Line Patterns in Free Groups

CHRISTOPHER H. CASHEN
NATAŠA MACURA

We study line patterns in a free group by considering the topology of the decomposition space, a quotient of the boundary at infinity of the free group related to the line pattern. We show that the group of quasi-isometries preserving a line pattern in a free group acts by isometries on a related space if and only if there are no cut pairs in the decomposition space. We also give an algorithm to detect such cut pairs.

20F65; 20E05

1 Introduction

Given a finitely generated free group F of rank greater than one and a word $w \in F$, the w-line at $g \in F$ is the set of elements $\{gw^m\}_{m\in\mathbb{Z}}$. A Cayley graph with respect to a free basis of F is a geometric model for F that is a tree. Up to translation and coarse equivalence, we may assume that w is cyclically reduced and not a power of another element, and in this case there is a unique geodesic in the tree that contains the vertices $\{gw^m\}_{m\in\mathbb{Z}}$.

The w-line at g is the same as the w line at h if and only if $\bar{h}g$ is a power of w; the w-lines are the cosets of $\langle w\rangle$ in F.

The *line pattern generated by* w is the collection of distinct w-lines. Similarly, if we take finitely many words w, as above, the line pattern generated by the collection is the union of the patterns generated by the individual words. We will denote the line pattern $\mathcal{L}$ when we do not wish to specify generators.

The main question is:

Question 1 *Let F and F' be finite rank free groups, possibly of different ranks. Consider collections of words $\{w_1, \dots, w_m\} \subset F$ and $\{w'_1, \dots, w'_n\} \subset F'$. Let $\mathcal{L}$ be the line pattern in F generated by $\{w_1, \dots, w_m\}$, and let $\mathcal{L}'$ be defined similarly for F'.*

[0]Published in *Geometry & Topology*, 15 (2011), no. 3, 1419–1475.

Is there a quasi-isometry $\phi\colon F \to F'$ that preserves the patterns, in the sense that there is some constant C so that for every line $l \in \mathcal{L}$ there is an $l' \in \mathcal{L}'$ such that the Hausdorff distance between $\phi(l)$ and l' is at most C, and vice versa?

A closely related question is:

Question 2 *Let F be a free group and $\mathcal{L}$ a line pattern in F. What is the group $\mathcal{QI}(F,\mathcal{L})$ of quasi-isometries of F that preserve the line pattern $\mathcal{L}$?*

In a pair of 1936 papers [21, 22], J. H. C. Whitehead gave an algorithm to answer the following question:

Given two finite (ordered) lists of words $(w_1,\dots,w_k)$ and $(w'_1,\dots w'_k)$ in a finite rank free group F, is there an automorphism ϕ of F such that for all i we have $\phi(w_i) = w'_i$?

Questions 1 and 2 may be viewed as geometric versions of Whitehead's question.

To motivate the statement of our results, it is instructive to consider line patterns in a different setting.

Line patterns in $\mathbb{H}^n$ for $n \geq 3$ have been studied by Schwartz [18]. His terminology is "symmetric pattern of geodesics". Let M be a compact hyperbolic orbifold of dimension $n \geq 3$. Pick any collection of closed geodesics in M. The lifts of these geodesics to the universal cover $\mathbb{H}^n$ form a line pattern; call it $\mathcal{L}$.

Theorem [18, Theorem 1.1]

$$\mathcal{QI}(\mathbb{H}^n,\mathcal{L}) \subset \mathrm{Isom}(\mathbb{H}^n)$$

This is an example of what we will call *pattern rigidity*. The hyperbolic orbifold case is special in that there is a canonical geometric model, $\mathbb{H}^n$, for $\pi_1 M$. Forgetting this for a moment, let Y be any geometric model for $\pi_1 M$. For example, Y could be a Cayley graph of $\pi_1 M$. We still get a line pattern $\mathcal{L}$ in Y, but it is not necessarily true that $\mathcal{QI}(Y,\mathcal{L}) \subset \mathrm{Isom}(Y)$. However, there is a quasi-isometry $\phi\colon Y \to \mathbb{H}^n$. Each line in $\mathcal{L}$ gets sent to a line in $\mathbb{H}^n$, so we get a line pattern $\phi(\mathcal{L})$ in $\mathbb{H}^n$. We have:

$$\phi\,\mathcal{QI}(Y,\mathcal{L})\phi^{-1} = \mathcal{QI}(\mathbb{H}^n,\phi(\mathcal{L})) \subset \mathrm{Isom}(\mathbb{H}^n)$$

In the free group situation we do not have a canonical space to take the place of $\mathbb{H}^n$ that works for every line pattern. For a given line pattern we will try to construct a space X and a quasi-isometry $\phi\colon F \to X$ such that pattern preserving quasi-isometries are conjugate into the isometry group of X:

$$\phi\,\mathcal{QI}(F,\mathcal{L})\phi^{-1} = \mathcal{QI}(X,\phi(\mathcal{L})) \subset \mathrm{Isom}(X)$$

A priori this would only give a quasi-action of $\mathcal{QI}(F,\mathcal{L})$ on X by maps bounded distance from isometries. We actually prove something stronger. We will say a line pattern $\mathcal{L}$ in F is *rigid* if there is a space X, a quasi-isometry $\phi\colon F \to X$, and an isometric action of $\mathcal{QI}(F,\mathcal{L})$ on X that agrees with conjugation by ϕ, up to bounded distance.

It is easy to see that not all patterns are rigid. A necessary condition is that the multiplicative quasi-isometry constants of $\mathcal{QI}(F,\mathcal{L})$ are bounded. Suppose $\mathcal{L}$ is contained in a proper free factor F' of F, so that $F = F' * F''$. Then $\mathcal{QI}(F'') \subset \mathcal{QI}(F,\mathcal{L})$ contains a sequence of quasi-isometries with unbounded constants, so the pattern is not rigid.

Another example where the lack of rigidity is apparent for algebraic reasons is the pattern generated by the word $ab\bar{a}\bar{b}$ in $F_2 = \langle a,b\rangle$. The automorphism group of F_2 preserves this line pattern, so again we have a sequence of pattern preserving quasi-isometries with unbounded constants.

However, algebraic considerations do not fully determine which patterns are rigid. Consider the pattern in F_2 generated by ab and $a\bar{b}$. There is only a finite group of outer automorphisms of F_2 that preserve this pattern, so all pattern preserving automorphisms are isometries, up to bounded distance. We might guess the pattern is rigid, but in fact it is quasi-isometrically equivalent to the $ab\bar{a}\bar{b}$ pattern, see Theorem 6.2.

Our main result shows that sufficiently complicated patterns are rigid. To make this precise, we use a topological space that is a quotient of the boundary at infinity of a tree for F. This space is called the *decomposition space* associated to the line pattern.

Main Theorem *Let $\mathcal{L}$ be a line pattern in a finitely generated, non-abelian free group, F. The following are equivalent:*

(1) *The line pattern is rigid.*

(2) *The decomposition space has no cut pairs.*

Remark We use the phrase "has no cut pairs" inclusively to mean that the space is connected, has no cut points and no cut pairs.

In the examples above, the pattern that is contained in a proper free factor would have a disconnected decomposition space. For the other two, the decomposition space is a circle.

Determining if the decomposition space is connected is essentially Whitehead's Algorithm, which is discussed in Section 3. The idea is to build a graph, the Whitehead

graph, associated to the line pattern. Connectivity of this graph is related to connectivity of the decomposition space, see Theorem 4.1.

In Section 4 we use generalizations of the Whitehead graph to identify finite cut sets in the decomposition space. In particular, Theorem 4.17 allows us to tell if there are cut pairs in the decomposition space.

We show in Section 5.1 that there is no possibility of rigidity if the decomposition space has cut pairs.

The proof that no cut pairs implies rigidity in Section 5.2 is similar in philosophy to the various geometric proofs of Stallings' Theorem, see Dunwoody [6], Gromov [7], Niblo [13] or Kapovich [8]. The idea in these proofs is to use minimal surfaces, or a combinatorial approximation thereof, to cut up a space into pieces. One then uses properties of the particular choice of surfaces to show that they are, or can be chosen to be, suitably independent, so that the complex dual to the cutting surfaces is a tree. The collection of minimal surfaces is in bijection with the edges of the resulting tree.

We do something similar with small cut sets in the decomposition space. With an appropriately chosen basis for F we can look at the tree $\mathcal{T}$ that is the Cayley graph for F with respect to this basis. We find cut sets in the decomposition space associated to the edges of this tree, which we call *edge cut sets*. The property of being an edge cut set depends on the choice of basis, so we study some topological properties enjoyed by edge cut sets and choose our definition of "small" to be all cut sets that also enjoy those properties. (See Section 5.2.1 for precise definitions.)

The cut sets we use have more complicated interactions than those in Stallings' Theorem, and in general the space dual to the collection of cut sets will not be a tree; it will be a cube complex quasi-isometric to a tree.

Working at infinity with topologically defined cut sets has the benefit that the cube complex we construct is canonical and inherits a canonical line pattern. $\mathcal{QI}(F,\mathcal{L})$ is conjugate to the group of isometries of the cube complex that preserve the line pattern, see Theorem 5.5.

This allows us to answer Questions 1 and 2 in the rigid case: Two line patterns in free groups are equivalent if and only if there is a pattern preserving isometry between the associated cube complexes. The free group F acts cocompactly by pattern preserving isometries on the cube complex, so $\mathcal{QI}(F,\mathcal{L})$ does as well. This allows us to give a description of $\mathcal{QI}(F,\mathcal{L})$ as a complex of groups. However, the vertex stabilizers will not, in general, be finitely generated groups.

Consideration of line pattern preserving quasi-isometries arises naturally in Geometric Group Theory. Work of Papasoglu [16] shows that group splittings of finitely presented groups over virtually cyclic subgroups are preserved by quasi-isometries. If a finitely presented, one-ended group has a non-trivial JSJ–decomposition over virtually cyclic subgroups, then each vertex group of the decomposition has a line pattern coming from the incident edge groups. The equivalence classes of these line patterns give quasi-isometry invariants for the group, and, in the rigid case, impose severe restrictions on quasi-isometries of the group.

In particular, the authors came upon this problem in the course of studying mapping tori of free group automorphisms. In the case of a linearly growing automorphism, the mapping torus has a JSJ-decomposition with vertex groups $F \times \mathbb{Z}$. Understanding the line patterns in the free factors of the vertex groups is a key step in the quasi-isometry classification of these mapping tori [5].

1.1 Remarks on Splittings

In a subsequent work [4], the first author has shown that cut pairs can be used to produce a JSJ-decomposition of the free group relative to the line pattern. The construction is analogous to Bowditch's JSJ-decomposition for hyperbolic groups [2], using the decomposition space in place of the boundary at infinity of the hyperbolic group. In this light the *shearing quasi-isometries* that we produce in Section 5.1 may be thought of as Dehn twists.

In the theory of JSJ-decompositions it is common to call vertices that do not split further *rigid vertices*, so it might seem natural to define a line pattern to be rigid if there are no relative splittings. The relative JSJ-decomposition shows that this is almost equivalent to the isometric action definition of rigidity used in the present work. That existence of a splitting implies non-rigidity is easy to see. The relative JSJ-decomposition shows that the converse is almost true; the converse fails only in the case that the line pattern is generated by the boundary words of a sphere with three boundary components. It can be shown this example does not split. However, we will see in Section 6.1 that the line pattern coming from the boundary words of a surface with boundary gives a decomposition space that is a circle, and therefore is non-rigid in the isometric action sense.

2 Preliminaries

In this section we recall relevant background material. Results are stated with references, and proofs are omitted.

2.1 Cut Sets and Cubings

If X is a topological space, a *cut set* is a subset $S \subset X$ such that $X \setminus S = \{x \in X \mid x \notin S\}$ is disconnected. A single point that is a cut set is a *cut point*; a pair of points that is a cut set is a *cut pair*, etc.

A cut set S is *minimal* if no proper subset of S is a cut set of X.

If S and S' are cut sets of X we say S' *crosses* S if $S' \setminus S$ has points in multiple components of $X \setminus S$. This is not a symmetric relation, but it is if we assume that S and S' are minimal.

A *cubing* is a simply connected, non-positively curved cube complex. Cubings can be used to encode the combinatorics of a collection of cut sets. Our treatment of cubings is based on work of Sageev [17].

Let $\{S_i\}_{i\in I}$ be a collection of closed, minimal cut sets of X so that for each i, $X \setminus S_i$ has exactly two connected components, A_i^0 and A_i^1. We will take the superscripts mod 2, so that the two components of $X \setminus S_i$ are A_i^ϵ and $A_i^{1+\epsilon}$ for $\epsilon \in \{0,1\}$. Let

$$\Sigma = \{A_i^0\}_{i\in I} \cup \{A_i^1\}_{i\in I}$$

Define a cube complex as follows. The vertices are the subsets V of Σ such that:

(1) For all $i \in I$ exactly one of A_i^0 or A_i^1 is in V.

(2) If $C \in V$ and $C' \in \Sigma$ with $C \subset C'$ then $C' \in V$.

Two vertices are connected by an edge if they differ by only one set in Σ.

One can identify Σ with 2^I. The i-th "coordinate" is either 0 for A_i^0 or 1 for A_i^1. Edges join vertices that differ in exactly one coordinate.

The vertices are the elements of 2^I that are "consistent" with the cut set structure in the sense that if for some i and j we have $A_i^1 \subset A_j^1$ then we do not have any vertices that are "1" in the i-th coordinate and "0" in the j-th coordinate. It is not consistent to be simultaneously in A_i^1 and A_j^0.

Informally, having ϵ in the i-th coordinate corresponds to being in A_i^ϵ. There is a subtlety here, though. An element of 2^I might be consistent without being realized as a component of $X \setminus \{S_i\}$. It is possible that there are vertices such that ϵ_i is the value of the i-th coordinate of the vertex, but $\cap_{i\in I} A_i^{\epsilon_i} = \emptyset$.

Remark There is a minor difference from Sageev's contruction. In his notation we would be considering $A_i = S_i \cup A_i^0$ and $A_i^c = A_i^1$. The nature of the cut sets we are interested in would make it problematic to include them in one of the components. There is only one place where this requires us to change Sageev's arguments, which we will point out shortly. Everywhere else, it is sufficient to replace a statement like:

$$A_i \subset A_j \implies A_j^c \subset A_i^c$$

with a statement like:

$$A_i^{\epsilon_i} \subset A_j^{\epsilon_j} \implies A_j^{1+\epsilon_j} \subset A_i^{1+\epsilon_i}$$

This statement follows easily from the fact that minimal cut sets are either mutually crossing or mutually non-crossing.

Edges in the complex correspond to changing one coordinate from 0 to 1, or vice versa. However, to maintain consistency not every coordinate can be changed:

Lemma 2.1 ([17, Lemma 3.2]) *If V is a vertex and $A_i^\epsilon \in V$ then*

$$W = (V \setminus \{A_i^\epsilon\}) \cup \{A_i^{1+\epsilon}\}$$

is a vertex if and only if A_i^ϵ is minimal in V, in the sense that A_i^ϵ does not contain any other $A_j^\delta \in V$.

It turns out in general that there are still too many vertices. The graph that has been constructed so far is not necessarily connected. This is where our construction differs from Sageev's. For both his construction and ours, the idea is to select a subcollection of the vertices, show that the subcollection belongs to a path connected subset of the graph, and then throw away everything not in that path component. Our construction will come later in Section 5. However, this is the only place in which Sageev uses the special properties of his chosen collection Σ. The rest of his arguments go through unchanged in our setting.

So assume that we have passed to a non-trivial path connected component of the original graph. Following Sageev again, one glues in one square (2 dimensional cube) whenever one sees the boundary of a square in the graph. One proceeds by induction to glue in an n–cube whenever one sees the boundary of an n–cube in the $(n-1)$–skeleton

of the complex. The result is a (possibly infinite dimensional) simply connected, non-positively curved cube complex, a cubing [17, Theorem 3.7].

There is an equivalence relation on the (directed) edges of a cubing. Two directed edges e and e' are equivalent if there is a finite sequence $e = e_0, e_1, \dots, e_k = e'$ such that for each i, e_i and e_{i+1} are opposite edges of some 2–cube, oriented in the same direction.

Equivalence classes of edges are called *combinatorial hyperplanes*. There is a corresponding idea of a *geometric hyperplane*. Consider an n–cube of the complex. It can be identified with a cube of side length 1 in $\mathbb{R}^n$ where the vertices have all coordinates in $\{\pm\frac{1}{2}\}$. Consider the edges that correspond to changing the n-th coordinate from $-\frac{1}{2}$ to $\frac{1}{2}$. These edges belong to a combinatorial hyperplane. The corresponding portion of a geometric hyperplane is the intersection of the n–cube with the coordinate hyperplane $\{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_n = 0\}$. Such pieces are then glued together for each cube with edges in the combinatorial hyperplane.

Theorem 2.2 [17, Theorem 4.10] *Suppose J is a geometric hyperplane in a cubing Y. Then J does not intersect itself and partitions Y into two connected components.*

We take the metric on the cubing to be the path metric on the 1-skeleton. The distance between two vertices is the minimal number of edges in an edge path joining them, and such a minimal edge path is called a *geodesic*.

A corollary of the preceding theorem is the following observation about geodesics: Let x and y be vertices in a cubing Y. If they are distance D apart, then a geodesic joining them must cross D geometric hyperplanes, one through the midpoint of each edge of the path. Each of these hyperplanes disconnects Y, with x and y in opposite components. Therefore, any geodesic from x to y must cross the same D hyperplanes. Conversely, the distance between x and y in Y is the number of hyperplanes separating them.

Fix a hyperplane. There is an $A_i^\epsilon \in \Sigma$ such that every directed edge e in the hyperplane joins a vertex V_e with a vertex $(V_e \setminus \{A_i^\epsilon\}) \cup \{A_i^{1+\epsilon}\}$. Furthermore, every edge of this form belongs to the hyperplane [17, Lemma 3.9].

Thus, we have a bijection between the set of geometric hyperplanes and the collection $\{S_i\}$ of cut sets. This is how the cubing encodes the collection of cut sets. Cut sets of X correspond to hyperplanes of Y. Distance in Y corresponds to being separated by a given number of cut sets. An n–cube in Y corresponds to a collection of n distinct, pairwise crossing cut sets S_i in X.

2.2 Graphs and Complexes of Groups

In this section we give a brief account of graphs and complexes of groups. The reader is referred to Bridson and Haefliger's book [3] for more detail.

A *graph of groups* is a construction that builds a group by amalgamating smaller groups. Start with a finite connected graph Γ, and associate to each vertex or edge γ a local group G_γ, along with injections $\phi_{e,v}\colon\ G_e \to G_v$ for each edge e and vertex v that is an endpoint of e.

The *fundamental group of the graph of groups* is then obtained by taking as generators all the vertex groups as well as one generator g_e for each edge e in the graph. The relations are:

(1) all the relations from the vertex groups,

(2) for each edge e with endpoints v and v', and for each $h \in G_e$,

$$g_e \phi_{e,v}(h) g_e^{-1} = \phi_{e,v'}(h),$$

(3) $g_e = 1$ for each edge e in a chosen maximal subtree of Γ.

The fundamental group does not depend on the choice of maximal subtree.

Associated to a graph of groups there is a simplicial tree $D\Gamma$ covering Γ called the Bass-Serre tree or the *development* of the graph of groups. The fundamental group of the graph of groups acts by isometries on $D\Gamma$, with vertex stabilizers equal to conjugates of the vertex groups in the graph of groups and edge stabilizers equal to conjugates of the edge groups.

Conversely, given a cocompact isometric action of a group G on a simplicial tree we get a graph of groups decomposition for G by taking the graph to be the quotient of the tree by the G action and choosing the local groups to be the vertex and edge stabilizers.

A *complex of groups* is a generalization of the graph of groups construction to higher dimensional complexes. In particular, a group acting cocompactly by isometries on a polyhedral complex can be given a complex of groups structure by associating to each cell in the quotient a group isomorphic to the stabilizer of the cell in the original complex.

Unlike in the graph of groups case, not every complex of groups is developable. That is, starting with a complex of groups Γ, there may not exist a complex X so that the fundamental group of the complex of groups acts on X with quotient Γ. However, if you start with a group acting on a polyhedral complex, then the resulting graph of

groups is developable: the development is just the polyhedral complex that you started with.

A developable complex of groups is *faithful* if no non-trivial element of the fundamental group of the complex of groups acts trivially on the development.

To ensure that the quotient is still a polyhedral complex, one should assume that if an element of the group leaves a cell invariant, then it fixes it pointwise. This is called an action *without inversions*. If this is not the case, it can be achieved by subdividing cells.

Lim and Thomas have worked out a covering theory for complexes of groups [9]. A particular result that will be of interest to us is:

Theorem 2.3 [9, Theorem 4] *Let X be a simply connected polyhedral complex, and let G be a subgroup of* $\mathrm{Aut}(X)$ *(acting without inversions) that induces a complex of groups Γ. Then there is a bijection between the set of subgroups of* $\mathrm{Aut}(X)$ *(acting without inversions) that contain G, and the set of isomorphism classes of coverings of faithful, developable complexes of groups by Γ.*

If G acts cocompactly on X then so does any subgroup H of $\mathrm{Aut}(X)$ containing G, and we get a covering of the compact quotient complexes. If the complex of groups coming from the H action has finite local groups then we get a finite covering, so G is a finite index subgroup of H.

2.3 Coarse Geometry

In this section and the next we establish the language and basic ideas of coarse geometry and trees. Again, see Bridson and Haefliger's book [3] for more detail.

Let (X, d_X) and (Y, d_Y) be metric spaces. Let A and B be subsets of X.

The (open) r–neighborhood of A is the set $N_r(A) = \{x \in X \mid d_X(x, A) < r\}$.

The *Hausdorff distance* between A and B is:

$$d_H(A, B) = \inf\{r \mid A \subset \overline{N_r(B)} \text{ and } B \subset \overline{N_r(A)}\}$$

We will use the common convention that some object is *r–[adjective]* if it has the property for the specified r, and is *[adjective]* if there exists some r such that the object is r–[adjective].

A and B are *r–coarsely equivalent* if $d_H(A, B) \leq r$.

A is *r–coarsely dense* in X if A is r–coarsely equivalent to X.

A map $\phi\colon X \to Y$ is a (λ, ϵ)-*quasi-isometric embedding* if there exist $\lambda \geq 1$ and $\epsilon \geq 0$ such that for all $x, x' \in X$:

$$\frac{1}{\lambda} d_X(x, x') - \epsilon \leq d_y(\phi(x), \phi(x')) \leq \lambda d_X(x, x') + \epsilon$$

If, in addition, the image of ϕ is ϵ-coarsely dense in Y, then ϕ is a (λ, ϵ)-*quasi-isometry*.

Maps ϕ and ψ from X to Y are *r–coarsely equivalent*, or are equivalent *up to r–bounded distance*, if for all $x \in X$, $d_Y(\phi(x), \psi(x)) \leq r$.

$\mathcal{QI}(X \to Y)$ is the set of quasi-isometries from X to Y modulo coarse equivalence.

Suppose A is r–coarsely dense in X and ϕ is a pseudo-map that assigns to each $a \in A$ a subset $\phi(a)$ in Y of diameter at most R. Suppose there are $\lambda \geq 1$ and $\epsilon \geq 0$ such that for all a and a' in A:

$$\frac{1}{\lambda} d_X(a, a') - \epsilon - R \leq \inf\{d_Y(y, y') \mid y \in \phi(a), y' \in \phi(a')\}$$

and

$$\sup\{d_Y(y, y') \mid y \in \phi(a), y' \in \phi(a')\} \leq \lambda d_x(a, a') + \epsilon + R$$

Then the pseudo-map ϕ determines a unique (up to coarse equivalence) extension to a $(\lambda, 2\lambda r + \epsilon + R)$–quasi-isometric embedding $\Phi\colon X \to Y$ such that for all $a \in A$, $\Phi(a) \in \phi(a)$. For each $x \in X$ choose a closest $a \in A$ and choose any $\Phi(x) \in \phi(a)$.

Suppose for some x we define $\Phi'(x)$ by choosing a different closest $a' \in A$ and $\Phi'(x) \in \phi(a')$. Then

$$\begin{aligned} d_Y(\Phi(x), \Phi'(x)) &\leq \sup\{d_Y(y, y') \mid y \in \phi(a), y' \in \phi(a')\} \\ &\leq \lambda d_X(a, a') + \epsilon + R \\ &\leq \lambda \cdot 2r + \epsilon + R, \end{aligned}$$

so Φ and Φ' are coarsely equivalent.

The fact that Φ is a quasi-isometric embedding follows easily.

If $\phi\colon X \to Y$ is a (λ, ϵ) quasi-isometry, consider the inverse pseudo-map that takes a point in $\phi(X)$ to its preimage in X. This preimage has diameter at most ϵ, and the image of ϕ is ϵ-coarsely dense in Y. We can therefore extend this pseudo-map to a $(\lambda, 2\epsilon(\lambda + \epsilon))$–quasi-isometry $\bar{\phi}\colon Y \to X$. The compositions $\phi \circ \bar{\phi}$ and $\bar{\phi} \circ \phi$ are coarsely equivalent to the identity maps in Y and X, respectively. We call $\bar{\phi}$ a *coarse inverse* of ϕ.

With this notion of inverse, the set $\mathcal{QI}(X) = \mathcal{QI}(X \to X)$ of quasi-isometries from X to itself, modulo coarse equivalence, becomes a group, the quasi-isometry group of X.

Let G be a finitely generated group and let $\mathcal{B}$ be a finite generating set. The *word metric* on G with respect to $\mathcal{B}$ is defined by setting $|g|$ to be the minimum length of a word equal to g in G written in terms of generators in $\mathcal{B}$ or their inverses.

The Cayley graph of G with respect to $\mathcal{B}$ is the graph with one vertex for each element of G and an edge $[g, g']$ connecting vertex g to vertex g' if $g' = gb$ for some $b \in \mathcal{B}$. Make this a metric graph by assuming that each edge has length one. The distance between two vertices g and g' is the length of the shortest edge path joining them. Thus, the distance from g to the identity vertex is the same as $|g|$ in the word metric. G acts on the Cayley graph by isometries via left multiplication.

While the Cayley graph depends on the choice of finite generating set, different choices yield quasi-isometric graphs. More generally, if G acts properly and cocompactly by isometries on a length space X, then X is quasi-isometric to G with (any) word metric. We call such a space X a *geometric model* of G.

2.4 Free Groups and Trees

Let F be the free group of rank n, with free generating set (free basis) $\mathcal{B} = \{a_1, \ldots, a_n\}$. For $g \in F$, let $\bar{g}$ denote g^{-1}.

Let $\mathcal{T} = \mathcal{C}_{\mathcal{B}}(F)$ be the Cayley graph of F with respect to $\mathcal{B}$. Since we have chosen a free generating set, $\mathcal{T}$ is a tree, a graph with no loops.

The tree has a boundary at infinity $\partial\mathcal{T}$ that is a Cantor set. Adding the boundary compactifies the tree; $\overline{\mathcal{T}} = \mathcal{T} \cup \partial\mathcal{T}$ is a compact topological space such that the induced topology on $\mathcal{T}$ agrees with the metric topology on $\mathcal{T}$. For any two points t and t' in $\overline{\mathcal{T}}$ there is a unique geodesic $[t, t']$ joining them.

Let v and w be vertices in $\mathcal{T}$. Define:

$$\text{shadow}^v(w) = \{x \in \overline{\mathcal{T}} \mid w \in [v, x]\}$$

Let $\text{shadow}^v_\infty(w) = \text{shadow}^v(w) \cap \partial\mathcal{T}$.

If $\xi \in \partial\mathcal{T}$ and $v \in \mathcal{T}$ let $v = v_0, v_1, \ldots$ be the vertices along $[v, \xi]$. The sets $\text{shadow}^v(v_i)$ give a neighborhood basis for ξ. The topology on $\overline{\mathcal{T}}$ is independent of the choice of v.

Since $\mathcal{T}$ is hyperbolic, any quasi-isometry $\phi\colon\ \mathcal{T} \to \mathcal{T}'$ extends to a homeomorphism $\partial\phi\colon\ \partial\mathcal{T} \to \partial\mathcal{T}'$.

2.5 Line Patterns and the Decomposition Space

Suppose $l = \{gw^m\}_{m\in\mathbb{Z}}$ is a line in the pattern. The line l has distinct endpoints at infinity:

$$l^+ = gw^\infty = \lim_{i\to\infty} gw^i$$

and

$$l^- = gw^{-\infty} = \lim_{i\to-\infty} gw^i$$

The lines in the pattern never have endpoints in common, so we can decompose $\partial\mathcal{T}$ into disjoint subsets that are either the pair of endpoints of a line from the pattern or a boundary point that is not the endpoint of a line.

Define the decomposition space $\mathbf{D}_\mathcal{L}$ (or just $\mathbf{D}$ when $\mathcal{L}$ is understood) associated to a line pattern $\mathcal{L}$ to be the space that has one point for each set in the decomposition of $\partial\mathcal{T}$, with the quotient topology.

Let $q\colon\ \partial\mathcal{T} \to \mathbf{D}$ be the quotient map. For $x \in \mathbf{D}$, $q^{-1}(x)$ is either a single point that is not the endpoint of any line in $\mathcal{L}$, or $q^{-1}(x) = \{l^+, l^-\}$ for some $l \in \mathcal{L}$. The former we call *bad points*, the latter, *good points*.

The quotient map q induces a bijection between $\mathcal{L}$ and the good points of $\mathbf{D}$, which we denote by q_*.

If $S \subset \mathbf{D}$ we will use the notation $\hat{S} = q^{-1}(S) \subset \partial\mathcal{T}$. Further, if S consists of good points we will use $\tilde{S}$ to be the collection of lines of $\mathcal{L}$ given by $q_*^{-1}(S)$.

The decomposition space is a perfect, compact, Hausdorff topological space.

A quasi-isometry ϕ from $\mathcal{T}$ to $\mathcal{T}'$ extends to a homeomorphism $\partial\phi\colon\ \partial\mathcal{T} \to \partial\mathcal{T}'$. In particular, if there are line patterns $\mathcal{L}$ in $\mathcal{T}$ and $\mathcal{L}' \in \mathcal{T}'$, and if ϕ is a pattern preserving quasi-isometry, then the homeomorphism $\partial\phi\colon\ \partial\mathcal{T} \to \partial\mathcal{T}'$ descends to a homeomorphism of the corresponding decomposition spaces.

3 Whitehead's Algorithm

In a pair of 1936 papers [21, 22], J. H. C. Whitehead gave an algorithm to answer the following question:

Given two finite (ordered) lists of words $(w_1, \ldots, w_k)$ and $(w'_1, \ldots w'_k)$ in a finite rank free group F, is there an automorphism ϕ of F such that for all i we have $\phi(w_i) = w'_i$?

A number of authors have refined Whitehead's Algorithm and applied it to related algebraic questions. Section I.4 of the book of Lyndon and Schupp [10] gives a version of Whitehead's Algorithm and some of the classical applications.

More recently, Stallings [19] and Stong [20] gave 3–manifold versions of Whitehead's Algorithm. In each of these papers the aim was to show that a version of Whitehead's Algorithm could be used to determine if, given a finite list of words $(w_1, \dots w_k)$ in F, there is a free splitting of F such that every w_i is conjugate into one of the free factors. Stallings calls this "algebraically separable". This algebraic question is then shown to be equivalent to a geometric question about whether or not a collection of curves in a handlebody has a property that Stallings calls "geometrically separable" and Stong calls "disk-busting".

In this section we review Whitehead's Algorithm. Our language is similar to that of Stallings and Stong, except that our group actions are on the left and path concatenations are on the right, while they use the opposite convention.

3.1 Whitehead Graphs

Let $\mathcal{B} = \{a_1, \dots, a_n\}$ be a free basis of F. Let $w \in F$ be a cyclically reduced word. The *Whitehead Graph of w with respect to $\mathcal{B}$* is the graph $\mathrm{Wh}_{\mathcal{B}}(*)\{w\}$ with $2n$ vertices labeled $a_1, \dots, a_n, \bar{a}_1, \dots \bar{a}_n$, and an edge between vertices v and v' for *each* occurrence of $\bar{v}v'$ in w (as a cyclic word). The graph depends on the choice of $\mathcal{B}$, and, of course, on w, but we will write $\mathrm{Wh}(*)$ when these are clear.

Remark At present the $(*)$ may be ignored; it will be explained in the next section.

For example, if $F = \langle a, b \rangle$ Figures 1-4 show some Whitehead Graphs.

Notice that $\mathrm{Wh}(*)\{a\}$ is disconnected; the vertices b and $\bar{b}$ are isolated.

$\mathrm{Wh}(*)\{ab^2\}$ is connected but becomes disconnect if the vertex b (or $\bar{b}$) is deleted; b and $\bar{b}$ are *cut vertices*.

$\mathrm{Wh}(*)\{ab\bar{a}\bar{b}\}$ is connected and has no cut vertices.

$\mathrm{Wh}(*)\{a^2ba^2\bar{b}^2\}$ is also connected with no cut vertices, and has multiple edges between vertices a and $\bar{a}$.

More generally, one can make a Whitehead graph representing finitely many words $w_1, \dots, w_m$. We call this Whitehead graph $\mathrm{Wh}(*)\{w_1, \dots, w_m\}$ or $\mathrm{Wh}(*)\{\mathcal{L}\}$, where $\mathcal{L}$ is the line pattern generated by $\{w_1, \dots, w_m\}$.

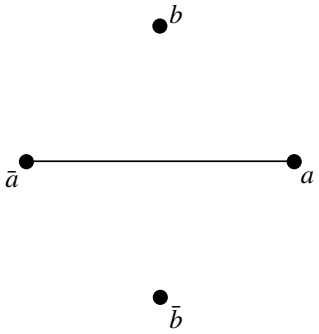

Figure 1: Wh$(*)\{a\}$

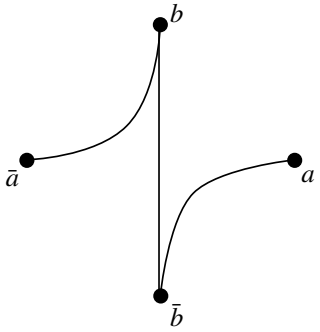

Figure 2: Wh$(*)\{ab^2\}$

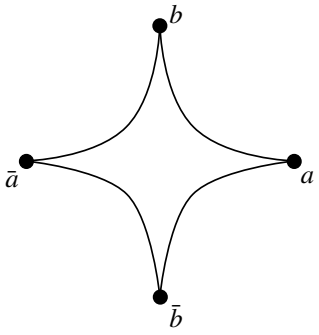

Figure 3: Wh$(*)\{ab\bar{a}\bar{b}\}$

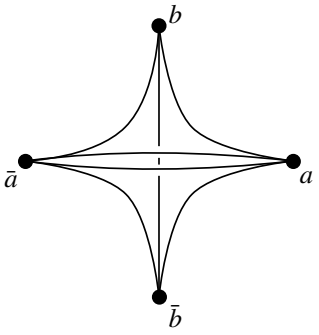

Figure 4: Wh$(*)\{a^2ba^2\bar{b}^2\}$

3.2 Whitehead Automorphisms

Applying $\phi \in \mathrm{Aut}(F)$ changes the Whitehead graph $\mathrm{Wh}_{\mathcal{B}}(*)\{w_1, \ldots, w_m\}$ to the Whitehead graph $\mathrm{Wh}_{\mathcal{B}}(*)\{[\phi(w_1)], \ldots, [\phi(w_m)]\}$, where $[\phi(w_i)]$ means choose a cyclically reduced word in the conjugacy class of $\phi(w_i)$.

An automorphism that permutes $\mathcal{B}$ or swaps a generator with its inverse gives an isomorphic Whitehead graph.

Definition 3.1 (Whitehead Automorphism) A *Whitehead automorphism* is an automorphism of the following form: Pick $x \in \mathcal{B} \cup \bar{\mathcal{B}}$, a generator or the inverse of a generator. Pick $Z \subset \mathcal{B} \cup \bar{\mathcal{B}}$ such that $x \in Z$ and $\bar{x} \notin Z$.

Define an automorphism $\phi_{x,Z}$ by defining $\phi_{x,Z}(x) = x$ and for the rest of the generators

$y \in \mathcal{B}$:

$$\phi_{x,Z}(y) = \begin{cases} y & \text{if } y \notin Z \text{ and } \bar{y} \notin Z \\ xy & \text{if } y \in Z \text{ and } \bar{y} \notin Z \\ y\bar{x} & \text{if } y \notin Z \text{ and } \bar{y} \in Z \\ xy\bar{x} & \text{if } y \in Z \text{ and } \bar{y} \in Z \end{cases}$$

We say that the automorphism $\phi_{x,Z}$ is the *Whitehead automorphism that pushes Z through x*.

To visualize what is happening, consider the rose with one vertex and one oriented loop for each element of $\mathcal{B}$. The fundamental group is F. The Whitehead automorphism $\phi_{x,Z}$ is the automorphism of the fundamental group induced by the homotopy equivalence that pushes one or both ends of the y–loop around the x–loop according to whether y or $\bar{y}$ or both are in Z, or leaves the y–loop alone if neither y nor $\bar{y}$ are in Z. See also Section 4.2.

Define the *complexity* of the collection $w_1, \dots, w_n$ to be the number of edges of $\mathrm{Wh}(*)\{w_1, \dots, w_m\}$. This is equivalent to the sum of the lengths of the w_i, and also half the sum of the valences of the vertices.

Comparing $\mathrm{Wh}_{\mathcal{B}}(*)\{w_1, \dots, w_m\}$ to $\mathrm{Wh}_{\mathcal{B}}(*)\{[\phi_{x,Z}(w_1)], \dots, [\phi_{x,Z}(w_m)]\}$ we see that the valences of vertices other than x and $\bar{x}$ do not change. The new valence of x and $\bar{x}$ is equal to the number of edges that go between Z and Z^c. Thus, the Whitehead automorphism reduces the complexity of the Whitehead graph exactly when there are fewer edges joining Z and Z^c than the valence of x.

Theorem 3.2 $\mathrm{Aut}(F)$ *is generated by:*

(1) *exchanges of a generator with its inverse*

(2) *permutations of the generators*

(3) *Whitehead automorphisms*

Proof This is clear since this set of automorphisms contains the Nielsen generators [14] for $\mathrm{Aut}(F)$. □

Whitehead's Algorithm is as follows: First, check if any Whitehead automorphisms reduce the complexity of the Whitehead graph. Repeat. Once you have reduced to minimal complexity, there are only finitely many graphs to consider. Build a graph with one vertex for each possible Whitehead graph with the given complexity, and an edge

between two vertices if one of the given generators of the automorphism groups takes one graph to the other. One can then show that an automorphism matching up the two lists of words exists if and only if the reduced Whitehead graphs for the two lists of words lie in the same connected component of this graph.

If $\{w_1, \ldots, w_m\}$ is a subset of a free basis then the minimal complexity Whitehead graph should have m disjoint edges.

If there is a free splitting $F = F' * F''$ with every w_i in F' or F'' then the minimal complexity Whitehead graph should be disconnected.

The presence of a cut vertex in the Whitehead graph indicates that the graph is not reduced. If x is a cut vertex, let Z be the union of $\{x\}$ and the vertices of a connected component of $\mathrm{Wh}(*)\{\mathcal{L}\} \setminus \{x\}$ not containing $\bar{x}$. The Whitehead automorphism $\phi_{x,Z}$ reduces complexity.

One application of Whitehead's Algorithm is that a word w is an element of a free basis of $F = F_n$ if and only if the minimal complexity Whitehead graph for w consists of a single edge and $2(n-1)$ isolated vertices.

More generally, the *width* of an element w is the rank of the smallest free factor of F containing w. The minimal complexity Whitehead graph for an element of width m in $F = F_n$ consists of $2(n-m)$ isolated vertices and a connected graph without cut vertices on the remaining vertices.

4 Whitehead Graphs and the Topology of the Decomposition Space

The decomposition space associated to a line pattern first appears in the literature in work of Otal [15], who proves that the decomposition space is connected if and only if there exists a basis $\mathcal{B}$ of F such that $\mathrm{Wh}_{\mathcal{B}}(*)$ is connected without cut vertices.

A similar theorem appears in the thesis of Reiner Martin [12], who references notes of Bestvina.

Theorem 4.1 [12, Theorem 49] *For any $w \in F - \{1\}$, the following are equivalent:*

(1) *w is contained in a proper free factor of F.*

(2) *The width of w is strictly less than the rank of F.*

(3) *There exists a disconnected Whitehead graph of w.*

(4) *The decomposition space associated to the pattern generated by w is disconnected.*

(5) *Every Whitehead graph for w with no cut vertices is disconnected.*

The goal of this section is to further explore the relationship between generalizations of the Whitehead graph and the topology of the decomposition space. In particular, we are interested in finite cut sets in the case that the decomposition space is connected.

Remark The theorem stated in [12] has an additional equivalent condition: for any basis there exists a generalized Whitehead graph that is disconnected. We will not make use of this. In our notation, Martin's generalized Whitehead graph is $\mathrm{Wh}_{\mathcal{B}}(N_r(*))\{w\}$ (see below).

4.1 Geometric Interpretation and Generalizations

Let $\mathcal{B}$ be a free basis for F_n, and let $\mathcal{T}$ be the corresponding Cayley graph. Let $\mathcal{L}$ be a line pattern in F_n. Let $\mathcal{X}$ be a connected subset of $\overline{\mathcal{T}}$ that includes any limit points of $\mathcal{X}$ in $\partial\mathcal{T}$.

Definition 4.2 (Whitehead Graph) The *Whitehead graph* $\mathrm{Wh}_{\mathcal{B}}(\mathcal{X})\{\mathcal{L}\}$ is a graph whose vertices are the connected components of $\overline{\mathcal{T}} \setminus \mathcal{X}$. Distinct vertices v_1 and v_2 are connected by an edge for each line $l \in \mathcal{L}$ with one endpoint in the component corresponding to v_1 and the other in the component corresponding to v_2.

Remarks The conditions on $\mathcal{X}$ ensure that for each component of $\overline{\mathcal{T}} \setminus \mathcal{X}$ there is a unique edge of $\mathcal{T}$ that connects the component to $\mathcal{X}$. Thus, vertices of a Whitehead graph are in bijection with the edges of $\mathcal{T}$ that have exactly one endpoint in $\mathcal{X}$.

When $\mathcal{X} = *$ is a single vertex the definition given here coincides with the classical definition of $\mathrm{Wh}_{\mathcal{B}}(*)\{\mathcal{L}\}$. Since $\mathcal{L}$ is equivariant we get the same graph for any choice of vertex.

We will suppress the $\mathcal{L}$ and $\mathcal{B}$ in the notation when these are clear from context.

The argument $\mathcal{X}$ in parentheses is a subset of $\overline{\mathcal{T}}$. In particular, we will often be interested in a geodesic segment $\mathcal{X}$ joining two points x and y in $\overline{\mathcal{T}}$. We will use closed interval notation $\mathcal{X} = [x, y]$ to denote such a segment. As the vertices of the tree are labeled by group elements one might also interpret $[x, y]$ to be the vertex labeled by the commutator of x and y. To avoid confusion, we will never use square brackets for commutators; square brackets always refer to a geodesic segment in $\overline{\mathcal{T}}$.

If $\mathcal{X}$ and $\mathcal{Y}$ have the same vertex set they result in the same Whitehead graph. It will sometimes be convenient to use notation like $[x, y)$ to mean the geodesic subsegment of $[x, y]$ including all but the last edge.

We will also give a combinatorial construction of our generalization of the Whitehead graph. However, the intuition that informs our arguments comes from the above geometric interpretation. One should visualize the Whitehead graph as the portion of the line pattern that passes through a subset of a tree, rather than as an abstract graph. Where appropriate, as in Figure 5 and Figure 6, we have included the relevant portions of the tree to aid in this visualization.

If $\mathcal{X}$ is a finite (finitely many vertices) connected subset of $\mathcal{T}$ we can build up Whitehead graphs $\mathrm{Wh}(\mathcal{X})\{\mathcal{L}\}$ in a combinatorial way by *splicing* together copies of $\mathrm{Wh}(*)\{\mathcal{L}\}$ for each of the vertices of $\mathcal{X}$. Splicing is a method of combining graphs. The term was coined by Manning in [11] where he uses splicing to construct Whitehead graphs of finite covers of a handlebody from the Whitehead graph of the base handlebody.

Let v be a vertex of a graph Γ and let v' be a vertex of a graph Γ' of valence equal to the valence of v. Given a bijection between edges of Γ incident to v and edges of Γ' incident to v', the *splicing map*, form a new graph whose vertices are the vertices of Γ and Γ' minus the vertices v and v'. Edges not incident to v or v' are retained in the new graph. Finally, for each pair of edges $[w, v]$ in Γ and $[w', v']$ in Γ' identified by the splicing map, add an edge $[w, w']$ in the new graph.

In other words, we have deleted v and v', leaving the edges incident to those vertices with "loose ends". The splicing map tells us how to splice a loose end at v to a loose end at v' to get an edge in the new graph.

For Whitehead graphs the splicing map is determined by the line pattern. Suppose $\mathcal{L}$ is some line pattern and we have chosen a basis. Suppose we have adjacent vertices g and ga in a Cayley graph, and corresponding Whitehead graphs $\mathrm{Wh}(g)\{\mathcal{L}\}$ and $\mathrm{Wh}(ga)\{\mathcal{L}\}$. We splice them together to build the Whitehead graph $\mathrm{Wh}([g, ga])\{\mathcal{L}\}$ on the (one edge) segment $[g, ga]$. The g vertex and ga vertex in $\mathcal{T}$ are adjacent across an a–edge, so the splicing vertices are the a–vertex of $\mathrm{Wh}(g)\{\mathcal{L}\}$ and the $\bar{a}$–vertex of $\mathrm{Wh}(ga)\{\mathcal{L}\}$. Each edge in $\mathrm{Wh}(g)\{\mathcal{L}\}$ incident to a corresponds to a length two subword of one of the generators of the line pattern of the form xa or $\bar{a}x$. Suppose an edge corresponds to a subword xa, and suppose the next letter is y, so there is a length three subword xay. We define the splicing map to identify the edge corresponding to this particular instance of the subword xa to the edge in $\mathrm{Wh}(ga)\{\mathcal{L}\}$ (incident to $\bar{a}$) corresponding to this particular instance of the subword ay.

We can make the splicing easier to visualize if we draw the Whitehead graphs with loose ends at the vertices. Figure 5 shows the Whitehead graph for the pattern generated by the words ab and $a\bar{b}$ in $F = F_2 = \langle a, b \rangle$, along with the underlying tree. The word ab will contribute an edge from $\bar{a}$ to b and an edge from $\bar{b}$ to a. The twists in the graph indicate the splicing maps.

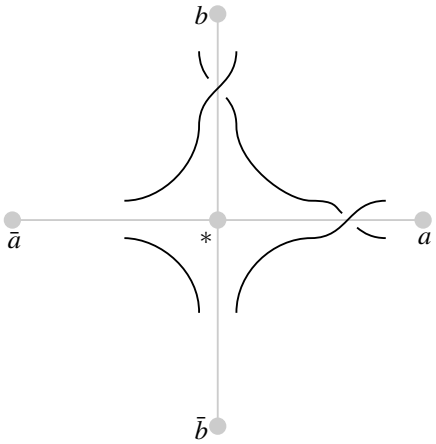

Figure 5: Wh$(*)\{ab, a\bar{b}\}$ (loose)

Let $*$ be the identity vertex. Take a copies of this graph at $*$ and at a and splice them together. We get the splicing map by considering the words. There is an ab–line at $*$. If the first letter is a, the previous letter was b, so we see an edge from $\bar{b}$ to a in the Whitehead graph at $*$. The next letter is b, so in the Whitehead graph at a we see an edge from $\bar{a}$ to b, and the twist in the graph indicates that these two edges should be spliced together.

Similarly, there is an $a\bar{b}$–line at $*$. It contributes an edge from b to a in the Whitehead graph at $*$, and this continues on to an edge from $\bar{a}$ to $\bar{b}$ in the Whitehead graph at a.

Note Unless noted otherwise, figures are drawn so that the splicing map is achieved by an orientation preserving isometry of the page.

The geometric and splicing constructions produce the same graph for sets $\mathcal{X}$ with finitely many vertices. We could try to take limits of the spliced graphs when $\mathcal{X}$ is infinite, but if $\mathcal{X} \subset \overline{\mathcal{T}}$ contains endpoints of some $l \in \mathcal{L}$, then splicing does not actually produce a graph.

If both endpoints of l are in $\overline{\mathcal{X}}$ then after finitely many splices there is an edge corresponding to l, but in the limit the edge grows to be an open interval not incident

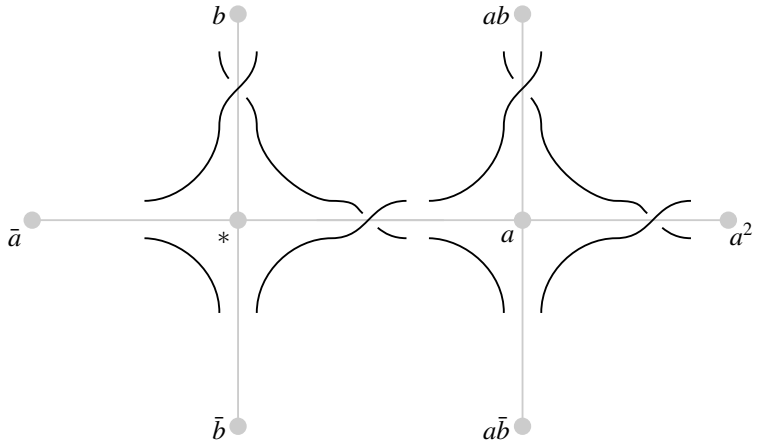

Figure 6: $\mathrm{Wh}(*)\{ab, a\bar{b}\}$ and $\mathrm{Wh}(a)\{ab, a\bar{b}\}$ (loose)

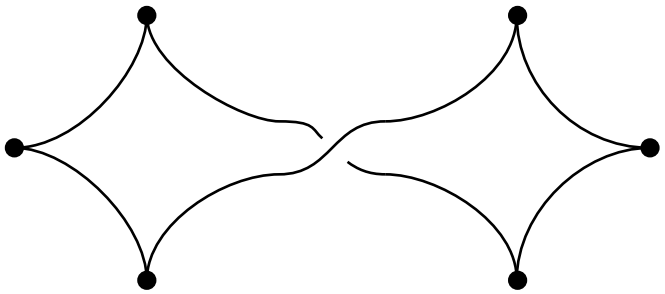

Figure 7: $\mathrm{Wh}([*, a])\{ab, a\bar{b}\}$

to any vertices; the vertices escape to infinity. This line does not occur if we follow the geometric definition, because it is not joining two different components of the complement of $\mathcal{X}$ in $\overline{\mathcal{T}}$. Similarly, if only one endpoint of l is in $\overline{\mathcal{X}}$ then splicing produces a graph $\mathcal{G}$ with a half line attached. If we throw out these "non-closed edges" we get the graph $\mathrm{Wh}(\mathcal{X})\{\mathcal{L}\}$ of the geometric definition.

Remark Stong [20] defines a generalized Whitehead graph that coincides with our definition, but, like Martin [12], only makes use of $\mathrm{Wh}_{\mathcal{B}}(N_r(*))\{w\}$.

4.2 Whitehead Automorphisms Revisited

Let us consider how application of a Whitehead automorphism changes the line pattern.

Suppose x, y and z are in $\mathcal{B} \cup \bar{\mathcal{B}}$ with $y \neq z$. Consider a Whitehead automorphism $\phi = \phi_{x,Z}$ (recall Definition 3.1). Let $l \in \mathcal{L}$ be a line that goes through vertices y, $*$ and z, where $*$ is the identity vertex. The line l is the geodesic that goes through vertices of the form $\{y(\bar{y}zu)^m\}_{m\in\mathbb{Z}}$, where u is some word in F that does not begin with $\bar{z}$ or end with y.

First suppose that $y, z \in Z$, $y, z \neq x$ and $\bar{y}, \bar{z} \notin Z$, so that $\phi(x) = x$, $\phi(y) = xy$ and $\phi(z) = xz$. Then $\phi(l)$ is the line that includes vertices of the form:

$$\{\phi(y(\bar{y}zu)^m)\}_{m\in\mathbb{Z}} = \{xy(\bar{y}\bar{x}xz\phi(u))^m\}_{m\in\mathbb{Z}} = \{xy(\bar{y}z\phi(u))^m\}_{m\in\mathbb{Z}}$$

Since u does not begin with $\bar{z}$ or end in y, the same is true for $\phi(u)$. Therefore, $\phi(l)$ goes through vertices xy, x and xz. The line l that went through $*$ has been "pushed through" the x edge to a line $\phi(l)$ that goes through x and not through $*$.

Using similar arguments one can show:

(1) $\phi(l)$ goes through x and not through $*$ if y and z are in Z.

(2) $\phi(l)$ goes through $*$ and not through x if y and z are in Z^c.

(3) $\phi(l)$ goes through both $*$ and x if exactly one of y or z is in Z.

4.3 Connectivity of the Decomposition Space

Fix a line pattern $\mathcal{L}$, and let $\mathbf{D} = \mathbf{D}_{\mathcal{L}}$ be the corresponding decomposition space. The next two lemmas use essentially the same ideas that go into the proofs of Theorem 4.1 in [12] and [15, Proposition 2.1].

Lemma 4.3 *If for some choice of basis $\mathcal{B}$ the graph $\mathrm{Wh}_{\mathcal{B}}(*)$ is disconnected, then $\mathbf{D}$ is disconnected.*

Proof Let $*$ be the identity vertex in $\mathcal{T}$. Let $\mathcal{B}$ be a free basis of F such that $\mathrm{Wh}_{\mathcal{B}}(*)$ is not connected. Vertices of $\mathrm{Wh}_{\mathcal{B}}(*)$ are in bijection with $\mathcal{B} \cup \bar{\mathcal{B}}$. There is some partition of $\mathcal{B} \cup \bar{\mathcal{B}}$ into subsets $\mathcal{A}$ and $\mathcal{A}'$ so that no lines of $\mathcal{L}$ connect $\mathcal{A}$ to $\mathcal{A}'$. Let

$$\hat{A} = \cup_{a\in\mathcal{A}} \mathrm{shadow}^*_\infty(a)$$

The sets $\hat{A}$ and $\hat{A}^c \subset \partial\mathcal{T}$ are both nonempty clopens, sets that are both open and closed. Since there are no lines of $\mathcal{L}$ with one endpoint in $\hat{A}$ and one in $\hat{A}^c$, their images in $\mathbf{D}$ are disjoint nonempty clopens, so $\mathbf{D}$ is disconnected. □

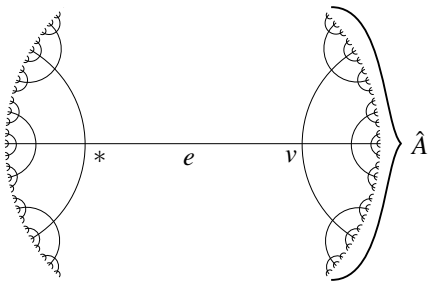

Figure 8: The boundary of the tree split into "halves".

Lemma 4.4 *Suppose there exists a free basis $\mathcal{B}$ of F such that* $\mathrm{Wh}_{\mathcal{B}}(*)$ *is connected without cut vertices. Let $\mathcal{T}$ be the Cayley graph of F corresponding to $\mathcal{B}$. Pick any edge e in $\mathcal{T}$. Let $*$ and v be the endpoints of e. Let $\hat{A} = \mathrm{shadow}^*_\infty(v)$. That is, $\hat{A}$ is the "half" of $\partial\mathcal{T}$ on the "v-side" of e. The set $A = q(\hat{A})$ is connected in* $\mathbf{D}$.

Proof Suppose there are open sets B and C of $\mathbf{D}$ such that $A \subset B \cup C$ and $A \cap B \cap C = \emptyset$. The set $\hat{A}$ is open in $\partial\mathcal{T}$, so $A' = \hat{A} \cap q^{-1}(B)$ and $A'' = \hat{A} \cap q^{-1}(C)$ are open. Assuming that A' is nonempty, we will show that A'' must be empty, which implies A is connected.

$\hat{A}$ is closed in $\partial\mathcal{T}$, so A' and A'' are closed. Compactness of $\partial\mathcal{T}$ implies A' and A'' are compact clopens. Since A' is compact and open, there are finitely many vertices $x_1, \ldots, x_a$ so that $A' = \cup_{i=1}^{a} \mathrm{shadow}^*_\infty(x_i)$

There is a similar finite collection $y_1, \ldots, y_b$ that determines A''.

Consider the convex hull $\mathcal{H}$ of $\{x_i\}_{i=1}^a \cup \{y_j\}_{j=1}^b \cup \{v\}$; it is a finite tree. Call the vertices of $\mathcal{H}$ other than $\{x_i\}_{i=1}^a \cup \{y_j\}_{j=1}^b \cup \{v\}$ the "interior vertices". Since $A' \cup A'' = \hat{A}$, the set $\mathcal{H}$ includes all edges incident to its interior vertices. Let $\mathcal{X}$ be the union of the set of interior vertices with $\{v\}$.

Construct the Whitehead graph $\mathrm{Wh}(\mathcal{X})$. It has $a + b + 1$ vertices corresponding to the x_i and y_j and the edge e of $\mathcal{T}$.

The graph is connected without cut vertices, since it can be constructed by splicing together finitely many copies of $\mathrm{Wh}(*)$, which is connected without cut vertices. In particular, the vertex e is not a cut vertex.

Assume A' is nonempty. If $x_1 = v$ then $A' = \hat{A}$, so $A'' = \emptyset$, and we are done. Otherwise, x_1 is a vertex of $\mathrm{Wh}(\mathcal{X})$ different from v. Since v is not a cut vertex, there are edges of $\mathrm{Wh}(\mathcal{X})$ incident to x_1 and not v. Such an edge corresponds to a line $l \in \mathcal{L}$ 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, and we already know that the image of the first endpoint is in B. This means that z must be in $\{x_2, \dots, x_a\}$. Since $\mathrm{Wh}(\mathcal{X})$ is connected and v is not a cut vertex, we conclude that all the vertices of $\mathrm{Wh}(\mathcal{X}) \setminus v$ belong to $\{x_1, \dots, x_n\}$, so $A'' = \emptyset$. Thus, A is connected in $\mathbf{D}$. □

Thus, if $\mathrm{Wh}(*)$ is connected without cut vertices, then for any edge e in $\mathcal{T}$ the boundaries of the two connected components of $\mathcal{T} \setminus e$ correspond to connected sets in the decomposition space. Since $\mathrm{Wh}(*)$ is connected without cut vertices there are at least two lines in $\mathcal{L}$ crossing e. This means that these two connected sets in the decomposition space have points in common.

Corollary 4.5 *Suppose* $\mathrm{Wh}(*)$ *has no cut vertices. Then the decomposition space is connected if and only if* $\mathrm{Wh}(*)$ *is connected.*

Proof of Theorem 4.1 Conditions (1) and (2) are equivalent by definition. The equivalence of (1) and (3) is a consequence of Whitehead's Algorithm.

(3) $\implies$ (4) is Lemma 4.3.

Corollary 4.5 implies the contrapositive of (4) $\implies$ (5).

It is always possible to eliminate cut vertices with Whitehead automorphisms, so (5) $\implies$ (3). □

4.4 Cut Sets in a Connected Decomposition Space

From now on, unless otherwise noted, we will assume that $\mathcal{L}$ is a fixed line pattern such that the decomposition space $\mathbf{D} = \mathbf{D}_{\mathcal{L}}$ is connected. Thus, we can choose a basis $\mathcal{B}$ such that $\mathrm{Wh}_{\mathcal{B}}(*)\{\mathcal{L}\}$ is connected without cut vertices. In particular, this is true for any basis for which $\mathrm{Wh}_{\mathcal{B}}(*)\{\mathcal{L}\}$ has minimal complexity, but for most of our arguments minimal complexity is not a necessary hypothesis. As $\mathcal{B}$ and $\mathcal{L}$ are fixed, we will suppress them in the Whitehead graph notation.

Our goal is to identify finite minimal cut sets.

Here is another corollary of Lemma 4.4:

Corollary 4.6 *Pick any edge* e *in* $\mathcal{T}$*. Let* $\tilde{S}$ *be the collection of (finitely many) lines of* $\mathcal{L}$ *that cross* e*. Then* $S = q_*(\tilde{S})$ *is a cut set in* $\mathbf{D}$*.*

Definition 4.7 (Edge Cut Set) A set $S = q_*(\tilde{S})$, such that $\tilde{S}$ is the set of lines of $\mathcal{L}$ crossing some edge e of $\mathcal{T}$ is called an *edge cut set*.

These cut sets will play a pivotal role in Section 5. They are the prototypes for the "small" cut sets that we will use to build a cubing.

It will turn out that for our purposes the most important cut sets come from finite subsets of $\mathcal{L}$. The edge cut sets mentioned above are an example of such sets. A finite collection of lines $\tilde{S} = \{l_1, \dots, l_k\} \subset \mathcal{L}$ gives a finite set $S = q_*(\tilde{S})$ of $\mathbf{D}$. We would like to be able to decide if such a set S is a cut set of $\mathbf{D}$. To do this, we consider generalized Whitehead graphs.

Let $\mathcal{X}$ be a compact connected subset of $\mathcal{T}$, and let l be a line of $\mathcal{L}$. If l intersects $\mathcal{X}$ then there is an edge of $\mathrm{Wh}(\mathcal{X})$ corresponding to l. Let $\mathrm{Wh}(\mathcal{X}) - l$ denote the graph that results from deleting the interior of the edge of $\mathrm{Wh}(\mathcal{X})$ corresponding to l. (If there is no such edge, because $l \cap \mathcal{X} = \emptyset$, then $\mathrm{Wh}(\mathcal{X}) - l$ is just $\mathrm{Wh}(\mathcal{X})$.) Similarly, $\mathrm{Wh}(\mathcal{X}) - \tilde{S}$ is the graph obtained from $\mathrm{Wh}(\mathcal{X})$ by deleting the interiors of any edges of $\mathrm{Wh}(\mathcal{X})$ coming from the lines of $\tilde{S}$. Given $\tilde{S}$, we would like to know the relationship between components of $\mathrm{Wh}(\mathcal{X}) - \tilde{S}$ and components of $\mathbf{D} \setminus S$. Much of the technical work of the next few sections is aimed at answering this question. Ultimately, in Lemma 4.20, we will show that, under certain conditions, given $\tilde{S}$ there is a choice of $\mathcal{X}$ so that components of $\mathrm{Wh}(\mathcal{X}) - \tilde{S}$ are in bijection with components of $\mathbf{D} \setminus S$.

There is an easy sufficient condition to see that a set $\tilde{S} = \{l_1, \dots, l_k\} \subset \mathcal{L}$ gives a cut set $S = q_*(\tilde{S})$ of $\mathbf{D}$:

Proposition 4.8 *Let $\tilde{S} = \{l_1, \dots, l_k\}$ be a finite collection of lines in $\mathcal{L}$. Let $\mathcal{X}$ be any compact, connected set in $\mathcal{T}$. If $\mathrm{Wh}(\mathcal{X}) - \tilde{S}$ is disconnected then $S = q_*(\tilde{S})$ is a cut set.*

One could prove this proposition with an argument similar to that of Corollary 4.6. We will prove a more general result in the next proposition. Before moving on, though, let us consider an example that shows that this proposition does not give a necessary condition for S to be a cut set. Clearly if $\mathcal{X}$ is disjoint from $\tilde{S}$ then $\mathrm{Wh}(\mathcal{X}) - \tilde{S}$ will not give any information about $\mathbf{D} \setminus S$, but in fact there are situations for which there is no compact $\mathcal{X}$ such that $\mathrm{Wh}(\mathcal{X}) - \tilde{S}$ adequately models $\mathbf{D} \setminus S$.

Consider the pattern $\mathcal{L}$ generated by the pair of words b and $ab\bar{a}\bar{b}$ in $F = \langle a, b \rangle$. The Whitehead graph (with loose ends) for this pattern is shown in Figure 9.

The graph $\mathrm{Wh}(*)\{b, ab\bar{a}\bar{b}\}$ is connected without cut vertices, so the decomposition space is connected. We claim that the endpoints of any b–line give a cut point in the

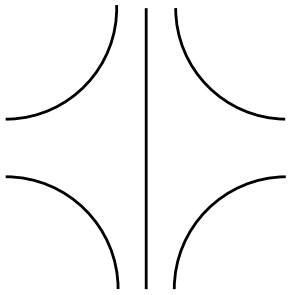

Figure 9: $\mathrm{Wh}(*)\{b, ab\bar{a}\bar{b}\}$ (loose)

decomposition space. For instance, the b–line through the identity vertex has endpoints b^∞, $b^{-\infty}$ in $\partial\mathcal{T}$. Let $B = q(b^\infty) = q(b^{-\infty})$.

Let $*$ be the identity vertex. Let

$$\hat{A} = \cup_{m\in\mathbb{Z}} \mathrm{shadow}^*_\infty(b^m a),$$

that is, $\hat{A}$ consists of all the boundary points ξ of $\mathcal{T}$ such that the first occurrence of a or $\bar{a}$ in the geodesic from the identity to ξ is an a.

Now $\hat{A}$ is open, and every line of $\mathcal{L}$ with one endpoint in $\hat{A}$ has both endpoints in $\hat{A}$. Let $A = q(\hat{A})$; the preimage is $\hat{A} = q^{-1}(A)$, so A is open in $\mathbf{D}$. Similarly, let A' be the image in $\mathbf{D}$ of the boundary points of $\mathcal{T}$ such that the first occurrence of a or $\bar{a}$ in the geodesic from the identity is an $\bar{a}$.

$\mathbf{D} = B \cup A \cup A'$, and $B = \overline{A} \setminus A = \overline{A'} \setminus A'$, so B is a cut point.

For any compact, connected $\mathcal{X}$, the Whitehead graph $\mathrm{Wh}(\mathcal{X})\{b, ab\bar{a}\bar{b}\}$ looks like a circle with a number of disjoint chords (see Figure 10).

The edges of the circle correspond to $ab\bar{a}\bar{b}$–lines, and the chords correspond to b–lines. This graph has no cut points, so deleting the interior of an edge does not disconnect it.

This example has shown that to decide if $\tilde{S}$ gives a cut set, it is not enough to delete the interiors of edges in a Whitehead graph over a compact set $\mathcal{X}$.

There will be several different notions of deleting parts of Whitehead graphs, so let us standardize notation.

We have already said that $\mathrm{Wh}(\mathcal{X}) - l$ is obtained from $\mathrm{Wh}(\mathcal{X})$ by deleting the interior of the edge corresponding to l, if such an edge exists. Similarly, obtain $\mathrm{Wh}(\mathcal{X}) - \tilde{S}$ by deleting the interiors of any edges corresponding to a line in $\tilde{S}$.

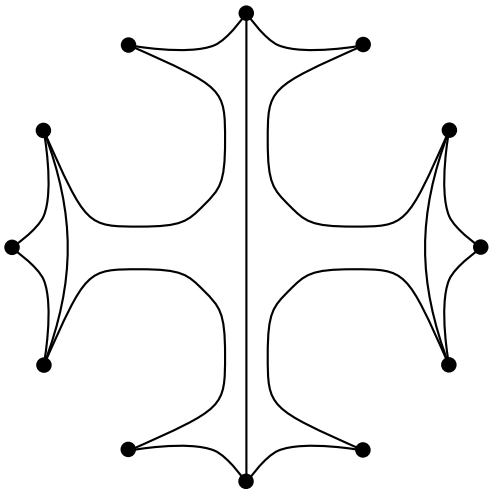

Figure 10: $\mathrm{Wh}(\overline{N_1(*)})\{b, ab\bar{a}\bar{b}\}$

Another option would be to delete not just the interior of an edge, but also its vertices. However, we do not want to lose other edges incident to these vertices. Let $\mathrm{Wh}(\mathcal{X}) - \ddot{l}$ be the graph obtained from $\mathrm{Wh}(\mathcal{X})$ by deleting the interior of the edge corresponding to l as well as the two vertices that are its endpoints, retaining loose ends at these vertices.

We will also be interested in comparing Whitehead graphs $\mathrm{Wh}(\mathcal{X})$ and $\mathrm{Wh}(\mathcal{Y})$ when $\mathcal{X} \subset \mathcal{Y}$. The simplest example is when e is an edge of $\mathcal{T}$ incident to exactly one vertex of $\mathcal{X}$. The edge e corresponds to a vertex in $\mathrm{Wh}(\mathcal{X})$. Define $\mathrm{Wh}(\mathcal{X}) \setminus e$ to be the graph obtained from $\mathrm{Wh}(\mathcal{X})$ by deleting this vertex, but retaining the incident edges as loose ends at e.

Similarly, if v is a vertex of $\mathcal{T}$ that is distance 1 from $\mathcal{X}$, then there is a unique edge e with one endpoint equal to v and the other in $\mathcal{X}$. Define $\mathrm{Wh}(\mathcal{X}) \setminus v = \mathrm{Wh}(\mathcal{X}) \setminus e$.

More generally, if $\mathcal{X} \subset \mathcal{Y} \in \overline{\mathcal{T}}$ then obtain $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{Y}$ from $\mathrm{Wh}(\mathcal{X})$ by deleting each vertex of $\mathrm{Wh}(\mathcal{X})$ that corresponds to an edge in $\mathcal{Y}$. Visualizing Whitehead graphs in the tree, $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{Y}$ is the portion of $\mathrm{Wh}(\mathcal{Y})$ that passes through the set $\mathcal{X}$.

Consider the line pattern $\mathcal{L}$ generated by $ab\bar{a}\bar{b}$ and b. Let l be the b–line through the identity vertex $*$. Let $\mathcal{X} = *$ and let $\mathcal{Y} = [\bar{b}, b]$. Figures 11–14 illustrate our different notions of deleting from $\mathrm{Wh}(\mathcal{X})$. The reader should note the relationships between Figures 11–14 and Figures 9 and 10.

Recall that the preimage in $\partial\mathcal{T}$ of a point in $\mathbf{D}$ is either a single point or the pair of endpoints of a line of $\mathcal{L}$. We call the former bad points and the latter good points.

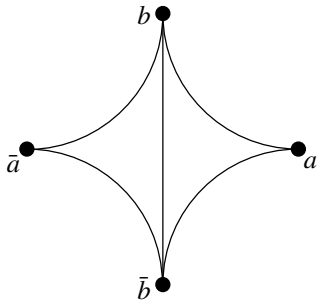

Figure 11: Wh(∗)

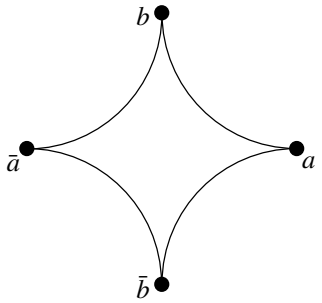

Figure 12: Wh(∗) − l

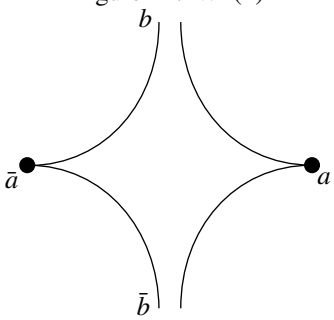

Figure 13: Wh(∗) − $\ddot{l}$

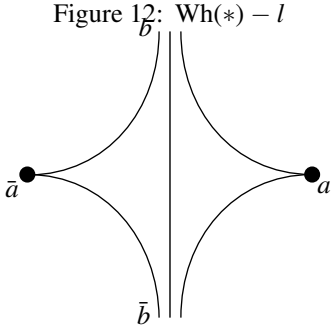

Figure 14: Wh(∗) \ $[\bar{b}, b]$

Lemma 4.9 (Hull Determines Connectivity) *Let S be a nonempty, finite subset of $\mathbf{D}$ that is not just a single bad point. Let $\mathcal{H}$ be the convex hull of $q^{-1}(S)$. There is a bijection between connected components of* Wh($\mathcal{H}$) *and connected components of $\mathbf{D} \setminus S$.*

Proof Components $\mathcal{A}_i$ of $\mathcal{T} \setminus \mathcal{H}$ are the kind of sets in Lemma 4.4. Therefore, $q(\partial \mathcal{A}_i)$ is connected in $\mathbf{D}$. The set $\partial \mathcal{A}_i$ is open in $\partial \mathcal{T}$. For a subcollection $\{\mathcal{A}_{i_j}\}$ corresponding to a connected component of Wh($\mathcal{H}$), we have that $q(\bigcup_j \partial \mathcal{A}_{i_j})$ is an open connected set in $\mathbf{D} \setminus S$, as in the proof of Lemma 4.4. The complement of this set in $\mathbf{D} \setminus S$ is either empty or is a union of sets of a similar form, corresponding to other connected components of $\mathcal{H}$. Thus, $q(\bigcup_j \partial A_{i_j})$ is closed, and is therefore a connected component of $\mathbf{D} \setminus S$. □

Pick any vertex $* \in \mathcal{T}$. If $\xi \in \mathbf{D}$ is a bad point, the previous argument applies if we take $\mathcal{H}$ to be the ray $[*, \xi]$. If Wh(∗) is connected without cut points then Wh($\mathcal{H}$) is connected. Therefore:

Corollary 4.10 *No bad point of* $\mathbf{D}$ *is a cut point.*

In a sense, Lemma 4.9 achieves our goal of relating the topology of the decomposition space to generalizations of the Whitehead graph. However, this generalized Whitehead graph is infinite. In the next two sections we show that the same information can be obtained from a finite portion of this Whitehead graph.

4.5 Identifying Cut Points and Cut Pairs

The previous corollary tells us that any cut point is a good point, so its preimage in $\partial\mathcal{T}$ is a pair of points. We have a similar situation if there is a cut pair consisting of two bad points; the preimage of such a set in $\partial\mathcal{T}$ is a pair of points. In both cases, the convex hull is a line.

Suppose $g \in F \setminus \{1\}$ is cyclically reduced with $\mathcal{H}^+ = g^\infty$ and $\mathcal{H}^- = g^{-\infty}$. Let $\mathcal{H}$ be the convex hull of these two points. Let $\mathcal{X} = [*, g)$ be the segment joining the identity vertex to the g vertex in $\mathcal{T}$, not including the vertex g.

We know, by Lemma 4.9, that the connected components of $\mathbf{D} \setminus q(\{\mathcal{H}^-, \mathcal{H}^+\})$ are in bijection with components of $\mathrm{Wh}(\mathcal{H})$. We can construct $\mathrm{Wh}(\mathcal{H})$ by splicing together g-translates of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$.

$\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ is $\mathrm{Wh}(\mathcal{X}) \setminus \{e, ge\}$ for some edges e incident to $*$ and ge incident to $g = g*$, so $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ has a collection of loose ends at e and at ge. The action of g identifies $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ with $\mathrm{Wh}(g\mathcal{X}) \setminus \mathcal{H}$, which has loose ends at ge and g^2e. The line pattern determines for us a splicing map for splicing the loose ends of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ at ge to the loose ends of $\mathrm{Wh}(g\mathcal{X}) \setminus \mathcal{H}$ at ge.

It is an easy consequence of the hypothesis that $\mathrm{Wh}(*)$ is connected without cut vertices that for any segment $[*, g^k) \subset \mathcal{H}$, every component of $\mathrm{Wh}([*, g^k)) \setminus \mathcal{H}$ contains a loose end at e (and a loose end at $g^k e$). Thus, the number of components of $\mathrm{Wh}(\mathcal{H})$ is bounded above by the number of components of $\mathrm{Wh}(*) \setminus \mathcal{H}$. To bound the number of connected components of $\mathrm{Wh}(\mathcal{H})$ below we need to know if distinct connected components of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ become connected when we splice on more translates.

Let P be a partition of the loose ends of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ at e that is at least as coarse as connectivity in $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$, ie, if two loose ends belong to the same connected component of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ then they belong in the same subset of the partition.

Let $|P|$ be the number of subsets in the partition; P is *nontrivial* if $|P| > 1$.

Since P is at least as coarse as connectivity, every vertex and edge in $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ is connected to loose ends in exactly one subset of P. Let P' be the partition of the loose ends of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ at ge such that two loose ends are in the same subset of the partition if and only if they are connected to loose ends at e in a common subset of the partition P.

The g action determines a partition gP of the loose ends of $\mathrm{Wh}(g\mathcal{X}) \setminus \mathcal{H}$ at ge by pushing forward the partition P.

We say the partition P is *compatible with the splicing map* if there is a bijection between subsets of the partitions P' and gP such that the splicing map splices edges in a subset of P' to edges in the corresponding subset of gP.

The trivial partition is always compatible with the splicing map, but this gives us no information. Another obvious partition to consider would be the partition that comes from connectivity in $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$. This is the partition in which two loose ends of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ at e belong to the same subset of the partition if and only if they belong to the same connected component of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$. Suppose this partition is compatible with the splicing map. This would mean that two loose ends of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ at ge in the same connected component of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$ must splice to two loose ends of $\mathrm{Wh}(g\mathcal{X}) \setminus \mathcal{H}$ at ge in the same connected component of $\mathrm{Wh}(g\mathcal{X}) \setminus \mathcal{H}$, so splicing introduces no new connectivity. In this case it follows that for all $k \geq 1$ the number of connected components of $\mathrm{Wh}([*, g^k)) \setminus \mathcal{H}$ is equal to the number of connected components of $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$.

However, this is not always the case. Splicing may introduce new connectivity. Compatibility of the partition controls how much new connectivity is introduced. If we have a partition compatible with the splicing map, then, after splicing, the partition P is at least as coarse as connectivity in $\mathrm{Wh}([*, g^2)) \setminus \mathcal{H}$. Moreover, P will still be compatible with the splicing map at g^2e, so we may continue by induction to show:

Proposition 4.11 *Suppose P is a partition that is compatible with the splicing map. Then for any segment $\mathcal{Y} = [*, g^k)$ of $\mathcal{H}$, the number of connected components of $\mathrm{Wh}(\mathcal{Y}) \setminus \mathcal{H}$ is greater than or equal to $|P|$.*

Any particular line $l \in \mathcal{L}$ overlaps with $\mathcal{H}$ for distance at most

$$|g| \times (2 + \text{maximum number of consecutive } g\text{'s in a generating word for } \mathcal{L}.)$$

This is infinite if and only if g is a generator of the line pattern. Thus, g is a generator of the line pattern if and only if in every $\mathcal{Y}$ there is a free edge (edge that is loose at

both ends) in $\mathrm{Wh}(\mathcal{Y}) \setminus \mathcal{H}$, since in this case $\mathcal{H} = \bar{l}$ where $l \in \mathcal{L}$ is the g–line through the identity.

Given a compatible partition, there are two cases to consider. If for some $\mathcal{Y}$ we have no free edges in $\mathrm{Wh}(\mathcal{Y}) \setminus \mathcal{H}$ then the number of connected components of $\mathrm{Wh}(\mathcal{H})$ is greater than or equal to $|P|$. In particular, if $|P| \geq 2$, then $q(\{\mathcal{H}^+, \mathcal{H}^-\})$ is a cut pair consisting of two bad points, a bad cut pair.

In this case we could have chosen the partition P so that one of the subsets is the singleton consisting of the loose end of the free edge. The partition P' also has a subset that is a singleton, consisting of the other loose end of the free edge. Such a partition has a *segregated free edge*.

We do not see the free edge in $\mathrm{Wh}(\mathcal{H})$, so in general we can only conclude that $\mathrm{Wh}(\mathcal{H})$ has at least $|P| - 1$ connected components. If $|P| - 1 \geq 2$ then $q(\{\mathcal{H}^+, \mathcal{H}^-\}) = q_*(l)$ is a cut point in $\mathbf{D}$.

Proposition 4.12 ($q(\{g^\infty, g^{-\infty}\})$ Cut Set Criterion) *Let $g \in F \setminus \{1\}$ be an element of the free group. With notation as above, let P be the finest partition that is compatible with the splicing map and at least as coarse as connectivity in $\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}$. Then:*

(1) *If P is trivial then $q(\{g^\infty, g^{-\infty}\})$ is not a cut set.*

(2) *If P is nontrivial and has no segregated free edge then $q(\{g^\infty, g^{-\infty}\})$ is a bad cut pair.*

(3) *If P has a segregated free edge and $|P| = 2$ then $q(\{g^\infty, g^{-\infty}\})$ is not a cut set.*

(4) *If P has a segregated free edge and $|P| > 2$ then $q(\{g^\infty, g^{-\infty}\})$ is a cut point.*

Proof If P is trivial then $\mathrm{Wh}([*, g^2)) \setminus \mathcal{H}$ is connected, so $q(\{g^\infty, g^{-\infty}\})$ is not a cut set. Similarly, if $|P| = 2$ and there is a segregated free edge then $\mathcal{H} = \bar{l}$ for $l \in \mathcal{L}$ and $\mathrm{Wh}([*, g^2]) - l$ is connected, so $q(\{g^\infty, g^{-\infty}\})$ is not a cut set.

In the other cases, $\mathrm{Wh}(\mathcal{H})$ has multiple components, so $q(\{g^\infty, g^{-\infty}\})$ is a cut set. □

The proposition tells us that given a g we can decide if $q(\{g^\infty, g^{-\infty}\})$ is a cut set. We call this a *periodic* cut set. Next we show that if there are cut points or cut pairs then there are periodic cut sets:

Proposition 4.13 *If $\mathbf{D}$ has cut points or cut pairs then there is some R depending on $\mathcal{L}$ and some g with $|g| \leq R$ such that $q(\{g^\infty, g^{-\infty}\})$ is a cut set.*

To identify cut points we just need to apply Proposition 4.12 to the generators of $\mathcal{L}$, so in this case it is sufficient to take R to be the length of the longest generator of $\mathcal{L}$. The work of proving Proposition 4.13 lies in finding an R that works for the cut pair case:

Lemma 4.14 *If $q(\{\mathcal{H}^+, \mathcal{H}^-\})$ is a cut pair then there is some R depending on $\mathcal{L}$ and some $g \in F \setminus \{1\}$ with $|g| \leq R$ such that $q(\{g^\infty, g^{-\infty}\})$ is a cut set.*

Note that $q(\{g^\infty, g^{-\infty}\})$ is either a cut point or a bad cut pair.

Proof Let $\mathcal{H}$ be the convex hull of $\{\mathcal{H}^+, \mathcal{H}^-\}$. We may assume that $\mathcal{H}$ contains the identity vertex $*$.

Use # to denote number of connected components.

Every connected component of $\mathrm{Wh}(*) \setminus \mathcal{H}$ contains an edge, so the number of components is at most the complexity of $\mathrm{Wh}(*)$.

For any segment $\mathcal{X}$ of $\mathcal{H}$ we have:

$$2 \leq \#\,\mathrm{Wh}(\mathcal{H}) \leq \#\left(\mathrm{Wh}(\mathcal{X}) \setminus \mathcal{H}\right) \leq \#\left(\mathrm{Wh}(*) \setminus \mathcal{H}\right) \leq \text{complexity of } \mathrm{Wh}(*)$$

Number the vertices of $\mathcal{H}$ consecutively with integers with $* = v_0$ and index increasing in the $\mathcal{H}^+$ direction. Number the edges of $\mathcal{H}$ so that e_i is incident to v_{i-1} and v_i. We consider these edges oriented in the direction of increasing index. An oriented edge of $\mathcal{T}$ comes with a label that is a generator or inverse of a generator of F.

The function $f(i) = \#(\mathrm{Wh}([\mathcal{H}^{-\infty}, v_i]) \setminus \mathcal{H})$ is nonincreasing and, for high enough i, stabilizes at $\#\,\mathrm{Wh}(\mathcal{H})$. Since we started with a cut pair, for high enough i there is no free edge in $\mathrm{Wh}([\mathcal{H}^{-\infty}, v_i]) \setminus \mathcal{H}$. After changing by an isometry and relabeling, if necessary, we may assume that $i = 0$ is "high enough" in the previous two statements.

Fix a numbering from 1 to $c = \#\,\mathrm{Wh}(\mathcal{H})$ on the components of $\mathrm{Wh}(\mathcal{H})$. At each v_i we get a partition P_i into c subsets of the loose ends of $\mathrm{Wh}(v_i) \setminus \mathcal{H}$ at e_i by connectivity in $\mathrm{Wh}(\mathcal{H})$. Similarly, we get a partition P'_i of the loose ends of $\mathrm{Wh}(v_i) \setminus \mathcal{H}$ at e_{i+1}. These partitions are at least as coarse as connectivity in $\mathrm{Wh}([v_i, v_j]) \setminus \mathcal{H}$ for any $j \geq i$.

By construction, the splicing map at e_{i+1} connecting loose ends of $\mathrm{Wh}(v_i) \setminus \mathcal{H}$ at e_{i+1} to loose ends of $\mathrm{Wh}(v_{i+1}) \setminus \mathcal{H}$ at e_{i+1} is compatible with the partitions P'_i and P_{i+1}.

For each edge pair (e_i, e_{i+1}) there is a corresponding label pair L_i that gives a nontrivial word of length two in F. There are $2n(2n-1)$ such words.

Let m be the number of partitions of (complexity of Wh$(*)$) things into c nonempty subsets. Consider the segment $[v_0, v_R]$, where $R = 2n(2n-1)m$. Some label pair appears at least m times. Let $\{i_j\}_{j=1}^{m}$ be a set of indices such that the L_{i_j} are the same.

Let $g_{j,k}$ be the element of F that takes v_{i_j} to v_{i_k}.

If we fix P_{i_1} we get a map of the elements $g_{1,k}$ into the set of all possible partitions by $g_{1,k} \mapsto g_{1,k}P_{i_1}$, so for some $1 \leq j < k$ we have $g_{1,j}$ and $g_{1.k}$ mapping to the same partition. Therefore, $g_{j,k}P_{i_j} = P_{i_k}$.

$g = g_{i,j}$ is then the desired element. □

Remark In the preceding proof we found an element g such that the g–action preserved a partition. We did not insist that the g–action also fixed the numbering of components of Wh$(\mathcal{H})$; these may be permuted. The proof may easily be modified to fix the numbering, at the expense of a larger bound on $|g|$.

Corollary 4.15 *Existence of a cut pair implies existence of a cut point or bad cut pair.*

Corollary 4.16 *With R as in the previous proposition, for any pair of points $\{\mathcal{H}^+, \mathcal{H}^-\}$ in $\partial\mathcal{T}$, if $\mathcal{X}$ is a segment of the convex hull $\mathcal{H}$ of $\{\mathcal{H}^+, \mathcal{H}^-\}$ of length greater than R, and if there are no cut pairs in the decomposition space, then one of the following is true:*

(1) Wh$(\mathcal{X}) \setminus \mathcal{H}$ *is connected.*

(2) Wh$(\mathcal{X}) \setminus \mathcal{H}$ *has two components, one of which is a free edge.*

Theorem 4.17 (Detecting Cut Pairs) *There is a finite algorithm for detecting cut pairs in the decomposition space*

Proof Given a list of words generating a line pattern, apply Whitehead automorphisms, if necessary, to eliminate cut vertices. If the graph becomes disconnected, stop; the decomposition space is disconnected by Corollary 4.5.

If it is possible to disconnect the Whitehead graph by deleting the interiors of two edges, stop; these two edges correspond to a cut pair. In particular, this happens if the Whitehead graph has any valence two vertices.

Use Proposition 4.12 to check if any of the generators of the line pattern give a cut point in the decomposition space. If so, stop.

Let R be the constant from Lemma 4.14. The idea now is to check segments of length R to see if we can find a disconnected Whitehead graph. There are a lot of these.

We streamline the process by only checking those long segments for which every sub-segment gives a disconnected Whitehead graph.

Let $\mathcal{X}_0 = \{*\}$.

We proceed by induction. Suppose $\mathcal{X}_i$ is defined.

Start with $\mathcal{X}_{i+1} = \mathcal{X}_i$. Consider pairs of points v and v' such that $d(v, \mathcal{X}_i) = d(v', \mathcal{X}_i) = 1$, such that $d(v, *) = d(v', *) = i+1$, and such that $* \in [v, v']$. If $\mathrm{Wh}(\mathcal{X}_i \cap [v, v']) \setminus [v, v']$ is not connected, add v and v' to $\mathcal{X}_{i+1}$.

Continue until stage $k = 1 + \lceil \frac{R}{2} \rceil$. Apply Corollary 4.16 and Proposition 4.13: there are pairs v and v' in $\mathcal{X}_k \setminus \mathcal{X}_{k-1}$ with $* \in [v, v']$ such that $\mathrm{Wh}(\mathcal{X}_{k-1} \cap [v, v']) \setminus [v, v']$ has more than one component that is not just a free edge if and only if there are cut pairs in the decomposition space.

□

Corollary 4.18 *If a Whitehead graph for a line pattern has the property that deleting any pair of vertices leaves at most one free edge and at most one other connected component, then the decomposition space has no cut pairs.*

Unfortunately, this corollary does not apply if a Whitehead graph has more than one edge between a pair of vertices. Indeed, consider the pattern in $F_2 = \langle a, b \rangle$ generated by the word $a^2ba^2\bar{b}^2$. The Whitehead graph in Figure 15 is reduced and contains the complete graph on the four vertices, but $q(\{a^\infty, a^{-\infty}\})$ is a cut pair, as is evident from Figure 16.

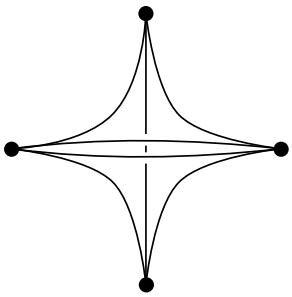

Figure 15: $\mathrm{Wh}(*)\{a^2ba^2\bar{b}^2\}$

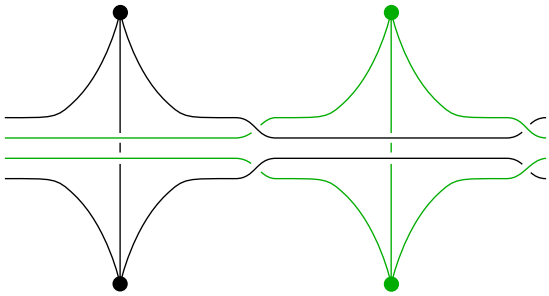

Figure 16: $\mathrm{Wh}([*, a])\{a^2ba^2\bar{b}^2\} \setminus [a^{-1}, a^2]$

4.6 Cut Sets When There are No Cut Pairs

Let S be a finite set in the decomposition space, and let $\mathcal{H}$ be the convex hull of $q^{-1}(S)$. Lemma 4.9 tells us that S is a cut set if and only if $\mathrm{Wh}(\mathcal{H})$ is disconnected. We will

pass to a finite subset of $\mathcal{H}$ whose Whitehead graph contains the same connectivity information.

Define the *core* $\mathcal{C}$ of $q^{-1}(S)$, to be the smallest closed, connected set such that $\mathcal{H} \setminus \mathcal{C}$ is a collection of disjoint infinite geodesic rays $\mathcal{R}_j\colon [1,\infty] \to \overline{\mathcal{T}}$. We use $\mathcal{R}_j(0)$ to denote the vertex of the core that is adjacent to $\mathcal{R}_j(1)$.

Let ξ be a point in $\partial\mathcal{T}$. If $q(\xi)$ is either a cut point or a bad point that is a member of a cut pair, it is not hard to see that there is a geodesic ray $\mathcal{R}$ with $\mathcal{R}(\infty) = \xi$ such that $\mathrm{Wh}(\mathcal{R}([1,\infty)) \setminus \mathcal{R}(0)$ is not connected.

Conversely, if there exists a geodesic ray $\mathcal{R}\colon [0,\infty] \to \overline{\mathcal{T}}$ with $\mathcal{R}(\infty) = \xi$ such that $\mathrm{Wh}(\mathcal{R}([1,\infty)) \setminus \mathcal{R}(0)$ is not connected, then, arguing as in the proof of Lemma 4.14, $q(\xi)$ is either a cut point or is a bad point that belongs to a cut pair.

If there are no cut points or cut pairs, then $\mathrm{Wh}(\mathcal{R}([1,\infty)) \setminus \mathcal{R}(0)$ is connected for any ray $\mathcal{R}$.

Proposition 4.19 *Suppose ξ is a point in $\partial\mathcal{T}$ such that $q(\xi)$ is a bad point that is not a member of a cut pair. Then $q(\xi)$ is not a member of any minimal finite cut set. In particular, if $\mathbf{D}$ has no cut pairs then no bad point belongs to any minimal finite cut set.*

Proof The assumption that $q(\xi)$ is not a member of a cut pair implies that for any ray $\mathcal{R}\colon [0,\infty] \to \overline{\mathcal{T}}$ with $\mathcal{R}(\infty) = \xi$, the Whitehead graph $\mathrm{Wh}(\mathcal{R}([1,\infty])) \setminus \mathcal{R}(0)$ is connected.

Let S be a finite cut set in $\mathbf{D}$ with $q(\xi) \in S$. Let $\mathcal{H}$ and $\mathcal{C}$ be the hull and core of $q^{-1}(S)$, respectively. Consider the ray $\mathcal{R}$ that is the component of $\mathcal{H} \setminus \mathcal{C}$ containing ξ.

Components of $\mathbf{D} \setminus S$ are in bijection with components of $\mathrm{Wh}(\mathcal{H})$, which, in turn, are in bijection with components of $\mathrm{Wh}(\mathcal{H} \setminus \mathcal{R}([1,\infty]))$, since $\mathrm{Wh}(\mathcal{R}([1,\infty])) \setminus \mathcal{R}(0)$ is connected. This is just the hull of $q^{-1}(S \setminus \{q(\xi)\})$.

Thus, $S \setminus \{q(\xi)\}$ is still a cut set, so S was not minimal. □

For a finite collection of lines $\tilde{S} = \{l_1, \ldots, l_k\} \subset \mathcal{L}$, the core is a finite tree. The convex hull minus the core is a collection of $2k$ disjoint rays:

$$\{\mathcal{R}_i^{\epsilon}\colon [1,\infty] \to \mathcal{T} \mid \lim_{t\to\infty} \mathcal{R}_i^{\epsilon}(t) = l_i^{\epsilon},\ \epsilon \in \{+,-\},\ i = 1 \ldots k\}$$

Lemma 4.20 *Let S be a finite set of good points of $\mathbf{D}$, none of which is a cut point. Components of $\mathbf{D} \setminus S$ are in bijection with components of $\mathrm{Wh}(\mathcal{C}) - \tilde{S}$*

Proof Suppose $\tilde{S} = \{l_1, \dots, l_k\}$.

For each i and ϵ, since $q_*(l_i)$ is not a cut point, $\mathrm{Wh}(\mathcal{R}_i^\epsilon([1,\infty])) \setminus \mathcal{R}_i^\epsilon(0)$ is connected.

$\mathrm{Wh}(\mathcal{H})$ is obtained from $\mathrm{Wh}(\mathcal{C}) - \{\ddot{l}_1, \dots, \ddot{l}_k\}$ by splicing on each $\mathrm{Wh}(\mathcal{R}_i^\epsilon([1,\infty])) \setminus \mathcal{R}_i^\epsilon(0)$.

This means to each deleted vertex of $\mathrm{Wh}(\mathcal{C}) - \{\ddot{l}_1, \dots, \ddot{l}_k\}$ we have spliced on a connected graph, so we might have just as well not deleted those vertices. □

In fact, we can use the argument of Lemma 4.20 to reduce the convex hull even further. If $\mathcal{C}$ is not just a vertex, then it has some valence one vertices that we call *leaves*. The edge connecting a leaf to the rest of the core is called the *stem*.

Suppose that for some leaf v of $\mathcal{C}$ every line of $\tilde{S}$ that goes through v goes through the stem of v. From $\mathrm{Wh}(v)$, delete the interiors of edges corresponding to the l_i and the vertex corresponding to the stem. The resulting graph is connected, so connected components of $\mathrm{Wh}(\mathcal{C}) - \tilde{S}$ are in bijection with connected components of $\mathrm{Wh}(\mathcal{C} \setminus \{v\}) - \tilde{S}$.

Thus, we may *prune* some leaves off of the core without changing the connectivity of the Whitehead graphs.

Prune off the leaves that do not have all of $\tilde{S}$ going through their stems, and repeat. The result is a well defined nonempty tree $p\mathcal{C}$, the *pruned core*. There are two possible outcomes: either every leaf contains a line of $\tilde{S}$ that does not go through the stem of that leaf, or the pruned core is a segment such that every line of $\tilde{S}$ runs over the entire segment. Edge cut sets (recall Definition 4.7) fall into the second category, and the segment is exactly the single edge defining the edge cut set. The segment could not be longer or else the Whitehead graph would have a cut vertex, contrary to hypothesis.

Remark Equivalently, the pruned core is the minimal subset of the tree that contains at least one point from every line of $\tilde{S}$ and all edges that are crossed by all of $\tilde{S}$.

Proposition 4.21 *An edge cut set is minimal if and only if it does not contain a cut point. In particular, if there are no cut points in the decomposition space, then edge cut sets are minimal.*

Proof Let $\tilde{S} = \{l_1, \dots, l_k\}$ be the set of lines of $\mathcal{L}$ going through an edge e of $\mathcal{T}$, so that $S = q_*(\tilde{S})$ is an edge cut set. We must have $k > 1$ or else $\mathrm{Wh}(*)$ has a cut vertex, contrary to hypothesis. If one of the points of S is a cut point, then this is a proper subset that is a cut set, so S is not minimal.

Conversely, suppose none of the points of S is a cut point. The pruned core $p\mathcal{C}$ of S is the single edge e defining S. There are at least two connected components of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$, lying on opposite sides of e. In fact, there are exactly two components, because otherwise $\mathrm{Wh}(*)$ has a cut vertex. By Lemma 4.20 these correspond to two connected components of $\mathbf{D} \setminus S$.

Each of the l_i has one endpoint in each component, so if any l_i is omitted from the set the two components will have a point in common, and we no longer have a cut set. Thus, S is minimal. □

Remark For a non-example consider Figure 9. The three lines crossing the b-edge give an edge cut set. However, one of these lines is the b–line through $*$, which already gives a cut point, so this edge cut set is not minimal.

Lemma 4.22 *Suppose there are no cut points in the decomposition space. If S is a cut set with pruned core $p\mathcal{C}$, then one of the following occurs:*

- *S is an edge cut set and $p\mathcal{C}$ is a single edge.*
- *$p\mathcal{C}$ is a single vertex.*
- *$p\mathcal{C}$ has leaves and through every leaf there is a line of $\tilde{S}$ that does not go through the stem.*

Proof Suppose there is a leaf for which every line of $\tilde{S}$ that goes through the leaf goes through the stem. Such a leaf would have been pruned off unless all of the lines of $\tilde{S}$ go through the stem. Thus, the pruned core is a segment such that all of the lines of $\tilde{S}$ run over the entire segment. This means that S is a subset of an edge cut set for each of the edges of the segment. Since S is a cut set and edge cut sets are minimal, S must be equal to the edge cut set for each of the edges of the segment. Since $\mathrm{Wh}(*)$ is connected without cut vertices, there is a unique such edge. □

Proposition 4.23 *If* $\mathbf{D}$ *has no cut pairs, the good points and bad points are topologically distinguished.*

Proof Every line in the pattern crosses edges, so every good point belongs to infinitely many edge cut sets, and edge cut sets are minimal.

Bad points are the points that do not belong to any minimal finite cut set. Good points are the points that do. □

Proposition 4.24 *Let S be a minimal finite cut set, none of whose elements are members of a cut pair. There are exactly two connected components of $\mathbf{D} \setminus S$.*

Proof By Proposition 4.19, S consists of good points. Let $\tilde{S} = \{l_1, \dots, l_k\} = q_*^{-1}(S)$. Components of $\mathbf{D} \setminus S$ are in bijection with components of $\mathrm{Wh}(\mathcal{C}) - \tilde{S}$. This is a finite graph, so $\mathbf{D} \setminus S$ has only finitely many components.

Let $A_1, \dots, A_m$ be a list of the components of $\mathbf{D} \setminus S$.

If $q_*(l_i)$ is not a limit point of A_j in $\mathbf{D}$ then A_j is still a connected component in $\mathbf{D} \setminus (S \setminus q_*(l_i))$. This contradicts minimality of S, so each point of S is a limit point in $\mathbf{D}$ of every A_j. This implies that for each i and j, at least one of the points l_i^+ and l_i^- is a limit point of $q^{-1}(A_j)$.

Now $\mathcal{H} \setminus \mathcal{C}$ is a collection of disjoint rays $\mathcal{R}_i^\epsilon$. The graph $\mathrm{Wh}(\mathcal{R}_i^\epsilon([1,\infty])) \setminus \mathcal{R}_i^\epsilon(0)$ is connected, so no l_i^+ or l_i^- is a limit point of more than one $q^{-1}(A_j)$.

Thus, there are exactly two components A_1 and A_2 of $\mathbf{D} \setminus S$, and each line l_i has one endpoint in $q^{-1}(A_1)$ and the other in $q^{-1}(A_2)$. □

Corollary 4.25 *Let S be a minimal finite cut set that is not an edge cut set, none of whose elements are members of a cut pair. For every vertex $v \in p\mathcal{C}$, the portion of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$ at v contains an edge from each component of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$.*

Proof If v is a leaf such that the portion of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$ at v belongs to a single component of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$ then v should have been pruned off.

If v is not a leaf, $p\mathcal{C} \setminus \{v\}$ has at least two components. If the Whitehead graph over one of those components sees only one component of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$ then it would have been possible to prune it off. Thus, every component of $p\mathcal{C} \setminus \{v\}$ must see two components of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$. There are only two components of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$, so both must be able to connect to all components of $p\mathcal{C} \setminus \{v\}$. In particular, they must connect through v. □

4.7 Indecomposable Cut Sets

In this section we will assume that the decomposition space has no cut pairs.

Our ultimate goal is to construct a cubing quasi-isometric to a bounded valence tree. For this purpose, we will need to choose a collection of cut sets in a such a way that there is a bound on the number of cut sets in the collection that cross any fixed cut set in the collection.

Cut sets with disjoint pruned cores do not cross, so we could control crossings if we could control the diameters of the pruned cores of the cut sets in some collection.

The following example shows that cut sets of a fixed size can have pruned cores with arbitrarily large diameter. We subsequently introduce the property of *indecomposability* to rule out this kind of bad behavior.

Let $\mathcal{L}$ be the line pattern in $F = \langle a, b \rangle$ generated by the words $ab\bar{a}\bar{b}$, a and b. The edge cut sets have size three. It can be shown that these are the only cut sets of size three and there are none smaller, see Section 6.2. It is also possible to find minimal cut sets of size four. Pick any two of the edge cut sets that share a line. The four lines of the symmetric difference are a minimal cut set. Figure 17 shows the line pattern. The two dashed lines indicate edge cut sets of size three. The four thickened lines make up the cut set of size four that is the symmetric difference. There is no bound on the size of the pruned core of such a cut set, nor on the number of such cut sets that cross each other. To see this, translate the two lines on the right by some high power of a, leaving the two lines on the left fixed. The result is still a four line cut set.

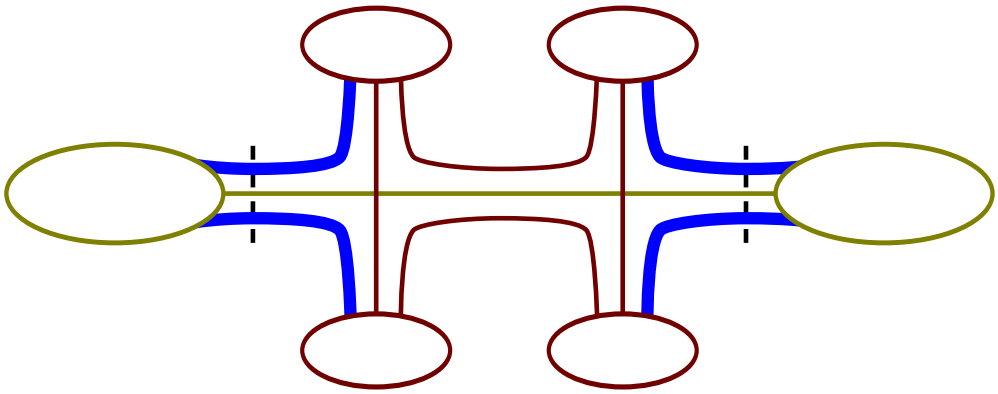

Figure 17: A problematic minimal cut set

We say that a minimal finite cut set $S \subset \mathbf{D}$ is *decomposable* if there are minimal cut sets Q and R such that:

(1) Q and R are non-crossing,

(2) $|Q| < |S|$ and $|R| < |S|$,

(3) $S = Q \Delta R = (Q \setminus R) \cup (R \setminus Q)$

A minimal finite cut set S is *indecomposable* if it is not decomposable.

The smallest cut sets in **D** are indecomposable since there are no smaller cut sets to decompose them into.

The following lemma is designed specifically to rule out the bad example described in Figure 17.

Lemma 4.26 *Suppose S is a finite minimal cut set and the pruned core $p\mathcal{C}$ of $\tilde{S}$ has an interior vertex v such that $\mathrm{Wh}(v) \setminus p\mathcal{C}$ has exactly two components, one of which is a free edge, and no lines of $\tilde{S}$ go through v. Then S is decomposable.*

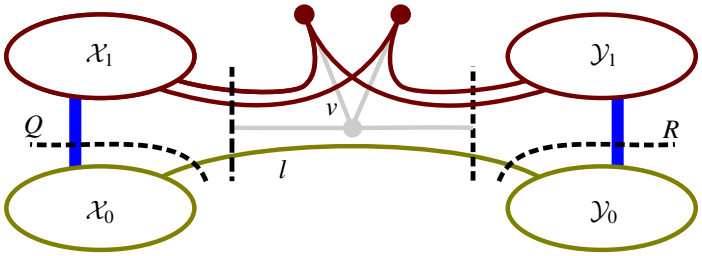

Figure 18: Schematic diagram of decomposable cut set

Proof Let l be the line of $\mathcal{L}$ corresponding to the free edge in $\mathrm{Wh}(v) \setminus p\mathcal{C}$. By Corollary 4.25, every component of $p\mathcal{C} \setminus \{v\}$ must see both components of $\mathrm{Wh}(p\mathcal{C}) \setminus \tilde{S}$. Thus, the line l enters every component of $p\mathcal{C} \setminus \{v\}$, which means that $p\mathcal{C} \setminus \{v\}$ has only two components. Let $\tilde{Q}$ be l together with the lines of $\tilde{S}$ on one side of $p\mathcal{C} \setminus \{v\}$, and let $\tilde{R}$ be l together with the lines of $\tilde{S}$ on the other side of $p\mathcal{C} \setminus \{v\}$.

Then $Q \cap R = q_*(l)$, and $S = Q \Delta R$.

Let A^0 and A^1 be the components of $\mathbf{D} \setminus S$. The line l does not belong to $\tilde{S}$, so we may assume that $q_*(l) \in A^0$. Let $\mathcal{X}$ and $\mathcal{Y}$ be the two components of $p\mathcal{C} \setminus \{v\}$. We may assume Q is on the $\mathcal{X}$ side.

Let $\mathcal{X}_0$ be the portion $\mathcal{X}$ corresponding to A^0, and define $\mathcal{X}_1$, $\mathcal{Y}_0$ and $\mathcal{Y}_1$ analogously, see Figure 18. The edge of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$ corresponding to l is the only connection between $\mathcal{X}_0$ and $\mathcal{Y}_0$, so $\tilde{Q}$ separates $\mathcal{X}_0$ from $\mathcal{X}_1 \cup \mathcal{Y}_0 \cup \mathcal{Y}_1$.

Thus, Q is a cut set. Moreover, Q is a minimal cut set since every edge corresponding to a line in $\tilde{Q}$ has one end in $\mathcal{X}_0$ and one end in $\mathcal{X}_1 \cup \mathcal{Y}_0 \cup \mathcal{Y}_1$.

By a similar argument, R is a minimal cut set.

Q and R are non-crossing because the only line of R that has an endpoint in $\mathcal{X}_0$ is $l = \tilde{Q} \cap \tilde{R}$.

Finally, as there are no cut pairs, we have:

$$3 \leq |Q| = |Q \setminus R| + |Q \cap R| = |Q \setminus R| + 1 \implies |Q \setminus R| \geq 2$$

Thus:

$$|S| = |Q \setminus R| + |R \setminus Q| > 1 + |R \setminus Q| = |R \cap Q| + |R \setminus Q| = |R|$$

Similarly, $|Q| < |S|$.

□

Theorem 4.27 (Edge cut sets are indecomposable) *Suppose we have chosen a free basis of F such that* Wh$(*)$ *is minimal complexity. Then edge cut sets are indecomposable.*

Proof Let e be an edge of $\mathcal{T}$. Let $\tilde{S}$ be the lines of $\mathcal{L}$ that cross e. Let $S = q_*(\tilde{S})$.

S is minimal by Proposition 4.21. Suppose S decomposes into Q and R. We must have $Q \cap R \neq \emptyset$, otherwise Q and R are proper subsets of S that are cut sets, contradicting minimality of S. Since Q and R do not cross and $S \setminus R = Q \setminus R$, it follows that S does not cross R. Thus, since they are minimal, R does not cross S. Therefore, $R \setminus S = R \cap Q$ is contained in one component of $\mathbf{D} \setminus S$. This means that $q_*^{-1}(Q \cap R) = \tilde{Q} \cap \tilde{R}$ is contained in one component of $\mathcal{T} \setminus e$. Let $*$ be the vertex of $\mathcal{T}$ incident to e on the $\tilde{Q} \cap \tilde{R}$ side. It follows that the cores of Q and R are contained in the component of $\mathcal{T} \setminus e$ containing $*$.

We define a version of pruned core relative to $*$. If $|R \setminus Q| = 1$ it is possible that $*$ is not in the core of R, so begin with the convex hull of $* \cup \mathcal{C}_R$. Now declare that $*$ is not a leaf, and prune leaves as in the construction of the pruned core. This time, however, we will also prune a leaf even if all of the lines of $\tilde{R}$ go through the stem. Call the resulting set the *partially pruned core* of $\tilde{R}$, and denote it $pp\mathcal{C}_{\tilde{R}}$. This partially pruned core is the minimal connected set that contains $*$ and at least one vertex of each line of $\tilde{R} \cap \tilde{Q}$, so, in fact, $pp\mathcal{C}_{\tilde{Q}} = pp\mathcal{C}_{\tilde{R}}$, and we may omit the subscripts.

$$|R \setminus Q| + |Q \cap R| = |R| < |S| = |Q \setminus R| + |R \setminus Q|$$

So $|Q \cap R| < |Q \setminus R|$. Similarly, $|Q \cap R| < |R \setminus Q|$.

There are two connected components of Wh$(pp\mathcal{C}) - \tilde{Q}$, call them component 0 and component 1. Since Q and S do not cross, everything on the side of e opposite $\tilde{Q} \cap \tilde{R}$ belongs to a single component.

First suppose $ppC = *$. Suppose the edge e oriented away from $*$ has label $x \in \mathcal{B} \cup \bar{\mathcal{B}}$; suppose the corresponding vertex in $\mathrm{Wh}(ppC) - \tilde{Q}$ is in component 1. Suppose the vertex corresponding to the edge labeled $\bar{x}$ is in component 0. Then the Whitehead automorphism that pushes the vertices in $\mathrm{Wh}(*)$ in component 1 through x changes the valence at x from $|S| = |Q \setminus R| + |R \setminus Q|$ to $|Q \cap R| + |R \setminus Q|$. Since $|Q \cap R| < |Q \setminus R|$ this contradicts the assumption that the Whitehead graph had minimal complexity.

Conversely, if the vertex $\bar{x}$ is in component 1 push

$$Z = \{x\} \cup \{\text{vertices of component } 0\}$$

through x and get a contradiction.

Now suppose ppC is not just $*$. Then there is some leaf $v \neq *$. Suppose the stem of v (oriented away from the leaf) has label $x \in \mathcal{B} \cup \bar{\mathcal{B}}$, and suppose the vertex in $\mathrm{Wh}(ppC) - \tilde{Q}$ corresponding to $\bar{x}$ is in component 1.

Figure 19 shows a schematic diagram of $\mathrm{Wh}(ppC)$.

The labeling in the diagram is as follows:

- $\mathcal{X}_0$ = the portion of component 0 on the v side of the stem.
- $\mathcal{X}_1$ = the portion of component 1 on the v side of the stem.
- $\mathcal{Y}_0$ = the portion of component 0 between the stem of v and e.
- $\mathcal{Y}_1$ = the portion of component 1 between the stem of v and e.
- $\mathcal{Z}$ = everything on the side of e opposite $\tilde{Q} \cap \tilde{R}$.
- lowercase letters represent the number of lines with endpoints in the specified regions with:
  - { a, b, c and h counting the lines of $\tilde{Q} \cap \tilde{R}$
  - { d and i counting the lines of $\tilde{Q} \setminus \tilde{R}$
  - { e and j counting the lines of $\tilde{R} \setminus \tilde{Q}$
  - { f and g counting the lines not in $\tilde{Q} \cup \tilde{R}$ crossing the stem.

Since $v \neq *$ is a leaf of ppC we must have $a > 0$ (this follows from Lemma 4.22) and $\mathcal{X}_0$ and $\mathcal{X}_1$ nonempty.

Minimality of Q implies that $\mathrm{Wh}(ppC) - \tilde{Q}$ has exactly two connected components. In the diagram they are $\mathcal{X}_0 \cup \mathcal{Y}_0$ and $\mathcal{X}_1 \cup \mathcal{Y}_1 \cup \mathcal{Z}$.

$\mathcal{X}_0 \cup \mathcal{Y}_0$ must belong to a connected component, so $\mathcal{Y}_0 = \emptyset$ if and only if $f = 0$. $\mathcal{Y}_0 = \emptyset$ also implies $c = h = i = 0$. Now $d + i = |Q \setminus R|$, so this would imply $d > 0$.

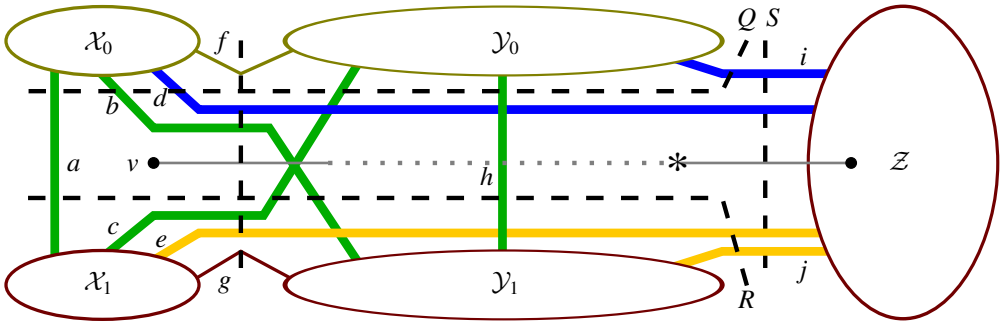

Figure 19: Schematic diagram for Wh(ppC)

R is also a minimal cut set, so Wh(ppC) $-\tilde{R}$ must have exactly two components. In the diagram they are $\mathcal{X}_0 \cup \mathcal{Y}_0 \cup \mathcal{Z}$ and $\mathcal{X}_1 \cup \mathcal{Y}_1$. Since $\mathcal{X}_1 \cup \mathcal{Y}_1$ is connected, $\mathcal{Y}_1 = \emptyset$ if and only if $g = 0$.

Thus, we have:

(1) $\mathcal{Y}_0 = \emptyset \iff f = 0 \implies d > 0,\ c = h = i = 0$

(2) $\mathcal{Y}_1 = \emptyset \iff g = 0 \implies e > 0,\ b = h = j = 0$

The Whitehead automorphism that pushes

$$Z = \{x\} \cup \{\text{vertices of Wh}(v) \text{ in component 0 of Wh}(ppC) - \tilde{Q}\}$$

through the stem changes the valence of vertex x from $b + c + d + e + f + g$ to $a + c + e + g$. By our minimal complexity assumption, we must therefore have $a \geq b + d + f$.

Now $|Q \setminus R| > |Q \cap R| \geq a + b + c \geq 2b + c + d + f$, which means that $|Q \setminus R| > d$, so $i > 0$.

We will now change Q and R to a new decomposing pair Q' and R' for S with strictly smaller partially pruned core.

Let $\tilde{Q}' \setminus \tilde{R}'$ be the lines of $\tilde{Q} \setminus \tilde{R}$ that do not pass through v.

Let $\tilde{R}'$ be the rest of $\tilde{S}$.

Let $\tilde{Q}' \cap \tilde{R}'$ be $\tilde{Q} \cap \tilde{R}$ minus the lines contributing to a and b plus the lines contributing to f.

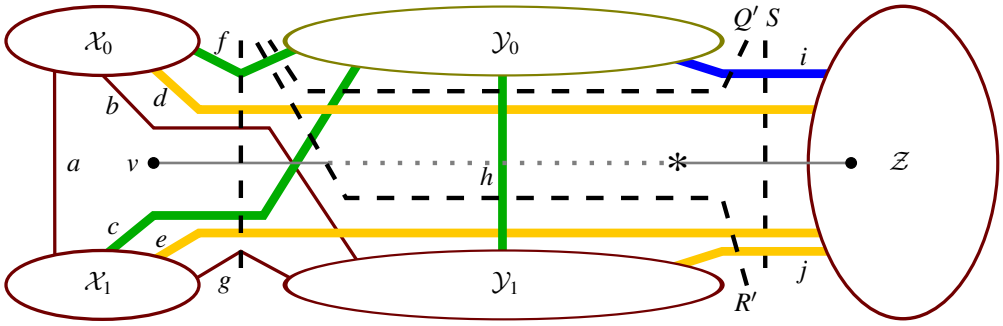

Figure 20: Schematic diagram for modified $\mathrm{Wh}(ppC)$

$$\begin{aligned}
|S| - |R'| &= |Q' \setminus R'| - |Q' \cap R'| \\
&= |Q \setminus R| - d - (|Q \cap R| - a - b + f) \\
&= |Q \setminus R| - |Q \cap R| + a + b - d - f \\
&\geq |Q \setminus R| - |Q \cap R| + (b + d + f) + b - d - f \\
&= |S| - |R| + 2b \geq |S| - |R| > 0
\end{aligned}$$

A similar computation shows $|S| - |Q'| \geq |S| - |Q| + 2(b+d) > 0$.

We must show that Q' and R' are non-crossing minimal cut sets.

The components of $\mathrm{Wh}(ppC) - \tilde{Q}'$ are $\mathcal{Y}_0$ and $\mathcal{X}_0 \cup \mathcal{X}_1 \cup \mathcal{Y}_1 \cup \mathcal{Z}$. To see that the latter is connected, note that $a > 0$, $e + j > 0$ and either $\mathcal{Y}_1 = \emptyset$ or $g > 0$.

Thus, Q' is a minimal cut set since $\mathrm{Wh}(ppC) - \tilde{Q}'$ has exactly two connected components and every line of $\tilde{Q}'$ goes from one component to the other.

By similar considerations, R' is a minimal cut set since $\mathrm{Wh}(ppC) - \tilde{R}'$ has components $\mathcal{Y}_0 \cup \mathcal{Z}$ and $\mathcal{X}_0 \cup \mathcal{X}_1 \cup \mathcal{Y}_1$.

That Q' and R' are non-crossing follows from the observation:

$$\mathcal{Y}_0 \cap (\mathcal{X}_0 \cup \mathcal{X}_1 \cup \mathcal{Y}_1) = \emptyset$$

We have not added anything to $\tilde{Q}' \cap \tilde{R}'$ except possibly some lines going through v, so the new partially pruned core is contained in the old one minus the vertex v.

If $\bar{x}$ is in component 0, repeat the argument with the roles of Q and R reversed and reach a similar conclusion.

Thus, by repeating this process, we can reduce the partially pruned core until we find some decomposing pair Q and R so that the partially pruned core is just $*$. We have already seen that that leads to a contradiction, so S is indecomposable. □

Theorem 4.28 (Indecomposables are bounded) *The diameter of the pruned core of an indecomposable cut set S is bounded in terms of $\mathcal{L}$ and $|S|$.*

Proof If $p\mathcal{C}$ is a point or a single edge we are done. Otherwise it is a tree with leaves. Each leaf contains a line from $\tilde{S}$ that does not go through its stem, so there are at most $|S|$ leaves.

Suppose $\mathcal{X}$ is a segment of $p\mathcal{C}$ that does not have any lines of $\tilde{S}$ going through it. By Corollary 4.25, at every vertex of $p\mathcal{C}$ there are edges of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$ from both components. Since S is indecomposable, by Lemma 4.26 it is not the case that one of these components is a free edge. Now apply Corollary 4.16 and conclude that there is a bound R on the length $\mathcal{X}$.

Similarly, if $\mathcal{X}$ is a segment of $p\mathcal{C}$ that meets exactly one of the l_i then it has length bounded by R.

It follows that the diameter of $p\mathcal{C}$ is at most $2R(|S| - 1)$. □

5 Rigidity

In the next two sections we prove the Main Theorem. In Section 5.1 we show that if there are cut pairs then the line pattern is not rigid. In Section 5.2 we prove the converse.

5.1 The Problem with Cut Pairs

If **D** has cut pairs then it has either a cut point or a bad cut pair, by Corollary 4.15. In either case, there is a cut set such that the preimage in $\partial\mathcal{T}$ is two points $\{g^{\infty}, g^{-\infty}\}$. The convex hull $\mathcal{H}$ of these two points is a line, and $\mathrm{Wh}(\mathcal{H})$ has multiple components, $A_1, \ldots, A_k$. For each i, let $\mathcal{X}_i$ be the union of components of $\mathcal{T} \setminus \mathcal{H}$ corresponding to A_i.

The action of g may permute these components, but $g^{k!}$ fixes them.

Let $\phi\colon\ \mathcal{T} \to \mathcal{T}$ be the quasi-isometry:

$$\phi(x) = \begin{cases} g^{k!}x & \text{if } x \in \mathcal{X}_1 \\ x & \text{otherwise} \end{cases}$$

This "shearing" quasi-isometry moves the $\mathcal{X}_1$ component along $\mathcal{H}$, fixing the rest of the tree.

It is not hard to see that ϕ^n is not bounded distance from an isometry for $n \neq 0$. Since F acts by isometries it follows that $F\phi^a$ and $F\phi^b$ are not the same coset of F in $\mathcal{QI}(F,\mathcal{L})$ when $a \neq b$, so F is an infinite index subgroup.

It is possible to show directly that ϕ could not be conjugate into an isometry group. Alternatively, notice that we can stack shearing quasi-isometries to produce a sequence of quasi-isometries with unbounded multiplicative quasi-isometry constants, see Figure 21. Take an element h of F such that $h\mathcal{H}$ is contained in the $\mathcal{X}_1$ component with $h\mathcal{X}_1 \subset \mathcal{X}_1$.

The desired sequence of quasi-isometries is (Φ_i), where:

$$\Phi_i = \phi^i \circ (h\phi^i\bar{h}) \circ (h^2\phi^i\bar{h}^2) \circ \cdots$$

That is, for any $x \in \mathcal{T}$ there exists some j such that $x \in h^j\mathcal{X}_1 \setminus h^{j+1}\mathcal{X}_1$. For $m > j$ $h^m\phi^i\bar{h}^m(x) = x$, so:

$$\Phi_i(x) = \phi^i \circ (h\phi^i\bar{h}) \circ (h^2\phi^i\bar{h}^2) \circ \cdots \circ (h^j\phi^i\bar{h}^j)(x)$$

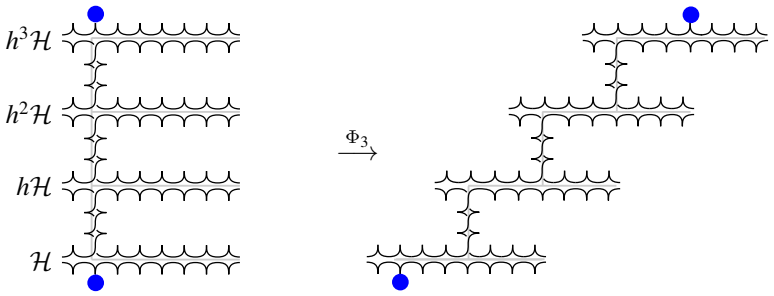

Figure 21: Shearing

5.2 Rigidity When There are No Cut Pairs

Let $\mathcal{L}$ be a line pattern in F. Choose a free basis $\mathcal{B}$ for F so that $\mathrm{Wh}(*) = \mathrm{Wh}_{\mathcal{B}}(*)\{\mathcal{L}\}$ has minimal complexity, and let $\mathcal{T}$ be the Cayley graph of F with respect to $\mathcal{B}$. Assume that $\mathbf{D} = \mathbf{D}_{\mathcal{L}}$ has no cut pairs. We will construct a cubing X, a quasi-isometry $\phi\colon \mathcal{T} \to X$, and an isometric action of $\mathcal{QI}(F,\mathcal{L})$ on X that agrees with $\phi\,\mathcal{QI}(F,\mathcal{L})\phi^{-1} \subset \mathrm{Isom}(X)$, up to bounded distance, completing the proof of the Main Theorem. The action of F on X will be cocompact, implying that $\mathcal{QI}(F,\mathcal{L})$ has a complex of groups structure.

5.2.1 Constructing the Cubing

We will now choose a collection of "small" cut sets and show that the corresponding cubing has the desired properties. The guiding principles in defining a suitable collection of cut sets are that:

(1) there are few enough cut sets so that the cubing is finite dimensional,

(2) there are enough cut sets so that the cubing is finite valence, and

(3) the cut sets are defined in terms of the topology of $\mathbf{D}$, so that homeomorphisms of $\mathbf{D}$ preserve the collection of cut sets, and hence the cubing.

Remark Notice that the cubing associated to the collection of edge cut sets is just the tree $\mathcal{T}$. However, the property of being an edge cut set depends on the choice of basis, so pattern preserving quasi-isometries do not necessarily preserve this cubing.

Let b be the maximum valence of a vertex in $\mathrm{Wh}(*)$. Let $\{S_i\}_{i\in I}$ be the collection of indecomposable cut sets of size at most b in $\mathbf{D}$. These are the "small" cut sets that we will use to build the cubing.

Remark We have chosen two topological properties enjoyed by the edge cut sets, size at most b and indecomposability, and defined our collection by these properties. These choices work. Other choices might also work. It is not strictly necessary that the chosen collection contains the edge cut sets, but this is convenient in showing that the cubing is quasi-isometric to $\mathcal{T}$, see Theorem 5.4.

For each $i \in I$, S_i is a finite minimal cut set, by definition, and, since there are no cut pairs, $\mathbf{D} \setminus S_i$ has exactly two connected components, by Proposition 4.24.

Let the two connected components of $\mathbf{D} \setminus S_i$ be A_i^0 and A_i^1. Let

$$\Sigma = \{A_i^0\}_{i\in I} \cup \{A_i^1\}_{i\in I}$$

Recall from Section 2.1 that from this information we define a graph as follows:

A vertex is a subset V of Σ such that:

(1) For all $i \in I$ exactly one of A_i^0 or A_i^1 is in V.

(2) If $C \in V$ and $C' \in \Sigma$ with $C \subset C'$ then $C' \in V$.

Two vertices are connected by an edge if they differ by only one set in Σ.

This gives a graph; it remains to select a path connected component of this graph to be the 1-skeleton of the cubing.

Define a *bad triple* $\bar{x} = \{x_1, x_2, x_3\}$ to be an unordered triple of distinct bad points in $\mathbf{D}$.

There are no bad points in minimal cut sets, so for any bad triple and any S_i, $\bar{x} \subset A_i^0 \cup A_i^1$. We let $\bar{x}$ decide democratically whether it will affiliate with A_i^0 or A_i^1: say $\bar{x} \in A_i^\epsilon$ if at least two of the x_j's are in A_i^ϵ.

Define $V_{\bar{x}} = \{A_i^\epsilon \in \Sigma \mid \bar{x} \in A_i^\epsilon\}$. This is a vertex of X. Define the 0-skeleton of the cubing to be the set $X^{(0)}$ of all vertices that are connected by a finite edge path to $V_{\bar{x}}$ for some bad triple $\bar{x}$.

The following lemma replaces Lemma 3.4 of [17].

Lemma 5.1 *For any bad triples $\bar{x}$ and $\bar{y}$, there are only finitely many S_i separating them.*

Proof Let $\bar{x} = \{x_1, x_2, x_3\}$ and $\bar{y} = \{y_1, y_2, y_3\}$ be bad triples.

The preimage $q^{-1}(\bar{x}) = \{q^{-1}(x_i)\}_{i=1,2,3}$ consists of three distinct points in $\partial\mathcal{T}$. The convex hull of three points in the boundary of a tree is a tripod. The core, as previously defined, is the unique vertex that is the branch point of the tripod. Denote this point $\mathcal{C}_{\bar{x}}$.

It is not hard to see that a cut set S_i separates $\bar{x}$ from $\bar{y}$ only if the pruned core $p\mathcal{C}$ of S_i intersects the finite geodesic segment joining $\mathcal{C}_{\bar{x}}$ and $\mathcal{C}_{\bar{y}}$ in $\mathcal{T}$.

By Theorem 4.28, there is a uniform bound a on the diameter of the pruned core of any S_i. Since $\mathcal{L}$ is locally finite, this means there is a uniform bound c on the number of S_i such that $* \in p\mathcal{C}_{S_i}$. If $\mathcal{Y}$ is any finite collection of vertices in $\mathcal{T}$, the number of S_i such that $p\mathcal{C}_{S_i} \cap \mathcal{Y} \neq \emptyset$ is at most $c|\mathcal{Y}|$.

Thus, the number of S_i separating $\bar{x}$ from $\bar{y}$ is at most $c \cdot \left(1 + d_{\mathcal{T}}(\mathcal{C}_{\bar{x}}, \mathcal{C}_{\bar{y}})\right)$. □

Add edges to the 0-skeleton as above to get the 1-skeleton $X^{(1)}$ of the cubing. With Lemma 5.1 replacing Lemma 3.4 of [17], the following theorem follows by the same proof as in [17]:

Theorem 5.2 [17, Theorem 3.3] *$X^{(1)}$ is connected.*

The rest of the construction of the cubing follows as in Section 2.1.

Remark We are forced to use this alternate way of choosing the vertices of the cubing because every good point in $\mathbf{D}$ belongs to infinitely many of the cut sets.

Also, Lemma 5.1 is false if one tries to use just bad points instead of bad triples. Two bad points are separated by infinitely many of the S_i.

Remark For a fixed vertex $v \in \mathcal{T}$, there are uncountably many bad triples $\bar{x}$ with $\mathcal{C}_{\bar{x}} = v$. However, these give only finitely many distinct vertices $V_{\bar{x}}$ in X, because the $V_{\bar{x}}$ can only differ in the finitely many coordinates i such that the pruned core of S_i contains v. Even this is an over count. If S_e is an edge cut set associated to an edge e incident to v, then every bad triple with $\mathcal{C}_{\bar{x}} = v$ lies in the same component of $\mathbf{D} \setminus S_e$. If our set of indecomposables is exactly the collection of edge cut sets then the cubing is isomorphic to the tree $\mathcal{T}$.

Notice X is defined in terms of the topology of $\mathbf{D}$, so we have:

Lemma 5.3 *Any homeomorphism of* $\mathbf{D}$ *induces an isomorphism of* X.

5.2.2 Estimates on the Cubing

Recall from the proof of Lemma 5.1, we have a bound a on the diameter of the pruned core of any S_i, and there is a c such that if $\mathcal{Y}$ is any finite collection of vertices in $\mathcal{T}$, the number of S_i such that $p\mathcal{C}_{S_i} \cap \mathcal{Y} \neq \emptyset$ is at most $c|\mathcal{Y}|$.

S_i and S_j are non-crossing if their pruned cores are disjoint, so we have a uniform bound $c(2n)^{\frac{a}{2}}$ on the number of S_j that cross a fixed S_i.

A k-cube in X corresponds to a collection of k pairwise crossing cut sets, so the cubing is finite-dimensional.

Pick a vertex $x \in X$. Let e and e' be edges incident to x. There are distinct hyperplanes H_e and $H_{e'}$ associated to these edges. Since e and e' are incident to a common vertex, there is no third hyperplane separating H_e from $H_{e'}$. Therefore, the valence of a vertex in X is bounded by the maximum size of a subcollection $\{S_i\}_{i\in J}$ of the indecomposable cut sets such that for any j and k in J, there is no $i \in I$ such that S_i separates S_j and S_k. If S_j and S_k have disjoint pruned cores then there is an edge cut set separating them, so the maximum size of the set J is at most $c(2n)^{\frac{a}{2}}$. Thus, X is uniformly locally finite.

A hyperplane H corresponds to an equivalence class of edges in X. The 1-neighborhood of H is the set of vertices that are endpoints of these edges. If k is the number of hyperplanes crossing H, then the 1-neighborhood of H has at most 2^{k+1} vertices and diameter at most $k+1$. Crossing hyperplanes correspond to crossing cut sets, so k is at most $c(2n)^{\frac{a}{2}}$.

5.2.3 The Rigidity Theorem

Theorem 5.4 *X is quasi-isometric to $\mathcal{T}$.*

Proof For each edge $e \in \mathcal{T}$ there is a corresponding edge cut set S_e. By construction, $S_e \in \{S_i\}$, so in the cubing X there is a corresponding hyperplane H_e. Define $\phi(e)$ to be the set of vertices in the 1-neighborhood of H_e. Recall from the previous section that this is a set of boundedly many vertices with bounded diameter.

$d_X(\phi(e), \phi(e'))$ is the number of hyperplanes separating H_e and $H_{e'}$. This is at least the number of edges separating e from e' in $\mathcal{T}$, which is $d_{\mathcal{T}}(e, e')$, and at most the number of $\{S_i\}$ such that $p\mathcal{C}_{S_i}$ meets the geodesic between e and e' in $\mathcal{T}$, which is bounded by $c \cdot d_{\mathcal{T}}(e, e')$. This shows that ϕ is a quasi-isometric embedding.

Suppose there is a vertex $x \in X$ not in the image of ϕ. This x has some incident edge, corresponding to some $S \in \{S_i\}$. The hypothesis that x is not in the image of ϕ implies that S does not cross any edge cut set, which means that $p\mathcal{C}_s$ is a single vertex $v \in \mathcal{T}$. There are boundedly many such S, and the distance from x to $\phi(\mathcal{T})$ is less than this bound, so ϕ is coarsely onto, hence a quasi-isometry. □

The quasi-isometry ϕ gives a collection of quasi-lines $\phi(\mathcal{L})$ in X. In fact, we can see this collection of quasi-lines directly from $\mathbf{D}$. Each good point in $\mathbf{D}$ belongs to infinitely many indecomposable cut sets. For $l \in \mathcal{L}$, the collection $\{S \mid S \text{ indecomposable}, |S| \leq b, q_*(l) \in S\}$ corresponds to a collection of hyperplanes in X. The union of these hyperplanes is coarsely equivalent to $\phi(l)$.

Theorem 5.5 (Rigidity Theorem) *For $i = 1, 2$, let F_i be a free group with line pattern $\mathcal{L}_i$. Let $\mathbf{D}_i$ be the decomposition space corresponding to $\mathcal{L}_i$ in F_i.*

Suppose, for each i, $\mathbf{D}_i$ has no cut pairs.

Let $\phi_i \colon F_i \to X_i$ be the quasi-isometry to the cube complex constructed above. Then:

$$\phi_2 \, \mathcal{QI}\{(F_1, \mathcal{L}_1) \to (F_2, \mathcal{L}_2)\} \phi_1^{-1} = \mathrm{Isom}\{(X_1, \phi_1(\mathcal{L}_1)) \to (X_2, \phi_2(\mathcal{L}_2))\}$$

Proof Elements of $\mathcal{QI}\{(F_1, \mathcal{L}_1) \to (F_2, \mathcal{L}_2)\}$ give homeomorphisms $\mathbf{D}_1 \to \mathbf{D}_2$. Homeomorphisms take indecomposable cut sets to indecomposable cut sets of the same size and preserve crossing and intersection. Therefore, we get isometries $X_1 \to X_2$ respecting the line patterns. □

The Rigidity Theorem answers Questions 1 and 2 for rigid patterns.

The free group acts on itself by pattern preserving isometries via left multiplication. Let $*$ be the identity vertex in $\mathcal{T}$. For any indecomposable cut set S, there is an element $g \in F$ such that $* \in g(p\mathcal{C}_S)$. There are only finitely many indecomposable cut sets of bounded size with $* \in p\mathcal{C}$, so F acts cocompactly on X. Therefore, $\mathcal{QI}(F, \mathcal{L}) \cong \mathrm{Isom}(X, \phi(\mathcal{L}))$ acts cocompactly on X. This gives an explicit presentation for $\mathrm{Isom}(X, \phi(\mathcal{L}))$ as a finite complex of groups. Moreover, the F action is already cocompact, so by applying Theorem 2.3 and the subsequent remarks, we have:

Corollary 5.6 *If $\mathcal{L}$ is a rigid line pattern and if $\mathcal{QI}(F, \mathcal{L})$ acts on X with finite stabilizers then F is a finite index subgroup of $\mathcal{QI}(F, \mathcal{L})$.*

One might hope that the finite stabilizers hypothesis in the preceding corollary is always true for rigid patterns, so that F is always finite index in $\mathcal{QI}(F, \mathcal{L})$. The example in Section 6.3 shows this is not true.

6 Examples

6.1 Whitehead Graph is the Circle

We will show in this section that when the Whitehead graph is a circle we get a quasi-isometrically flexible line pattern.

Theorem 6.1 *For a line pattern $\mathcal{L}$ in F_n, the following are equivalent:*

(1) *Every Whitehead graph $\mathrm{Wh}_{\mathcal{B}}(*)$ that has no cut vertex is a circle.*

(2) *Some Whitehead graph $\mathrm{Wh}_{\mathcal{B}}(*)$ is a circle.*

(3) *$\mathcal{L}$ is generated by the boundary words of a compact surface with boundary with fundamental group F_n.*

(4) ***D** is a circle.*

(5) *Every minimal cut set of **D** has size two.*

Otal [15, Theorem 2] proved essentially the same theorem using essentially the same argument.

Proof Clearly (1) $\Longrightarrow$ (2), because Whitehead automorphisms will eliminate cut vertices.

If some Whitehead graph $\mathrm{Wh}_{\mathcal{B}}(*)$ is a circle, find a planar embedding of it. Fill the interior of this circle with a $2n$–gon. Blow up each vertex of the $2n$–gon to a segment to give a $4n$–gon. Identify these newly blown up sides in pairs using the splicing maps. The result is a compact surface with fundamental group F_n such that the words labeling boundary components generate $\mathcal{L}$. Thus (2) $\Longrightarrow$ (3).

We can give this surface a hyperbolic metric so that the universal cover is just $\mathcal{T}$ fattened, and the boundary components are horocycles that are in bijection with the lines of $\mathcal{L}$. This gives us a homeomorphism between the decomposition space and $S^1 = \partial\mathbb{H}^2$. Thus (3) $\Longrightarrow$ (4).

(4) $\Longrightarrow$ (5) is a topological property of circles.

Now, suppose every minimal cut set of $\mathbf{D}$ has size two. Since $\mathbf{D}$ is connected, every Whitehead graph for $\mathcal{L}$ is connected. Let $\mathcal{B}$ be a free basis such that $\mathrm{Wh}_{\mathcal{B}}(*)$ has no cut vertices. (Such a basis always exists, by Whitehead's Algorithm.) The edges incident to a vertex of $\mathrm{Wh}_{\mathcal{B}}(*)$ correspond to an edge cut set. This is a minimal cut set by Proposition 4.21, so by hypothesis has size two. Therefore, $\mathrm{Wh}_{\mathcal{B}}(*)$ is a finite, connected, valence two graph, hence, a circle. Thus, (5) $\Longrightarrow$ (1). □

Theorem 6.2 *Let F and F' be free groups, possibly of different rank. Let $\mathcal{L}$ and $\mathcal{L}'$ be line patterns in F and F', respectively. Suppose $\mathbf{D}_{\mathcal{L}}$ is a circle. There is a pattern preserving quasi-isometry from F to F' if and only if $\mathbf{D}_{\mathcal{L}'}$ is also a circle.*

Proof The "only if" direction is clear, as a pattern preserving quasi-isometry induces a homeomorphism of decomposition spaces.

Suppose both $\mathbf{D}_{\mathcal{L}}$ and $\mathbf{D}_{\mathcal{L}'}$ are circles. By Theorem 6.1, there exist free bases $\mathcal{B}$ of F and $\mathcal{B}'$ of F' such that $\mathrm{Wh}_{\mathcal{B}}(*)\{\mathcal{L}\}$ and $\mathrm{Wh}_{\mathcal{B}'}(*)\{\mathcal{L}'\}$ are circles.

As in the proof of Theorem 6.1 we can associate each pattern with the boundary curves of the universal cover of a surface with boundary. It is a theorem of Behrstock and Neumann [1] that there are many boundary preserving quasi-isometries of such surfaces. □

We saw some circle patterns in the Introduction. Let $F = F_2 = \langle a, b\rangle$.

Let $\mathcal{L}_1$ be the line pattern generated by the word $ab\bar{a}\bar{b}$.

Let $\mathcal{L}_2$ be the line pattern generated by the words ab and $a\bar{b}$.

For each of these $\mathrm{Wh}(*)\{\mathcal{L}_i\}$ is a circle, so the two patterns are quasi-isometrically equivalent.

This example also shows that neither the number of generators of a line pattern nor the widths of the generators are quasi-isometry invariants.

6.2 Whitehead Graph is the Complete Graph

Let K_{2n} be the complete graph on $2n$ vertices, the graph consisting of $2n$ vertices with exactly one edge joining each pair of vertices.

Suppose $\mathcal{L}$ is a line pattern in $F = F_n$ so that for some free basis $\mathcal{B}$, $\mathrm{Wh}_{\mathcal{B}}(*)\{\mathcal{L}\} = K_{2n}$.

The decomposition space $\mathbf{D}$ has no cut pairs.

Suppose S is a minimal finite cut set of $\mathbf{D}$ that is not an edge cut set. $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$ has two components. The portion of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$ at a leaf contains vertices from both components.

The portion of $\mathrm{Wh}(p\mathcal{C}) - \tilde{S}$ at a leaf is a graph obtained from K_{2n} be deleting a vertex, corresponding to the stem of the leaf, and interiors of some number of edges coming from lines of $\tilde{S}$ that go through the leaf but not through the stem. The result is a disconnected graph with at least one vertex in each of the components. Thus, we have partition of $2n-1$ vertices into two subsets, and we must delete all the edges between them. The subsets have sizes m and $2n-1-m$, for some $1 \leq m \leq 2n-2$, and the number of edges between them is $m(2n-1-m) \geq 2n-2$. There are at least two leaves, so $|S| \geq 4n-4 > 2n-1$. The edge cut sets have size $2n-1$, so our construction of a cubing uses only the edge cut sets. Thus, the cubing is just the tree $\mathcal{T}$.

Suppose there are two different $2n$–valent trees $\mathcal{T}$ and $\mathcal{T}'$, each with a line pattern whose Whitehead graph is K_{2n}. Choose any vertices $* \in \mathcal{T}$ and $*' \in \mathcal{T}'$, and define $\phi(*) = *'$. Choose any vertex $v \in \mathcal{T}$ adjacent to $*$ and any vertex $v' \in \mathcal{T}'$ adjacent to $*'$, and define $\phi(v) = v'$. Choose any bijection $\tilde{\phi}$ between the lines crossing $[*, v]$ and the lines crossing $[*', v']$. We claim that these choices determine a pattern preserving isometry ϕ. In particular, suppose we have chosen to identify lines l and l'. The line l goes through one other vertex w adjacent to $*$, and the line l' goes through one other vertex w' adjacent to $*'$. Define $\phi(w) = w'$, and continue to extend ϕ in this manner.

In this case it is easy to compute:

$$\mathcal{QI}(F, \mathcal{L}) \cong \mathrm{Sym}(2n) *_{\mathrm{Sym}(2n-1)} (\mathrm{Sym}(2n-1) \times \mathrm{Sym}(2))$$

Here, $\mathrm{Sym}(2n)$ is the symmetric group on $2n$ objects, stabilizing a vertex of the tree and permuting the incident edges, and $\mathrm{Sym}(2n-1) \times \mathrm{Sym}(2)$ is the stabilizer of an edge of $\mathcal{T}$, which should be subdivided to avoid edge inversions.

The preceding discussion proves:

Theorem 6.3 *Suppose $\mathcal{L}$ is a line pattern in $F = F_n$ such that $\mathrm{Wh}_{\mathcal{B}}(*)\{\mathcal{L}\} = K_{2n}$. Suppose that $F' = F_m$ is another free group, possibly of different rank, with line pattern $\mathcal{L}'$.*

There is a pattern-preserving quasi-isometry $F \to F'$ if and only if $\mathbf{D}_{\mathcal{L}'}$ has the following properties:

(1) *There are no cut sets of size less than $2n-1$.*

(2) *The collection of cut sets of size $2n-1$ yields a cubing that is a $2n$-valent tree.*

(3) *The induced line pattern in the cubing restricts to the complete graph K_{2n} in the star of any vertex.*

For example, the line pattern $\mathcal{L}$ in $F = F_2$ with basis $\mathcal{B} = \{a, b\}$ generated by a, b, and $ab\bar{a}\bar{b}$ has Whitehead graph $\mathrm{Wh}_{\mathcal{B}}(*)\{\mathcal{L}\} = K_4$.

Compare this to the line pattern $\mathcal{L}'$ in $F' = F_3$ with basis $\mathcal{B}' = \{x, y, z\}$ generated by y, zx, $z\bar{x}\bar{y}$ and $xy\bar{z}$. The Whitehead graph $\mathrm{Wh}_{\mathcal{B}'}(*)\{\mathcal{L}'\}$ looks like two copies of K_4 spliced together, see Figure 22.

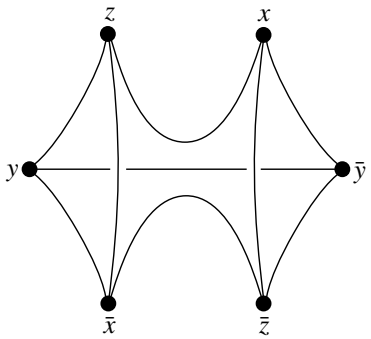

Figure 22: $\mathrm{Wh}_{\{x,y,z\}}(*)\{y, zx, z\bar{x}\bar{y}, xy\bar{z}\}$

It is not hard to show that the smallest cut sets are the obvious ones of size three. These yield a cubing that is a 4-valent tree, essentially blowing up each vertex of F_3 into a pair of vertices.

This pattern is quasi-isometric to the K_4 pattern in F_2.

6.3 A Rigid Example for which the Free Group is not Finite Index in the Group of Pattern Preserving Quasi-isometries

Consider the line pattern in $F = \langle a, b \rangle$ generated by the words a, b, and $ab\bar{a}b\bar{a}b\bar{a}\bar{b}$. Let $\mathcal{T}$ be the Cayley graph of F with respect to $\{a, b\}$.

It is easy to check that Whitehead graph in Figure 23 is reduced and the decomposition space has no cut pairs, so the pattern is rigid.

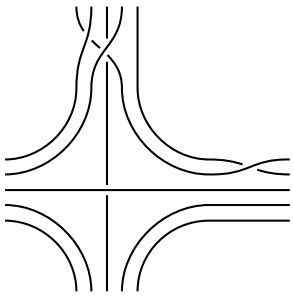

Figure 23: Wh$(*)\{a, b, ab\bar{a}b\bar{a}b\bar{a}\bar{b}\}$ (loose)

The edge cut sets have size five. Deleting any vertex of the Whitehead graph leaves a graph that requires at least three more edges to be deleted to disconnect the graph. Thus, any other cut sets have size at least six. As the edge cut sets are the only cut sets of size less than or equal to five, the rigid cube complex is just the tree $\mathcal{T}$.

We will show that F is not a finite index subgroup of $\mathcal{QI}(F, \mathcal{L})$. Not only are the vertex stabilizers in $\mathcal{QI}(F, \mathcal{L})$ not finite, they are not even finitely generated.

Define an isometry ϕ of $\mathcal{T}$ piecewise as follows. First, note that the automorphism α of F that exchanges a with $\bar{a}$ preserves the pattern. It inverts a, fixes b, and takes $ab\bar{a}b\bar{a}b\bar{a}\bar{b}$ to a cyclic permutation of itself. To the branch of the tree consisting of words beginning with b, apply the automorphism α. To each branch of the tree beginning with $a^n b$ for some n, apply the automorphsim $a^n \circ \alpha \circ \bar{a}^n$. Fix the rest of the tree.

The isometry ϕ is built piecewise from pattern preserving automorphisms of F. It fixes the "bottom half" of $\mathcal{T}$, fixes the b–line through a^n for each n, and reflects each branch beginning with $a^n b$ through the b–line through a^n.

There are lines of the pattern that pass through multiple pieces of the domain of ϕ, so we check that the ϕ is defined consistently for these lines. As illustrated in Figure 24, the only lines shared by the bottom half of the tree and the vertical branches are the fixed b–lines (green lines in the figure are fixed). The reflections in adjacent vertical

branches agree on the two lines they share (the two thickened blue lines are exchanged). Therefore, ϕ pieces together to give a pattern preserving isometry.

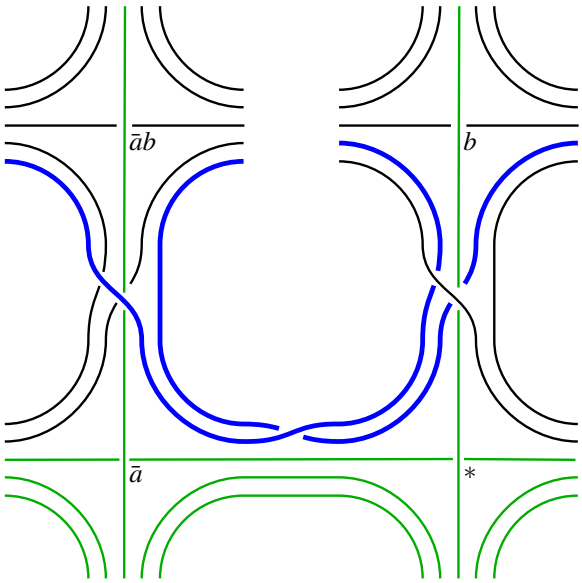

Figure 24: $\mathrm{Wh}([\bar{a}b, b])\{a, b, ab\bar{a}b\bar{a}\bar{b}\}$ (loose)

Thus, for any n, $b^n \circ \phi \circ \bar{b}^n$ is a pattern preserving isometry that fixes every line in the n-neighborhood of the identity vertex, but is not the identity map. It follows that the stabilizer of the identity vertex is not finitely generated.

6.4 A Cube Complex That is Not a Tree

Finally, we give an example of a rigid line pattern for which our argument does not produce a cube complex that is a tree.

Consider the line pattern in $F_3 = \langle a, b, c\rangle$ generated by the four words $\bar{a}bc$, $\bar{a}cb$, $\bar{a}b^3$ and $\bar{a}c^3$. The Whitehead graph (with loose ends), is shown in Figure 25.

The reader may verify that this is a minimal Whitehead graph and there are no cut points or cut pairs in the decomposition space. In fact, the smallest cut sets are the edge cut sets of size four corresponding to the a–edges. These are the only cut sets of size four.

The other edge cut sets have size five, so we construct a cube complex using indecomposable cut sets of size four and five. Figure 26 depicts the Whitehead graph (along with

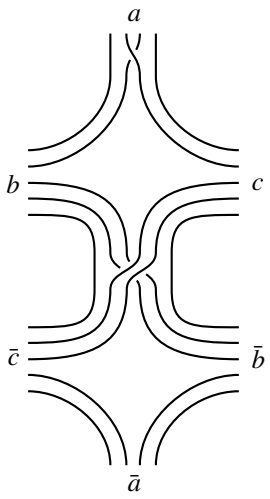

Figure 25: Wh(∗){$\bar{a}bc, \bar{a}cb, \bar{a}b^3, \bar{a}c^3$} (loose)

portions of the Whitehead graph over two neighboring vertices) with the 1–skeleton of the cube complex overlaid.

Note that every cut set of size five is crossed by another cut set of size five. However, the edge cut sets are still topologically distinguished! They are the cut sets of size five that are crossed minimally (once) by another cut set of size five. The other cut sets of size five are crossed by either two or five other cut sets of size five.

Had we said, "build the cube complex associated to the cut sets of size four and those of size five that are crossed by exactly one other cut set of size five" we would have recovered the tree as the cube complex.

In every example we know, it is possible, after computing the cube complex, to pick out a topologically distinguished collection of cut sets whose associated cube complex is a tree. We do not know if this is true in general.

References

[1] Jason A. Behrstock and Walter D. Neumann, *Quasi-isometric classification of graph manifold groups*, Duke Math. J. **141** (2008), 217–240, doi:10.1215/S0012-7094-08-14121-3, arXiv:math.GT/0604042.

[2] Brian H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186, doi:10.1007/BF02392898.

[3] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der mathematischen Wissenschaften, vol. 319, Springer, Berlin, 1999.

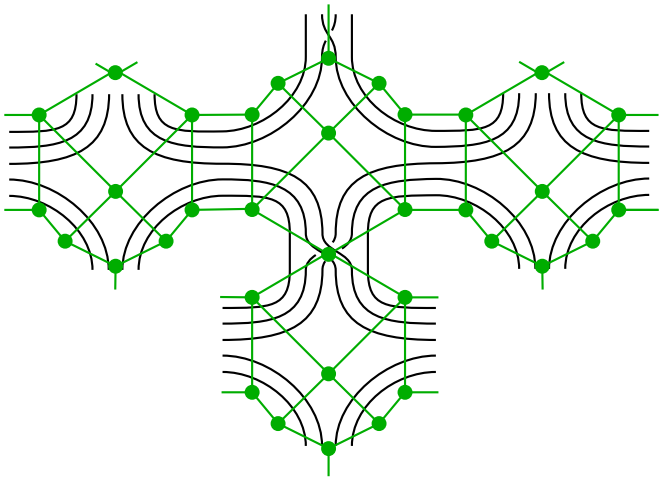

Figure 26: Whitehead graph with cube complex

[4] Christopher H. Cashen, *Splitting line patterns in free groups*, preprint, arXiv:1009.2492, 2010.

[5] Christopher H. Cashen and Nataša Macura, *Quasi-isometries of mapping tori of linearly growing free group automorphisms*, in preparation, 2011.

[6] Martin J Dunwoody, *The accessibility of finetly presented groups*, Invent. Math. **81** (1985), no. 3, 449–457, doi:10.1007/BF01388581.

[7] M. Gromov, *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263. MR919829(89e:20070)

[8] Michael Kapovich, *Energy of harmonic functions and Gromov's proof of Stallings' theorem*, preprint, arXiv:0707.4231, 2007.

[9] Seonhee Lim and Anne Thomas, *Covering theory for complexes of groups*, J. Pure Appl. Algebra **212** (2008), no. 7, 1632–1663, doi:10.1016/j.jpaa.2007.10.012, arXiv:math/0605303.

[10] Roger C. Lyndon and Paul E. Schupp, *Combinatorial group theory*, Classics in Mathematics, Springer-Verlag, Berlin, 2001, Reprint of the 1977 edition. MR1812024(2001i:20064)

[11] Jason Fox Manning, *Virtually geometric words and Whitehead's algorithm*, Math. Res. Lett. **17** (2010), no. 5, 917–925, arXiv:0904.0724.

[12] Reiner Martin, *Non-uniquely ergodic foliations of thin type, measured curents and automorphisms of free groups*, Ph.D. thesis, UCLA, 1995.

[13] Graham A. Niblo, *A geometric proof of Stallings' theorem on groups with more than one end*, Geom. Dedicata **105** (2004), 61–76, doi:10.1023/B:GEOM.0000024780.73453.e4. MR2057244(2005e:20060)

[14] Jakob Nielsen, *Die isomorphismengruppe der freien gruppen*, Math. Ann. **91** (1924), 169–209, doi:10.1007/BF01556078.

[15] Jean-Pierre Otal, *Certaines relations d'equivalence sur l'ensemble des bouts d'un groupe libre*, J. London Math. Soc. **s2-46** (1992), no. 1, 123–139, doi:10.1112/jlms/s2-46.1.123. MR1180888(93h:20041)

[16] Panos Papasoglu, *Quasi-isometry invariance of group splittings*, Ann. of Math. (2) **161** (2005), no. 2, 759–830, doi:10.4007/annals.2005.161.759.

[17] Michah Sageev, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. (3) **71** (1995), no. 3, 585–617, doi:10.1112/plms/s3-71.3.585. MR1347406(97a:20062)

[18] Richard Evans Schwartz, *Symmetric patterns of geodesics and automorphisms of surface groups*, Invent. Math. **128** (1997), no. 1, 177–199, doi:10.1007/s002220050139.

[19] John R. Stallings, *Whitehead graphs on handlebodies*, Geometric group theory down under (Canberra, 1996), de Gruyter, Berlin, 1999, pp. 317–330. MR1714852(2001i:57028)

[20] Richard Stong, *Diskbusting elements of the free group*, Math. Res. Lett. **4** (1997), no. 2, 201–210. MR1453054(98h:20049)

[21] J. H. C. Whitehead, *On certain sets of elements in a free group*, Proc. London Math. Soc. (2) **41** (1936), no. 2097, 48–56.

[22] ———, *On equivalent sets of elements in a free group*, Ann. of Math. (2) **37** (1936), no. 4, 782–800, http://www.jstor.org/stable/1968618. MR1503309

Department of Mathematics, University of Utah
Salt Lake City, UT 84112, USA

Department of Mathematics, Trinity University
San Antonio, TX 78212, USA

cashen@math.utah.edu, nmacura@trinity.edu

http://www.math.utah.edu/~cashen