simple closed curve in Σ_1 intersects one of the curves of the multicurve. This means that Σ_1 does not contribute any relative splittings of F . The fundamental group of Σ_1 is an example of what will be called a *rigid* vertex group.
- (2) The subsurface Σ_2 is a sphere with four holes. It is ‘empty’, in the sense that there are no curves of the multicurve in its interior. Thus, any essential, non-peripheral simple closed curve in Σ_2 yields a cyclic splitting of F rel $[\underline{w}]$. However, as in the previous example, these splittings are not universal. The fundamental group of Σ_2 is an example of what will be called a *QH-surface* vertex group².
- (3) The annulus A_3 has a core curve representing $[\langle a \rangle] \in [\underline{w}]$, so $\Delta(a^\infty) = \Delta(\bar{a}^\infty)$ is a cut point of \mathcal{D} .
- (4) The core curve of annulus A_4 is not homotopic to a curve of the multicurve. The points $\Delta((dc\bar{d}\bar{c}a)^\infty)$ and $\Delta((dc\bar{d}\bar{c}a)^{-\infty})$ are distinct points in \mathcal{D} . They form a cut pair in \mathcal{D} .

²These are the torsion-free examples of the ‘quadratically hanging’ vertex groups of Rips and Sela [18].

Observe the following features of this decomposition, and compare Theorem 4.25. The decomposition is bipartite: there is a collection of annuli and a collection of more complicated subsurfaces, and each subsurface is adjacent only to members of the opposite collection. Among the more complicated subsurfaces there are those that are filled (the rigid vertices) and those that are empty (the QH-surface vertices). Every splitting of F rel $[w]$ comes from either the core curve of one of the annuli or an essential, non-peripheral simple closed curve in one of the empty subsurface pieces. \diamond

1.1 The Decomposition Space and the rJSJ

We are interested in cut sets of \mathcal{D} . If \mathcal{D} is connected and $\mathcal{D} \setminus \{x\}$ is not connected, then x is called a *cut point*. Similarly, if \mathcal{D} is connected and $\{x_0, x_1\}$ is a pair of points, neither of which is a cut point, such that $\mathcal{D} \setminus \{x_0, x_1\}$ is not connected, then $\{x_0, x_1\}$ is called a *cut pair*. We call $(F, [w])$ *rigid* if \mathcal{D} is connected with no cut points and no cut pairs³. We call $(F, [w])$ a *QH-surface* if there exists a compact surface with boundary, Σ , such that $F = \pi_1(\Sigma)$ and $[w] = [\partial\Sigma]$. In this case we write $(F, [w]) \sim (\Sigma, [\partial\Sigma])$.

1.1.1 Induced Multiclasses

Definition 1.5 Let G be a non-cyclic vertex group of a graph of groups decomposition Γ of F rel $[w]$ with cyclic edge stabilizers. Define the *induced multiclass* in G , denoted $\text{Ind}_\Gamma^G([w])$, to be the set of distinct G -conjugacy classes of G -maximal cyclic subgroups that either contain the image of an edge injection into G or are contained in an F -maximal cyclic subgroup whose conjugacy class is in $[w]$.

In Example 1.4, the induced multiclass in $\pi_1(\Sigma_1)$ is:

$$\{[\langle a \rangle], [\langle c \rangle], [\langle d \rangle], [\langle \bar{a}c \rangle], [\langle d\bar{c}\bar{d} \rangle], [\langle d\bar{c}\bar{d}a \rangle]\}$$

These classes come from the four curves of the multicurve in the interior of Σ_1 , plus the two boundary curves, one of which was a member of the multicurve, and one of which was not.

Similarly, $\text{Ind}_\Gamma^{\pi_1(\Sigma_2)} = \{[\langle bab \rangle], [\langle e \rangle], [\langle d\bar{c}\bar{d}cab\bar{a}\bar{b}\bar{e} \rangle], [\langle d\bar{c}\bar{d}a \rangle]\}$. All of these classes correspond to boundary curves of Σ_2 , three of which were members of the multicurve, and one of which was not.

³The fact that $(F, [w])$ is quasi-isometrically rigid if and only if \mathcal{D} is connected with no cut points and no cut pairs is the main result of [4]. Since we will not be concerned with quasi-isometric rigidity in this paper, we take these condition on \mathcal{D} to be the definition of rigidity.

The vertex corresponding to G is said to be *rigid* or *QH* if $(G, \text{Ind}_1^G([\underline{w}]))$ is rigid or is a QH-surface, respectively. It will turn out that all non-cyclic vertices of the rJSJ are either rigid or QH.

1.1.2 Outline of the Construction of the rJSJ

Suppose that F does not split freely relative to $[\underline{w}]$, which is equivalent to supposing that \mathcal{D} is connected, by Theorem 2.17. Suppose that $\{x_0, x_1\}$ and $\{y_0, y_1\}$ are cut pairs such that y_0 and y_1 lie in different complementary components of $\{x_0, x_1\}$. In this case, we say $\{y_0, y_1\}$ *crosses* $\{x_0, x_1\}$. If there does not exist a cut pair crossing $\{x_0, x_1\}$ then we say $\{x_0, x_1\}$ is *uncrossed*.

Proposition 4.6 generalizes a construction of Bowditch to show that if F does not split freely rel $[\underline{w}]$, and if $(F, [\underline{w}])$ is neither rigid nor a QH-surface, then \mathcal{D} contains a cut point or an uncrossed cut pair.

Proposition 4.10 says that an uncrossed cut pair is rational⁴, that is, it is stabilized by an infinite cyclic subgroup of F . Moreover, there are only finitely many conjugacy classes of stabilizers of uncrossed cut pairs.

We show in Proposition 4.13 that the collection of cut points and uncrossed cut pairs in \mathcal{D} has the structure of a simplicial tree. Otal proved this for the cut points. We generalize his proof to work simultaneously with the cut points and uncrossed cut pairs. Since F acts by homeomorphism on \mathcal{D} , and since there are finitely many conjugacy classes of stabilizers of cut points and uncrossed cut pairs, we get a cocompact F -action on this cut point/uncrossed cut pair tree. Theorem 4.25 says the rJSJ is the quotient graph of groups of this action.

A consequence of Proposition 4.13 is that every uncrossed cut pair of \mathcal{D} corresponds to a cyclic splitting of F rel $[\underline{w}]$.

The main work is proving Proposition 4.6, Proposition 4.10, and Proposition 4.13. Verifying that the resulting graph of groups satisfies the desired properties of the rJSJ is routine.

⁴ ∂F can be thought of as the set of infinite, freely reduced words in the generators of F and their inverses. The points whose expressions as such are eventually periodic are commonly called ‘rational’, in analogy to decimal representations of rational numbers. Being a rational point of ∂F is equivalent to being fixed by an infinite cyclic subgroup. We extend the terminology to call a pair of points *rational* if they are fixed by an infinite cyclic subgroup.

1.1.3 Otal, Bowditch, and Cut Points

Otal [17] showed that the cut points of \mathcal{D} have a simplicial tree structure. Bowditch [3] proved that the cut points of the boundary of a relatively hyperbolic group have a simplicial tree structure.

It can be shown, see Manning [14], that \mathcal{D} is equivariantly homeomorphic to the Bowditch boundary of F relative to $[w]$, so Otal's result is an early special case of Bowditch's result. For our purposes these results are not sufficient: The cut point tree does not see all the universal relative cyclic splittings, because it misses the ones coming from uncrossed cut pairs. Our methods treat cut points and uncrossed cut pairs in a unified way.

1.2 Virtual Geometricity

In Section 5 we apply the Relative JSJ-Decomposition Theorem to characterize virtual geometricity.

$[w] \subset F$ is *geometric* if it can be represented by an embedded multicurve in the boundary of a handlebody with fundamental group F .

Otal's main result in [17], suitably reinterpreted, is that in the case that the rJSJ is trivial, $[w]$ is geometric if and only if the corresponding decomposition space is planar. Furthermore, planarity of the decomposition space can be deduced from the Whitehead graph of $[w]$.

$[w]$ is *virtually geometric* if there is a finite index subgroup G of F such that the 'lift' of $[w]$ to G is geometric. The *lift of $[w]$ to G* , for $[w] = \{[\langle w_1 \rangle], \dots, [\langle w_k \rangle]\}$, is the multiclass of G that contains every conjugacy class of maximal cyclic subgroup of G that is conjugate in F into one of the $\langle w_i \rangle$, see Section 2.3.

We use the rJSJ to reduce virtual geometricity to geometricity of the induced multiclass in the vertex groups:

Characterization of Virtual Geometricity (Theorem 5.9) *For a multiclass in a free group, the following are equivalent:*

- (1) *The multiclass is virtually geometric.*
- (2) *The decomposition space is planar.*
- (3) *For every non-cyclic vertex group of the rJSJ, the induced multiclass is geometric.*

Thus, virtually geometric multiclasses are those that are built from geometric pieces.

When there are no uncrossed cut pairs we use the rJSJ to explicitly construct a finite index subgroup and handlebody that demonstrate virtual geometricity.

When there are uncrossed cut pairs we first pinch them to cut points and then apply the previous construction. The pinching is done in such a way as to preserve planarity of the decomposition space, using a technical fact, Proposition 3.24, that the closure of a complementary component of a cut pair in \mathcal{D} is arc-connected.

1.3 Acknowledgements

I thank Jason Manning, who noticed that the decomposition space of [4] was the same as the space considered by Otal, and brought the question of virtual geometricity to my attention. I also thank him for helpful comments on an earlier version of this paper.

I thank the anonymous referee for their diligence.

This work was partially supported by the Agence Nationale de la Recherche (ANR) grant ANR-2010-BLAN-116-01 GGAA, the European Research Council (ERC) grant of Goulnara ARZHANTSEVA, grant agreement #259527, and the Austrian Science Fund (FWF):M1717-N25

2 Preliminaries

2.1 Definitions and Notation

The *degree* of a homomorphism from the integers into a free group is the index of its image in the maximal cyclic subgroup containing the image.

A nontrivial element $g \in F$ is *indivisible* if it is not a proper power.

Let \underline{b} be a basis of F . The Cayley graph of F with respect to \underline{b} is a tree \mathcal{T} . Assign each edge length one; F acts isometrically on \mathcal{T} by left multiplication. We will use $\mathbf{1}$ to denote the vertex corresponding to the identity element of F .

The tree \mathcal{T} has a Gromov boundary at infinity $\partial\mathcal{T}$ that is identified with ∂F . This boundary compactifies the tree: $\overline{\mathcal{T}} = \mathcal{T} \cup \partial\mathcal{T}$ is a compact topological space whose topology on \mathcal{T} agrees with the metric topology. For $x, y \in \overline{\mathcal{T}}$, there exists a unique geodesic $[x, y]$ connecting them. By a *simplicial geodesic* we shall mean an isometric

embedding $\phi: [a, b] \hookrightarrow \overline{\mathcal{T}}$ taking integers to vertices, with $a, b \in \mathbb{Z} \cup \{\pm\infty\}$. We use the notation $\phi: [a, b] \rightarrow [x, y]$ to indicate a simplicial geodesic with $\phi(a) = x$ and $\phi(b) = y$. There is unique such simplicial geodesic if a or b is finite.

For a fixed vertex u , a basis for the topology of $\partial\mathcal{T}$ is given by the sets

$$\text{Shadow}^u(v) = \{\xi \in \partial\mathcal{T} \mid v \in [u, \xi]\} \quad \text{for } v \in \mathcal{T} \setminus \{u\}.$$

The resulting topology does not depend on the choice of u .

2.2 Cut Pairs

Recall that a minimal cut set is a subset $Y \subset X$ such that $X \setminus Y$ is not connected but $X \setminus Z$ is connected for every proper subset Z of Y .

Lemma 2.1 *Let Y be a closed minimal cut set of connected, locally connected space X . Every complementary component limits to every point of Y .*

Proof $X \setminus Y$ is locally connected, so components are proper, non-empty clopens. If $y \in Y$ is not a limit point of a component C then C is still a proper, non-empty clopen in $X \setminus (Y \setminus y)$, contradicting minimality of Y . \square

Definition 2.2 If Y is a cut set of X then Z *crosses* Y if there are points $z_0, z_1 \in Z$ in different components of $X \setminus Y$.

Recall that a cut pair is a minimal cut set of size two, and a cut pair is said to be uncrossed if no cut pair crosses it. The following lemma is easily verified:

Lemma 2.3 *Let X be a connected, locally connected space in which cut points and cut pairs have finitely many complementary components.*

- (1) *Crossing is a symmetric relation among cut pairs.*
- (2) *A cut pair cannot cross a cut point.*
- (3) *A cut pair with at least three complementary components is uncrossed.*


2.3 A Multiclass Lifted to a Finite Index Subgroup

Definition 2.4 The *lift* of $[w]$ to a finite index subgroup G is the multiclass $[w]_G$ consisting of the distinct G -conjugacy classes of G -maximal cyclic subgroups that are contained in an F -maximal cyclic subgroup whose conjugacy class is in $[w]$.

The inclusion $\iota: G \hookrightarrow F$ extends to a homeomorphism $\partial\iota: \partial G \rightarrow \partial F$ that takes $\partial[w]_G$ to $\partial[w]$, inducing a G -equivariant homeomorphism $\mathcal{D}_{[w]_G} \rightarrow \mathcal{D}_{[w]}$.

Example 2.5 Let $F = \langle a, b \rangle$. Let $[w] = \{[\langle a \rangle], [\langle b \rangle]\}$. Let G be the index 2 subgroup $G = \langle a^2, b, ab\bar{a} \rangle < F$. The lift of $[w]$ to G is $[w]_G = \{[\langle a^2 \rangle], [\langle b \rangle], [\langle ab\bar{a} \rangle]\}$.

$$\begin{aligned} [w]_G &= \{[\langle a^2 \rangle], [\langle b \rangle], [\langle ab\bar{a} \rangle]\} \\ G &= \pi_1(\tilde{H}) = \langle a^2, b, ab\bar{a} \rangle \end{aligned}$$



$$\begin{aligned} [w] &= \{[\langle a \rangle], [\langle b \rangle]\} \\ F &= \pi_1(H) = \langle a, b \rangle \end{aligned}$$




Figure 6: Lifting to a finite index subgroup

We can visualize the situation by taking F to be the fundamental group of a handlebody H and picking curves representing $[\langle a \rangle]$ and $[\langle b \rangle]$. Then G corresponds to a 2-fold cover \tilde{H} of H , and $[w]_G$ is represented by the curves in \tilde{H} covering the chosen curves in H , as in Figure 6. \diamond

2.4 Normalization

See Serre [20] for an introduction to graphs of groups and Bass-Serre theory.

Suppose Γ is a graph of groups decomposition of F with cyclic edge groups. It will be convenient to normalize Γ . We describe a sequence of moves that change the graph of groups description without changing the group itself or the conjugacy classes of non-cyclic vertex groups. If e is an edge of Γ let $\eta(e, 0)$ and $\eta(e, 1)$ denote the initial and terminal vertices of e , respectively. Let $\phi_{e,i}: G_e \hookrightarrow G_{\eta(e,i)}$ be the edge injection of an edge group into a vertex group.

First, if there is an edge incident to two non-cyclic vertex groups, subdivide it by adding a vertex with stabilizer equal to the stabilizer of the edge group.

Second, for each edge e let $G_e = \langle z_e \rangle$ and let $\langle z_{e,i} \rangle$ be the maximal cyclic subgroup of $G_{\eta(e,i)}$ containing $\phi_{e,i}(G_e)$. Since no nontrivial element is conjugate to a power of itself in the free group, it is possible to choose the generators z_e and $z_{e,i}$ so that for all e and i the map $\phi_{e,i}$ takes z_e to a positive power of $z_{e,i}$ and so that if $\bar{g} \langle z_{e,i} \rangle g = \langle z_{e',i'} \rangle$ for some $g \in G_{\eta(e,i)} = G_{\eta(e',i')}$ then $\bar{g} z_{e,i} g = z_{e',i'}$.

Third, if an edge group maps into a non-maximal cyclic subgroup of a non-cyclic vertex group, we may un-collapse an edge as in Figure 7.

$$G \xrightarrow{g^p} \dots \longrightarrow G \xrightarrow{g} z \langle z \rangle \xrightarrow{z^p} \dots$$

Figure 7: Un-collapse an edge.

Fourth, consider two edges e and e' incident to a non-cyclic vertex $v = \eta(e, i) = \eta(e', i')$. Suppose $\phi_{e,i}(G_e)$ and $\phi_{e',i'}(G_{e'})$ are distinct and conjugate in G_v . Choose $g \in G_v$ such that $\bar{g} \phi_{e,i}(G_e) g = \phi_{e',i'}(G_{e'})$. Replace the edge map $\phi_{e,i}$ with $\bar{g} \phi_{e,i} g$.

Fifth, fold all the edges together that map into a common maximal cyclic subgroup in a given non-cyclic vertex, as in Figure 8.

$$\dots \xrightarrow{y^p} \langle y \rangle \xrightarrow{y} \dots \xrightarrow{g} G \xrightarrow{g} z^q \langle z \rangle \longrightarrow G \xrightarrow{g} z^q \langle z \rangle \xrightarrow{z^{pq}} \dots$$

Figure 8: Fold two edges.

Folding is always possible when two edges map into a common maximal cyclic subgroup of a non-cyclic vertex because we are in the free group. Consider the possible obstructions:

- If $\langle y \rangle = \langle z \rangle$ then the two edges of Figure 8 form a loop corresponding to a stable letter t conjugating y to $y^{\pm q}$. This would mean $\langle t, y \rangle$ is a Baumslag-Solitar subgroup, but free groups do not contain such subgroups.
- If $\langle y \rangle \neq \langle z \rangle$ we could imagine the situation depicted in Figure 9 with $r > 1$ and $q > 1$. In this case, $\langle y, z \rangle$ is a virtually free-by-cyclic subgroup, but free groups do not contain such subgroups.

$$\langle y \rangle \xrightarrow{y^r} \dots \xrightarrow{g} G \xrightarrow{g} z^q \langle z \rangle$$

Figure 9: Cannot occur in a free group.

Similarly, since F is free, every edge group maps onto a maximal cyclic subgroup in one of its two vertex groups. Otherwise we would find a Baumslag-Solitar or virtually free-by-cyclic subgroup. Therefore, for any edge that is incident to two cyclic vertices, one of the inclusions of the edge group into the vertex groups is an isomorphism, and we can collapse the edge, as in Figure 10.

$$\cdots \xrightarrow{y^p} \langle y \rangle^y \xrightarrow{z^q} \langle z \rangle \longrightarrow \cdots \xrightarrow{z^{pq}} \langle z \rangle$$

Figure 10: Collapse an edge.

We are left with a new graph of groups decomposition of F that is bipartite: vertex groups are either maximal cyclic subgroups or are non-cyclic. Cyclic vertex groups are adjacent only to non-cyclic vertex groups, and vice versa. For each edge, the edge map to the incident non-cyclic vertex group maps the edge group onto a maximal cyclic subgroup of the vertex group. Furthermore, for each non-cyclic vertex group G and each maximal cyclic subgroup C of G there is at most one incident edge whose edge group maps into a conjugate of C .

Remark Another normalization that is often applied to a graph of groups is to make them *reduced*. This means that if there is a non-loop edge e incident to a vertex $\eta(e, i)$ such that the edge inclusion $\phi_{e,i}$ is an isomorphism, then the edge e should be collapsed. We do not assume that Γ is reduced, because in some cases doing so would ruin the ‘bipartite’ condition.

2.5 Whitehead Graphs

Our tool for understanding the topology of the decomposition space associated to a multiclass is the generalized Whitehead graph of the multiclass. This machinery was developed in [4].

2.5.1 Classical Whitehead Graph

Let $[\langle w \rangle]$ be a conjugacy class of maximal cyclic subgroups of F . The (classical) Whitehead graph $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[\langle w \rangle]\}$ of $[\langle w \rangle]$ with respect to a basis \underline{b} of F is a graph with $2n$ vertices labeled with the elements of \underline{b} and their inverses. Let w be a freely and cyclically reduced word in \underline{b}^{\pm} that generates a representative of $[\langle w \rangle]$. One edge of $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[\langle w \rangle]\}$ joins vertex x to vertex y for each occurrence of $\bar{x}y$ in w , thought of as a cyclic word. This definition extends to a multiclass by adding edges for each class of the multiclass. We will see a geometric interpretation and examples in Section 2.5.3.

Let $||[\langle w \rangle]||_{\underline{b}}$ be the minimal \underline{b} -length of a generator of a representative of $[\langle w \rangle]$. The *complexity* of the Whitehead graph is the number of edges, which is equal to $\sum_{[\langle w \rangle] \in [\underline{w}]} ||[\langle w \rangle]||_{\underline{b}}$. A Whitehead graph $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[\underline{w}]\}$ is *minimal* if its complexity is minimal among the complexities of $\mathfrak{W}_{\phi^{-1}(\underline{b})}(\mathbf{1})\{[\underline{w}]\}$ for $\phi \in \text{Aut}(F)$.

Whitehead's Algorithm [22] picks a basis \underline{b} for which $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[\underline{w}]\}$ is minimal. The proof shows that there is a finite set of 'Whitehead automorphisms' so that if $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[\underline{w}]\}$ is not minimal then there exists a Whitehead automorphism ϕ that strictly reduces the complexity. The algorithm checks if any Whitehead automorphism reduces the complexity, and repeats this process until no reducing Whitehead automorphism exists. An important observation in the proof is that if a Whitehead graph $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[\underline{w}]\}$ is connected and has a cut vertex, then it is not minimal.

An easy extension of Whitehead's methods yields the following:

Proposition 2.6 *The following are equivalent:*

- (1) *Some Whitehead graph for $[\underline{w}]$ is not connected.*
- (2) *Every minimal Whitehead graph for $[\underline{w}]$ is not connected.*
- (3) *F splits freely rel $[\underline{w}]$.*

It is easy to see that if F splits freely rel $[\underline{w}]$ then \mathcal{D} is not connected:

Corollary 2.7 *If there is a basis \underline{b} such that $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[\underline{w}]\}$ is not connected, then \mathcal{D} is not connected.*

The converse is also true, see Theorem 2.17.

2.5.2 Standing Assumption

From now on, unless otherwise noted, we assume that $[\underline{w}] = \{[\langle w_1 \rangle], \dots, [\langle w_k \rangle]\}$ is fixed and \underline{b} is a basis of F such that $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[\underline{w}]\}$ is connected without cut vertices. Let \mathcal{T} denote the Cayley tree of F with respect to \underline{b} .

Having fixed a reference basis, we simplify notation by considering the *multiword* $\underline{w} = \{w_1, \dots, w_k\}$, where the w_i are cyclically reduced and generate non-conjugate maximal cyclic subgroups. Similarly, by choosing representatives we pass from the induced multiclass in a vertex group of a splitting to an induced multiword, and from a lifted multiclass in a finite index subgroup to a lifted multiword.

We drop $[\underline{w}]$ and \underline{b} from the notation unless they are necessary for clarity.

Remark There is no loss of generality from these assumptions. If F splits freely relative to $[\underline{w}]$ then first pass to a maximal relative free splitting and then deal with the factors separately. If there is no such free splitting then Proposition 2.6 says that to ensure the no-cut-vertex assumption it suffices to choose the basis that gives the minimal complexity Whitehead graph. However, minimality is not necessary.

2.5.3 Generalized Whitehead Graph and Friends

For each $\{(\bar{f}\bar{w}if)^\infty, (\bar{f}wif)^\infty\} \in \partial \underline{w}$ there is a unique bi-infinite geodesic with endpoints $(\bar{f}\bar{w}if)^\infty$ and $(\bar{f}wif)^\infty$. The stabilizer $\langle \bar{f}wif \rangle$ acts cocompactly with translation length $|w_i|$.

Definition 2.8 The *line pattern generated by \underline{w}* is the set \mathbf{L} ($=\mathbf{L}_{\underline{w}}$) of ‘lines’, bi-infinite geodesics \mathcal{L} in $\bar{\mathcal{T}}$, with endpoints $\mathcal{L}^-, \mathcal{L}^+ \in \partial \mathcal{T}$ such that $\{\mathcal{L}^-, \mathcal{L}^+\} \in \partial \underline{w}$.

Definition 2.9 Let \mathcal{X} be a connected subset of $\bar{\mathcal{T}}$ with $\mathcal{X} \cap \mathcal{T} \neq \emptyset$. The *Whitehead graph over \mathcal{X}* , denoted $\mathfrak{W}(\mathcal{X})$ ($=\mathfrak{W}_{\underline{b}}(\mathcal{X})\{\underline{w}\}$) is a graph with one vertex for each component of $\bar{\mathcal{T}} \setminus \bar{\mathcal{X}}$ and one edge \mathfrak{E} joining vertices \mathfrak{V} and \mathfrak{V}' for each $\mathcal{L} \in \mathbf{L}$ with one endpoint in \mathfrak{V} and the other in \mathfrak{V}' .

Additionally, each vertex and edge carries a piece of data that records whence it came: Let $\mathcal{T}_{\mathfrak{V}}$ denote the component of $\bar{\mathcal{T}} \setminus \bar{\mathcal{X}}$ corresponding to vertex \mathfrak{V} , and let $\mathcal{L}_{\mathfrak{E}}$ be the line in \mathbf{L} corresponding to edge \mathfrak{E} .

Recall that $\mathbf{1}$ denotes the vertex of \mathcal{T} corresponding to the identity element of F , so the notation $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{\underline{w}\}$ of the previous section refers to the Whitehead graph over the vertex $\mathbf{1}$.

We have partially defined functions from $\bar{\mathcal{T}}$ and \mathbf{L} to $\mathfrak{W}(\mathcal{X})$:

Definition 2.10 If $x \in \bar{\mathcal{T}} \setminus \bar{\mathcal{X}}$ define $\Theta_{\mathcal{X}}(x)$ to be the vertex of $\mathfrak{W}(\mathcal{X})$ corresponding to the component of $\bar{\mathcal{T}} \setminus \bar{\mathcal{X}}$ containing x .

If $\mathcal{L} \in \mathbf{L}$ such that $\mathcal{L} \cap \mathcal{X} \neq \emptyset$ and $\partial \mathcal{L} \cap \partial \bar{\mathcal{X}} = \emptyset$, define $\Theta_{\mathcal{X}}(\mathcal{L})$ to be the edge of $\mathfrak{W}(\mathcal{X})$ contributed by \mathcal{L} .

We shorten the notation to Θ when \mathcal{X} is apparent.

The point of this extra data is to consider Whitehead graphs not just as abstract graphs, but as pictures of \mathbf{L} in \mathcal{T} .

Figure 11 depicts the line pattern $\mathbf{L}_{\{ab, ab\bar{a}\bar{b}\}}$ in $F = \langle a, b \rangle$. It also shows closeups of the lines that pass through the vertex b and through the edge $[1, a]$. The Whitehead graphs $\mathfrak{W}(b)\{ab, ab\bar{a}\bar{b}\}$ and $\mathfrak{W}([1, a])\{ab, ab\bar{a}\bar{b}\}$ in Figure 12 show the corresponding Whitehead graphs.

Notice that the Whitehead graph over \mathcal{X} looks like the closeup of the line pattern passing through \mathcal{X} , with lines exiting \mathcal{X} through a common edge of \mathcal{T} pinched to a vertex. It will sometimes be convenient *not* to do this pinching, but at the same time remember the incidence of the edges. For this we introduce the following formalism:




Figure 11: Line pattern $\mathbf{L}_{\{ab, ab\bar{a}\bar{b}\}}$ with closeups.

Definition 2.11 A graph with loose ends at v_1, \dots, v_k is a graph Γ with a specified subset of vertices $\{v_1, \dots, v_k\}$ that have been marked ‘deleted’. An edge e of Γ incident to a deleted v_i is said to *have a loose end at v_i* .

A *component* of a graph with loose ends is an equivalence class of edges and undeleted vertices given by the incidence relation.

Definition 2.12 For connected sets $\mathcal{X} \subset \mathcal{Y} \subset \bar{\mathcal{T}}$ with $\mathcal{X} \cap \mathcal{T} \neq \emptyset$, let $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$ denote $\mathfrak{W}(\mathcal{X})$ with loose ends at each vertex $\mathfrak{V} \in \mathfrak{W}(\mathcal{X})$ such that $\mathcal{T}_{\mathfrak{V}} \cap \mathcal{Y} \neq \emptyset$.

In other words, to construct $\mathfrak{W}(\mathcal{X})$ we look at the lines of \mathbf{L} passing through \mathcal{X} and pinch off a vertex for each edge e of $\mathcal{T} \setminus \mathcal{X}$ incident to \mathcal{X} . For $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$, we do the same, except that we do not pinch a vertex if $e \in \mathcal{Y}$. Instead we will have loose ends at a deleted vertex corresponding to such an e . See Figure 13 for an example, and compare to Figure 12. The utility of this definition is that it will let us build up large Whitehead graphs from smaller pieces. The idea, which will be made precise in Section 2.5.5, is that if $\mathcal{Y} = \coprod_i \mathcal{X}_i$ is a connected subset of \mathcal{T} written as a disjoint union of connected pieces \mathcal{X}_i then we can build up $\mathfrak{W}(\mathcal{Y})$ from the various ‘graphs with loose ends’ $\mathfrak{W}(\mathcal{X}_i) \odot \mathcal{Y}$ by ‘splicing loose ends’.





Figure 12: Whitehead graphs


 Figure 13: Whitehead graph with loose ends $\mathfrak{W}([1, a])\{ab\bar{a}b, ab\} \odot [\bar{b}, ab]$

Definition 2.13 For connected sets $\mathcal{X} \subset \mathcal{Y} \subset \bar{\mathcal{T}}$ with $\mathcal{X} \cap \mathcal{T} \neq \emptyset$, let \mathfrak{C} be a component of $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$.

\mathfrak{C} has a loose end at \mathfrak{V} if it contains an edge with a loose end at \mathfrak{V} . \mathfrak{C} has an end at $\xi \in \partial\mathcal{T}$ if ξ is a limit point of $\bigcup_{\mathfrak{C} \in \mathfrak{C}} \mathcal{L}_{\mathfrak{C}}$.

Define $\partial\mathcal{T}_{\mathfrak{C}} = \coprod_{\mathfrak{V} \in \mathfrak{C}} \partial\mathcal{T}_{\mathfrak{V}}$.

2.5.4 Splicing

Manning [13] gave a construction for combining two graphs called *splicing*. Let Γ_0 and Γ_1 be graphs with vertices $\gamma_i \in \Gamma_i$ of the same valence. Let a *splice map* σ be a bijection between edges incident to γ_0 and edges incident to γ_1 . The result of splicing Γ_0 and Γ_1 at γ_0 and γ_1 by σ is defined to be a graph whose vertices are the union of vertices of Γ_0 and Γ_1 , minus γ_0 and γ_1 . Edges not incident to γ_0 or γ_1 are retained. If $e_0 = [u, \gamma_0]$ and $e_1 = \sigma(e_0) = [\gamma_1, v]$, then add an edge $[u, v]$ in the new graph. In the above terminology, take the graph Γ_0 with loose ends at γ_0 and the graph Γ_1 with loose ends at γ_1 and ‘tie up’ the loose ends by matching them using the given splice map σ .

2.5.5 Cutting Whitehead Graphs into Pieces and Splicing them Together

The following lemma is the motivation for the splicing construction:

Lemma 2.14 For connected sets $\mathcal{X} \subset \mathcal{Y} \subset \overline{\mathcal{T}}$ with $\mathcal{X} \cap \mathcal{T} \neq \emptyset$, suppose $\mathcal{X}_0 \amalg \mathcal{X}_1 = \mathcal{X} \setminus e$ for some edge e of \mathcal{X} in \mathcal{T} . Then $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$ is obtained from $\mathfrak{W}(\mathcal{X}_0) \odot \mathcal{Y}$ and $\mathfrak{W}(\mathcal{X}_1) \odot \mathcal{Y}$ by discarding any edges \mathfrak{E} with $\partial \mathcal{L}_{\mathfrak{E}} \cap \overline{\mathcal{X}} \neq \emptyset$ and splicing remaining loose ends at $\Theta_{\mathcal{X}_0}(e)$ in $\mathfrak{W}(\mathcal{X}_0) \odot \mathcal{Y}$ to loose ends at $\Theta_{\mathcal{X}_1}(e)$ in $\mathfrak{W}(\mathcal{X}_1) \odot \mathcal{Y}$.

In this situation we say $\mathfrak{W}(\mathcal{X}_i) \odot \mathcal{Y}$ includes into $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$ via splicing. See Figure 14.

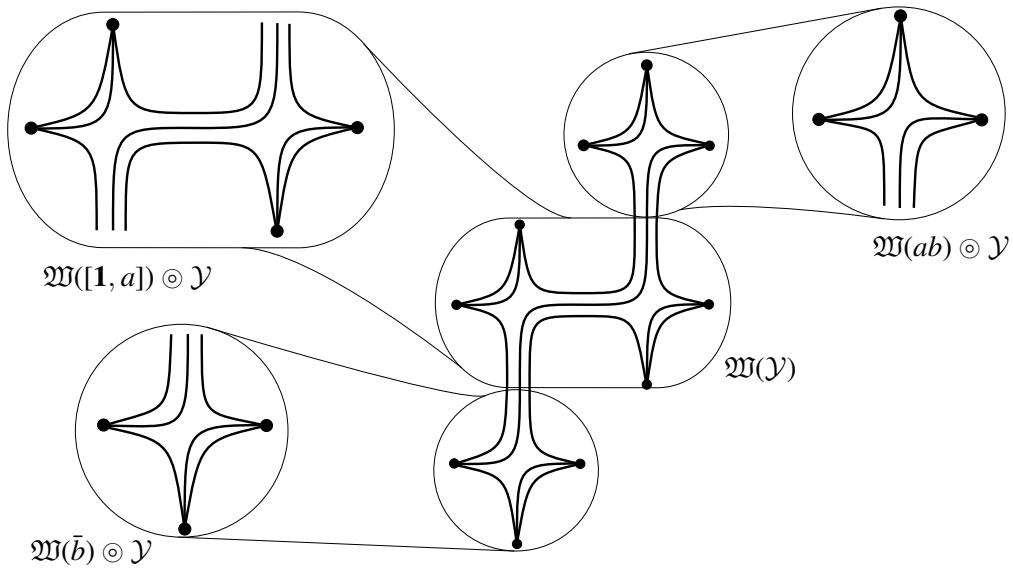


Figure 14: Inclusion via splicing for $F = \langle a, b \rangle$, $\underline{w} = \{ab\bar{a}\bar{b}, ab\}$, $\mathcal{Y} = [\bar{b}, ab]$

Manning states a similar result in the case that $\mathcal{X} = \mathcal{Y}$ is bounded. The proof is immediate from the definitions. The essential observation is that there is a natural splicing map defined by sending an edge $\mathfrak{E} \in \mathfrak{W}(\mathcal{X}_0) \odot \mathcal{Y}$ with a loose end at $\Theta_{\mathcal{X}_0}(e)$ to the unique edge $\mathfrak{E}' \in \mathfrak{W}(\mathcal{X}_1) \odot \mathcal{Y}$ with a loose end at $\Theta_{\mathcal{X}_1}(e)$ such that $\mathcal{L}_{\mathfrak{E}} = \mathcal{L}_{\mathfrak{E}'}$.

Manning notes that splicing two graphs that are connected without cut vertices yields a graph that is connected without cut vertices. $\mathfrak{W}(\mathbf{1})$ is connected without cut vertices, by our standing assumption, so, by induction on the number of vertices of \mathcal{X} :

Lemma 2.15 ([13]) *If \mathcal{X} is a bounded connected subset of \mathcal{T} then $\mathfrak{W}(\mathcal{X})$ is connected without cut vertices.*

2.6 A First Connectivity Lemma

Recall that $\Delta: \partial F \rightarrow \mathcal{D}$ is the quotient map.

Lemma 2.16 ([17, Prop 2.1][15, Theorem 49]) *Let e be an edge in \mathcal{T} with endpoints u and v . Let $\mathcal{A} = \text{Shadow}^u(v)$. The set $A = \Delta(\mathcal{A})$ is connected in \mathcal{D} .*

We include the proof for completeness:

Proof Suppose $B, C \subset \mathcal{D}$ are open with $A \subset B \cup C$ and $A \cap B \cap C = \emptyset$. Let $\mathcal{B} = \mathcal{A} \cap \Delta^{-1}(B)$ and $\mathcal{C} = \mathcal{A} \cap \Delta^{-1}(C)$. Compactness of $\partial\mathcal{T}$ implies \mathcal{B} and \mathcal{C} are compact clopens. Since \mathcal{B} is compact and open, there are finitely many vertices x_1, \dots, x_a so that $\mathcal{B} = \bigcup_{i=1}^a \text{Shadow}^u(x_i)$. Assume x_1, \dots, x_a contains as few points as possible, ie, the shadows of a proper subcollection of the x_i have union a proper subset of \mathcal{B} . There is a similar finite collection y_1, \dots, y_b that determines \mathcal{C} .

Assume \mathcal{B} is non-empty. If $x_1 = v$ then $\mathcal{B} = \mathcal{A}$, so $\mathcal{C} = \emptyset$, and we are done. Otherwise, consider the convex hull \mathcal{H} of $\{x_i\}_{i=1}^a \cup \{y_j\}_{j=1}^b \cup \{u\}$; it is a finite tree with leaves $\{x_i\}_{i=1}^a \cup \{y_j\}_{j=1}^b \cup \{u\}$. Let $\mathcal{X} = \mathcal{H} \setminus (\{x_i\}_{i=1}^a \cup \{y_j\}_{j=1}^b \cup \{u\})$. By Lemma 2.15, $\mathfrak{W}(\mathcal{X})$ is connected without cut vertices.

Consider the vertex $\Theta(x_1) \in \mathfrak{W}(\mathcal{X})$. Since $\Theta(u)$ is not a cut vertex, there are edges of $\mathfrak{W}(\mathcal{X})$ incident to $\Theta(x_1)$ and not $\Theta(u)$. Such an edge corresponds to a line $\mathcal{L} \in \mathbf{L}$ with one endpoint in the shadow of x_1 and the other endpoint in the shadow of z for some $z \in \{x_i\}_{i=2}^a \cup \{y_j\}_{j=1}^b$. In the decomposition space these two endpoints are identified, so they must both be in \mathcal{B} , hence $z \in \{x_i\}_{i=2}^a$. Since $\mathfrak{W}(\mathcal{X})$ is connected and $\Theta(u)$ is not a cut vertex, we conclude that all the vertices of $\mathfrak{W}(\mathcal{X})$ except $\Theta(u)$ belong to $\{\Theta(x_1), \dots, \Theta(x_n)\}$, so $\mathcal{C} = \emptyset$. Thus, A is connected. \square

Together with Proposition 2.6, this implies:

Theorem 2.17 *Let $[w]$ be an arbitrary multiclass (not necessarily satisfying the standing assumption of Section 2.5.2). F admits a free splitting relative to $[w]$ if and only if $\mathcal{D}_{[w]}$ is not connected.*

Proof Choose a basis \underline{b} for F such that $[w]$ is Whitehead minimal. By Proposition 2.6, F splits freely rel $[w]$ if and only if $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[w]\}$ is not connected. If $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[w]\}$ is not connected then \mathcal{D} is not connected by Corollary 2.7.

Suppose $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[w]\}$ is connected. Pick an edge $e \in \mathcal{T}$ with vertices u and v . Let $A = \text{Shadow}^u(v) \subset \partial\mathcal{T}$ and $A' = \text{Shadow}^v(u) \subset \partial\mathcal{T}$. Note $\partial\mathcal{T} = A \amalg A'$. Whitehead minimality implies that $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[w]\}$ has no cut vertex, so Lemma 2.16 applies. Moreover, $\Delta(A)$ and $\Delta(A')$ have a point in common, since $\mathfrak{W}_{\underline{b}}(\mathbf{1})\{[w]\}$ connected implies there is some line in \mathbf{L} crossing e . Thus, $\mathcal{D} = \Delta(A) \cup \Delta(A')$ is a union of connected sets with non-empty intersection, so it is connected. \square

3 The Topology of the Decomposition Space

In this section we prove some results for later use. We begin, in Section 3.1 by recalling Moore's Decomposition Theorem, which will be needed in Section 5. This theorem requires a technical hypothesis called 'upper semi-continuity'.

In Section 3.2 we explicitly construct a neighborhood basis for \mathcal{D} . It follows quickly that the decomposition we consider is upper semi-continuous and that \mathcal{D} is metrizable. Furthermore, when \mathcal{D} is connected it is locally connected and arc connected.

In the two remaining subsections we prove technical results that will be used in Section 4 to show that when the rJSJ is non-trivial there do exist cut points or uncrossed cut pairs of \mathcal{D} , and that uncrossed cut pairs are rational.

The purpose of Section 3.3 is to make rigorous the idea that we can 'see' a neighborhood of $x \in \mathcal{D}$ by looking at components of a Whitehead graph over a ray in \mathcal{T} tending towards a point of $\Delta^{-1}(x)$ in ∂F . Precisely, given a point $x \in \mathcal{D}$ and a neighborhood N' of x there exists a connected open neighborhood $N \subset N'$ of x such that for any simplicial geodesic $\phi: [0, \infty] \rightarrow \overline{\mathcal{T}}$ with $\phi(\infty) \in \Delta^{-1}(x)$ one of the following holds:

- $\Delta^{-1}(x) \notin \partial \underline{w}$ and there is a bijection between the components of $N \setminus \{x\}$ and components of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$.
- $\Delta^{-1}(x) \in \partial \underline{w}$ and there is a bijection between the half of the components of $N \setminus \{x\}$ and components of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$.

In the second case we only see half the components of $N \setminus \{x\}$ because $|\Delta^{-1}(x)| = 2$. The other half of the components correspond to components of a Whitehead graph over a ray tending towards the other point in $\Delta^{-1}(x) \in \partial F$.

There are two main results of Section 3.4. One, Proposition 3.25, is a refinement of a rational approximation result from [4]. The other, Lemma 3.32, says that if \mathcal{H} is the convex hull in $\overline{\mathcal{T}}$ of the preimage of a cut pair of \mathcal{D} then at every point x of \mathcal{H} the Whitehead graph $\mathfrak{W}(x) \odot \mathcal{H}$ has at least two components.

3.1 Aside on General Decomposition Spaces

We recall some results about more general decomposition spaces.

Let X be a topological space. Let $X = \coprod_{i \in I} X_i$ be a decomposition of X as a disjoint union. The *non-degenerate* elements are the non-singleton X_i . The *decomposition space* is the index set I with the quotient topology induced by $X_i \mapsto i$.

Definition 3.1 A decomposition $X = \coprod_{i \in I} X_i$ is *upper semi-continuous* if for each $i \in I$ the set X_i is compact and for each open set $U \supset X_i$ there is an open set $V \supset X_i$ such that for all $j \in I$ if $X_j \cap V$ is non-empty then $X_j \subset U$.

Lemma 3.2 ([11, Theorem 3-33]) *The decomposition space of an upper semi-continuous decomposition of a compact Hausdorff space is Hausdorff.*

The following theorem will be used in Section 5.

Theorem 3.3 (Moore’s Decomposition Theorem [16]) *An upper semi-continuous decomposition of the 2-sphere into connected, non-separating sets has decomposition space homeomorphic to the 2-sphere.*

3.2 Basic Topology of the Decomposition Space

We now return our attention to the decomposition space \mathcal{D} of ∂F associated to \underline{w} .

Let $\mathbf{L}|_{\mathcal{X}} = \{\mathcal{L} \in \mathbf{L} \mid \mathcal{L} \cap \mathcal{X} \neq \emptyset\}$. If \mathcal{X} is bounded, $\mathbf{L}|_{\mathcal{X}}$ is finite.

3.2.1 Neighborhood Basis

The goal of this section is to describe a neighborhood basis of \mathcal{D} . A set \mathcal{A} in $\partial \mathcal{T}$ is *saturated* if $\mathcal{A} = \Delta^{-1}(\Delta(\mathcal{A}))$. The Δ -image of an open saturated set is open. First we build open saturated neighborhoods of points $\xi \in \partial \mathcal{T}$. The construction proceeds in stages. The idea is to start with a basic neighborhood of ξ in $\partial \mathcal{T}$. This neighborhood might not be saturated, so we add some points to make it saturated. The resulting set might not be open, so add basic neighborhoods of all of the newly added points, and repeat.

Let M be the ‘maximum overlap’ in \mathbf{L} : the maximum length of $\mathcal{L} \cap \mathcal{L}'$ for distinct lines \mathcal{L} and \mathcal{L}' in \mathbf{L} . This is finite, since \underline{w} consists of finitely many words.

Let ξ_0 be a point in $\partial \mathcal{T}$, and let $r \in \mathbb{N}$.

Let $\phi_0: [0, \infty) \rightarrow [\mathbf{1}, \xi_0]$ be the simplicial geodesic. If ξ_0 is the endpoint of a line $[\xi_0, \xi_1] \in \mathbf{L}$ let $r_0 = \max\{r, 1 + d_{\mathcal{T}}(\mathbf{1}, [\xi_0, \xi_1])\}$. Otherwise, there is at most one line of \mathbf{L} with one endpoint in $\text{Shadow}^1(\phi_0(r + M + 1))$ and the other in the complement of $\text{Shadow}^1(\phi_0(r))$. If there are none, let $r_0 = r$. If there is one, \mathcal{L} , let $r_0 > r$ be the minimal number such that $\mathcal{L} \cap [\mathbf{1}, \xi_0] \subset \phi_0([0, r_0 + M])$. Then no line of \mathbf{L}

has one endpoint in $\text{Shadow}^1(\phi_0(r_0 + M + 1))$ and the other in the complement of $\text{Shadow}^1(\phi_0(r_0))$, since such a line would overlap \mathcal{L} by more than M .

Consider two neighborhoods of ξ_0 , an ‘inner neighborhood’ $\mathcal{S}_0 = \text{Shadow}^1(\phi_0(r_0 + M + 1))$ contained in an ‘outer neighborhood’ $\mathcal{O}_0 = \text{Shadow}^1(\phi_0(r_0))$. Let \mathbf{L}'_0 be the lines of \mathbf{L} with exactly one endpoint in \mathcal{S}_0 and both endpoints in \mathcal{O}_0 .

If ξ_0 does not belong to an element of $\partial_{\underline{W}}$ then \mathbf{L}'_0 accounts for all lines with exactly one endpoint in \mathcal{S}_0 and we define ‘stage 0’ to be $\{\mathcal{S}_0\}$. Otherwise $\mathcal{L} = [\xi_0, \xi_1]$ is the unique line with one endpoint in \mathcal{S}_0 and one outside of \mathcal{O}_0 . Let $\phi_1: [0, \infty] \rightarrow [\mathbf{1}, \xi_1]$ be the simplicial geodesic. Let $\mathcal{S}_1 = \text{Shadow}^1(\phi_1(r_0 + M + 1))$ and $\mathcal{O}_1 = \text{Shadow}^1(\phi_1(r_0))$. Let $\mathbf{L}'_1 = (\mathbf{L}|_{\phi_1(r_1+M)} \cap \mathbf{L}|_{\phi_1(r_1+M+1)}) \setminus \{\mathcal{L}\}$. Define ‘stage 0’ to be $\{\mathcal{S}_0, \mathcal{S}_1\}$. Note in this case $\Delta(\mathcal{L}^-) = \Delta(\mathcal{L}^+) \in \Delta(\mathcal{S}_0) \cap \Delta(\mathcal{S}_1)$.

Suppose now we have constructed stage k . Let $r_{k+1} = r_k + M + 1$. Let I be the indices of the stage k sets. For $i \in I$ and $\mathcal{L}_{i,j} \in \mathbf{L}'_i$, let $\phi_{i,j}: [0, \infty] \rightarrow [\mathbf{1}, \xi_{i,j}]$ be the simplicial geodesic, where $\xi_{i,j}$ is the endpoint of $\mathcal{L}_{i,j}$ that is not in \mathcal{S}_i . Define an inner neighborhood $\mathcal{S}_{i,j} = \text{Shadow}^1(\phi_{i,j}(r_{k+1} + M + 1))$ and outer neighborhood $\mathcal{O}_{i,j} = \text{Shadow}^1(\phi_{i,j}(r_{k+1}))$ of $\xi_{i,j}$.

For fixed i and all j we have $\mathcal{O}_{i,j} \subset \mathcal{O}_i$ and $\mathcal{O}_{i,j}$ is disjoint from \mathcal{S}_i and from every $\mathcal{O}_{i,j'}$ for $j' \neq j$.

Let $\mathbf{L}'_{i,j}$ be the lines of \mathbf{L} with exactly one endpoint in $\mathcal{S}_{i,j}$ and with both endpoints in $\mathcal{O}_{i,j}$. We know that there is one line of \mathbf{L} with an endpoint in $\mathcal{S}_{i,j}$ and the other not in $\mathcal{O}_{i,j}$ — it is $\mathcal{L}_{i,j}$, and its other endpoint is in \mathcal{S}_i . There can be no other such line, for it would overlap $\mathcal{L}_{i,j}$ by more than M .

Call \mathcal{S}_i the *predecessor* of $\mathcal{S}_{i,j}$. Let ‘stage $k + 1$ ’ be the sets $\mathcal{S}_{i,j}$ produced from the sets of stage k . By construction, the stage $k + 1$ sets are disjoint from each other and from all sets in the previous stages. This is true for their Δ images in \mathcal{D} as well, except that $\Delta(\mathcal{S}_{i,j}) \cap \Delta(\mathcal{S}_i) = \Delta(\mathcal{L}_{i,j}^-) = \Delta(\mathcal{L}_{i,j}^+)$.

Repeat this construction for as many stages as possible, potentially infinitely many, to produce a collection $\{\mathcal{S}_\tau\}_\tau$ of disjoint open sets. Let $\mathcal{S}(r) = \cup_\tau \mathcal{S}_\tau$. By construction, $\mathcal{S}(r)$ is an open saturated neighborhood of ξ_0 , so its image $S_{\xi_0}(r) = \Delta(\mathcal{S}(r))$ is an open neighborhood of $\Delta(\xi_0)$.

Definition 3.4 For $x \in \mathcal{D}$ pick $\xi_0 \in \Delta^{-1}(x)$ and define $N(x, r) = S_{\xi_0}(r)$.

By construction, $N(x, r)$ does not depend on the choice of $\xi_0 \in \Delta^{-1}(x)$.

It is immediate from the construction that:

Proposition 3.5 $\{N(x, r) \mid r \in \mathbb{N}\}$ is a neighborhood basis for x .

3.2.2 Upper Semi-continuity

Proposition 3.6 *The decomposition of $\partial\mathcal{T}$ whose non-degenerate elements are ∂_w is an upper semi-continuous decomposition.*

Proof This follows directly from Section 3.2.1 since the sets S_{ξ_0} are saturated. \square

Corollary 3.7 *\mathcal{D} is Hausdorff.*

Proposition 3.8 *\mathcal{D} is metrizable.*

Proof $\Delta: \partial\mathcal{T} \rightarrow \mathcal{D}$ is a continuous map from a compact space to a Hausdorff space, by Corollary 3.7, so it is a closed map. The codomain is first-countable by Proposition 3.5. The domain is metrizable, so a theorem of Stone [21, Theorem 1] says \mathcal{D} is metrizable. \square

3.2.3 Connectivity

Lemma 3.9 *For every $x \in \mathcal{D}$ and $r \in \mathbb{N}$, the set $N(x, r)$ is connected.*

Proof Lemma 2.16 says that for each of the sets S_τ in the construction of $N(x, r)$, the set $\Delta(S_\tau)$ is connected in \mathcal{D} . If there are two stage 0 sets their Δ -images have a point in common, and the Δ -image of every set from a higher stage has a point in common with the Δ -image of its predecessor. Thus $N(x, r)$ is connected. \square

Proposition 3.10 (cf [17, Proposition 2.1]) *\mathcal{D} is connected and locally connected.*

Proof It is locally connected by Lemma 3.9.

Let u and v be neighboring vertices in \mathcal{T} . Let $\mathcal{A} = \text{Shadow}^u(v)$. Then $\mathcal{A}^c = \partial\mathcal{T} \setminus \mathcal{A} = \text{Shadow}^v(u)$. By Lemma 2.16, $\Delta(\mathcal{A})$ and $\Delta(\mathcal{A}^c)$ are each connected. $\mathfrak{W}(\mathbf{1})$ is connected, so there exists $\mathcal{L} \in \mathbf{L}|_u \cap \mathbf{L}|_v$, and $\Delta(\mathcal{L}^-) = \Delta(\mathcal{L}^+) \in \Delta(\mathcal{A}) \cap \Delta(\mathcal{A}^c)$. \square

An *arc* is an embedded path. A space is *arc-connected* if any two points can be joined by an arc. A space is *Peano* if it is compact, Hausdorff, connected, locally connected, and metrizable.⁵

⁵Alternatively, a space is Peano if it is a continuous image of the unit interval. These definitions are equivalent for Hausdorff spaces by the Hahn-Mazurkiewicz Theorem.

Proposition 3.11 \mathcal{D} is Peano.

Proof \mathcal{D} is compact. It is Hausdorff by Corollary 3.7, connected and locally connected by Proposition 3.10, and metrizable by Proposition 3.8. \square

Theorem 3.12 ([23, Theorem 31.2]) *A Peano space is arc-connected.*

Corollary 3.13 \mathcal{D} is arc-connected.

3.3 Whitehead Graphs and Cut Sets of the Decomposition Space

Lemma 3.14 *If $x \in \mathcal{D}$ such that $|\Delta^{-1}(x)| = 1$ then x is not a cut point.*

Proof Let $\xi = \Delta^{-1}(x)$. Let $\phi: [0, \infty) \rightarrow [\mathbf{1}, \xi]$ be the simplicial geodesic. For all $i \in \mathbb{N}$ the set $\Delta(\text{Shadow}^{\phi(i)}(\phi(i-1)))$ is connected, by Lemma 2.16. $\mathcal{D} \setminus \{x\}$ is an increasing union of such sets, so it is connected. \square

Proposition 3.15 (cf [4, Lemma 4.9]) *Let $S \subset \mathcal{D}$ be closed with $|\Delta^{-1}(S)| > 1$, and let \mathcal{H} be the convex hull of $\Delta^{-1}(S)$. Then $\mathcal{C} \mapsto \Delta(\partial\mathcal{T}_{\mathcal{C}})$ is a bijection between components of $\mathfrak{W}(\mathcal{H})$ and components of $\mathcal{D} \setminus S$.*

Proof Let \mathcal{C} be a component of $\mathfrak{W}(\mathcal{H})$. Let \mathfrak{V} be a vertex of \mathcal{C} . By Lemma 2.16, $\Delta(\partial\mathcal{T}_{\mathfrak{V}})$ is connected in \mathcal{D} . If vertices \mathfrak{V} and \mathfrak{V}' are joined by an edge \mathcal{E} in \mathcal{C} then, by definition, $\mathcal{L}_{\mathcal{E}} \in \mathbf{L}$ is a line with one endpoint in $\partial\mathcal{T}_{\mathfrak{V}}$ and the other in $\partial\mathcal{T}_{\mathfrak{V}'}$, so $\Delta(\partial\mathcal{T}_{\mathfrak{V}})$ and $\Delta(\partial\mathcal{T}_{\mathfrak{V}'})$ have a point in common. This implies $\Delta(\partial\mathcal{T}_{\mathcal{C}})$ is connected. It is also open, since $\partial\mathcal{T}_{\mathcal{C}}$ is open and saturated.

Since S is closed, for every $\xi \in \partial\mathcal{T} \setminus \Delta^{-1}(S)$ there is a vertex $v \in \mathcal{H}$ such that $[v, \xi] \cap \mathcal{H} = \{v\}$. Therefore, $\Theta_{\mathcal{H}}(\xi)$ is a vertex of $\mathfrak{W}(\mathcal{H})$ with $\xi \in \partial\mathcal{T}_{\Theta_{\mathcal{H}}(\xi)}$.

Letting \mathcal{C} range over all components of $\mathfrak{W}(\mathcal{H})$, we get disjoint, connected, open subsets of the form $\Delta(\partial\mathcal{T}_{\mathcal{C}})$, whose union is all of $\mathcal{D} \setminus S$. \square

Lemma 3.16 *Let $S \subset \mathcal{D}$ be a closed minimal cut set, and let \mathcal{H} be the convex hull of $\Delta^{-1}(S)$. For all $x \in S$, each component \mathcal{C} of $\mathfrak{W}(\mathcal{H})$ has an end at a point in $\Delta^{-1}(x)$.*

Proof \mathcal{C} has an end at a point in $\Delta^{-1}(x)$ if and only if $\Delta(\partial\mathcal{T}_{\mathcal{C}})$ has x as a limit point, which it must, by Lemma 2.1. \square

Definition 3.17 Let $\phi: [0, l] \rightarrow \overline{\mathcal{T}}$ be a simplicial geodesic for some $l \in \mathbb{N} \setminus \{1\}$. An edge path $\mathfrak{E}_0, \dots, \mathfrak{E}_k$ in $\mathfrak{W}(\phi([1, l-1])) \odot \phi([0, l])$ is *non-backtracking* if the path in \mathcal{T} obtained by concatenating the segments $\mathcal{L}_{\mathfrak{E}_i} \cap \phi([0, l])$ does not backtrack.

Lemma 3.18 For $l \in \mathbb{N} \setminus \{1\}$, let $\phi: [0, l] \rightarrow \mathcal{T}$ be a simplicial geodesic. Let \mathfrak{E} be an edge of $\mathfrak{W}(\phi([1, l-1])) \odot \phi([0, l])$. There exists a non-backtracking edge path $\mathfrak{E} = \mathfrak{E}_0, \dots, \mathfrak{E}_k$ with $\mathcal{L}_{\mathfrak{E}_k} \in \mathbf{L}|_{\phi(0)}$.

Proof Let r be the smallest integer such that $\phi(r) \in \mathcal{L}_{\mathfrak{E}_0}$. If $r \leq 0$ we are done. Suppose not. By minimality of r , the edge $\Theta_{\phi(r)}(\mathcal{L}_{\mathfrak{E}_0})$ is incident to an undeleted vertex $\mathfrak{V} \in \mathfrak{W}(\phi(r)) \odot \phi([0, l])$.

$\mathfrak{W}(\phi(r))$ is connected without cut vertices, so the vertex $\Theta_{\phi(r)}(\phi(r+1))$ is not a cut vertex. Thus, there is an edge path $\mathfrak{E}'_1, \dots, \mathfrak{E}'_j$ in $\mathfrak{W}(\phi(r))$ connecting \mathfrak{V} to $\Theta_{\phi(r)}(\phi(r-1))$ that does not go through the vertex $\Theta_{\phi(r)}(\phi(r+1))$. Choosing the shortest such path guarantees that $\mathcal{L}_{\mathfrak{E}'_i} \cap \phi([0, l]) = \phi(r)$ for all $i < j$, and $\mathcal{L}_{\mathfrak{E}'_j} \cap \phi([0, l]) = \phi([r', r])$ for some $r' < r$. Extend the existing edge path by edges $\mathfrak{E}_i = \Theta_{\phi([1, l-1])}(\mathcal{L}_{\mathfrak{E}'_i})$. This gives a non-backtracking edge path beginning with \mathfrak{E} that reaches closer to $\phi(0)$. Proceed by induction. \square

The number of lines in $\mathbf{L}|_e$ for an edge $e \in \mathcal{T}$ corresponding to basis element b is equal to the valence of the b -vertex in $\mathfrak{W}(\mathbf{1})$.

Corollary 3.19 The maximum valence of a vertex in $\mathfrak{W}(\mathbf{1})$ is an upper bound for the number of components of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$.

Proposition 3.20 $\mathcal{D} \setminus S$ has finitely many components for every finite $S \subset \mathcal{D}$.

Proof Let $\{\xi_1, \dots, \xi_k\} = \Delta^{-1}(S)$, and suppose $\mathcal{D} \setminus S$ is not connected. Then $k \geq 2$, by Lemma 3.14. Let \mathcal{H} be the convex hull of $\Delta^{-1}(S)$. If $k = 2$, let \mathcal{X} be an arbitrary vertex in \mathcal{H} ; otherwise, let \mathcal{X} be the convex hull of the branch points of \mathcal{H} . For $1 \leq i \leq k$, let $\phi_i: [0, \infty] \hookrightarrow \overline{\mathcal{T}}$ be the simplicial geodesic ray with $\phi_i([0, \infty]) \cap \mathcal{X} = \phi(0)$ and $\phi(\infty) = \xi_i$. By Lemma 2.14, $\mathfrak{W}(\mathcal{H})$ is obtained by splicing the $\mathfrak{W}(\phi_i([1, \infty])) \odot \mathcal{H}$ to $\mathfrak{W}(\mathcal{X}) \odot \mathcal{H}$. The former have finitely many components by Corollary 3.19, and the latter is finite, so $\mathfrak{W}(\mathcal{H})$ has finitely many components. By Proposition 3.15, $\mathcal{D} \setminus S$ has finitely many components. \square

Lemma 3.21 Let $\phi: [-\infty, \infty] \rightarrow [\overline{g}^\infty, g^\infty]$ be a simplicial geodesic for some cyclically reduced $g \in F$. Let $\mathcal{L}, \mathcal{L}' \in \mathbf{L}|_{\phi((0,1))}$. Then $\Theta_{\phi([-\infty, \infty])}(\mathcal{L})$ and $\Theta_{\phi([-\infty, \infty])}(\mathcal{L}')$ are in the same component of $\mathfrak{W}(\phi([-\infty, \infty]))$ if and only if $\Theta_{\phi([1, \infty])}(\mathcal{L})$ and $\Theta_{\phi([1, \infty])}(\mathcal{L}')$ are in the same component of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$.

Proof $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$ includes into $\mathfrak{W}(\phi([-\infty, \infty]))$ via splicing, so the ‘if’ direction is clear.

For the converse, consider the finitely many lines in $\mathbf{L}|_{\phi((0,1))}$. For every pair \mathcal{L}_i and \mathcal{L}_j that contribute edges to a common component of $\mathfrak{W}(\phi([-\infty, \infty]))$, choose an edge path $\mathfrak{P}_{i,j}$ connecting them.

Define a g -action on $\mathfrak{W}(\phi([-\infty, \infty]))$ by sending edge \mathfrak{E} to edge $\Theta_{\phi([-\infty, \infty])}(g\mathcal{L}_{\mathfrak{E}})$. The g -action permutes components, and by replacing g with a suitable power we may assume the permutation is trivial. Let m be sufficiently large so that for all i and j and all $\mathfrak{E} \in \mathfrak{P}_{i,j}$, we have $g^m\mathcal{L}_{\mathfrak{E}} \cap \phi([-\infty, \infty]) \subset \phi([1, \infty])$.

By Lemma 3.18, there is a non-backtracking edge path connecting $\Theta_{\phi([1, \infty])}(\mathcal{L})$ to an edge \mathfrak{E} such that $\mathcal{L}_{\mathfrak{E}} \in \mathbf{L}|_{g^m\phi(1)}$, and a similarly defined edge \mathfrak{E}' for \mathcal{L}' . Since the g -action is trivial on components, $\Theta_{\phi([-\infty, \infty])}(\mathcal{L}_{\mathfrak{E}})$ and $\Theta_{\phi([-\infty, \infty])}(\mathcal{L}_{\mathfrak{E}'})$ are in the same component of $\mathfrak{W}(\phi([-\infty, \infty]))$. Therefore, they are connected in $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$ by one of the $g^m\mathfrak{P}_{i,j}$. Concatenating these three paths connects $\Theta_{\phi([1, \infty])}(\mathcal{L})$ and $\Theta_{\phi([1, \infty])}(\mathcal{L}')$ in $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$. \square

Lemma 3.22 *For every $x \in \mathcal{D}$ and every neighborhood N' of x and every simplicial geodesic $\phi: [0, \infty] \rightarrow \overline{\mathcal{T}}$ with $\Delta(\phi(\infty)) = x$ there exists a connected open neighborhood $N \subset N'$ of x such that:*

$$\#(N \setminus \{x\}) = |\Delta^{-1}(x)| \lim_{r \rightarrow \infty} \#(\mathfrak{W}(\phi([r, \infty])) \odot \phi([0, \infty]))$$

Here, $|\cdot|$ denotes cardinality and $\#(\cdot)$ denotes number of components.

Note that x is a cut point in \mathcal{D} if and only if:

$$\lim_{r \rightarrow \infty} \#(\mathfrak{W}(\phi([r, \infty])) \odot \phi([0, \infty])) > 1 \text{ and } |\Delta^{-1}(x)| = 2$$

Proof As r increases, $\#(\mathfrak{W}(\phi([r, \infty])) \odot \phi([0, \infty]))$ is a non-decreasing sequence of integers that by Corollary 3.19 is bounded above. Thus, for any sufficiently large r_1 we may assume that the limit has been achieved and that $\Delta(\text{Shadow}^{\phi(0)}(\phi(r_1))) \subset N'$. Let $\mathfrak{C}_1, \dots, \mathfrak{C}_k$ be the components. For each i , let $\mathcal{L}_{i,1}, \dots, \mathcal{L}_{i,m_i}$ be the lines of \mathbf{L} that cross the edge $\phi((r_1 - 1, r_1))$ and satisfy $\Theta_{\phi([r_1, \infty])}(\mathcal{L}_{i,j}) \in \mathfrak{C}_i$.

For each $\mathcal{L}_{i,j}$ take a connected open neighborhood $N_{i,j}$ of $\Delta(\mathcal{L}_{i,j}^+)$ as in Section 3.2.1. We can take these neighborhoods small enough so that they are disjoint and contained in N' . For each i define $N_i = \Delta(\partial\mathcal{T}_{\mathfrak{C}_i}) \cup \bigcup_j N_{i,j}$. The N_i are disjoint, connected, open sets that all have x as a limit point, and x is the only limit point that any two have in common.

If $\Delta^{-1}(x) = \phi(\infty)$ then let $N = \{x\} \cup \bigcup_i N_i$ and we are done. Otherwise, let ϕ' be a simplicial geodesic ray converging to the other point in $\Delta^{-1}(x)$, and repeat the construction to produce connected open sets N'_1, \dots, N'_k . The number k of components is the same since there is a group element acting cocompactly on $[\phi'(\infty), \phi(\infty)]$. Let $N = \{x\} \cup \bigcup_i N_i \cup \bigcup_i N'_i$. \square

Proposition 3.23 *Let $\phi: [-\infty, \infty] \rightarrow [h\bar{w}^\infty, hw^\infty]$ be a simplicial geodesic for some $w \in \underline{w}$. Let $x_0 = \Delta(hw^\infty)$, and suppose x_0 is not a cut point.*

$$\lim_{r \rightarrow \infty} \#(\mathfrak{W}(\phi([r, \infty])) \odot \phi([0, \infty])) = \#(\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])) = 1$$

Furthermore, for every neighborhood N' of x_0 there exists an open connected neighborhood $N \subset N'$ of x_0 such that $N \setminus \{x_0\}$ has precisely two components.

Proof The first statement is true because $\langle w \rangle$ acts cocompactly on $\phi((-\infty, \infty))$ and x_0 is not a cut point. The second statement is an application of Lemma 3.22. \square

The next fact will be used in Section 5, but it is convenient to prove it now:

Proposition 3.24 *If A is the closure of a complementary component of a cut point or cut pair in \mathcal{D} , then A is arc-connected.*

Proof Let x be the cut point or one of the points of the cut pair. Let N' be a neighborhood of x in \mathcal{D} . Construct open connected sets $N_i \subset N'$ as in Lemma 3.22. Let $N = \{x\} \cup \bigcup_{N_i \subset A} N_i$. Then $N \subset A \cap N'$ is a connected set containing x that is open in A . Therefore, A is locally connected at x .

A is arc-connected just as in Corollary 3.13, since all of the other necessary properties are inherited from \mathcal{D} . \square

3.4 Cut Point and Cut Pair Detection and Approximation

Proposition 3.25 *If $x \in \mathcal{D}$ is a cut point then the stabilizer of x is conjugate to $\langle w \rangle$ for some $w \in \underline{w}$. In particular, there are finitely many orbits of cut points.*

Proof By Lemma 3.14, $\Delta^{-1}(x)$ consists of two points, so $\Delta^{-1}(x) = \{\bar{h}\bar{w}^\infty, \bar{h}w^\infty\}$ for some $h \in F$ and $w \in \underline{w}$. The stabilizer of x is therefore $\bar{h}\langle w \rangle h$. \square

Lemma 3.26 (Rational Approximation cf [4, Lemma 4.12]) *Let $\phi: [0, \infty] \rightarrow \bar{\mathcal{T}}$ be a simplicial geodesic. There exist elements $g, h \in F$ and $a \in \underline{b} \cup \bar{b}$ such that:*

- (1) the oriented edges $[h, ha]$ and $[gh, gha]$ belong to $\phi([0, \infty]) \cap [\bar{g}^\infty, g^\infty]$,
- (2) components of $\mathfrak{W}([ha, gh]) \odot [h, gha]$ that are in different components of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$ are in different components of $\mathfrak{W}([\bar{g}^\infty, g^\infty])$, and
- (3) for each line $\mathcal{L} \in \mathbf{L}|_h \cap \mathbf{L}|_{ha}$, the lines \mathcal{L} and $g\mathcal{L} \in \mathbf{L}|_{gh} \cap \mathbf{L}|_{gha}$ contribute edges to the same component of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$.

Moreover, g is conjugate to an element whose word length with respect to \underline{b} is bounded in terms of the rank of F and the maximum valence among vertices in $\mathfrak{W}(\mathbf{1})$, independent of ϕ .

Proof $n = \text{rank}(F)$. Let x be the maximum valence of $\mathfrak{W}(\mathbf{1})$. Let y be the x -th Bell number, the number of distinct partitions of x items into nonempty subsets. Let $z = 1 + (2n)^{y+2}$. Along any directed segment \mathcal{X} of $\phi([1, \infty])$ of length z there is some $a \in \underline{b} \cup \bar{\underline{b}}$ such that there are at least $y + 2$ many directed a -edges in the segment. Fix the first of these, $e = [g_0, g_0a]$.

Fix a numbering of the components of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$. The set $\mathbf{L}|_e$ of lines of \mathbf{L} that contain e is finite. Fix a numbering of them $1, \dots, k$. Of course, $k \leq x$, by Corollary 3.19. Partition them into subsets according to which component of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$ the corresponding edge belongs.

Consider an element $g' \in F$ such that the oriented edge $g'e$ is in \mathcal{X} . There is a bijection $\mathbf{L}|_e \rightarrow \mathbf{L}|_{g'e} : \mathcal{L} \mapsto g'\mathcal{L}$. Push forward the numbering of $\mathbf{L}|_e$ to $\mathbf{L}|_{g'e}$, and consider the partition of $1, \dots, k$ according to which component of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$ each line of $\mathbf{L}|_{g'e}$ belongs to. So g' gives a new partition of $1, \dots, k$.

There are at least $y + 1$ such elements g' , but at most y distinct partitions of the numbers $1, \dots, k$, so there exist g_1 and g_2 such that the oriented edges g_1e and g_2e are edges of $\phi([1, \infty])$ (say, with g_2e between g_1e and $\phi(\infty)$) and for each line $\mathcal{L} \in \mathbf{L}$ containing g_1e , the corresponding line $g_2\bar{g}_1\mathcal{L}$ containing g_2e is in the same component of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$.

The desired elements are $h = g_1g_0$ and $g = g_2\bar{g}_1$. The word length of $\bar{h}gh$ is the distance from g_1e to g_2e , which is at most z . \square

Corollary 3.27 *With notation as above, $\mathcal{D} \setminus \Delta(\{\bar{g}^\infty, g^\infty\})$ has at least as many components as $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$.*

Corollary 3.28 *If \mathcal{D} does not have any cut points or cut pairs with more than two complementary components, then for every simplicial geodesic $\phi: [0, \infty] \rightarrow \bar{\mathcal{T}}$ there are at most two components of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$.*

Lemma 3.29 *Let $\phi: [-\infty, \infty] \rightarrow \overline{\mathcal{T}}$ be a simplicial geodesic such that $\Delta(\phi(\infty))$ and $\Delta(\phi(-\infty))$ are distinct points and neither is a cut point. If, for some increasing sequence (r_i) of non-negative integers, there are at least two components of every $\mathfrak{W}(\phi([-r_i, r_i])) \odot \phi([-r_i - 1, r_i + 1])$, then $\Delta(\{\phi(-\infty), \phi(\infty)\})$ is a cut pair.*

Proof $\mathfrak{W}(\phi([-\infty, \infty]))$ has at least two components. Let \mathcal{H} be the convex hull of $\Delta^{-1}(\Delta(\{\phi(-\infty), \phi(\infty)\}))$. By Proposition 3.15 we must show $\mathfrak{W}(\mathcal{H})$ has at least two components. We are done if $\mathcal{H} = \phi([-\infty, \infty])$. Otherwise, we obtain $\mathfrak{W}(\mathcal{H})$ from $\mathfrak{W}(\phi([-\infty, \infty]))$ by deleting at most two vertices and splicing on connected graphs at the deleted vertices. Thus, $\mathfrak{W}(\mathcal{H})$ has the same number of components as $\mathfrak{W}(\phi([-\infty, \infty]))$. \square

Lemma 3.30 *Let \mathcal{H} be the convex hull of a subset of $\partial\mathcal{T}$, and let \mathcal{Y} be a connected subset of \mathcal{H} . Let \mathfrak{C} be a component of $\mathfrak{W}(\mathcal{Y}) \odot \mathcal{H}$, and let $\mathcal{X} = \mathcal{Y} \cap \bigcup_{\mathfrak{C} \in \mathfrak{C}} \mathcal{L}_{\mathfrak{C}}$.*

- (1) *For every leaf v of \mathcal{X} there exists an edge $\mathfrak{C} \in \mathfrak{C}$ with $v \in \mathcal{L}_{\mathfrak{C}}$ and an edge $e \in \mathcal{H} \setminus \mathcal{Y}$ incident to v such that \mathfrak{C} has a loose end at $\Theta_{\mathcal{Y}}(e)$.*
- (2) *The sum of the number of distinct ends and loose ends of \mathfrak{C} is at least two.*

Proof Let v be a leaf of \mathcal{X} . There is at most one edge e_0 of \mathcal{X} incident to v .

Let \mathfrak{C} be an edge of \mathfrak{C} such that $v \in \mathcal{L}_{\mathfrak{C}}$. Let $e_1 \neq e_0$ be an edge of $\mathcal{L}_{\mathfrak{C}}$ incident to v . Since v is a leaf of \mathcal{X} , e_1 is not an edge of \mathcal{Y} , so there exists a vertex $\Theta_{\mathcal{Y}}(e_1) \in \mathfrak{W}(\mathcal{Y})$. If $e_1 \in \mathcal{H} \setminus \mathcal{Y}$ then we are done: \mathfrak{C} has a loose end at $\Theta_{\mathcal{Y}}(e_1)$. If there are no such edges then, since $\mathfrak{W}(v)$ is connected without cut vertices, it would mean that v is a leaf of \mathcal{H} , but \mathcal{H} has no leaves. This proves (1).

(2) follows directly from (1) unless \mathcal{X} is a single vertex. If $\mathcal{X} = v$ then repeat the above argument with $e_0 = e_1$ to find a second edge $e_2 \neq e_1$ such that \mathfrak{C} also has an edge with a loose end at $\Theta_{\mathcal{Y}}(e_2)$. \square

Lemma 3.31 *Let $S \subset \mathcal{D}$ be a closed minimal cut set that is not a cut point and contains a point x such that $|\Delta^{-1}(x)| = 2$. For $\epsilon \in \pm$, let $\psi_x^\epsilon: [0, \infty] \rightarrow \overline{\mathcal{T}}$ be a simplicial geodesic ray such that $\psi_x^\epsilon([1, \infty])$ contains no branch point of the convex hull \mathcal{H} of $\Delta^{-1}(S)$ and such that $\{\psi_x^-(\infty), \psi_x^+(\infty)\} = \Delta^{-1}(x)$. Then for $\epsilon \in \pm$ there is a unique component $\mathfrak{C}_\epsilon \subset \mathfrak{W}(\mathcal{H})$ such that $\psi_x^\epsilon([1, \infty]) \cap \bigcup_{\mathfrak{C} \in \mathfrak{C}_\epsilon} \mathcal{L}_{\mathfrak{C}} \neq \emptyset$, and these are the only two components of $\mathfrak{W}(\mathcal{H})$. Hence, S has exactly two complementary components.*

Proof By Proposition 3.23, each $\mathfrak{W}(\psi_x^\epsilon([1, \infty])) \odot \phi([0, \infty])$ is connected. Since $\psi_x^\epsilon([1, \infty])$ contains no branch point of \mathcal{H} , it follows that $\mathfrak{W}(\psi_x^\epsilon([1, \infty])) \odot \phi([0, \infty]) = \mathfrak{W}(\psi_x^\epsilon([1, \infty])) \odot \mathcal{H}$ includes into $\mathfrak{W}(\mathcal{H})$ via splicing. Therefore, there is a single component $\mathfrak{C}_\epsilon \subset \mathfrak{W}(\mathcal{H})$ containing the image of $\mathfrak{W}(\psi_x^\epsilon([1, \infty])) \odot \phi([0, \infty])$.

If C is a component of $\mathcal{D} \setminus S$ then by Lemma 3.16 the corresponding component \mathfrak{C} of $\mathfrak{W}(\mathcal{H})$ has an end at $\psi_x^\epsilon(\infty)$ for at least one ϵ , so \mathfrak{C} is either \mathfrak{C}_- or \mathfrak{C}_+ . \square

Lemma 3.32 *Let $\{x, y\}$ be a cut pair in \mathcal{D} . Let \mathcal{H} be the convex hull of $\Delta^{-1}(\{x, y\})$. Let $\phi: [0, l] \rightarrow \mathcal{H}$ be a simplicial geodesic of finite length $l \geq 2$. Then $\mathfrak{W}(\phi([1, l-1])) \odot \phi([0, l])$ has at least two components.*

Proof If $\Delta^{-1}(\{x, y\})$ is two points then ϕ can be extended to be a simplicial geodesic $\phi: [-\infty, \infty] \rightarrow \mathcal{H}$. It is an easy consequence of Lemma 3.18 that if $\mathfrak{W}(\phi([1, l-1])) \odot \phi([0, l])$ is connected then so is $\mathfrak{W}(\mathcal{H})$, so suppose $|\Delta^{-1}(\{x, y\})| > 2$.

Let $\mathcal{X} = \phi([1, l-1])$ and $\mathcal{Y} = \phi([0, l])$.

For $z \in \{x, y\}$, if $|\Delta^{-1}(z)| = 2$ then for $\epsilon \in \pm$ let $\psi_z^\epsilon: [0, \infty] \rightarrow \overline{\mathcal{T}}$ be the simplicial geodesic rays such that $\{\psi_z^-(\infty), \psi_z^+(\infty)\} = \Delta^{-1}(z)$ and such that $\psi_z^\epsilon(0)$ is the only branch point of \mathcal{H} in the image of ψ_z^ϵ . Let $\mathcal{L}_z \in \mathbf{L}$ be the line such that $\{\mathcal{L}_z^-, \mathcal{L}_z^+\} = \Delta^{-1}(z)$.

We make three preliminary claims:

Claim 3.32.1 *Let $\mathcal{Z} \subset \mathcal{H}$ be connected. Let $z \in \{x, y\}$ with $|\Delta^{-1}(z)| = 2$. Then $\Theta_{\mathcal{Z}}(\mathcal{L}_z)$ is the only possible component of $\mathfrak{W}(\mathcal{Z}) \odot \mathcal{H}$ that, for both $\epsilon \in \pm$, contains edges \mathfrak{C}_ϵ with $\mathcal{L}_{\mathfrak{C}_\epsilon} \cap \psi_z^\epsilon([1, \infty]) \neq \emptyset$.*

Proof of Claim. Every component of $\mathfrak{W}(\mathcal{Z}) \odot \mathcal{H}$ that is not equal to $\Theta_{\mathcal{Z}}(\mathcal{L}_x)$ or $\Theta_{\mathcal{Z}}(\mathcal{L}_y)$ includes into a component of $\mathfrak{W}(\mathcal{H})$ via splicing. The claim follows from Lemma 3.31. \diamond

Claim 3.32.2 *Let $\mathcal{Z} \subset \mathcal{H}$ be connected. No component of $\mathfrak{W}(\mathcal{Z}) \odot \mathcal{H}$ has loose ends at three distinct deleted vertices.*

Proof of Claim. If \mathcal{Z} does not contain a branch point of \mathcal{H} then the claim is trivial because $\mathfrak{W}(\mathcal{Z}) \odot \mathcal{H}$ has at most two deleted vertices.

If there is only one branch point, or if there are two and \mathcal{Z} contains both, or if \mathcal{Z} is disjoint from one of \mathcal{L}_x or \mathcal{L}_y , then the claim follows from Claim 3.32.1.

The remaining possibility is that $\mathcal{L}_x \cap \mathcal{L}_y$ contains an edge and \mathcal{Z} contains exactly one branch point of \mathcal{H} . Then $\mathfrak{W}(\mathcal{Z}) \odot \mathcal{H}$ has at most three deleted vertices, which, without loss of generality, we may assume are $\Theta_{\mathcal{Z}}(\psi_x^+(\infty))$, $\Theta_{\mathcal{Z}}(\psi_y^+(\infty))$, and $\Theta_{\mathcal{Z}}(\psi_x^-(\infty)) = \Theta_{\mathcal{Z}}(\psi_y^-(\infty))$.

Suppose \mathcal{C} is a component with loose ends at all three deleted vertices. Let $\mathcal{Z}' = \mathcal{Z} \cup (\mathcal{L}_x \cap \mathcal{L}_y)$. Consider the component \mathcal{C}' of $\mathfrak{W}(\mathcal{Z}') \odot \mathcal{H}$ containing the image of \mathcal{C} under splicing $\mathfrak{W}(\mathcal{Z}) \odot \mathcal{H}$ to $\mathfrak{W}((\mathcal{L}_x \cap \mathcal{L}_y) \setminus \mathcal{Z}) \odot \mathcal{H}$. Since \mathcal{C} had loose ends at $\Theta_{\mathcal{Z}}(\psi_x^+(\infty))$ and $\Theta_{\mathcal{Z}}(\psi_y^+(\infty))$, so does \mathcal{C}' . Since \mathcal{C} had a loose end at $\Theta_{\mathcal{Z}}(\psi_x^-(\infty)) = \Theta_{\mathcal{Z}}(\psi_y^-(\infty))$, the set $\bigcup_{\mathcal{E} \in \mathcal{C}'} \mathcal{L}_{\mathcal{E}}$ contains an edge in $\mathcal{L}_x \cap \mathcal{L}_y$. Thus, $\mathcal{Z}' \cap \bigcup_{\mathcal{E} \in \mathcal{C}'} \mathcal{L}_{\mathcal{E}}$ contains a leaf $v \in \mathcal{L}_x \cap \mathcal{L}_y$ with $v \neq \psi_x^+(0) = \psi_y^+(0)$. By Lemma 3.30, we conclude that $v = \psi_x^-(0) = \psi_y^-(0)$ and that \mathcal{C}' has a loose end at either $\Theta_{\mathcal{Z}'}(\psi_x^-(1))$ or $\Theta_{\mathcal{Z}'}(\psi_y^-(1))$. In either case this contradicts Claim 3.32.1. \diamond

Claim 3.32.3 Let $\mathcal{Z} \subset \mathcal{H}$ be a bounded, connected set such that for $z \in \{x, y\}$ and $\epsilon \in \pm$ we have four distinct vertices $\Theta_{\mathcal{Z}}(\psi_z^\epsilon(\infty))$. For $\epsilon, \epsilon' \in \pm$, let $P^{\epsilon, \epsilon'}$ be the set of components of $\mathfrak{W}(\mathcal{Z}) \odot \mathcal{H}$ that contain an edge with a loose end at $\Theta_{\mathcal{Z}}(\psi_x^\epsilon(\infty))$ and an edge with a loose end at $\Theta_{\mathcal{Z}}(\psi_y^{\epsilon'}(\infty))$. Then one of the following is true:

- $P^{+,+}$ and $P^{-,-}$ are non-empty and $P^{+,-}$ and $P^{-,+}$ are empty.
- $P^{+,-}$ and $P^{-,+}$ are non-empty and $P^{+,+}$ and $P^{-,-}$ are empty.

Proof of Claim. Lemma 3.30 and Claim 3.32.1 imply that every component except $\Theta_{\mathcal{Z}}(\mathcal{L}_x)$ and $\Theta_{\mathcal{Z}}(\mathcal{L}_y)$ belongs to one of the $P^{\epsilon, \epsilon'}$. Claim 3.32.2 implies that no component belongs to more than one of the $P^{\epsilon, \epsilon'}$. If fewer than two of the $P^{\epsilon, \epsilon'}$ are non-empty then $\mathfrak{W}(\mathcal{Z})$ is not connected without cut vertices, which we know it is.

Proposition 3.23 implies that each $\mathfrak{W}(\psi_z^\epsilon([1, \infty])) \odot \psi_z^\epsilon([0, \infty])$ is connected. Therefore, by splicing, we see that all of the components in $P^{\epsilon_1, \epsilon_2}$ and $P^{\epsilon_3, \epsilon_4}$ include into a common component of $\mathfrak{W}(\mathcal{H})$ if $\epsilon_1 = \epsilon_3$ or $\epsilon_2 = \epsilon_4$. $\mathfrak{W}(\mathcal{H})$ has at least two components, so the claim follows. \diamond

We now proceed with the proof of the lemma.

If $\mathcal{Y} \subset \mathcal{L}_z$ for $z \in \{x, y\}$ then $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$ has one component that is the single edge $\Theta_{\mathcal{X}}(\mathcal{L}_z)$ with two loose ends, and at least one other component containing vertices, so at least two components.

Otherwise, if \mathcal{X} does not contain a branch point of \mathcal{H} then \mathcal{X} separates $\Delta^{-1}(x)$ from $\Delta^{-1}(y)$ and $\mathfrak{W}(\mathcal{X}) \odot \mathcal{H} = \mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$. If this is connected then at most one component of $\mathfrak{W}(\mathcal{H})$ has ends in $\Delta^{-1}(x)$ and $\Delta^{-1}(y)$, contradicting Lemma 3.16.

If $\mathfrak{W}(\mathcal{X}) \odot \mathcal{H}$ has four deleted vertices then partition the components into parts $\Theta_{\mathcal{X}}(\mathcal{L}_x)$, $\Theta_{\mathcal{X}}(\mathcal{L}_y)$, and the non-empty $P^{\epsilon, \epsilon'}$ of Claim 3.32.3. Assume, without loss of generality, that $P^{+,+}$ and $P^{-,-}$ are non-empty. $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$ is obtained from $\mathfrak{W}(\mathcal{X}) \odot \mathcal{H}$ by un-deleting two vertices $\Theta_{\mathcal{X}}(\psi_x^\epsilon(\infty))$ and $\Theta_{\mathcal{X}}(\psi_y^{\epsilon'}(\infty))$.

If $\epsilon' = -\epsilon$ then the two components of $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$ are:

$$\Theta_{\mathcal{X}}(\psi_x^\epsilon(\infty)) \cup \Theta_{\mathcal{X}}(\mathcal{L}_x) \cup \bigcup_{\mathfrak{C} \in P^{\epsilon, \epsilon}} \mathfrak{C}$$

and

$$\Theta_{\mathcal{X}}(\psi_y^{-\epsilon}(\infty)) \cup \Theta_{\mathcal{X}}(\mathcal{L}_y) \cup \bigcup_{\mathfrak{C} \in P^{-\epsilon, -\epsilon}} \mathfrak{C}$$

If $\epsilon' = \epsilon$ then the components of $P^{-\epsilon, -\epsilon}$ remain separate components in $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$, distinct from the component containing the un-deleted vertices.

If $\mathfrak{W}(\mathcal{X}) \odot \mathcal{H}$ has three deleted vertices then partition the components of $\mathfrak{W}(\mathcal{X}) \odot \mathcal{H}$ into three parts according to their two loose ends. To get $\mathfrak{W}(\mathcal{X}) \odot \mathcal{Y}$ we un-delete one vertex, which combines two of the parts into a single component but leaves the other part alone. \square

4 Splittings

Armed with the machinery of Section 3, we are now prepared to construct the relative JSJ decomposition.

Otal [17] makes the following observation:

Lemma 4.1 *If F splits over $\langle g \rangle$ relative to \underline{w} then $\Delta(\{\bar{g}^\infty, g^\infty\})$ is a cut point or cut pair in \mathcal{D} .*

Since a rigid decomposition space has no cut points or cut pairs:

Corollary 4.2 *If (F, \underline{w}) is rigid there are no cyclic splittings of F relative to \underline{w} .*

We will prove a converse in Theorem 4.14. This takes care of one case for which the rJSJ is trivial. We saw another trivial case in Example 1.3, in which the decomposition space is a circle. We explore this case in Section 4.1. In particular, a circle has no cut points or uncrossed cut pairs.

In the language of Guirardel and Levitt [9], a subgroup of F is *universally elliptic rel \underline{w}* if it is elliptic in every cyclic splitting of $F \text{ rel } \underline{w}$. A graph of groups decomposition is a JSJ decomposition if all the splittings are over universally elliptic subgroups and the decomposition is maximal with respect to this property. Our goal is to show that cut points and uncrossed cut pairs correspond to universally elliptic relative cyclic splittings. The first step is to show that the stabilizer of a cut point or uncrossed cut pair is an infinite cyclic group over which F splits rel \underline{w} . For cut points this was already noted by Otal. We show that uncrossed cut pairs have infinite cyclic stabilizers in Section 4.2, and in Section 4.3 we construct a simplicial tree with a cocompact F -action whose edge stabilizers are the stabilizers of cut points and uncrossed cut pairs of \mathcal{D} .

In Section 4.4 we prove the stabilizers of cut points and uncrossed cut pairs of \mathcal{D} are exactly the maximal cyclic universally elliptic subgroups of $F \text{ rel } \underline{w}$ over which F splits rel \underline{w} . In Section 4.5 we conclude that the splitting we have constructed is the rJSJ.

4.1 Crossing Pairs and the Circle

In this subsection we give criteria for the decomposition space to be a circle.

Lemma 4.3 ([17, Theorem 2], [4, Theorem 6.1]) *The following are equivalent:*

- (1) (F, \underline{w}) is a QH-surface.
- (2) Some Whitehead graph for \underline{w} is a circle.
- (3) Every Whitehead graph for \underline{w} with no cut vertex is a circle.
- (4) \mathcal{D} is a circle.
- (5) Every minimal cut set of \mathcal{D} is a cut pair.

Proofs of the following two lemmas are elementary and are left to the reader.

Lemma 4.4 (Cut Pair Exchange) *Suppose $\{x_0, x_1\}$ and $\{y_0, y_1\}$ are crossing cut pairs in \mathcal{D} . Then $\{x_0, y_0\}$ is a cut pair.*

Lemma 4.5 *Suppose $\{x, y\}$ and $\{y, z\}$ are cut pairs of \mathcal{D} and for every neighborhood N' of y there exists a connected neighborhood $N \subset N'$ of y such that $N \setminus \{y\}$ has exactly two components. Then $\{x, z\}$ is a cut pair.*

The following proposition is a generalization of a construction of Bowditch for boundaries of hyperbolic groups [2].

Proposition 4.6 \mathcal{D} is a circle if and only if all of the following are satisfied:

- (1) \mathcal{D} is connected.
- (2) \mathcal{D} has no cut points.
- (3) \mathcal{D} has cut pairs.
- (4) Every cut pair in \mathcal{D} is crossed by a cut pair.

Proof A circle satisfies these conditions. We prove the converse.

Define an equivalence relation on \mathcal{D} by $x \sim y$ if $x = y$ or if $\{x, y\}$ is a cut pair. Transitivity follows from Claim 4.6.1 and Lemma 4.5.

Claim 4.6.1 For every cut pair $\{x_0, x_1\}$ and every neighborhood N' of x_0 there exists a connected neighborhood $N \subset N'$ of x_0 such that $N \setminus \{x_0\}$ has precisely two components.

Proof of Claim. Every cut pair is crossed, so by Lemma 2.3 every cut pair has precisely two complementary components. By Corollary 3.28, for every simplicial geodesic $\phi: [0, \infty] \rightarrow \overline{\mathcal{T}}$ there are at most two components of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$. From this and the fact that there are no cut points, Lemma 3.22 gives the desired neighborhood. \diamond

Claim 4.6.2 Equivalence classes are closed.

Proof of Claim. If $[x]$ is a single point we are done. Otherwise, suppose $(y_i) \rightarrow y$ for $x \neq y_i \in [x]$. For some $\eta \in \Delta^{-1}(y)$ there exists a subsequence of (y_i) and a choice of $\eta_i \in \Delta^{-1}(y_i)$ so that $(\eta_i) \rightarrow \eta$ in $\partial\mathcal{T}$. Choose $\xi \in \Delta^{-1}(x)$ and let $\phi: [-\infty, \infty] \rightarrow [\xi, \eta]$ be a simplicial geodesic. Passing to a further subsequence of (η_i) , there are positive integers r_i such that $r_{i+1} > r_i + 1$ and $[\xi, \eta_i] \cap [\xi, \eta] = \phi([-\infty, r_i + 1])$. Since $x \neq y_i \in [x]$, each $\{x, y_i\}$ is a cut pair, so there are at least two components of $\mathfrak{W}(\phi([-r_i, r_i])) \odot \phi([-r_i - 1, r_i + 1])$, by Lemma 3.32. By Lemma 3.29, $\Delta(\{\phi(-\infty), \phi(\infty)\}) = \{x, y\}$ is a cut pair, so $y \in [x]$. \diamond

Claim 4.6.3 All of \mathcal{D} is in one equivalence class.

Given the claim, every point of \mathcal{D} is a member of a cut pair, and it follows from Lemma 4.3 that \mathcal{D} is a circle.

Proof of Claim. We have assumed that a cut pair exists, so there is an equivalence class $[x]$ consisting of more than one point. Suppose that $[x]$ is not all of \mathcal{D} .

Let U be a component of $\mathcal{D} \setminus [x]$. Since \mathcal{D} is locally connected by Proposition 3.10, and since $[x]$ is closed by Claim 4.6.2, U is open in \mathcal{D} . Since \mathcal{D} is connected without cut points, U has at least two limit points in $[x]$. Pick distinct points y_0 and y_1 in $\overline{U} \cap [x]$. Since they are in $[x]$, these points are a cut pair, and $\mathcal{D} \setminus \{y_0, y_1\}$ has exactly two components, A_0 and A_1 . Assume $U \subset A_0$.

Let $\{z_0, z_1\}$ be a cut pair crossing $\{y_0, y_1\}$ with complementary components B_0 and B_1 . Assume $z_0 \in A_0$, $z_1 \in A_1$, $y_0 \in B_0$ and $y_1 \in B_1$.

By Lemma 4.4, z_0 and z_1 are in $[x] \subset \mathcal{D} \setminus U$. Thus, U is contained in B_ϵ , where ϵ is either 0 or 1. Since $U \subset A_0$, we have $U \subset A_0 \cap B_\epsilon$.

Now, $\{y_\epsilon, z_0\}$ is a cut pair whose components are $C_0 = A_0 \cap B_\epsilon$ and $C_1 = A_1 \cup B_{1-\epsilon} \cup \{y_{1-\epsilon}\} \cup \{z_1\}$. However, U , and hence C_0 , has $y_{1-\epsilon} \in C_1$ as a limit point, which is a contradiction. Thus, $[x] = \mathcal{D}$. \diamond \square

Corollary 4.7 *If \mathcal{D} is not rigid and not a circle then \mathcal{D} contains cut points or uncrossed cut pairs.*

4.2 Uncrossed Cut Pairs

In this section we show that uncrossed cut pairs have infinite cyclic stabilizers. A priori, the preimage in ∂F of a pair of points in \mathcal{D} could be as many as four points. This first step is to rule out that possibility.

Lemma 4.8 *The preimage of an uncrossed cut pair is two points.*

Proof Let $\{x_0, x_1\} \subset \mathcal{D}$ be an uncrossed cut pair, and suppose $|\Delta^{-1}(x_0)| = 2$. Then there is an $h \in F$ and a $w \in \underline{w}$ such that $\Delta^{-1}(x_0) = \{h\bar{w}^\infty, hw^\infty\}$. Replacing $\{x_0, x_1\}$ by $\{\bar{h}x_0, \bar{h}x_1\}$, we may assume h is trivial. Let \mathcal{H} be the convex hull of $\Delta^{-1}(\{x_0, x_1\})$. Let $\phi: [-\infty, \infty] \rightarrow [\bar{w}^\infty, w^\infty]$ be the simplicial geodesic with $\phi(0) = \mathbf{1}$. Let p be large enough so that $\phi([-\infty, -p-1]) \cup \phi([p+1, \infty])$ contains no branch point of \mathcal{H} .

By Lemma 3.31, for each $\epsilon \in \pm$ there is a unique component $\mathcal{C}_\epsilon \subset \mathfrak{W}(\mathcal{H})$ such that $\cup_{\mathcal{E} \in \mathcal{C}_\epsilon} \mathcal{L}_\mathcal{E}$ meets $\phi([\epsilon \cdot (p+1), \epsilon \cdot \infty])$, and these are the only two components.

We reach a contradiction by exhibiting a cut pair crossing $\{x_0, x_1\}$. For each $\epsilon \in \pm$, we have $w^{\epsilon \cdot (2p+1)}x_1 \subset \Delta(\partial\mathcal{T}_{\mathcal{C}_\epsilon})$, so $\{w^{2p+1}x_1, w^{-2p-1}x_1\}$ crosses $\{x_0, x_1\}$.

$\{w^{2p+1}x_1, w^{-2p-1}x_1\}$ is a cut pair by Lemma 4.5: $w^{2p+1}\{x_1, x_0\} = \{w^{2p+1}x_1, x_0\}$ and $w^{-2p-1}\{x_0, x_1\} = \{x_0, w^{-2p-1}x_1\}$ are both cut pairs, and, by Lemma 3.22, for every neighborhood N' of x_0 there exists a connected neighborhood $N \subset N'$ such that $N \setminus \{x_0\}$ has exactly two components. \square

Lemma 4.9 *If $\Delta(\{\xi, g^\infty\})$ is a cut pair, for some $g \in F$ and $\xi \in \partial\mathcal{T}$ with $\xi \neq \bar{g}^\infty$, then it is a crossed cut pair.*

Proof We assume, without loss of generality, that g is cyclically reduced. Let \mathcal{H} be the convex hull of $\{\bar{g}^\infty, g^\infty, \xi\}$.

There is a g -action on $\mathfrak{W}([\bar{g}^\infty, g^\infty])$ given by $\mathfrak{E} \mapsto \Theta_{[\bar{g}^\infty, g^\infty]}(g\mathcal{L}_{\mathfrak{E}})$ for an edge \mathfrak{E} . This action permutes the components of $\mathfrak{W}([\bar{g}^\infty, g^\infty])$. Replacing g by a suitable power, we may assume the components are fixed.

Let $\psi: [0, \infty]$ be the simplicial geodesic ray with $\psi(\infty) = \xi$ and $\psi([0, \infty]) \cap [\bar{g}^\infty, g^\infty] = \psi(0)$. Let $\phi: [-\infty, \infty] \rightarrow [\bar{g}^\infty, g^\infty]$ be the simplicial geodesic with $\phi(0) = \psi(0)$.

Let $\mathcal{L} \in \mathbf{L}|_{\phi((0,1))}$ be a line such that $\Theta_{\mathcal{H}}(\mathcal{L})$ belongs to a component of $\mathfrak{W}(\mathcal{H})$ with an end at $\psi(\infty)$. Then there exists an edge path $\mathfrak{P}: \Theta_{[\bar{g}^\infty, g^\infty]}(\mathcal{L}) = \mathfrak{E}_0, \dots, \mathfrak{E}_k$ in $\mathfrak{W}([\bar{g}^\infty, g^\infty])$ with \mathfrak{E}_k incident to $\Theta_{[\bar{g}^\infty, g^\infty]}(\psi(\infty))$. For all sufficiently large m , we have $\phi([-\infty, \infty]) \cap \bigcup_{\mathfrak{E} \in g^m \mathfrak{P}} \mathcal{L}_{\mathfrak{E}} \subset \phi([1, \infty])$. Since the g -action preserves components, \mathfrak{E}_0 and $g^m \mathfrak{E}_0$ are in the same component. By Lemma 3.18 and Lemma 3.21, there is an edge path in $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$ that connects $\Theta_{\phi([1, \infty])}(\mathcal{L}_{\mathfrak{E}_0})$ to $\Theta_{\phi([1, \infty])}(\mathcal{L}_{g^m \mathfrak{E}_0})$. By concatenating $g^m \mathfrak{P}$, we see that $\Theta_{\phi([1, \infty])}(\mathcal{L}_{\mathfrak{E}_0})$ is in the $\Theta_{\phi([1, \infty])}(g^m \psi(\infty))$ component of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([0, \infty])$. Since this is true for all sufficiently large m , every line in $\mathbf{L}|_{\phi((0,1))}$ that contributes an edge to a component of $\mathfrak{W}(\mathcal{H})$ with an end at $\psi(\infty)$ contributes an edge to the same component of $\mathfrak{W}(\phi([1, \infty])) \odot \phi([1, \infty])$, so there is only one component, \mathfrak{C}_+ , of $\mathfrak{W}(\mathcal{H})$ with ends at $\phi(\infty)$ and $\psi(\infty)$. The same argument in the $\phi(-\infty)$ direction shows there is only one component, \mathfrak{C}_- , of $\mathfrak{W}(\mathcal{H})$ with ends at $\phi(-\infty)$ and $\psi(\infty)$. Since $\Delta(\{\psi(\infty), \phi(\infty)\})$ is a cut pair, $\mathfrak{C}_+ \neq \mathfrak{C}_-$.

It follows from Lemma 3.30 that for $\epsilon \in \pm$, $\mathcal{H} \cap \bigcup_{\mathfrak{E} \in \mathfrak{C}_\epsilon} \mathcal{L}_{\mathfrak{E}} = [\psi(\infty), \phi(\epsilon \cdot \infty)]$, so $\Theta_{\mathcal{H}}(g^{\epsilon-1} \psi(\infty)) \in \mathfrak{C}_\epsilon$. The map $\mathfrak{E} \mapsto \Theta_{[\psi(\infty), \phi(\infty)]}(\mathcal{L}_{\mathfrak{E}})$ sends \mathfrak{C}_+ onto a component of $\mathfrak{W}([\psi(\infty), \phi(\infty)])$ not containing the vertex $\Theta_{[\psi(\infty), \phi(\infty)]}(\phi(-\infty))$, and sends \mathfrak{C}_- into the component of $\mathfrak{W}([\psi(\infty), \phi(\infty)])$ containing $\Theta_{[\psi(\infty), \phi(\infty)]}(\phi(-\infty))$, so $\Theta_{[\psi(\infty), \phi(\infty)]}(\bar{g}\psi(\infty))$ and $\Theta_{[\psi(\infty), \phi(\infty)]}(g\psi(\infty))$ are in different components of $\mathfrak{W}([\psi(\infty), \phi(\infty)])$. Moreover, for $\mathcal{L} \in \mathbf{L}|_{[\bar{g}\phi(0), g\phi(0)]}$ we have that $\Theta_{g\mathcal{H}}(\mathcal{L}) \in g\mathfrak{C}_-$ if and only if $\Theta_{\bar{g}\mathcal{H}}(\mathcal{L}) \in \bar{g}\mathfrak{C}_+$. It follows that $\mathfrak{W}([\bar{g}\psi(\infty), g\psi(\infty)])$ has a component not containing $\Theta_{[\bar{g}\psi(\infty), g\psi(\infty)]}(\phi(-\infty))$ and $\Theta_{[\bar{g}\psi(\infty), g\psi(\infty)]}(\phi(\infty))$. Therefore, $\Delta(\{\bar{g}\psi(\infty), g\psi(\infty)\})$ is a cut pair, and it crosses $\Delta(\{\psi(\infty), \phi(\infty)\})$. \square

Proposition 4.10 *Uncrossed cut pairs are rational: For every uncrossed cut pair S there exists a non-trivial element $f \in F$ such that $\langle f \rangle$ is the stabilizer of S .*

Proof Let $S = \{x_0, x_1\}$. By Lemma 4.8, $|\Delta^{-1}(x_i)| = 1$ for both $i \in \{0, 1\}$. Let $\phi: [-\infty, \infty] \rightarrow [\Delta^{-1}(x_0), \Delta^{-1}(x_1)]$ be a simplicial geodesic.

For increasing r , the sequence $\#\mathfrak{W}(\phi([r+1, \infty])) \odot \phi([r, \infty])$ is a non-decreasing sequence of integers, bounded above by the maximum valence of $\mathfrak{W}(\mathbf{1})$. Assume r is large enough so that the sequence has achieved its maximum.

Let g, h , and a be the elements provided by Lemma 3.26 for $\phi: [r, \infty] \rightarrow \overline{\mathcal{T}}$. There is some $r' > r$ such that $\phi(r') = h$ and some $r'' > r' + 1$ such that $\phi(r'') = gha$. Because of our choice of r , Lemma 3.26 (3) implies that for $\mathcal{L} \in \mathbf{L}|_{\phi([r', r'+1])}$, the edges $\Theta_{\phi([r'+1, \infty])}(\mathcal{L})$ and $\Theta_{\phi([r'+1, \infty])}(g\mathcal{L})$ are in the same component of $\mathfrak{W}(\phi([r'+1, \infty])) \odot \phi([r', \infty])$.

If $\phi(\infty) \neq g^\infty$ and $\phi(-\infty) \neq \bar{g}^\infty$ then, since $\Delta(\{\bar{g}^\infty, g^\infty\})$ is a cut point or cut pair by Corollary 3.27 and S is uncrossed, $\Theta_{\phi([-\infty, \infty])}(\bar{g}^\infty)$ and $\Theta_{\phi([-\infty, \infty])}(g^\infty)$ are in the same component of $\mathfrak{W}(\phi([-\infty, \infty]))$. Let $\mathcal{L} \in \mathbf{L}|_{\phi([r', r'+1])}$ be a line such that $\Theta_{\phi([-\infty, \infty])}(\mathcal{L})$ is in a component of $\mathfrak{W}(\phi([-\infty, \infty]))$ not containing $\Theta_{\phi([-\infty, \infty])}(\bar{g}^\infty)$ and $\Theta_{\phi([-\infty, \infty])}(g^\infty)$. Then $\Theta_{\phi([r'+1, \infty])}(\mathcal{L})$ and $\Theta_{\phi([r'+1, \infty])}(g\mathcal{L})$ are in the same component of $\mathfrak{W}(\phi([r'+1, \infty])) \odot \phi([r', \infty])$, which, in light of the inclusion via splicing, implies that $\Theta_{\phi([-\infty, \infty])}(\mathcal{L})$ and $\Theta_{\phi([-\infty, \infty])}(g\mathcal{L})$ are in the same component of $\mathfrak{W}(\phi([-\infty, \infty]))$. This implies that $\Theta_{\phi([-\infty, \infty])}(g\phi(-\infty))$ is in a component of $\mathfrak{W}(\phi([-\infty, \infty]))$ not containing the vertex $\Theta_{\phi([-\infty, \infty])}(g^\infty)$. However, $\Theta_{\phi([-\infty, \infty])}(g\phi(\infty))$ is in the component of $\mathfrak{W}(\phi([-\infty, \infty]))$ containing $\Theta_{\phi([-\infty, \infty])}(g^\infty)$, and $\Delta(\{g\phi(-\infty), g\phi(\infty)\})$ is a cut pair. This contradicts the hypothesis that S is uncrossed.

In the two cases that $[\bar{g}^\infty, g^\infty]$ and $\phi([-\infty, \infty])$ share exactly one endpoint, Lemma 4.9 gives a contradiction to the hypothesis that S is uncrossed. Therefore $\phi(\infty) = g^\infty$ and $\phi(-\infty) = \bar{g}^\infty$, and we take f to be an indivisible root of g . \square

Proposition 4.11 \mathcal{D} has finitely many orbits of uncrossed cut pairs.

Proof The element g in Proposition 4.10 is conjugate to a word of bounded length provided by Lemma 3.26. There are finitely many such conjugacy classes. \square

Lemma 4.12 If \mathcal{D} is not rigid or a circle then there is an indivisible element $g \in F$ such that $\Delta(\{\bar{g}^\infty, g^\infty\})$ is a cut set that is not crossed by any cut pair.

Proof By Corollary 4.7 there exist cut points or uncrossed cut pairs. For cut points apply Proposition 3.25; for uncrossed cut pairs apply Lemma 4.8. \square

4.3 The Splitting Criterion

Proposition 4.13 is a generalization of a construction of Otal [17], who proves it in the case that $\{S_i\}_{i \in I}$ is a single orbit of cut points. The main change is Claim 4.13.1, which replaces Otal's Lemma 3.3.

Proposition 4.13 *Consider a non-empty, F -invariant collection of disjoint cut sets $\{S_i\}_{i \in I}$ in \mathcal{D} satisfying the following conditions:*

- (1) *For each $i \in I$ there is an indivisible $g_i \in F$ such that $S_i = \Delta(\{\bar{g}_i^\infty, g_i^\infty\})$.*
- (2) *The cut sets are pairwise non-crossing.*
- (3) *The set $\{S_i\}_{i \in I}$ is a union of finitely many F -orbits.*

Then F splits as a graph of groups $\text{rel } w$ with cyclic edge stabilizers. The vertex set is bipartite, with Type 1 vertices stabilized by maximal cyclic subgroups generated by the g_i and Type 2 vertices stabilized by non-cyclic subgroups.

Proof For each $i \in I$ there is a partial ordering $<_i$ on $\{S_j\}_{j \in I \setminus \{i\}}$ defined by $S_j <_i S_k$ if S_j separates S_i from S_k , that is, if S_i and S_k are in different complementary components of S_j . Since the cut sets are pairwise non-crossing, the complementary component of S_j containing S_k is well defined.

Claim 4.13.1 *There exist $<_i$ -minimal elements.*

Proof of Claim. Suppose $S_i, S_j,$ and S_k are elements of $\{S_i\}_{i \in I}$. Let \mathcal{X} be $[\bar{g}_i^\infty, g_i^\infty] \cap [\bar{g}_k^\infty, g_k^\infty]$, if this intersection is non-empty. Otherwise, let \mathcal{X} be the geodesic segment in \mathcal{T} connecting $[\bar{g}_i^\infty, g_i^\infty]$ to $[\bar{g}_k^\infty, g_k^\infty]$. Suppose $[\bar{g}_j^\infty, g_j^\infty]$ does not intersect \mathcal{X} . Then there exists an edge $e \in \mathcal{T}$ incident to, but not contained in, $[\bar{g}_j^\infty, g_j^\infty]$, such that e separates $[\bar{g}_j^\infty, g_j^\infty]$ from at least three of the points $\bar{g}_i^\infty, g_i^\infty, \bar{g}_k^\infty,$ and g_k^∞ . It then follows from Lemma 2.16 that S_i and S_k contain points in a common component of $\mathcal{D} \setminus S_j$, so S_j does not separate S_i from S_k in \mathcal{D} .

We conclude that for fixed S_i and S_k the only S_j such that $S_j <_i S_k$ belong to the finite set of those for which $[\bar{g}_j^\infty, g_j^\infty]$ intersects \mathcal{X} . \diamond

Define a graph on which F acts without inversions as follows. The graph has two types of vertices. There is a Type 1 vertex v_i for each S_i . Given a Type 1 vertex v_i , S_i has finitely many complementary components $C_{i,1}, \dots, C_{i,m_i}$. For each i, j , consider the subset $\{v_i\} \cup \{v_k \mid S_k \text{ is } <_i \text{-minimal, and } S_k \in C_{i,j}\}$. Define this subset to be a Type 2 vertex. Define adjacency by inclusion.

Since the S_i are cut sets, this graph is a tree. The quotient of this tree by the F -action contains one vertex of Type 1 for each orbit of cut set, and some finite number of adjacent Type 2 vertices. The stabilizers of the Type 1 vertices are the groups $\langle g_i \rangle$, so we have a cyclic splitting of F .

The generators of the line pattern must be conjugate into the vertex groups, otherwise we would have a line in the pattern crossing from one component of some $\mathfrak{W}([\bar{g}_j^\infty, g_j^\infty])$ to another, which is absurd. \square

Combining Proposition 4.13 and Lemma 4.12 gives us a splitting theorem:

Theorem 4.14 (Splitting Theorem) *If \mathcal{D} is not rigid then either (F, \underline{w}) is a three-holed sphere or F splits over \mathbb{Z} relative to \underline{w} .*

Proof If \mathcal{D} is neither rigid nor a circle then Lemma 4.12 provides a g so that the translates of $\Delta(\{\bar{g}^\infty, g^\infty\})$ satisfy the hypotheses of Proposition 4.13, which gives a relative splitting.

If \mathcal{D} is a circle then, by Lemma 4.3, (F, \underline{w}) is a surface with boundary. Either this is a three-holed sphere, or there exists an essential, non-peripheral simple closed curve in the surface, which gives a relative splitting. \square

4.4 Refining Splittings

Theorem 4.14 tells us when we can split F rel \underline{w} . In this subsection we determine when a splitting can be refined.

Definition 4.15 Let Γ be a graph of groups decompositions of F rel \underline{w} with cyclic edge stabilizers. Define the *augmented multiword* $\text{Aug}_\Gamma(\underline{w})$ to be a multiword in F obtained by choosing generators of representatives of the distinct conjugacy classes of maximal cyclic subgroups of F containing the elements of \underline{w} and the generators of each of the edge groups of Γ . The choices can be, and are, made so that $\underline{w} \subset \text{Aug}_\Gamma(\underline{w})$.

Definition 4.16 Define the *augmentation map*, α_Γ , to be the quotient map $\alpha_\Gamma: \mathcal{D}_\underline{w} \rightarrow \mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$. Note that $\Delta_{\text{Aug}_\Gamma(\underline{w})} = \alpha_\Gamma \circ \Delta_\underline{w}$.

Lemma 4.17 *Let G be a non-cyclic vertex group of Γ . The decomposition space $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$ of G corresponding to $\text{Ind}_\Gamma^G(\underline{w})$ embeds naturally into the decomposition space $\mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$ of F corresponding to $\text{Aug}_\Gamma(\underline{w})$.*

Proof The inclusion $\iota: G \hookrightarrow F$ extends to an embedding $\partial\iota: \partial G \hookrightarrow \partial F$. The equivalence relation on ∂G coming from $\text{Ind}_\Gamma^G(\underline{w})$ is the restriction to $\partial\iota(\partial G)$ of the equivalence relation on ∂F coming from $\text{Aug}_\Gamma(\underline{w})$. \square

Lemma 4.18 *Let Γ be a graph of groups decomposition of F rel \underline{w} with cyclic edge groups. Let G be a non-cyclic vertex of Γ . The decomposition space $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$ of G with respect to $\text{Ind}_\Gamma^G(\underline{w})$ is connected.*

Proof By Proposition 3.10, if $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$ is not connected there is a free splitting of G rel $\text{Ind}_\Gamma^G(\underline{w})$. This gives a free splitting of F rel $\text{Aug}_\Gamma(\underline{w})$, which implies that $\mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$ is not connected. This is not possible since $\mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$ is a quotient of $\mathcal{D}_{\underline{w}}$, which, by our standing assumption, is connected. \square

Lemma 4.19 *Let Γ be a graph of groups decomposition of F rel \underline{w} with cyclic edge groups. Let G be a non-cyclic vertex of Γ such that $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$ is rigid. Let $\langle g \rangle$ be the stabilizer of an edge incident to G . Then $S = \Delta_{\underline{w}}(\{\bar{g}^\infty, g^\infty\})$ is either a cut point or uncrossed cut pair of $\mathcal{D}_{\underline{w}}$.*

Proof Suppose not. S is a cut set, so by Lemma 2.3 it is a cut pair with exactly two complementary components, B_0 and B_1 . Assume $\Delta(\partial G) \subset B_0$.

Since S is crossed then there is a cut pair R of $\mathcal{D}_{\underline{w}}$ crossing S . Let r be the point of R in B_0 , and let A_0 and A_1 be the complementary components of R .

Let $\{S_i\}_{i \in I}$ be the collection of cut points and cut pairs corresponding to the edges of Γ . S is one of these, so set $S = S_{i_0}$. Set:

$$J = \{j \neq i_0 \mid S_j \text{ crosses } R \text{ and } S_j \subset B_0\}$$

If $J = \emptyset$ then $\alpha_\Gamma(r) \in \mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$, and:

$$\alpha_\Gamma(\bar{A}_0) \cap \alpha_\Gamma(\bar{A}_1) \cap \mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})} = \alpha_\Gamma(\{S_{i_0}, r\})$$

If $J \neq \emptyset$ then there is a $<_{i_0}$ -minimal element S_{i_1} of $\{S_j\}_{j \in J}$, and:

$$\alpha_\Gamma(\bar{A}_0) \cap \alpha_\Gamma(\bar{A}_1) \cap \mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})} = \alpha_\Gamma(\{S_{i_0}, S_{i_1}\})$$

In either case we get a cut pair of $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$, contradicting rigidity. \square

Definition 4.20 *An uncrossed collection in $\mathcal{D}_{\underline{w}}$ is a non-empty union of orbits of cut points and uncrossed cut pairs.*

There are finitely many such orbits by Proposition 3.25 and Proposition 4.11, so, given an uncrossed collection $\{S_i\}_{i \in I}$, Proposition 4.13 provides a corresponding graph of groups decomposition Γ . For the remainder of this section, fix an uncrossed collection $\{S_i\}_{i \in I}$ and corresponding graph of groups Γ with Bass-Serre tree $\text{BS}(\Gamma)$.

Lemma 4.21 (Universality of Uncrossed Splittings) *The stabilizer of a cut point or cut pair of $\mathcal{D}_{\underline{w}}$ is elliptic in Γ .*

Proof Let S be a cut point or cut pair. By considering the F -action on the convex hull of $\Delta^{-1}(S)$ it is clear that the stabilizer of S is either trivial or a maximal cyclic subgroup. If its stabilizer is trivial we are done, so assume its stabilizer is $\langle g \rangle$.

If g is not elliptic then it has an axis in $\text{BS}(\Gamma)$. A Type 1 vertex on this axis corresponds to a cut set S_i separating the two points of S . By Lemma 2.3, this would mean S crosses S_i , contradicting the hypothesis that the S_i are uncrossed. \square

Lemma 4.22 *If S is a cut point or cut pair of $\mathcal{D}_{\underline{w}}$ then $\alpha_\Gamma(S)$ is a cut point or cut pair of $\mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$. If R is a cut point or cut pair of $\mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$ then $\alpha_\Gamma^{-1}(R)$ is a cut point or cut pair of $\mathcal{D}_{\underline{w}}$.*

Proof α_Γ identifies points of uncrossed cut pairs, so $\alpha_\Gamma(S)$ fails to be a cut set only if some uncrossed cut pair crosses S . That is impossible, since by Lemma 2.3 crossing is a symmetric relation.

Conversely, let R be a cut point or cut pair of $\mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$. It is clear that $\alpha_\Gamma^{-1}(R)$ is a cut set in $\mathcal{D}_{\underline{w}}$; we just need to show that it consists of at most two points.

Suppose a point $r \in R$ has preimage $\alpha_\Gamma^{-1}(r)$ consisting of two points. Then r is stabilized by a conjugate of an element $h \in \text{Aug}_\Gamma(\underline{w}) \setminus \underline{w}$, so $\alpha_\Gamma^{-1}(r) = S_i$ for some i . By the first part of the lemma, $\alpha_\Gamma(S_i) = \{r\}$ is a cut point, so $R = \{r\}$, and $\alpha_\Gamma^{-1}(R) = S_i$ is a cut pair. \square

Lemma 4.23 *Let G be a non-cyclic vertex group of Γ . Let g be an element of G such that $\Delta_{\text{Ind}_\Gamma^G(\underline{w})}(\{\bar{g}^\infty, g^\infty\})$ is a cut point or uncrossed cut pair in $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$. Then $\Delta_{\underline{w}}(\{\bar{g}^\infty, g^\infty\})$ is a cut point or uncrossed cut pair, respectively, of $\mathcal{D}_{\underline{w}}$ that is not in the uncrossed collection.*

Proof There is a Type 2 vertex $\{v_j\}_{j \in J}$ of $\text{BS}(\Gamma)$ corresponding to G . For each $j \in J$, the corresponding cut set S_j becomes a point in $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$, and each complementary component of the image of $\Delta_{\text{Ind}_\Gamma^G(\underline{w})}(\{\bar{g}^\infty, g^\infty\})$ in $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$ contains some of these

S_j points, so $\alpha_\Gamma^{-1}(\Delta_{\text{Ind}_\Gamma^G(\underline{w})}(\{\bar{g}^\infty, g^\infty\})) = \Delta_{\underline{w}}(\{\bar{g}^\infty, g^\infty\})$ is a cut point or cut pair in $\mathcal{D}_{\underline{w}}$ separating some of these S_j . The Type 2 vertices are subsets $\{v_j\}_{j \in J}$ such that for $j_0, j_1 \in J$ there is no $i \in I$ such that S_i separates S_{j_0} from S_{j_1} . Thus, $\Delta_{\underline{w}}(\{\bar{g}^\infty, g^\infty\})$ is a cut set that is not in $\{S_i\}_{i \in I}$.

Since $\Delta_{\underline{w}}(\{\bar{g}^\infty, g^\infty\})$ is not one of the S_i , its cardinality is the same as that of the image in $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$. Thus, they are either both cut points or both cut pairs.

It remains only to show that if $\Delta_{\text{Ind}_\Gamma^G(\underline{w})}(\{\bar{g}^\infty, g^\infty\})$ is an uncrossed cut pair then $\Delta_{\underline{w}}(\{\bar{g}^\infty, g^\infty\})$ is uncrossed. Suppose $\Delta_{\underline{w}}(\{\bar{g}^\infty, g^\infty\})$ is crossed by some other cut pair R of $\mathcal{D}_{\underline{w}}$. R cannot cross any of the S_i , so $R \subset \Delta(\partial G)$, which implies $\alpha_\Gamma(R)$ and $\alpha_\Gamma(\Delta_{\text{Ind}_\Gamma^G(\underline{w})}(\{\bar{g}^\infty, g^\infty\}))$ are crossing cut pairs in the embedded copy of $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$ in $\mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$. Thus we get crossing cut pairs of $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$, contrary to hypothesis. \square

Lemma 4.24 (Refinement Lemma) *Let G be a non-cyclic vertex of Γ . If $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$ is neither a circle nor rigid then there is a refinement Γ' of Γ obtained by splitting G rel $\text{Ind}_\Gamma^G(\underline{w})$. This is a splitting over an uncrossed collection containing $\{S_i\}_{i \in I}$.*

Proof By Lemma 4.18, $\mathcal{D}_{\text{Ind}_\Gamma^G(\underline{w})}$ is connected. If it is not a circle and not rigid, then by Lemma 4.12 there is an element $g \in G$ such that $\Delta_{\text{Ind}_\Gamma^G(\underline{w})}(\{\bar{g}^\infty, g^\infty\})$ is a cut point or uncrossed cut pair that is not in the uncrossed collection. By Lemma 4.23, $\Delta_{\underline{w}}(\{\bar{g}^\infty, g^\infty\})$ is a cut point or uncrossed cut pair of $\mathcal{D}_{\underline{w}}$. Thus, we can add the orbit of $\Delta_{\underline{w}}(\{\bar{g}^\infty, g^\infty\})$ to the set $\{S_i\}_{i \in I}$ to get a larger uncrossed collection, and hence a graph of groups decomposition Γ' refining Γ . \square

4.5 The Decomposition Theorem

Theorem 4.25 (Relative JSJ-Decomposition Theorem) *There exists a canonical relative JSJ-decomposition (rJSJ), a graph of groups decomposition Γ of F relative to \underline{w} with cyclic edge groups, satisfying the following conditions:*

- (1) *If there is more than one vertex, the graph is bipartite. Cyclic vertex groups are adjacent only to non-cyclic vertex groups, and vice-versa. Furthermore, if G is a non-cyclic vertex group the incident edge groups map onto G -maximal cyclic subgroups of G in distinct G -conjugacy classes. Finally, the sum of the degrees of the edge inclusions at any cyclic vertex group is at least 2.*
- (2) *Γ is universal: if F splits over a cyclic subgroup relative to \underline{w} then the cyclic subgroup is conjugate into one of the vertex groups.*
- (3) *Γ is maximal: it cannot be refined and still satisfy these conditions.*

Moreover, the rJSJ is characterized by splitting F over the stabilizers of cut points and uncrossed cut pairs in $\mathcal{D}_{\underline{w}}$. There are three mutually exclusive possibilities:

- (a) (F, \underline{w}) is rigid. $\mathcal{D}_{\underline{w}}$ has no cut points or cut pairs. The rJSJ is trivial.
- (b) (F, \underline{w}) is a QH-surface. $\mathcal{D}_{\underline{w}}$ is a circle. The rJSJ is trivial.
- (c) The rJSJ is nontrivial. For every non-cyclic vertex group G we have that $(G, \text{Ind}_{\Gamma}^G(\underline{w}))$ is either rigid or a QH-surface. Stabilizers of cut points and uncrossed cut pairs are conjugate to the cyclic vertex groups.

Consequently, if F splits over $\langle g \rangle$ relative to \underline{w} then $\langle g \rangle$ is conjugate into the stabilizer of one of the cyclic vertices or one of the QH-surface vertices of the rJSJ.

Remark Conditions (2) and (3) are standard requirements for a JSJ decomposition. In general there is not a canonical JSJ decomposition satisfying these condition, but a whole deformation space of JSJ decompositions [9]. A particular JSJ decomposition can be chosen from this deformation space by applying the normalizations from Section 2.4. Condition (1) says these normalizations have been performed.

Proof If (F, \underline{w}) is rigid or is a three-holed sphere then there are no splittings rel \underline{w} . The rJSJ is trivial, and we are done.

If (F, \underline{w}) is a QH-surface other than a three-holed sphere then splittings rel \underline{w} come from essential, non-peripheral simple closed curves on the surface. For any such curve we can find another intersecting it, giving us an incompatible splitting. Therefore, no such splitting is universal, so the rJSJ is trivial, and we are done.

If we are not in either of these cases then by Corollary 4.7 there exists a cut point or an uncrossed cut pair in $\mathcal{D}_{\underline{w}}$.

Take the uncrossed collection $\{\mathcal{S}_i\}_{i \in I}$ consisting of all cut points and uncrossed cut pairs, and let Γ be the graph of groups provided by Proposition 4.13.

The tree $\text{BS}(\Gamma)$ is canonically defined by the topology of the decomposition space, and the F -action is induced by the F -action on \mathcal{D} , so the resulting graph of groups decomposition is canonical.

We will show Γ satisfies conditions (1)-(3). Conversely, we will show that any Γ' satisfying conditions (1)-(3) has Bass-Serre tree $\text{BS}(\Gamma')$ equivariantly isomorphic to $\text{BS}(\Gamma)$, so Γ and Γ' are equivalent graph of groups decompositions.

Condition (1) says that the graph of groups is normalized as in Section 2.4. Using the facts that F is free and that the cyclic vertex groups of Γ are maximal cyclic subgroups of F it is easy to see that Γ satisfies these conditions.

Uncrossed splittings are universal by Lemma 4.21, so Γ satisfies condition (2).

If for some non-cyclic vertex G of Γ the pair $(G, \text{Ind}_\Gamma^G(\underline{w}))$ is neither rigid nor a QH-surface then by Lemma 4.24 there is a refinement of Γ coming from a larger uncrossed collection. This is absurd; we have already included all cut points and uncrossed cut pairs in our uncrossed collection. Together with condition (2) this implies that any refinement of Γ must come from splitting a QH-surface vertex group. The resulting splitting would not be universal, because if there is a way to split the surface then there is always an incompatible way to split it. Thus, condition (3) is satisfied.

Now suppose Γ' is another graph of groups decomposition of $F \text{ rel } \underline{w}$ satisfying conditions (1)-(3). Condition (c) must be satisfied or else it would be possible to refine Γ' in a universal way, contradicting maximality.

Consider a cyclic vertex group $\langle g \rangle$ of Γ' . We would like to show $S = \Delta(\{\bar{g}^\infty, g^\infty\})$ is a cut point or uncrossed cut pair in $\mathcal{D}_{\underline{w}}$. The number of components of $\mathcal{D}_{\underline{w}} \setminus S$ is equal to the sum of the degrees of the edge maps into the vertex group. By condition (1), this is at least two, so S is a cut point if $g \in \underline{w}$ or cut pair otherwise.

If the sum of the degrees of the edge maps is greater than two, or if the sum of the degrees is equal to two and one of the adjacent non-cyclic vertices is rigid, then S is an uncrossed cut pair, by Lemma 4.19.

Otherwise, either the vertex separates two QH-surfaces glued along boundary curves or it is adjacent to one QH-surface and the edge maps into the cyclic vertex group with degree 2. In the first case, the cyclic vertex can be removed by gluing together the two QH-surfaces to give a larger QH-surface. In the second case, the cyclic vertex can be removed by gluing a Möbius strip to the corresponding QH-surface along their boundary curves. In either case, the new surface contains an essential, non-peripheral simple closed curve that intersects the curve we just glued along. This would provide a splitting of $F \text{ rel } \underline{w}$ incompatible with Γ' , contradicting universality. Thus, each cyclic vertex group of Γ' is the stabilizer of a cut point or uncrossed cut pair. Furthermore, the cyclic vertex groups account for all the cut points and uncrossed cut pairs, since Γ' is maximal.

$\text{BS}(\Gamma)$ and $\text{BS}(\Gamma')$ are both equivariantly isomorphic to the tree constructed in Proposition 4.13, hence, to each other. \square

5 Virtually Geometric Multiwords

A *handlebody* is a 3-manifold obtained by gluing 1-handles to a 3-ball. These are commonly imagined as thickened graphs, although we do not assume orientability.

A multiword $\underline{w} = \{w_1, \dots, w_k\}$ in $F = F_n$ is *geometric* if there exists a handlebody H with fundamental group F such that the conjugacy classes of the w_i can be represented by an embedded multicurve in the boundary of H . The multiword is *virtually geometric* if it becomes geometric upon lifting to a finite index subgroup of F . (Recall Definition 2.4)

Example 5.1 $\underline{w} = \{\bar{a}^2 \bar{b} a b\}$ in $F_2 = \langle a, b \rangle$




Figure 15: $\mathfrak{W}(\mathbf{1})\{\bar{a}^2 \bar{b} a b\}$




Figure 16: Corresponding handlebody for $\bar{a}^2 \bar{b} a b$

In Figure 15 we have a Whitehead graph with vertices blown up to discs and a numbering around each vertex. Following the edges according to the numbering reads off the word $\bar{a}^2 \bar{b} a b$. Embed this graph on the surface of a three-ball.

The numbering around the b -disc can be read $3 - 5$ going counterclockwise, and the numbering around the \bar{b} -disc can be read $3 - 5$ going clockwise, so we can glue these discs together, matching the numbering, to make an orientable b -handle.

The numbering around both the a -disc and \bar{a} -disc is $1 - 4 - 2$ going counterclockwise. We can glue these discs together and match the numbering, but we must do it in a non-orientable way. In drawing the corresponding handlebody in Figure 16 we leave the a -disc and \bar{a} -disc apart, but one should imagine that they have been identified to create a non-orientable a -handle. \diamond

Similarly, any multiword with a Whitehead graph that is planar and valence at most three is geometric. The same argument shows, more generally:

Proposition 5.2 ([7]) *If there exists a Whitehead graph of \underline{w} with a planar embedding such that the cyclic orderings of edges incident to inverse vertices are consistent (either the same or opposite), then \underline{w} is geometric.*

In fact, more is true. If a multiword is geometric then every minimal Whitehead graph of \underline{w} has a planar embedding such that the cyclic orderings of edges incident to inverse vertices are consistent.

Thus, there is an algorithm to determine geometricity: The Whitehead graph is finite, so it is possible to check by brute force to see if there exists a planar embedding that respects cyclic ordering around the vertices.

These claims follow from work of Zieschang [7] (see also Berge's *Documentation for the program Heegaard* [1]).

There is a positive algorithm for checking virtual geometricity of \underline{w} by enumerating subgroups of successively larger index, computing the lift of \underline{w} , and then checking for geometricity. There is not an obvious bound on how large an index is necessary for a given multiword, so there is not a corresponding negative algorithm. Gordon and Wilton [8] even asked whether every one element multiword is virtually geometric. Manning [13] answered in the negative by showing that the word $w = bbaaccabc$ in $F_3 = \langle a, b, c \rangle$ is not virtually geometric.

Otal [17] notes that a free splitting of the free group corresponds to a connected sum of the corresponding handlebodies, so, as usual, we will confine our attention to the case that F does not split freely rel \underline{w} .

5.1 Rigid Multiwords and Geometricity

Lemma 5.3 ([17, Proposition 0]) *If \underline{w} is geometric then \mathcal{D} is planar.*

We include the proof for completeness:

Proof Realize \underline{w} by an embedded multicurve on the surface of a handlebody, which lifts to a collection of disjoint arcs on the boundary surface of the universal cover of the handlebody. The universal cover is a thickened tree, and may be compactified by including the Cantor set boundary of the tree. The resulting space is a 3–ball with a collection of disjoint arcs in the bounding 2–sphere. By Moore's Decomposition Theorem (Theorem 3.3), collapsing each of these arcs to a point sends the 2–sphere to the 2–sphere, and the image of the Cantor set in the quotient is the decomposition space, embedded, non-surjectively, into S^2 . \square

Inclusion of a finite index subgroup induces a homeomorphism of decomposition spaces, so:

Corollary 5.4 *If \underline{w} is virtually geometric then \mathcal{D} is planar.*

Theorem 5.5 (cf [17, Theorem 1]) *Let (F, \underline{w}) be rigid. The following are equivalent:*

- (1) *The multiword \underline{w} is geometric.*
- (2) *The decomposition space $\mathcal{D}_{\underline{w}}$ is planar.*
- (3) *Every minimal Whitehead graph for \underline{w} has an embedding in the plane with consistent cyclic orderings of edges incident to inverse vertices.*

Proof Lemma 5.3 shows (1) implies (2).

(2) implies (3) is the content of [17, Lemma 4.4]. The hypotheses for this lemma are that every element of the multiword is ‘indecomposable’ and that the decomposition space embeds into S^2 in such a way that closures of the complementary regions intersect pairwise in at most one point. The first hypothesis is unnecessarily strong. Indecomposability of each element of the multiword is only used to prove that the decomposition space has no cut points. The second hypothesis is satisfied if the decomposition space has no cut pairs. Therefore, rigidity is a sufficient hypothesis.

(3) implies (1) by Proposition 5.2. □

Corollary 5.6 *Virtual geometricity implies geometricity for rigid multiwords.*

5.2 Non-rigid Multiwords and Virtual Geometricity

In this section we prove the Characterization of Virtual Geometricity (Theorem 5.9). We first (Theorem 5.7) prove the theorem in the special case that the decomposition space has no uncrossed cut pairs. Given Theorem 5.7, the proof of Theorem 5.9 amounts to showing that if the decomposition space is planar then uncrossed cut pairs can be pinched to cut points without making the space non-planar.

Theorem 5.7 *Assume that $\mathcal{D}_{\underline{w}}$ is connected with no uncrossed cut pairs. Let Γ be the JSJ decomposition of F relative to \underline{w} . The following are equivalent:*

- (1) *The multiword \underline{w} is virtually geometric.*
- (2) *The decomposition space $\mathcal{D}_{\underline{w}}$ is planar.*
- (3) *$\text{Ind}_{\Gamma}^G(\underline{w})$ is geometric for every non-cyclic vertex group G of Γ .*

Proof Corollary 5.4 shows (1) implies (2).

Let G be a non-cyclic vertex group of Γ . Since there are no uncrossed cut pairs, $\underline{w} = \text{Aug}_{\Gamma}(\underline{w})$, so Lemma 4.17 shows that the decomposition space $\mathcal{D}_{\text{Ind}_{\Gamma}^G(\underline{w})}$ embeds into $\mathcal{D}_{\underline{w}}$. Thus, if $\mathcal{D}_{\text{Ind}_{\Gamma}^G(\underline{w})}$ is non-planar then $\mathcal{D}_{\underline{w}}$ is non-planar as well. Therefore,

(2) implies that $\mathcal{D}_{\text{Ind}_{\Gamma}^G(\underline{w})}$ is planar. If $(G, \text{Ind}_{\Gamma}^G(\underline{w}))$ is a QH–surface then $\text{Ind}_{\Gamma}^G(\underline{w})$ is geometric. If $(G, \text{Ind}_{\Gamma}^G(\underline{w}))$ is rigid and $\mathcal{D}_{\text{Ind}_{\Gamma}^G(\underline{w})}$ is planar Theorem 5.5 says $\text{Ind}_{\Gamma}^G(\underline{w})$ is geometric. Thus, (2) implies (3).

Now assume (3). From a graph of groups we may build a corresponding graph of spaces [19]: For each vertex group choose a vertex space with fundamental group isomorphic to the vertex group. For each edge group choose a space with fundamental group isomorphic to the edge group, and let the edge space be the product of that space with the unit interval. Use the edge injections of the graph of groups to define attaching maps of edge spaces to the corresponding vertex spaces. The resulting space will have fundamental group isomorphic to the fundamental group of the graph of groups.

For each non-cyclic vertex group, the induced multiword is geometric, so we can choose the vertex space to be a handlebody with an embedded multicurve in the boundary representing the induced multiword.

For the edge spaces we use annuli. Later we will want to thicken them to make the resulting graph of spaces a 3–manifold.

For the moment we will also make a geometricity assumption on the cyclic vertex groups. Suppose one of the following possibilities are true for each cyclic vertex group $\langle g \rangle$:

- There are some number k of incident edges and each edge injection is degree one. In this case we choose the vertex space to be a solid torus with $k + 1$ disjoint curves on the boundary, one representing the element g and k to be attaching curves to which we will glue boundary curves of annulus edge spaces.
- There are some number k of incident edges and the degrees of the edge injections are all two except for possibly one of degree one. In this case we choose the vertex space to be a solid Klein bottle, and again we have $k + 1$ disjoint curves on the boundary representing g and the attaching curves.

The resulting graph of spaces has fundamental group F and has an embedded multicurve representing \underline{w} such that the multicurve is disjoint from the edge spaces. It is not yet a 3–manifold with boundary; we need to fatten the annuli. To see if this is possible, consider for each boundary component of each annulus a small tubular neighborhood of the attaching curve in the boundary of the corresponding handlebody. If for each annulus the two neighborhoods are either both annuli or both Möbius strips then the annuli may be fattened to make the graph of spaces a 3–manifold. Now, a fattened annulus is composed of a 1–handle and a 2–handle, so this does not explicitly give the resulting space a handlebody structure. However, a graph of aspherical spaces is

aspherical [19], and a compact aspherical 3–manifold with free fundamental group is a handlebody [10], so this space really is a handlebody, and \underline{w} is geometric.

Thus, assuming (3), there are two possible obstructions to geometricity:

- (a) The degrees of the edge injections into some cyclic vertex group are not of one of the two forms described above.
- (b) Some annulus cannot be fattened because one boundary neighborhood is an annulus and the other is a Möbius strip.

Claim 5.7.1 These obstructions vanish in a finite index subgroup of F .

Lift \underline{w} to this finite index subgroup, and then apply the graph of spaces construction to see (3) \implies (1).

Proof of Claim. There are finitely many elements $g_i \in \underline{w}$ such that $\Delta(\{\bar{g}_i^\infty, g_i^\infty\})$ is a cut point in \mathcal{D}_w .

From the proof of Proposition 4.13, an edge injection of degree greater than one into a cyclic vertex group $\langle g_i \rangle$ occurs when the g_i –action permutes some components of $\mathcal{D}_w \setminus \Delta(\{\bar{g}_i^\infty, g_i^\infty\})$. There are only finitely many components, so there exists some minimal positive power a_i of g_i such that the $g_i^{a_i}$ –action fixes each complementary component. Additionally, if some edge incident to the $\langle g_i \rangle$ vertex attaches to a handlebody around a non-orientable handle, and if a_i is odd, then consider $g_i^{2a_i}$.

Marshall Hall’s Theorem implies there exists a finite index subgroup H_i of F in which $g_i^{a_i}$ (alternatively, $g_i^{2a_i}$) generates a free factor. Let H be the finite index subgroup $\cap_i H_i$. If we apply the Relative JSJ-Decomposition Theorem to H we get a graph of groups covering the graph of groups decomposition for F . By construction, the smallest power of g_i in H is a multiple of $g_i^{a_i}$, so all edge inclusions are degree one. This takes care of obstruction (a), and we can choose all the cyclic vertex spaces to be solid tori.

Furthermore, we can take the vertex spaces to be handlebodies finitely covering the original handlebodies. If some attaching curve in the original decomposition ran along a Möbius strip then it runs along an even covering of the Möbius strip in the covering handlebodies. Thus, all attaching curves have annulus neighborhoods, which takes care of obstruction (b). ◇ □

We would now like to show that if \mathcal{D}_w is planar then $\mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$ is planar. To get $\mathcal{D}_{\text{Aug}_\Gamma(\underline{w})}$ from \mathcal{D}_w we must pinch uncrossed cut pairs to points. To make sure planarity

is preserved we first embed \mathcal{D}_w in a sphere and then find an upper semi-continuous collection of arcs so that collapsing the arcs achieves the pinching of the uncrossed cut pairs. It will suffice to find such a collection of arcs for an arbitrary vertex group of the rJSJ:

Lemma 5.8 *If \mathcal{D}_w is planar then for each non-cyclic vertex group G of the rJSJ, the decomposition space $\mathcal{D}_{\text{Ind}_F^G(w)}$ is planar.*

Proof Let $\chi: \mathcal{D}_w \hookrightarrow S^2$ be an embedding. Let $\{S_i\}_{i \in I}$ be the collection of uncrossed cut pairs of \mathcal{D}_w in $\Delta(\partial G)$. Let $S_i = \{x_{i,0}, x_{i,1}\}$. Since G is a vertex group of the rJSJ, for each $i \in I$ all of $\Delta(\partial G) \setminus S_i$ is contained in a single complementary component C_i of S_i , for otherwise we could use Lemma 4.24 to find a refinement of the rJSJ. There is at least one other complementary component C'_i of S_i . By Proposition 3.24 and Lemma 2.1, there exists an arc $A_i \subset \overline{C'_i}$ connecting the two points of S_i . Note that $A_i \cap \Delta(\partial G) = S_i$. Recalling the terminology of Section 3.1:

Claim 5.8.1 The sets $\chi(A_i)$ are the non-degenerate elements of an upper semi-continuous decomposition of S^2 .

Assuming the claim, Moore's Decomposition Theorem (Theorem 3.3) then says that the quotient of the sphere obtained by collapsing each of these arcs to a point is again the sphere. The image of $\chi(\Delta(\partial G))$ in this quotient is $\mathcal{D}_{\text{Ind}_F^G(w)}$ embedded, non-surjectively, in S^2 . Thus, $\mathcal{D}_{\text{Ind}_F^G(w)}$ is planar, and the lemma is proven.

Proof of Claim 5.8.1. Essentially the proof is that each S_i separates $A_i \setminus S_i$ from all of the other arcs, so arcs can only be close at their endpoints, and we know the endpoints are well behaved because the boundary pattern of a multiword gives an upper semi-continuous decomposition of the boundary of the free group.

Formally, fix an $i \in I$ and let U be an open neighborhood of $\chi(A_i)$ in S^2 . We will produce an open neighborhood V of $\chi(A_i)$ such that $\chi(A_j) \cap V \neq \emptyset$ implies $\chi(A_j) \subset U$ for all $j \in I$.

By Lemma 4.8, for each j and k the set $\Delta^{-1}(x_{j,k})$ is a single point; call it $\xi_{j,k}$. Let U' be an open neighborhood of $\chi(A_i)$ in S^2 such that $\overline{U'} \subset U$.

Claim 5.8.2 $J = \{j \in I \mid \chi(S_j) \subset U' \text{ and } \chi(A_j) \not\subset U'\}$ is finite.

Proof of Claim. Suppose not. For each $j \in J$ choose a point $\chi(y_j) \in \chi(A_j) \setminus U$. There is some $\sigma: \mathbb{N} \hookrightarrow J$ such that $(\chi(y_{\sigma(k)}))$ is a convergent sequence, converging to a point $\chi(y) \in \chi(\mathcal{D}_{\underline{w}}) \setminus U$. By Proposition 3.10, y has a neighborhood basis in $\mathcal{D}_{\underline{w}}$ of connected neighborhoods $N(y, r)$. Thus, given r there is a K so that for all $k > K$ we have $y_{\sigma(k)} \in N(y, r)$. But $S_{\sigma(k)}$ is a cut set in $\mathcal{D}_{\underline{w}}$ separating $y_{\sigma(k)}$ from all of the other y_j , so if $y_{\sigma(k)}$ and y_j are both contained in the connected set $N(y, r)$ then so is at least one of the points of $S_{\sigma(k)}$. Thus, y is a limit point of $\cup_{j \in J} S_j$, so $\chi(y)$ is a limit point of $\chi(\cup_{j \in J} S_j)$. This is impossible, since $\cup_{j \in J} \chi(S_j) \subset U' \subset \overline{U'}$ and $\{\chi(y_j)\}_{j \in J} \subset S^2 \setminus U$ are contained in disjoint closed sets. \diamond

Let $U'' = U' \setminus (\cup_{j \in J} \chi(A_j))$. This is an open neighborhood of $\chi(A_i)$ contained in U . By Proposition 3.6, the decomposition of $\partial\mathcal{T}$ whose non-degenerate elements are elements of the boundary pattern for $\text{Aug}_{\Gamma}(\underline{w})$ is upper semi-continuous. Thus, for the neighborhood $\mathcal{U}'' = \Delta^{-1}(\chi^{-1}(\chi(\mathcal{D}_{\underline{w}}) \cap U'')) \subset \partial\mathcal{T}$ of $\{\xi_{i,0}, \xi_{i,1}\}$ there exists an open neighborhood $\mathcal{V} \subset \mathcal{U}''$ of $\{\xi_{i,0}, \xi_{i,1}\}$ so that if for some j and k we have $\xi_{j,k} \in \mathcal{V}$ then $\xi_{j,1-k} \in \mathcal{U}''$. As in Section 3.2.1, we may assume \mathcal{V} is saturated and $\Delta(\mathcal{V})$ has two components, one containing $x_{i,0}$ and the other containing $x_{i,1}$. Since χ is an embedding, there exists an open set $V' \subset S^2$ such that $V' \cap \chi(\mathcal{D}_{\underline{w}}) = \chi(\Delta(\mathcal{V}))$.

Let $W = S^2 \setminus \chi(C_i)$, and let $V = U'' \cap (V' \cup W)$.

Suppose, for some $j \neq i$, that $\chi(A_j) \cap V \neq \emptyset$. Since C_i is the complementary component of S_i containing A_j , we have $\chi(A_j) \cap W = \emptyset$, so $A_j \cap \Delta(\mathcal{V}) \neq \emptyset$. On the other hand, S_j separates S_i from $A_j \setminus S_j$, and $\Delta(\mathcal{V})$ consists of connected neighborhoods of the two points of S_i , so if $A_j \cap \Delta(\mathcal{V}) \neq \emptyset$ then one of the $x_{j,k}$ must be in $\Delta(\mathcal{V})$. By definition of \mathcal{V} , this means that $\{\xi_{j,0}, \xi_{j,1}\} \subset \mathcal{U}''$. By definition of \mathcal{U}'' , this means that $\chi(A_j) \subset U$. \diamond \square

Theorem 5.9 (Characterization of Virtual Geometricity) *Assume that $\mathcal{D}_{\underline{w}}$ is connected. Let Γ be the JSJ decomposition of F relative to \underline{w} . The following are equivalent:*

- (1) *The multiword \underline{w} is virtually geometric.*
- (2) *The decomposition space $\mathcal{D}_{\underline{w}}$ is planar.*
- (3) *For every non-cyclic vertex group G of Γ , the induced multiword $\text{Ind}_{\Gamma}^G(\underline{w})$ is geometric.*

Thus, virtually geometric multiwords are those that are built from geometric pieces.

Proof Corollary 5.4 shows (1) implies (2).

If the decomposition space is planar then Lemma 5.8 shows that the induced decomposition spaces of each of the vertex groups in the rJSJ are planar. Theorem 5.5 and the fact that QH–surface multiwords are always geometric show that all of the induced multiwords are geometric. Thus, (2) implies (3).

By Theorem 5.7, if the induced multiword in each non-cyclic vertex group of the rJSJ is geometric then $\text{Aug}_\Gamma(\underline{w})$ is virtually geometric. This implies \underline{w} is virtually geometric, since \underline{w} is a subset of $\text{Aug}_\Gamma(\underline{w})$, so (3) implies (1). \square

5.3 Examples

5.3.1 Baumslag’s Word

$w = \bar{a}^2 \bar{b} \bar{a} b \bar{a} b$ in $F_2 = \langle a, b \rangle$ is known as *Baumslag’s word*. In response to a question of Gordon and Wilton, Manning showed, by brute force, that this word becomes geometric in an orientable handlebody with fundamental group an index four subgroup of F_2 .

The rJSJ for $F = \langle a, b \rangle \cong \langle a, b, c \mid c = \bar{b}ab \rangle$ is shown in Figure 17.




Figure 17: rJSJ-Decomposition of $\langle a, b \rangle$ for $\bar{a}^2 \bar{b} \bar{a} b \bar{a} b$ ($c = \bar{b}ab$)

The word w becomes $\bar{a}^2 \bar{c} a c$ when rewritten in the rank two vertex group, so the induced multiword is $\{\bar{a}^2 \bar{c} a c, a, c\}$. One can check that this multiword is rigid, so this is the rJSJ. (Checking rigidity takes some work, using techniques of [4].)

In Figure 18 we have a reduced Whitehead graph/Heegaard diagram for the induced multiword that shows it is geometric. Figure 19 shows a (non-orientable) handlebody with embedded multicurve representing $\{\bar{a}^2 \bar{c} a c, a, c\}$.

The obstruction to geometricity of w is that the curve representing c runs around a orientable handle, while the curve representing a does not. We cannot achieve the conjugation of a to c by a fattened annulus.

To correct this problem, pass to the index two subgroup $G = \langle A = a^2, b, B = ab\bar{a} \rangle$.




Figure 18: Whitehead graph/Heegaard diagram Figure 19: Corresponding non-orientable handlebody for $\{a^2 \bar{c} a c, a, c\}$

After applying the automorphism that sends B to bB and fixes b and A , the image of w^2 is $\bar{A}(\bar{b}A\bar{b})BAB\bar{A}\bar{B}^2(\bar{b}A\bar{b})$.

The splitting over $\langle A \rangle$ is an HNN extension with b conjugating A to C , and the induced multiword in the vertex group $\langle A, B, C \rangle$ is $\{A, C, \bar{A}\bar{C}B\bar{A}B\bar{A}\bar{B}^2C\}$, which is geometric in a non-orientable handlebody, as seen in Figure 20.




Figure 20: A non-orientable handlebody for $\{A, C, \bar{A}\bar{C}B\bar{A}\bar{B}^2C\}$

This time we can build a 3-manifold graph of spaces because we only need to conjugate words that run around orientable handles. Gluing on a fattened annulus conjugating A to C gives a non-orientable handlebody with fundamental group isomorphic to G for which the image of w^2 is geometric.

5.3.2 Baumslag-Solitar Words

Another interesting family of examples is given by the Baumslag-Solitar words $w_{p,q} = \bar{a}^q \bar{b} a^p b$ in $F_2 = \langle a, b \rangle$. We will assume that $0 < p \leq q$. Gordon and Wilton [8] have shown that $w_{p,q}$ is virtually geometric when p and q are relatively prime.

The decomposition space associated to $w_{p,q}$ is connected without cut points. The pair $\Delta(\{\bar{a}^\infty, a^\infty\})$ is a cut pair. $\mathfrak{W}([\bar{a}^\infty, a^\infty])$ has components $\mathfrak{C}_j^+ = \{\Theta(a^i b) \mid i \equiv j \pmod p\}$ for $0 \leq j < p$ that are cyclically permuted by the a -action, and components $\mathfrak{C}_j^- = \{\Theta(a^i \bar{b}) \mid i \equiv j \pmod q\}$ for $0 \leq j < q$ that are cyclically permuted by the a -action.

The case $p = q = 1$ is special; in this case the Whitehead graph is a circle, which implies the decomposition space is a circle and the word is geometric.

Otherwise, the number of complementary components is $p+q > 2$, so that $\Delta(\{\bar{a}^\infty, a^\infty\})$ is an uncrossed cut pair. The rJSJ for this case is shown in Figure 21.




Figure 21: rJSJ-Decomposition of $\langle a, b \rangle$ for $\bar{a}^q \bar{b} a^p b$

The rank two vertex group is $\langle A = a^q, C = \bar{b} a^p b \rangle$, and the induced multiword in this vertex group is $\{A, C, \bar{A}C\}$. The Whitehead graph for this multiword is a circle, which implies the vertex decomposition space is a circle and the induced multiword is geometric. Thus, Theorem 5.9 says $w_{p,q}$ is at least virtually geometric.

The cyclic vertex group has edge inclusions of degrees p and q .

If $p = 1$ and $q = 2$ we can make this geometric by using a solid Klein bottle for the cyclic vertex space. (We saw the non-cyclic vertex space for this example back in Example 5.1.)

If $p = q$ the word is also geometric, because two disjoint degree p curves fit into the boundary of a solid torus.

In all other cases, the word $w_{p,q}$ is not geometric. Virtual geometricity can be verified by passing to the index $m = \text{lcm}(p, q)$ subgroup:

$$G = \langle A, B_0, B_1, \dots, B_{m-1} \mid A = a^m, B_i = a^i \bar{b} a^i \rangle$$

References

- [1] John Berge, *Documentation for the program Heegaard*, circa 1990, available at <http://www.math.uic.edu/~t3m>.

- [2] Brian H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186.
- [3] Brian H. Bowditch, *Peripheral splittings of groups*, Trans. Amer. Math. Soc. **353** (2001), no. 10, 4057–4082.
- [4] Christopher H. Cashen and Nataša Macura, *Line patterns in free groups*, Geom. Topol. **15** (2011), no. 3, 1419–1475.
- [5] Christopher H. Cashen and Jason F. Manning, *virtuallygeometric*, 2014, computer program, https://bitbucket.org/christopher_cashen/virtuallygeometric.
- [6] Christopher H. Cashen and Jason F. Manning, *Virtual geometricity is rare*, LMS J. Comput. Math. **18** (2015), no. 1, 444–455.
- [7] Heiner Cišang, *Simple path systems on full pretzels*, Mat. Sb. (N.S.) **66 (108)** (1965), 230–239, translated in AMS Translations. Series 2. Vol. 92: Seventeen papers on topology and differential geometry.
- [8] Cameron Gordon and Henry Wilton, *On surface subgroups of doubles of free groups*, J. Lond. Math. Soc. (2) **82** (2010), no. 1, 17–31.
- [9] Vincent Guirardel and Gilbert Levitt, *JSJ decompositions: definitions, existence, uniqueness. I: The JSJ deformation space*, preprint (2009), arXiv:0911.3173.
- [10] John Hempel, *3-Manifolds*, Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, 1976, Ann. of Math. Studies, No. 86.
- [11] John G. Hocking and Gail S. Young, *Topology*, second ed., Dover Publications, Inc., New York, 1988.
- [12] Olga Kharlampovich and Alexei G. Myasnikov, *Effective JSJ decompositions*, Groups, languages, algorithms, Contemp. Math., vol. 378, Amer. Math. Soc., Providence, RI, 2005, pp. 87–212.
- [13] Jason Fox Manning, *Virtually geometric words and Whitehead’s algorithm*, Math. Res. Lett. **17** (2010), no. 5, 917–925.
- [14] Jason Fox Manning, *The Bowditch boundary of (G, \mathcal{H}) when G is hyperbolic*, preprint (2015), arXiv:1504.03630.
- [15] Reiner Martin, *Non-uniquely ergodic foliations of thin type, measured currents and automorphisms of free groups*, Ph.D. thesis, UCLA, 1995.
- [16] R. L. Moore, *Concerning upper semi-continuous collections of continua*, Trans. Amer. Math. Soc. **27** (1925), no. 4, 416–428.
- [17] Jean-Pierre Otal, *Certaines relations d’équivalence sur l’ensemble des bouts d’un groupe libre*, J. London Math. Soc. **s2-46** (1992), no. 1, 123–139.
- [18] E. Rips and Z. Sela, *Structure and rigidity in hyperbolic groups. I*, Geom. Funct. Anal. **4** (1994), no. 3, 337–371.

- [19] Peter Scott and Terry Wall, *Topological methods in group theory*, Homological group theory (Proc. Sympos., Durham, 1977), London Math. Soc. Lecture Note Ser., vol. 36, Cambridge Univ. Press, Cambridge, 1979, pp. 137–203.
- [20] Jean-Pierre Serre, *Trees*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003, Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.
- [21] Arthur H. Stone, *Metrizability of decomposition spaces*, Proc. Amer. Math. Soc. **7** (1956), 690–700.
- [22] J. H. C. Whitehead, *On equivalent sets of elements in a free group*, Ann. of Math. (2) **37** (1936), no. 4, 782–800.
- [23] Stephen Willard, *General topology*, Dover Publications, Inc., Mineola, NY, 2004, Reprint of the 1970 original [Addison-Wesley, Reading, MA].

Fakultät für Mathematik, Universität Wien
1090 Wien, Österreich

`christopher.cashen@univie.ac.at`

`http://www.mat.univie.ac.at/~cashen`