

Moduli of graphs and automorphisms of free groups*

Marc Culler¹ and Karen Vogtmann²

¹ Department of Mathematics, Rutgers University, New Brunswick, NJ, USA

² Department of Mathematics, Columbia University, New York, NY 10027, USA

0. Introduction

This paper represents the beginning of an attempt to transfer, to the study of outer automorphisms of free groups, the powerful geometric techniques that were invented by Thurston to study mapping classes of surfaces. Let F_n denote the free group of rank n . We will study the group $\text{Out}(F_n)$ of outer automorphisms of F_n by studying its action on a space X_n which is analogous to the Teichmüller space of hyperbolic metrics on a surface; the points of X_n are metric structures on graphs with fundamental group F_n . We begin by making this notion precise.

By a *graph* we shall mean a connected 1-dimensional CW-complex. The 0-cells will be called *nodes* and the 1-cells *edges*. The *valence* of a node x is the number of oriented edges which terminate at x , i.e. the minimum number of components of an arbitrarily small deleted neighborhood of x . An *\mathbb{R} -graph* is a graph endowed with a metric such that each edge is locally isometric to an interval in \mathbb{R} and such that the distance between two points is the length of the shortest edge-path joining them. An \mathbb{R} -graph is said to be *minimal* if it is not homotopy equivalent to any proper subgraph. Throughout this paper we will consider only \mathbb{R} -graphs which are minimal and have no nodes of valence 2. (A minimal \mathbb{R} -graph cannot have nodes of valence 1).

Fix a (topological) graph R_0 with one node and n edges, and choose an identification $F_n \cong \pi_1(R_0)$. If G is an \mathbb{R} -graph, then a homotopy equivalence $g: R_0 \rightarrow G$ is called a *marking* on G . We define two markings $g_1: R_0 \rightarrow G_1$ and $g_2: R_0 \rightarrow G_2$ to be *equivalent* if there exists an isometry $i: G_1 \rightarrow G_2$ making the following diagram commute up to (free) homotopy:

$$\begin{array}{ccc}
 & & G_1 \\
 & \nearrow^{g_1} & \downarrow i \\
 R_0 & & G_2 \\
 & \searrow_{g_2} &
 \end{array}$$

* Partially supported by research grants from the National Science Foundation

An equivalence class of markings will be called a *marked graph*, and the notation (g, G) will be used for the equivalence class containing the marking $g: R_0 \rightarrow G$.

Observe that if G is an \mathbb{R} -graph then the universal cover \tilde{G} is a simplicial \mathbb{R} -tree in the sense of Morgan and Shalen [9]. A marking $g: R_0 \rightarrow G$ determines an isometric action of $F_n = \pi_1(R_0)$ on \tilde{G} . Two marked graphs are equivalent if and only if the associated actions are conjugate by an isometry between the \mathbb{R} -trees.

A marking $g: R_0 \rightarrow G$ determines a real-valued length function on F_n . If $w \in F_n$ is a word, and g_* is the induced map on π_1 , then the length $l(w)$ of w is the length of the shortest loop in G in the free homotopy class determined by $g_*(w)$. Since the length of a word is determined by a free homotopy class, we have $l(w) = l(x^{-1}wx)$, i.e. l is actually a length function on conjugacy classes of words in F_n . It is clear that equivalent marked graphs induce the same length function on F_n .

This length function l is a special case of those considered in [9], which are defined as follows. Given an isometric action of a group Γ on an \mathbb{R} -tree and an element x of Γ which does not have a fixed point, there is an invariant line in the tree upon which x acts by translation. The length of x is defined to be the translation distance along this axis. Elements with fixed points are defined to have length zero. The action is called *minimal* if there are no invariant proper sub-trees, and is said to be *abelian* if every element of $[\Gamma, \Gamma]$ has length zero. It is shown in [2] that non-abelian minimal actions on \mathbb{R} -trees are determined up to conjugacy by the associated length function. If G is a minimal \mathbb{R} -graph then a marking $g: R_0 \rightarrow G$ determines a minimal action of F_n on \tilde{G} , since otherwise G would have a proper subgraph with isomorphic fundamental group. The action is free and hence non-abelian. Thus the length function associated with an equivalence class of marked graphs is unique.

By the *volume* of an \mathbb{R} -graph we mean the sum of the lengths of the edges. We now define the space X_n to be the set of all marked graphs for which the underlying minimal \mathbb{R} -graph has fundamental group of rank n and has volume 1. Let \mathcal{C} denote the set of all conjugacy classes in F_n . By the remarks above, X_n may be embedded in $\mathbb{R}^{\mathcal{C}}$ by sending a marked graph to its length function; the maps $(g, G) \rightarrow l(c)$ for $c \in \mathcal{C}$ give coordinate maps for this embedding. Clearly the origin is not in the image of this map. Moreover, the uniqueness statement above together with our condition that the volume of an \mathbb{R} -graph be 1 imply that distinct elements of X_n do not have length functions which are positive scalar multiples of each other. Thus, in fact, we obtain an embedding $\theta: X_n \rightarrow \mathbb{P}^{\mathcal{C}}$ where $\mathbb{P}^{\mathcal{C}}$ is the infinite dimensional projective space associated to $\mathbb{R}^{\mathcal{C}}$. We give X_n the topology induced by the map θ .

There is a natural right action of the automorphism group $\text{Aut}(F_n)$ on X_n given as follows: any automorphism α can be realized by a map $A: R_0 \rightarrow R_0$. If $(g, G) \in X_n$, then $(g, G) \cdot \alpha = (g \circ A, G)$. This action is well defined, and the group $\text{Inn}(F_n)$ of inner automorphisms acts trivially; thus we have an action of $\text{Out}(F_n) = \text{Aut}(F_n)/\text{Inn}(F_n)$ on X_n . Note that the stabilizer of any marked graph (g, G) is isomorphic to the group of isometries of G , and is hence finite.

It is also shown in [2] that the subspace of $\mathbb{IP}^{\mathcal{G}}$ consisting of all length functions which arise from isometric actions of a given finitely generated group on \mathbb{R} -trees is compact, as is the smaller space of length functions associated to actions for which the edge stabilizers do not have non-abelian free subgroups. (An edge stabilizer is a subgroup which fixes each point of a non-degenerate embedded arc in the \mathbb{R} -tree.) The actions associated to points of X_n are free, and hence the edge stabilizers are trivial. Therefore the closure \bar{X}_n of X_n in $\mathbb{IP}^{\mathcal{G}}$ is compact, and the ideal points of this compactification are length functions arising from isometric actions, with cyclic edge stabilizers, of F_n on (possibly non-simplicial) \mathbb{R} -trees. The action of $\text{Out}(F_n)$ on X_n clearly extends to an action on \bar{X}_n .

Note that this situation is very similar to that which arises in Thurston's classification of surface automorphisms [11, 3]. We are hoping that the same program could succeed here. In particular, we would like to consider the action of a single outer automorphism on \bar{X}_n , to find a fixed point for this homeomorphism, and to analyze the structure of the automorphism in terms of the fixed point. Finding such a fixed point requires some knowledge of the algebraic topology of \bar{X}_n . Our main theorem is a major step in this direction.

Theorem. *The space X_n is contractible.*

It seems likely that to prove a structure theorem for outer automorphisms of F_n one will have to understand the entire space \bar{X}_n of real-valued length functions. However, the result which we will prove here really only involves the combinatorial structure of a simplicial spine of the space X_n . Specifically, we show that there is an equivariant deformation retraction of X_n onto a contractible simplicial complex of dimension $2n - 3$.

The action of $\text{Out}(F_n)$ on X_n and hence on this subcomplex has finite stabilizers and finite quotient. The group $\text{Out}(F_n)$ is known to have torsion-free subgroups of finite index; hence these subgroups have finite classifying spaces of dimension at most $2n - 3$. Thus we obtain the following.

Corollary. *$\text{Out}(F_n)$ is of type VFL and $\text{vcd}(\text{Out}(F_n)) = 2n - 3$.*

We remark that the complex mentioned above was independently studied by Gersten, who simultaneously found a somewhat different proof that it is contractible. The equality in the corollary is his observation; it is an immediate consequence of the fact that $\text{Out}(F_n)$ contains free abelian subgroups of rank $2n - 3$. An example of such a subgroup is the image in $\text{Out}(F_n)$ of the subgroup in $\text{Aut}(F_n)$ generated by the automorphisms $\alpha_i: x_i \rightarrow x_1 x_i$ and $\beta_i: x_i \rightarrow x_i x_1$ for $1 < i \leq n$.

1. A deformation retract of X_n

Fix n and set $F = F_n$ and $X = X_n$.

1.1. Definition of the retract K

Each marked graph (g, G) in X determines a family of marked graphs, obtained by varying the metric on G . That is, the family consists of all marked graphs (g_1, G_1) in X for which there exists a homeomorphism $h: G \rightarrow G_1$ making the diagram

$$\begin{array}{ccc} & & G \\ & \nearrow g & \downarrow h \\ R_0 & & \\ & \searrow g_1 & G_1 \end{array}$$

homotopy commutative. Clearly the marked graph (g_1, G_1) is uniquely determined by the lengths of the edges $h(e)$, as e ranges over the edges of G . We may thus parametrize the marked graphs in this family as points in the positive cone of a Euclidean space for which the coordinates are indexed by the edges of the graph G . Since marked graphs in X have volume 1, this parameter space is actually an open k -simplex σ , where $k+1$ is the number of edges of G . The map $\theta: \sigma \rightarrow \mathbb{P}^{\mathcal{E}}$, which sends a point of σ to the projectivized length function associated with the corresponding marked graph, is injective by the uniqueness result mentioned above and has linear coordinate functions. Thus θ is a homeomorphism onto its image in $\mathbb{P}^{\mathcal{E}}$ and has a continuous extension to a map $\bar{\theta}$ from the closure $\bar{\sigma}$ of σ into $\mathbb{P}^{\mathcal{E}}$. Each open face of $\partial\bar{\sigma}$ is obtained by assigning length zero to certain edges of G . The map $\bar{\theta}$ will map the open face into X if and only if the graph G' , obtained by collapsing these edges of G , is homotopy equivalent to G . In this case the restriction of $\bar{\theta}$ to the open face is exactly the parametrization of the family of length functions associated to G' as described above.

Thus X is the union of a set of disjoint open simplices. The closure of any of these simplices in X is homeomorphic to a closed simplex minus a collection of closed faces, and will be called an *ideal simplex*. The faces which are missing from the ideal simplex determined by (g, G) correspond to subsets of the edges of G whose union contains a non-trivial loop. A maximal ideal simplex corresponds to a graph G with the maximal number of edges; since such a maximal graph has $3n-3$ edges, we have $\dim(X) = 3n-4$.

Note that if two of these ideal simplices meet, then their intersection is an (ideal) face of each simplex. Thus there exists a simplicial complex X^* in which X embeds as a union of open simplices. This complex is constructed from a disjoint union of closed simplices, containing a closed k -simplex for each ideal k -simplex in X , by identifying a closed k -simplex with a face of a closed $k+1$ -simplex whenever the same relation holds among the corresponding ideal simplices in X .


At this point it is convenient to replace our space X by a deformation retract Y which is somewhat simpler. The subspace Y consists of those points of X represented by marked graphs (g, G) with the property that G has no separating edges. If σ is an ideal simplex of X represented by a marked graph (g, G) , where G has separating edges, then σ meets Y in the face obtained by assigning length zero to each separating edge. The deformation retraction can

be defined by uniformly shrinking the lengths of the separating edges to zero while uniformly increasing the lengths of the other edges to preserve the volume. This deformation retraction is equivariant with respect to the action of $\text{Out}(F)$. The construction above realizes Y as a union of open simplices in a subcomplex Y^* of X^* , and this retraction extends to a deformation retraction of X^* onto Y^* .

Let $\partial Y^* = Y^* - Y$. Let Y' denote the barycentric subdivision of Y^* , and let K be the maximal full subcomplex of Y' which is disjoint from ∂Y^* . Then K is a deformation retract of Y : The deformation retraction is performed by collapsing every simplex τ in Y' to the face of τ which is contained in K . Again, this retraction can be done equivariantly. The action of $\text{Out}(F)$ extends to a simplicial action on K .

A vertex of K is the barycenter of a simplex in Y^* ; i.e. it is a marked graph (g, G) such that all of the edges of G have the same length. Therefore by retracting Y onto K we are ignoring the metric structure on the marked graphs and concentrating on the combinatorial structure of the space Y .

Suppose that (g, G) and (g', G') are vertices of K . The open simplex in X determined by (g, G) is a face of that determined by (g', G') if and only if G is obtained from G' by collapsing a set of edges and g equals the composition of g' with the collapsing map. We will say that a vertex (g, G) of K is obtained from (g', G') by *blowing down* an edge e of G' if there is a cellular homotopy equivalence $d: G' \rightarrow G$ which collapses the edge e , and $g' \simeq d \circ g$. Thus a $k+1$ -tuple $((g_0, G_0), \dots, (g_k, G_k))$ of vertices of K forms a k -simplex if G_i can be obtained from G_{i-1} by a composition d_i of edge collapses, and the diagram



is homotopy commutative. In other words, K is homeomorphic to the geometric realization of the category of equivalence classes of marked graphs, in which an arrow is a sequence of blowings down.

We can determine the dimension of K by noting that blowing down an edge of a marked graph decreases the number of nodes by 1 (otherwise the Euler characteristic would increase). Since a maximal graph has $2n-2$ nodes, we can do at most $2n-3$ blowings down, i.e. $\dim K = 2n-3$.

1.2. Labelled graphs

A vertex (g, G) of K can be represented by a labelled graph. Each edge in the complement of a maximal tree in G is assigned a label, where a label consists of an orientation and a word in the standard basis of $F = \pi_1(R_0)$. The word assigned to a labelled edge is defined as follows. Choose a homotopy inverse to g which collapses the maximal tree to the node of R_0 . Each labelled edge is mapped to a loop in R_0 ; the edge is assigned the word which corresponds to the based homotopy class of this loop.




Fig. 1

Note that there are many (in fact infinitely many) labelled graphs which represent the same vertex of K . This is the case, for example, with all of the labelled graphs shown in Fig. 1.

Different choices of a maximal tree will result in different labelled graphs that represent the same vertex (graphs #2 and #3), as will different choices for the inverse to g (#3 and #4). In the first case the words in the labels will be replaced by their images under one of a finite number of automorphisms of F , while in the second they will be replaced by their images under an inner automorphism. Of course, if the labels are permuted by an automorphism of the graph (#1 and #2), it will still describe the same vertex of K . In particular, if an edge e of G is attached to a single node, then the orientation can be omitted from the label on that edge.

1.3. Roses

We define the *level* of a vertex (g, G) of K to be the number of nodes of G . A vertex of level 1 will also be called a *rose*. By the above remarks, the roses are in one-to-one correspondence with conjugacy classes of unoriented, unordered bases of F , and the action of $\text{Out}(F)$ on the set of roses is transitive. It is clear that every vertex v of K is contained in the star of at least one rose, e.g. the rose obtained by collapsing all of the unlabelled edges of a labelled graph representing v . We summarize these remarks as:

1.3.1. Gertrude Stein Lemma. The complex K is the union of the stars of the roses. The stars of any two roses are homeomorphic. \square

We will prove our theorem by showing that the complex K is contractible. The strategy of the proof is this. By the Gertrude Stein Lemma, K is the union of the stars of the roses. Suppose that $W \subset \mathcal{C}$ is an arbitrary finite set of conjugacy classes in F . We define an integer-valued function $\|\cdot\|_W$ on roses by setting

$$\|\rho\|_W = n \sum_{w \in W} l(w), \quad (1.3.2)$$

where l is the length function on F associated to ρ . The factor of n accounts for the fact that the marked graph representing ρ has edges of length $1/n$. We show that this function behaves like a non-singular Morse function on K . More specifically, let ' $<$ ' be an arbitrary partial ordering on the set of roses in K which respects $\|\cdot\|_W$ in the sense that if $\|\rho\|_W < \|\rho'\|_W$, then $\rho < \rho'$. Then for each rose ρ for which $\|\rho\|_W$ is not minimal,

$$st(\rho) \cap \bigcup_{\rho' < \rho} st(\rho')$$

is contractible. Thus K is homotopy equivalent to the subcomplex K_{\min} which is the union of stars of roses ρ for which $\|\rho\|_W$ is minimal. By choosing W carefully we can arrange that K_{\min} be the star of a single rose, which implies that K is contractible. This strategy reduces the proof to a local analysis; we need only understand the full subcomplex of $st(\rho)$ given above.

1.4. Examples

For $n=2$, the space Y can be easily described; it is homeomorphic to the upper half-plane, and its decomposition into ideal triangles is isomorphic to the decomposition which is equivariant with respect to the action of $GL(2, \mathbb{Z})$ via the full group of isometries of the hyperbolic plane. (It is well known that $\text{Out}(F_2)$ is isomorphic to $GL(2, \mathbb{Z})$.) We will use labelled graphs to represent the marked graphs that determine the open simplices in Y . Note that, up to homeomorphism, there are only two graphs with fundamental group of rank 2 that have no vertices of valence 1 or 2 and no separating edges. These are shown in Fig. 2.

Suppose that (g, G) is a marked graph which is the barycenter of a 2-simplex σ in Y_2 . Then G must be homeomorphic to the graph in Fig. 2 which has three edges, and each edge must have length $1/3$. With respect to the length function associated with (g, G) , exactly six conjugacy classes in F_2 have length $2/3$. Moreover there exists a basis $\{u, v\}$ of F_2 such that the set $\{u, u^{-1}, v, v^{-1}, uv, v^{-1}u^{-1}\}$ is a complete set of representatives for these classes. Thus (g, G) is described by any of the labelled graphs in Fig. 3.

The barycenters of the 1-faces of σ are described by the labelled graphs shown in Fig. 4.




Fig. 2




Fig. 3




Fig. 4




Fig. 5

The group F_2 has the property that any automorphism which induces the identity on its abelianization is inner. It follows that the six-element set above is unique up to the action of the inner automorphism group of F_2 . If we now fix a basis of $F_2/[F_2, F_2]$, we may associate to each primitive element x of F_2 an element of $\mathbb{Q} \cup \{\infty\}$ by taking the ratio of the coordinates of its image in the abelianization. This number determines the set $\{x, x^{-1}\}$ up to conjugacy. Thus the six elements $\{u, u^{-1}, v, v^{-1}, uv, v^{-1}u^{-1}\}$ give rise to three extended rational numbers a/b , c/d , and $(a+c)/(b+d)$ such that $ad - bc = \pm 1$. Conversely these numbers characterize the six elements up to the action of $\text{Inn}(F_2)$.

Thus each ideal 2-simplex in Y is uniquely described by a set of three rational numbers as above. It contains all of its open 1-faces, each of which is determined by two of the three numbers and consequently is contained in exactly one other 2-simplex. The vertices of each 2-simplex are missing. It follows that Y is homeomorphic to the upper half-space by a homeomorphism that takes the ideal 2-simplices in Y to those in the Farey diagram of the modular group (Fig. 5).




Fig. 6

The complex K in this case is a tree. The vertices are the barycenters of the ideal simplices and there is an edge joining the barycenter of each 2-simplex to that of each of its faces. Thus K is the usual tree for $GL(2, \mathbb{Z})$. (See Bass and Serre [1].) The space X_2 contains, in addition, ideal 2-simplices determined by marked graphs which are described by the labelled graphs of the type shown in Fig. 6. These ideal simplices are disjoint. Each is missing its vertices and two of its open 1-faces, and is adjoined to Y along the third. There is exactly one such simplex attached to each ideal 1-simplex in Y . Thus X_2 is constructed from Y by attaching a “fin” along each ideal 1-simplex.

For $n \geq 3$, the space Y_n is not as well behaved as Y_2 . To begin with, it is not a manifold; a codimension 1 ideal simplex may be a face of either two or three maximal simplices. The complex K is also much more complicated. When $n = 3$, for example, the link of a level 1 vertex is a 2-torus with 12 disks attached, forming a complex which is homotopy equivalent to a wedge of eleven 2-spheres.

2. The star of a rose

We will now concentrate on the star in K of a single rose ρ . Recall that a vertex v of K is joined to ρ by an edge if and only if v and ρ are represented respectively by markings $g: R_0 \rightarrow G$ and $r: R_0 \rightarrow R$ for which there exist a sequence of edge collapses $G \rightarrow G_{l-1} \rightarrow \dots \rightarrow G_2 \rightarrow R$ making the following diagram commute up to homotopy.

$$\begin{array}{c}
 G \rightarrow G_{l-1} \rightarrow \dots \rightarrow G_2 \rightarrow R \\
 \swarrow \quad \searrow \\
 \quad \quad R_0
 \end{array}$$

Let $d: G \rightarrow R$ denote the composition of these edge collapses. Then d is a homotopy equivalence which collapses a subgraph T of G to the node of R . It follows that T is a tree. Since the number of edges of G that are not in T equals the number of edges of R , we see that T is a maximal tree in G .

The commutativity of the above diagram together with the fact that d is a homotopy equivalence means that the marking g is determined up to homotopy by d and r . Thus we temporarily ignore the markings and study the process of collapsing edges, and the inverse process of blowing up edges. It is helpful to introduce a combinatorial description of these processes.

2.1. Combinatorial graphs

A nice definition of a combinatorial graph is that coined by Gersten [4]: A graph is a set G together with an involution $x \rightarrow \bar{x}$ and a retraction t from G onto the fixed set $N(G)$ of the involution. The elements of $N(G)$ are the nodes of G and the elements of $E(G) = G - N(G)$ are the oriented edges. To realize G as a CW-complex, one takes $N(G)$ as the set of 0-cells and, for each orbit $\{e, \bar{e}\} \subset E(G)$ adds a 1-cell with boundary $\{t(e), t(\bar{e})\}$. Thus the involution interchanges the orientations of an edge, and t assigns to each oriented edge its terminal node. The combinatorial graphs form a category in which the morphisms are functions between the underlying sets which respect the retractions and involutions. These correspond to cellular maps which may collapse 1-cells. The combinatorial structure on a graph can also be specified, as by Bass and Serre [1], by giving the sets $N(G)$ and $E(G)$, the involution on $E(G)$ and the terminal vertex map $t: E(G) \rightarrow N(G)$.

In the sequel we assume that the reader can easily translate between the combinatorial and topological structures on graphs. We therefore may sometimes intermix the two points of view without explicitly stating which approach we are using.

2.2. Collapsing trees

Let (r, R) be a rose. By the discussion above, the vertices of the star of (r, R) in K correspond to surjective morphisms $d: G \rightarrow R$ for which the inverse of the node of R is a maximal tree in the graph G . We will show that such morphisms correspond to certain “complete” collections of subsets of $E(R)$. In fact, the star of (r, R) is isomorphic to a poset of complete collections, partially ordered by inclusion. A complete collection of subsets of $E(R)$ has a simple geometric description in terms of Venn diagrams, which we will use throughout the rest of the paper.

Definition. Two subsets A and B of $E(R)$ will be said to be *compatible* if one of the four sets $A \cap B$, $\bar{A} \cap B$, $A \cap \bar{B}$ and $\bar{A} \cap \bar{B}$ is empty, where \bar{X} denotes the complement of X in $E(R)$. If A and B are not compatible, we say A *crosses* B . A collection of subsets of $E(R)$ is *complete* if it consists of pairwise compatible subsets of $E(R)$ and is closed under set-theoretic complement.

Proposition 2.2.1. *There is a one-to-one correspondence between complete collections of subsets of $E(R)$ and morphisms $d: G \rightarrow R$ for which the inverse image of the node of R is a maximal tree in G .*

Proof. Let x denote the node of R . Let $d: G \rightarrow T$ be a morphism with $T = d^{-1}(x)$ a maximal tree in G . We will define, for each (directed) edge e of T , a subset e_* of $E(R)$ such that the collection $\{e_* : e \in E(T)\}$ is complete.

An edge e of T separates T into two components. Let T_e denote the subtree of $T - \{e\}$ which contains the terminal vertex $t(e)$ of e . We define the subset e_* of $E(R)$ to be the set of all edges of R of the form $d(f)$, where f is an edge in

$E(G) - E(T)$ and the terminal vertex of f is contained in T_e . Note that d restricts to a bijection between $E(G) - E(T)$ and $E(R)$.

Let e and f be two edges of T . Since $T - \{e, f\}$ has at most three components, it follows that one of the four sets

$$e_* \cap f_*, \bar{e}_* \cap f_*, e_* \cap \bar{f}_*, \bar{e}_* \cap \bar{f}_*$$

is empty. Thus e_* and f_* are compatible. Note also that \bar{e}_* is the set-theoretic complement of e_* . This shows that the set $\{e_* : e \in E(T)\}$ is complete.

Conversely, given a complete collection I of subsets of $E(R)$, we define a graph G and a morphism $d: G \rightarrow R$, which collapses a maximal tree T in G to the node of R , so that $I = \{e_* : e \in E(T)\}$. We first define the tree T by taking $E(T) = I$ with the complement operator as involution. Take $N(T)$ to be the set of equivalence classes in I under the equivalence relation generated by defining $i \sim \bar{j}$ whenever i is a maximal proper subset of j . Set $t(i)$ equal to the equivalence class of i . Thus oriented reduced edge-paths in T correspond to strictly ascending chains of sets in I , which implies that T is a tree. Next, for each oriented edge e of R add an oriented edge e' to G and define $t(e')$ to be equal to $t(i)$, where i is the maximal element of I not containing e . It is straightforward to check that t is well-defined (since the sets in I are compatible) and that $I = \{e_* : e \in E(T)\}$. \square

The process by which the graph G is constructed from I will be called *blowing up*, and the graph G will be denoted R^I . Each edge of R corresponds to a unique edge of R^I under this construction, and it will be convenient to abuse notation by referring to both by the same name.

The process of blowing up can be viewed geometrically, in terms of a Venn diagram for the sets of I . Regard the elements of $E(R)$ as points in the plane. For each pair $\{i, \bar{i}\} \subset I$ draw a simple closed curve which separates the elements of i from those of \bar{i} . The condition that the sets in I be compatible means that these curves can be taken to be disjoint. The tree T has a node for each connected component of the complement of this family of curves, and an edge joining two nodes if they lie in adjacent components. In addition, for each edge e of $E(R)$ there is an edge in $G = R^I$ joining the node in the component of e to that in the component of \bar{e} . We illustrate this by an example, shown in Fig. 7.




Fig. 7

In this example, we have

$$E(R) = \{x, \bar{x}, y, \bar{y}, z, \bar{z}, w, \bar{w}\}$$

$$I = \{\{x, y\}, \{\bar{x}, \bar{y}, z, \bar{z}, w, \bar{w}\}, \{\bar{x}, \bar{w}\},$$

$$\{x, y, \bar{y}, z, \bar{z}, w\}, \{\bar{y}, \bar{w}\}, \{x, \bar{x}, y, z, \bar{z}, w\},$$

$$\{\bar{y}, z, \bar{z}\}, \{x, \bar{x}, y, w, \bar{w}\}\}.$$

The graph G is shown on the right. The solid lines represent the maximal tree T , and the dotted lines are the edges in R^I which come from the edges of R .

Recall that we are considering only graphs G with no separating edges and no vertices of valence 2. Let I be a complete collection of subsets of $E(R)$. A subset i in the collection I corresponds to a separating edge of R^I if and only if i is invariant under the involution on $E(R)$. In terms of the Venn diagram, a non-separating edge corresponds to a simple closed curve in the diagram which encloses some element $e \in E(R)$ but excludes \bar{e} . The graph R^I will have no valence 2 vertices if and only if every subset in the collection I has more than one element.

These observations lead us to the following definition.

Definition. An *ideal edge* of a graph R with one node is a subset of $E(R)$ such that

- (1) the set i has at least two elements but less than $2n - 1$ elements;
- (2) there is some $e \in E(R)$ with $e \in i$ and $\bar{e} \in \bar{i}$.

Let $\rho = (r, R)$ be a rose. We have seen that each vertex of $st(\rho)$ can be represented uniquely by a marked graph (g_I, R^I) , where I is a complete collection of ideal edges of R , and g_I is the composition of r with a homotopy inverse for the collapsing map $R^I \rightarrow R$. Note that the collection I can be described independently of the choice of R as a set of conjugacy classes in F (cf. 1.3).

If i is in I , then the graph obtained from R^I by blowing down the edge corresponding to i is R^J , where $J = I - \{i, \bar{i}\}$. In fact, the vertices (g_I, R^I) and (g_J, R^J) of $st(\rho)$ are connected by an edge if and only if $I \subset J$ or $J \subset I$. The following proposition follows easily from these remarks.

Proposition 2.2.2. *The star of the rose (r, R) in K is isomorphic to the geometric realization of the poset of complete collections of ideal edges of R , ordered by inclusion. \square*


3. Reductivity of ideal edges

3.1. Whitehead moves

Let $\rho = (r, R)$ be a rose, and $i \subset E(R)$ an ideal edge of R . Define $D(i)$ to be the set of edges $e \in E(R)$ such that $e \in i$ and $\bar{e} \in \bar{i}$. Note that $D(i)$ is not empty. If $e \in D(i)$, then when e is regarded as an edge of the graph $R^{(i, \bar{i})}$, we have $t(e) \neq t(\bar{e})$; thus we can collapse the edge $\{e, \bar{e}\}$ of $R^{(i, \bar{i})}$ to obtain a new rose,

denoted ρ_e^i . The operation $\rho \rightarrow \rho_e^i$ will be called the *Whitehead move* (i, e) . Note that $(i, e) = (i, \bar{e}) = (\bar{i}, e) = (\bar{i}, \bar{e})$.

A Whitehead move determines an edge-path of length 2 in the complex K , as is shown by the diagram



In fact, we will show that any two roses in K can be joined by a path which is given by Whitehead moves. This is essentially a restatement of the fact that the Whitehead automorphisms generate $\text{Aut}(F_n)$ (cf. [12]).

We recall briefly the definition of a Whitehead automorphism of a free group F_n with basis $\{x_1, \dots, x_n\}$. Let A be a subset of $L = \{x_1^{\pm 1}, \dots, x_n^{\pm 1}\}$ for which there exists some letter $a \in L$ such that $a \in A$ and $a^{-1} \notin A$. Then the Whitehead automorphism (A, a) is the automorphism which fixes a and whose action on $L - \{a, a^{-1}\}$ is given by

$$\begin{aligned} x &\rightarrow axa^{-1} && \text{if } x \in A \text{ and } x^{-1} \in A \\ x &\rightarrow xa^{-1} && \text{if } x \in A \text{ and } x^{-1} \notin A \\ x &\rightarrow ax && \text{if } x \notin A \text{ and } x^{-1} \in A \\ x &\rightarrow x && \text{if } x \notin A \text{ and } x^{-1} \notin A. \end{aligned}$$

The relation between Whitehead moves and Whitehead automorphisms is described by the following lemma.

Lemma 3.1.1. *Let ρ_0 be the rose $(1, R_0)$ in K . Then a rose ρ is obtained from ρ_0 by a Whitehead move if and only if $\rho = \rho_0 \cdot \omega$ for some Whitehead automorphism ω given in terms of the usual geometric basis for $\pi_1(R_0)$. In the notation established above, we have $(\rho_0)_a^A = \rho_0 \cdot (A, a)$, where A is a subset of letters labelling the oriented edges of R_0 .*

Proof. This lemma is contained in the work of Hoare [5], whose “cut-and-paste” operation, when applied to the cointial graph associated to the standard basis, corresponds to a Whitehead move. The reader is referred to [5] for a detailed proof of the lemma. We include here an example which illustrates the idea of the proof. In this example, we apply a Whitehead move to ρ_0 in the case $n=5$ (Fig. 8). Here the geometric basis of $\pi_1(R_0)$ is $\{a, x, y, z, w\}$ and the set A is $\{a, x, \bar{y}, z, \bar{z}\}$. Note that if ρ_0 is represented by a labelled graph and $\alpha \in \text{Out}(F_n)$, then the rose $\rho_0 \cdot \alpha$ is described by the labelled graph whose labels are the images under α of those on ρ_0 . \square

Proposition 3.1.2. *K is path-connected.*

Proof. Let $\rho_0 = (1, R_0)$. As we noted in (1.2), $\text{Aut}(F_n)$ acts transitively on roses. Thus if ρ is any rose, we have $\rho = \rho_0 \cdot \varphi$ for some automorphism $\varphi \in \text{Aut}(F_n)$. Since the Whitehead automorphisms generate $\text{Aut}(F_n)$ we have




Fig. 8

$\rho = \rho_0 \cdot \omega_1 \cdot \dots \cdot \omega_k$, where each ω_i is a Whitehead automorphism. The proposition is proved by induction on k , the case $k=1$ being a consequence of Lemma 3.1.1. Thus we assume that there is a path in K joining ρ_0 to $\rho_0 \cdot \omega_1 \cdot \dots \cdot \omega_{k-1}$ and show that there is a path from ρ_0 to $\rho_0 \cdot \omega_1 \cdot \dots \cdot \omega_k$. By Lemma 3.1.1 there is a path joining ρ_0 to $\rho_0 \cdot \omega^{-1}$. Thus we obtain a path σ from $\rho_0 \cdot \omega^{-1}$ to $\rho_0 \cdot \omega_1 \cdot \dots \cdot \omega_{k-1}$. The automorphism ω_k acts on K by a homeomorphism which maps σ to a path joining ρ_0 to $\rho_0 \cdot \omega_1 \cdot \dots \cdot \omega_k$. This proves the proposition. \square

3.2. Whitehead moves and the norm

It is necessary for us to understand the behavior of the norm $\|\cdot\|_W$ (1.3.2) with respect to Whitehead moves. Our analysis is similar to that done by Higgins and Lyndon [6] and McCool [7] for Whitehead automorphisms. For the following definition, we assume that we have fixed a rose $\rho = (r, R)$, and a finite subset W of the set \mathcal{C} of conjugacy classes in F .

Definition. Let i be an ideal edge of R , and $e \in D(i)$. The *reductivity* of the Whitehead move (i, e) is $\text{red}(i, e) = \|\rho\|_W - \|\rho_e^i\|_W$. The *reductivity* of the ideal edge i is

$$\text{red}(i) = \max_{e \in D(i)} \text{red}(i, e).$$

The ideal edge i is *reductive* if $\text{red}(i) \geq 0$, *zero-reductive* if $\text{red}(i) = 0$ and *strictly reductive* if $\text{red}(i) > 0$. If $\text{red}(i) \geq \text{red}(j)$ for all ideal edges j of R , then i is *maximally reductive*.

It turns out that the reductivity is an intrinsic invariant of an ideal edge. This is made clear by the following interpretation. For any marking $g: R_0 \rightarrow G$ and any oriented edge $e \in E(G)$, let $|e|_g$ denote the number of occurrences of the oriented edge e or of \bar{e} in the reduced edge-cycles that represent the conjugacy classes $g_*(w)$ for $w \in W$. An edge collapse which is a homotopy equivalence clearly maps reduced edge-cycles to reduced edge-cycles. Thus if $g': R_0 \rightarrow G'$ is obtained from $g: R_0 \rightarrow G$ by collapsing an edge $\{f, \bar{f}\}$ and if e is an oriented edge in $E(G) - \{f, \bar{f}\}$ then

$$|e|_g = |e'|_{g'}$$

where e' is the image of e under the edge collapse $G \rightarrow G'$. If (g, G) is the marked graph obtained by blowing up an ideal edge i of R , then $\text{red}(i, e) = |i|_g - |e|_g$. Moreover, $|i|_g$ can be expressed entirely in terms of the oriented edges of R . We consider the reduced edge-cycles in R which represent the elements of W . Then $|i|_g$ equals the number of pairs of consecutive oriented edges e, f in these edge-cycles with $e \in i$ and $\bar{f} \in \bar{i}$ or with $e \in \bar{i}$ and $\bar{f} \in i$.

Since these quantities are intrinsic, we drop the subscript g . Thus given a rose ρ , a representation (r, R) of ρ , and a finite set $W \subset \mathcal{C}$, we establish the following suggestive notation (which is borrowed from Higgins and Lyndon).

Notation. For two ideal edges i and j of R , $i \cdot j$ is the number of occurrences in the reduced cycles representing W of e followed by f , where either $e \in i$ and $\bar{f} \in j$ or $e \in j$ and $\bar{f} \in i$. For an ideal edge i , $|i| \equiv i \cdot \bar{i}$. For $e \in E(R)$, $|e| \equiv \{e\} \cdot \overline{\{e\}}$ is the number of occurrences of the oriented edge e or of \bar{e} in these reduced cycles.

Definition. Let $W \subset \mathcal{C}$ be a finite set of cyclic words in F , and $\rho = (r, R)$ a rose. The *star graph* $S_W(\rho)$ is the graph whose set of nodes is $E(R)$ and which has an (unoriented) edge joining e to f for each occurrence of e and f as consecutive oriented edges in the reduced cycles representing $r_*(w)$ for $w \in W$.

Note that $|e|$ can be interpreted as the valence of the node of the star graph representing e . We can also interpret $|i|$ for an ideal edge i geometrically using the star graph and the Venn diagram representing i . To do this we map the star graph into the plane so that

- (1) the nodes go to the points in our Venn diagram;
- (2) each edge of the star graph intersects the simple closed curve corresponding to $\{i, \bar{i}\}$ at most once; and
- (3) no double point of the map lies on the simple closed curve. It is clear that such a map exists, although it will not usually be an imbedding. Then $|i|$ is the number of intersections of the star graph with the curve.

3.3. The factorization Lemma

If ρ and ρ' are two roses whose stars in K have a non-empty intersection, then ρ' can be obtained from ρ by blowing up some ideal edges and then collapsing an equal number of non-ideal edges. The next lemma estimates $\|\rho\|_W - \|\rho'\|_W$.

Factorization Lemma 3.3.1. *Let ρ and ρ' be roses represented by (r, R) and (r', R') respectively. Suppose that ρ' is constructed from ρ by first blowing up the set $I = \{i_1, \dots, i_k\}$ of pairwise compatible ideal edges and then blowing down the set of ordinary edges $O = \{e_1, \dots, e_k\}$. Then $\|\rho\|_W - \|\rho'\|_W \leq \sum_{\alpha=1}^k \text{red}(i_\alpha)$.*

Proof. We will construct a permutation $\sigma \in \Sigma_k$ such that for each $\alpha = 1, \dots, k$, $e_\alpha \in D(i_{\sigma(\alpha)})$ or $\bar{e}_\alpha \in D(i_{\sigma(\alpha)})$. After renaming the e_α , we may assume $e_\alpha \in D(i_{\sigma(\alpha)})$. Then

$$\begin{aligned} \|\rho\|_W - \|\rho'\|_W &= \sum_{\alpha=1}^k |i_\alpha| - \sum_{\alpha=1}^k |e_\alpha| \\ &= \sum_{\alpha=1}^k (|i_{\sigma(\alpha)}| - |e_\alpha|) \\ &= \sum_{\alpha=1}^k \text{red}(i_{\sigma(\alpha)}, e_\alpha) \leq \sum_{\alpha=1}^k \text{red}(i_{\sigma(\alpha)}) = \sum_{\alpha=1}^k \text{red}(i_\alpha). \end{aligned}$$

Consider the graph $G = R^I$ constructed by blowing up the ideal edges in I . Considering edges of R as edges of G , I and O respectively comprise the sets of edges of maximal trees T_I and T_O in G . As such, the collections of cosets $\{[i_1], \dots, [i_k]\}$ and $\{[e_1], \dots, [e_k]\}$ form bases for the rational vector space C_1/Z_1 of 1-chains modulo the 1-cycles in the chain complex of G . Therefore we may write

$$[e_\alpha] = \sum_{\beta=1}^k \varepsilon_{\alpha\beta} [i_\beta]$$

in C_1/Z_1 , where the matrix $[\varepsilon_{\alpha\beta}]$ is non-singular. In fact the coefficient $\varepsilon_{\alpha\beta}$ is 0, 1 or -1 and is non-zero if and only if $\{i_\beta, \bar{i}_\beta\}$ is an edge in the unique reduced path in T_I that joins the endpoints of e_α . Thus $\varepsilon_{\alpha\beta} \neq 0$ if and only if either $e_\alpha \in D(i_\beta)$ or $\bar{e}_\alpha \in D(i_\beta)$. Since $\text{Det}([\varepsilon_{\alpha\beta}]) \neq 0$ there must exist some $\sigma \in \Sigma_k$ with $\prod \varepsilon_{\alpha\sigma(\alpha)} \neq 0$. Hence $\varepsilon_{\alpha\sigma(\alpha)} \neq 0$ for all α , so $e_\alpha \in D(i_{\sigma(\alpha)})$. \square

3.4. Diagrams

The lemmas in the next two sections are concerned with the reductivities of two ideal edges i and j which cross, and of the four associated ideal edges $i \cap j$, $i \cap \bar{j}$, $\bar{i} \cap j$ and $\bar{i} \cap \bar{j}$. We will give the arguments in terms of geometric properties of the Venn diagrams for i and j .

Consider a rose ρ represented by a marked graph (r, R) and two ideal edges i and j which cross. Choose distinct points in the plane corresponding to the oriented edges of R . Represent i and j by Jordan curves as in the earlier discussion of Venn diagrams. These curves can be taken to meet transversely in two points. Realize each edge of the star graph S_W of our finite set $W \subset \mathcal{C}$ as an arc joining the points corresponding to the appropriate ideal edges of R . Each of these arcs can be taken to intersect the Jordan curves transversely in at most one point and not to pass through the points of intersection of the two Jordan curves. There are four arcs of the Jordan curves that join the two intersection points. These are labelled α, β, γ and δ in the diagram in Fig. 9.




Fig. 9

Let a, b, c and d be the respective numbers of geometric intersections of edges of the star graph with α, β, γ and δ . In the arguments that follow we will refer to the diagram above and assume that $\alpha, \beta, \gamma, \delta, a, b, c$ and d are as defined here.

Observe that $|i| = a + c$ and $|j| = b + d$. Also observe that any edge of S_W joining points in $i \cap j$ with points in $(\overline{i \cap j})$ will meet β or γ . However, edges of S_W that join points in $\overline{i \cap j}$ to points in $i \cap j$ may also meet β and γ . Thus

$$|i \cap j| \leq b + c.$$

Similarly,

$$|\overline{i \cap j}| \leq c + d$$

and

$$|i \cap \overline{j}| \leq a + b$$

$$|\overline{i \cap \overline{j}}| \leq a + d.$$

4. Existence of a reductive edge

4.1. The Higgins-Lyndon Lemma

The object of the lemmas in this section is to show that if ρ is a rose which does not have minimal norm then $st(\rho)$ has a non-empty intersection with $st(\rho')$ for some rose ρ' with $\|\rho'\|_W < \|\rho\|_W$. This amounts to showing that if $\rho = (r, R)$, then R has a strictly reductive ideal edge. Actually we prove a somewhat stronger statement, which is needed for the inductive argument in Sect. 6. We show that, if I is a complete set of zero-reductive edges, then there exists a strictly reductive ideal edge which is compatible with each ideal edge in I .

We begin with a lemma which was proved in a slightly different form by Higgins and Lyndon [6]. Our proof is given in terms of the Venn diagrams described in (3.4).

Higgins-Lyndon Lemma 4.1.1. *Let ρ be a rose represented by (r, R) and let i and j be ideal edges of R which cross. Let e and f be oriented edges of R with $e \in D(i)$ and $f \in D(j)$. Then, for some choice of g equal to either e or f and k equal to one of $i \cap j, i \cap \overline{j}, \overline{i \cap j}, \overline{i \cap \overline{j}}$, we have*

$$\text{red}(k, g) \geq \min(\text{red}(i, e), \text{red}(j, f))$$

with strict inequality unless $\text{red}(i, e) = \text{red}(j, f)$.




Fig. 10

Proof. We may assume that $\text{red}(i, e) \geq \text{red}(j, f)$. We will then assume that for each of the choices for k and g , $\text{red}(k, g) \leq \text{red}(j, f)$ and $\text{red}(k, g) < \text{red}(i, e)$. These assumptions will lead to a contradiction, so we conclude that either $\text{red}(k, g) > \text{red}(j, f)$ or $\text{red}(k, g) \geq \text{red}(i, e) \geq \text{red}(j, f)$. The latter inequality will be strict unless $\text{red}(i, e) = \text{red}(j, f)$.

By replacing j by \bar{j} if necessary we may assume $e \in j$. We then have eight cases to consider, since either $\bar{e} \in i \cap j$ or $\bar{e} \in \bar{i} \cap \bar{j}$, either $f \in i \cap j$ or $f \in \bar{i} \cap \bar{j}$ and either $\bar{f} \in i \cap \bar{j}$ or $\bar{f} \in \bar{i} \cap \bar{j}$. In each of these cases, by our assumptions, if $e \in k$ and $\bar{e} \in \bar{k}$ then $\text{red}(k, e) < \text{red}(i, e)$, and if $f \in k$ and $\bar{f} \in \bar{k}$ then $\text{red}(k, f) < \text{red}(j, f)$. These inequalities imply certain inequalities involving a, b, c and d for each of the four choices for k (cf. 3.4 for the definitions of a, b, c and d). We show that the inequalities obtained this way are inconsistent, giving us our contradiction.

To illustrate, suppose $\bar{e} \in i \cap j, f \in i \cap j$ and $\bar{f} \in i \cap \bar{j}$. The diagram for i and j is then as shown in Fig. 10.

The inequalities we obtain are:

$$\begin{aligned} |e| - (b + c) &\leq |e| - |i \cap j| = \text{red}(i \cap j, e) < \text{red}(i, e) \\ &= |e| - (a + c) \Rightarrow a + c < b + c \Rightarrow a < b. \end{aligned}$$

$$\begin{aligned} |f| - (b + c) &\leq |f| - |i \cap j| = \text{red}(i \cap j, f) \leq \text{red}(j, f) \\ &= |f| - (b + d) \Rightarrow b + d \leq b + c \Rightarrow d \leq b. \end{aligned}$$

$$\begin{aligned} |f| - (a + b) &\leq |f| - |i \cap \bar{j}| = \text{red}(i \cap \bar{j}, f) \leq \text{red}(j, f) \\ &= |f| - (b + d) \Rightarrow b + d \leq a + b \Rightarrow d \leq a. \end{aligned}$$

$$\begin{aligned} |e| - (c + d) &\leq |e| - |\bar{i} \cap j| = \text{red}(\bar{i} \cap j, e) < \text{red}(i, e) \\ &= |e| - (a + c) \Rightarrow a + c < c + d \Rightarrow a < d. \end{aligned}$$

Obviously the last two inequalities are contradictory.

We omit the complete derivations in the other seven cases. Instead we supply a table in Fig. 11, in which we give the Venn diagram and the inconsistent inequalities obtained in each case. In each diagram the ideal edge i is represented by the circle on the left, and j by the circle on the right. This completes the proof of the Higgins-Lyndon Lemma. \square

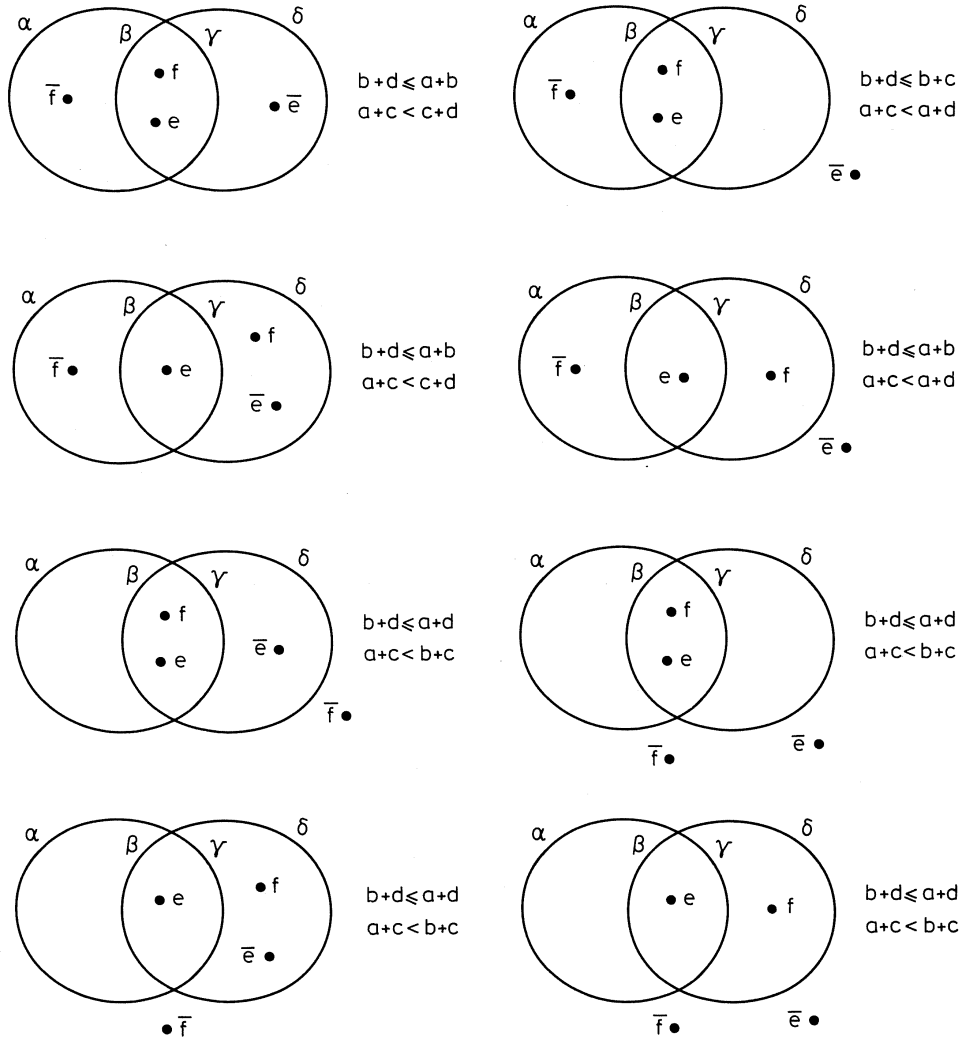


Fig. 11

4.2. Existence Theorems

In [6] the lemma above is used to prove the following.

Proposition 4.2.1. (Higgins-Lyndon). *Let ρ be a rose, and ρ' and ρ'' be roses obtained from ρ by Whitehead moves. Assume $\|\rho'\|_W < \|\rho\|_W \cong \|\rho''\|_W$. Then there exists a sequence of Whitehead moves leading from ρ' to ρ'' such that for each intermediate rose σ we have $\|\sigma\|_W < \|\rho\|_W$.*

Proof. Suppose $\rho' = \rho_e^i$ and $\rho'' = \rho_f^j$. If i and j are compatible, then $\|\rho_{e,f}^{i,j}\|_W = \|\rho_f^j\|_W - \text{red}(i, e) < \|\rho\|_W$, and $\rho_{e,f}^{i,j}$ is connected to each of ρ' and ρ'' by Whitehead moves.

If i crosses j , then let k and g be the edges given by applying the Higgins-Lyndon Lemma to the Whitehead moves (i, e) and (j, f) . Note that $\|\rho_e^k\|_W < \|\rho\|_W$. If $g = e$ then ρ_e^k is connected to ρ_f^j by a Whitehead move and,




Fig. 12

since k is compatible with j , we may apply the preceding argument with ρ' replaced by ρ_e^k . If $g = f$ then ρ_e^k is connected to ρ_f^j by a Whitehead move, and we may apply the argument above with ρ'' replaced by ρ_e^k since i and k are compatible. \square

We remark that the paths from ρ' to ρ'' corresponding to these sequences of Whitehead moves are homotopic in K . In fact the homotopy can be realized by a sequence of moves in which the path is pushed across certain 2-cells in K . Shown in Fig. 12 are three 2-cells embedded in K which correspond to the three cases of the proof. The union of the 2-cells of these types is a 2-dimensional subcomplex of K which is the universal cover of the complex constructed by McCool in [7].

By repeated applications of the preceding proposition to an arbitrary path of Whitehead moves, Higgins and Lyndon also show the following

Proposition 4.2.2. *Let ρ and ρ' be roses with $\|\rho\|_W$ minimal. Then there is a sequence $\rho = \rho_0, \rho_1, \dots, \rho_m = \rho'$ of roses such that ρ_i is obtained from ρ_{i-1} by a Whitehead move for $i = 1, \dots, m$ and such that*

$$\|\rho_0\|_W = \|\rho_1\|_W = \dots = \|\rho_k\|_W < \|\rho_{k+1}\|_W < \dots < \|\rho_m\|_W$$

for some $k \in \{0, 1, \dots, m\}$. \square

Since each application of Proposition 4.2.1 produces a homotopic path, this implies that $\pi_1(K) \cong \pi_1(K_{\min})$.

We will make use of the following corollary of Proposition 4.2.2.

Corollary 4.2.3. *Let $\rho = (r, R)$ be a rose. If $\|\rho\|_W$ is not minimal, then R has a strictly reductive ideal edge. \square*

In fact, using the Higgins-Lyndon Lemma we can strengthen this corollary to provide the following fundamental result:

Existence Theorem 4.2.4. *Let R be a rose represented by (r, R) . Let I be a (possibly empty) complete set of zero-reductive ideal edges. Then there exists a strictly reductive ideal edge which is compatible with each ideal edge in I .*

Proof. Let j be a strictly reductive ideal edge of R which is compatible with the maximum number of ideal edges in I . If j crosses i for some $i \in I$ then by the

Higgins-Lyndon Lemma one of $j \cap i$, $j \cap \bar{i}$, $\bar{j} \cap i$ and $\bar{j} \cap \bar{i}$ will be a strictly reductive ideal edge. Each of these is compatible with any ideal edge that is compatible with both i and j , and hence with more of the ideal edges in I than j . This contradicts the maximality assumption, so j must be compatible with every ideal edge in I . \square

for We remark that in the case where W is the set of conjugacy classes of primitive elements in F then the existence result 4.2.3 was proved by Whitehead in [12] as a consequence of his result that the star graph for such a set W has a cut vertex. *ab*

5. Outermost intersections

The main result of this section is the Pushing Lemma, which is a refinement of the Higgins-Lyndon Lemma. It is used in the proof of our theorem to make possible what is more or less an outermost curve argument applied to Venn diagrams.

5.1. The minimax property

We continue to work with a rose ρ that is represented by a marking $r: R_0 \rightarrow R$ and with a fixed finite subset W of \mathcal{C} . Let I be a (possibly empty) complete set of zero-reductive ideal edges. An ideal edge of R will be said to have the *minimax property* with respect to I if it is compatible with each ideal edge in I , has maximal reductivity among all ideal edges of R that are compatible with I , and has no proper subsets with maximal reductivity. Let i have the minimax property and choose an oriented edge $e \in i$ such that $\text{red}(i, e) = \text{red}(i)$. Let j be an ideal edge of R which crosses i . Replacing j by \bar{j} if necessary, we may assume that $e \in j$. We define the *outer slices* of j along i to be the two ideal edges $i \cap \bar{j}$ and $\bar{i} \cap \bar{j}$. These edges are illustrated in the Venn diagram of Fig. 13.

The Pushing Lemma states that one of the outer slices has reductivity at least as large as that of j .

Pushing Lemma 5.1.1. *Let I be a complete set of zero-reductive ideal edges. Let i be an ideal edge of R with the minimax property relative to I , and let e be an oriented edge of R such that $e \in i$, $\bar{e} \in \bar{i}$ and $\text{red}(i, e) = \text{red}(i)$. Let j be an ideal edge*




Fig. 13




Fig. 14

compatible with I that crosses i . Then for one of the outer slices j' of j along i we have

$$\text{red}(j') \geq \text{red}(j).$$

Proof. We assume that $e \in j$ (Fig. 14 shows the standard Venn diagram) and that f is an oriented edge with $f \in j, f \in \bar{j}$ and $\text{red}(j, f) = \text{red}(j)$.

Observe that since $i \cap j$ is an ideal edge with $i \cap j$ properly contained in i , $e \in i \cap j$ and $\bar{e} \in (\overline{i \cap j})$, it follows from the minimax property of i that $\text{red}(i \cap j, e) < \text{red}(i, e)$. This implies that

$$|e| - (b + c) \leq \text{red}(i \cap j, e) < \text{red}(i) = |e| - (a + c).$$

Hence $a < b$, so if $\bar{f} \in \bar{i} \cap \bar{j}$ then

$$\text{red}(\bar{i} \cap \bar{j}) \geq \text{red}(\bar{i} \cap \bar{j}, \bar{f}) \geq |f| - (a + d) > |f| - (b + d) = \text{red}(j),$$

so the outer slice $\bar{i} \cap \bar{j}$ has the desired property.

Assume now that $\bar{f} \in i \cap \bar{j}$.

If $\bar{e} \in \bar{i} \cap j$ then, by the maximality of $\text{red}(i)$, we have

$$|e| - (c + d) \leq \text{red}(\bar{i} \cap j, \bar{e}) \leq \text{red}(i) = \text{red}(i, e) = |e| - (a + c).$$

Hence $a \leq d$ and, using the assumption $\bar{f} \in i \cap \bar{j}$ we have

$$\text{red}(i \cap \bar{j}) \geq \text{red}(i \cap \bar{j}, \bar{f}) \geq |f| - (a + b) \geq |f| - (b + d) = \text{red}(j),$$

i.e. the outer slice $i \cap \bar{j}$ has the desired property.

Finally, assume that $\bar{e} \in \bar{i} \cap \bar{j}$. This leaves two cases to consider: Either $f \in i \cap j$ or $f \in \bar{i} \cap j$.

Case I. $f \in i \cap j$ (Fig. 15)




Fig. 15

The minimax condition implies that $\text{red}(i \cap j, f) < \text{red}(i, e)$. Hence

$$|f| - (b + c) \leq \text{red}(i \cap j, f) < \text{red}(i, e) = |e| - (a + c).$$

Thus $|f| - b < |e| - a$, so

$$\text{red}(\bar{i} \cap \bar{j}, \bar{e}) \geq |e| - (a + d) > |f| - (b + d) = \text{red}(j).$$

Case II. $f \in \bar{i} \cap j$ (Fig. 16)




Fig. 16

By the minimax property we have

$$|\bar{f}| - (a + c) = \text{red}(i, \bar{f}) \leq \text{red}(i) = |e| - (a + c),$$

which implies $|f| < |e|$. As observed at the beginning of the argument, we have $a < b$. Thus

$$\text{red}(\bar{i} \cap \bar{j}, e) \geq |e| - (a + d) > |f| - (a + d) > |f| - (b + d) = \text{red}(j).$$

This completes the proof of the Pushing Lemma. \square

6. The contractibility of K

In this section we prove that the complex K (and hence the space X) is contractible. Recall that K can be described as the geometric realization of the category of equivalence classes of marked graphs (g, G) , with $\pi_1(G) \cong F_n$, in which the arrows are sequences of collapses (cf. 1.1).

6.1. Statement of the theorem

By the Gertrude Stein Lemma (1.2.1), the complex K is the union of the stars of its roses:

$$K = \bigcup_{\rho \in K} \text{st}(\rho).$$

Let W be a nonempty finite set of conjugacy classes in F_n , and ' $<$ ' a total ordering of the roses in K subordinate to the norm $\|\cdot\| = \|\cdot\|_W$ (cf. 1.2). Define

$$K_{\min} = \bigcup_{\substack{\|\rho\| \\ \text{minimal}}} \text{st}(\rho).$$

We prove the theorem in the following form:

Theorem 6.1.1. *The complex K_{\min} is a deformation retract of K .*

We then show that for an appropriate choice of W , K_{\min} consists of the star of the single rose ρ_0 . Thus we obtain

Corollary 6.1.2. *The complex K is contractible.*

As mentioned in the introduction this implies

Corollary 6.1.3. *The group $\text{Out}(F_n)$ is of type VFL and has virtual cohomological dimension $2n - 3$.*

Observe that 6.1.1 and 6.1.2 together imply that K_{\min} is contractible for any choice of a finite subset W of \mathcal{C} . Moreover, K_{\min} is invariant under the action of the subgroup $\text{Out}_W(F)$ which consists of all outer automorphisms that permute the conjugacy classes in the set W . By showing that the quotient of K_{\min} under this action is finite we can conclude that $\text{Out}_W(F)$ is of type VFL. This result was also obtained by Gersten, and is an extension of McCool's theorem [8] that these groups are finitely presented. We remark that the mapping class groups of bounded surfaces are of this form.

Corollary 6.1.4. *For any finite set W of conjugacy classes in F_n the group $\text{Out}_W(F_n)$ is of type VFL.*

Proof. $\text{Out}_W(F_n)$ acts on the contractible complex K_{\min} with finite stabilizers, so it suffices to show $K_{\min}/\text{Out}_W(F_n)$ is finite.

For any rose ρ , there are only a finite number of subsets $W' \subset \mathcal{C}$ with the property that $\|\rho\|_W = \|\rho\|_{W'}$, since there are only a finite number of graphs on $2n$ vertices with the sum of the valences of the nodes equal to twice $\|\rho\|_W$. Fix a rose ρ_0 in K_{\min} . Let $W = W_0, W_1, \dots, W_s$ be all subsets W' of \mathcal{C} with $\|\rho_0\|_W = \|\rho_0\|_{W'}$ which are images of W under an outer automorphism of F . Choose automorphisms $\alpha_0 = \text{id}, \alpha_1, \dots, \alpha_s$ such that $\alpha_i(W) = W_i$. Note that for any $\alpha \in \text{Out}(F_n)$ and any rose ρ , we have $\|\rho \cdot \alpha\|_W = \|\rho\|_{\alpha(W)}$. Thus

$$\|\rho_0 \cdot \alpha_i\|_W = \|\rho_0\|_{W_i} = \|\rho_0\|_W,$$

i.e. $\rho_0 \cdot \alpha_i$ is rose in K_{\min} for each i .

If ρ is any rose in K_{\min} , choose an automorphism α with $\rho_0 \cdot \alpha = \rho$. Then $\|\rho_0\|_W = \|\rho\|_W = \|\rho_0\|_{\alpha(W)}$, so $\alpha(W) = W_i$ for some i . Thus $W = \alpha^{-1} \circ (\alpha_i(W))$, i.e. $\alpha^{-1} \circ \alpha_i$ stabilizes W ; since $\rho \cdot (\alpha^{-1} \circ \alpha_i) = \rho_0 \cdot \alpha_i$, we have shown that ρ is equivalent modulo $\text{Out}_W(F_n)$ to one of the finite number of roses $\{\rho_0 \cdot \alpha_i\}$. Therefore $K_{\min}/\text{Out}_W(F_n)$ is finite. \square

6.2. Proof of the theorem

We prove that K deformation retracts to K_{\min} by showing that for each rose ρ , the subcomplex

$$K_{\leq \rho} = \bigcup_{\rho' \leq \rho} st(\rho')$$

deformation retracts to K_{\min} .

Let ρ be a rose whose norm is not minimal. Then $K_{\leq \rho} = st(\rho) \cup K_{< \rho}$. If ρ is the first rose in the ordering which is not minimal, then $K_{< \rho} = K_{\min}$. Therefore

we may assume by induction that $K_{<\rho}$ retracts onto K_{\min} . Since $st(\rho)$ is contractible, in order to show that $K_{\leq\rho}$ retracts to K_{\min} it suffices to show that $st(\rho) \cap K_{<\rho}$ is contractible.

Note that the Existence Theorem (4.2.4) shows that $st(\rho) \cap K_{<\rho}$ is non-empty.

We have reduced our problem to working inside the star of a single rose $\rho = (r, R)$. By Proposition 2.2.2, we may identify $st(\rho)$ with the geometric realization of the poset of complete collections of ideal vertices, where the partial ordering is given by inclusion. We will make this identification throughout this section.

The following lemma, which is standard (see [10]) is useful in dealing with complexes of this type.

Poset Lemma 6.2.1. *Let $f: P \rightarrow P$ be a poset map from the poset P to itself, such that $p \leq f(p)$ for all $p \in P$. Then f induces a deformation retraction from the geometric realization of P to the geometric realization of the image $f(P)$.*

We begin by defining certain subcomplexes of $st(\rho)$:

Definition. The *upper star* $st_+(v)$ of a vertex v of K is the subcomplex of the star of v spanned by vertices of level greater than or equal to the level of v . (Recall that the *level* of a graph is the number of nodes). The *upper link* $lk_+(v)$ is the subcomplex of the link of v spanned by vertices of level greater than the level of v .

The *upper star* $st_+(L)$ of a subcomplex L of K is the union (not the intersection) of the upper stars of all vertices of L .

Let I be a complete set of ideal edges of ρ , and $v = \rho^I = (g_I, R^I)$ the associated vertex of K . Then $st_+(v)$ is the realization of the poset of all complete sets J of ideal edges which contain I . The following subcomplexes of $st_+(v)$ will be important:

Define

$$R(I) = \{J \supset I : J \text{ is complete and contains a reductive ideal edge}\}$$

$$S(I) = \{J \supset I : J \text{ is complete and contains a strictly reductive ideal edge}\}$$

$$TR(I) = \{J \supset I : J \text{ is complete and consists entirely of reductive ideal edges}\}$$

$$TS(I) = \{J \supset I : J \text{ is complete, consists entirely of reductive ideal edges, and contains a strictly reductive ideal edge}\}$$

Define $R(v)$, $S(v)$, $TR(v)$, and $TS(v)$ to be the geometric realizations of $R(I)$, $S(I)$, $TR(I)$, and $TS(I)$ respectively.

Lemma 6.2.2. *Let I be a complete set of zero-reductive ideal edges of R , and $v = \rho^I$. Then $TS(v)$ is contractible.*

Proof. Since every ideal edge in I is zero-reductive, the Existence Theorem says that there exists an ideal edge m which is compatible with each ideal edge in I

and has the minimax property. We will retract $TS(v)$ to the vertex of $TS(v)$ represented by $I \cup \{m, \bar{m}\}$ by using the Pushing Lemma (Sect. 5).

For any subcomplex L of $st(v)$, let $C(L, m)$ denote the set of all ideal edges j such that j crosses m and $j \in J$ for some vertex ρ^J of L . Define the *complexity* $c(L)$ to be the number of elements of $C(L, m)$.

Claim. Let L' be a subcomplex of $TS(v)$ such that

- (i) $c(L') = 0$
 - (ii) $st_+(L') \cap TS(v) \subset L'$
- and (iii) $I \cup \{m, \bar{m}\} \in L'$.

Then L' is contractible.

Proof. By condition (i), m is compatible with each ideal edge in each vertex J of L' . This together with the fact that L' is in the star of I implies that the map sending J to $J \cup I \cup \{m, \bar{m}\}$ is a well-defined poset map on L' . Condition (ii) guarantees that the image of this map is contained in L' , and thus by the Poset Lemma gives a deformation retraction from L' onto its image. Since every vertex in the image contains $I \cup \{m, \bar{m}\}$, which is also in the image by condition (iii), the map sending J to $I \cup \{m, \bar{m}\}$ now gives a deformation retraction of the image to the vertex $I \cup \{m, \bar{m}\}$. \square

If $c(TS(v)) = 0$, we are done by the claim.

If $c(TS(v)) > 0$, we will retract $TS(v)$ to a subcomplex L' satisfying hypotheses (ii) and (iii) of the claim which has strictly lower complexity. We can repeat this operation, reducing the complexity at each stage. We end up with a subcomplex of complexity zero, which can be retracted to $I \cup \{m, \bar{m}\}$ by the claim. This will prove the lemma.

Notation. Choose $e \in m$ such that $\text{red}(m, e) = \text{red}(m)$. Orient each ideal edge i so that $e \in i$.

Let $C' \subset C(TS(v), m)$ be the set of ideal edges $i \in C(TS(v), m)$ with $\#(m \cap \bar{i})$ minimal. Choose $a \in C'$ such that a is *outermost*, i.e. $\#a \geq \#c'$ for all $c' \in C'$. By the Pushing Lemma, one of the sets $a \cup \bar{m}$ or $m \cup a$ is an ideal edge with reductivity $\geq \text{red}(a)$; denote this ideal edge by a_0 . Define a map from $TS(v)$ to itself by sending J to $J \cup \{\bar{a}_0, a_0\}$ if $a \in J$, and J to itself if $a \notin J$.

Claim. This map is a well-defined poset map.

Proof. We must show that a_0 is compatible with J for every J which contains a . Fix such a J , and let $b \in J$. Since b is compatible with a but not with a_0 , we have either $b \subset a$ or $a \supset b$. If $b \subset a$, it does not cross a_0 . If $b \supset a$, and b crosses a_0 , then b must cross m and $\#(m \cap \bar{b}) \leq \#(m \cap \bar{a})$. But $a \in C'$, so $\#(m \cap \bar{b}) \geq \#(m \cap \bar{a})$. Thus $\#(m \cap \bar{b}) = \#(m \cap \bar{a})$, and $b \in C'$. But this contradicts the assumption that a is outermost, so b must be compatible with a_0 . \square

This poset map retracts $TS(v)$ onto a subcomplex. We now retract this subcomplex by the poset map which sends J to itself if a is not in J , and J to $J - \{a, \bar{a}\}$ if $a \in J$. The image of this map is a subcomplex L' of $TS(v)$ with $C(L', m) = C(TS(v), m) - \{a, \bar{a}\}$, i.e. with strictly smaller complexity.

One checks easily that this procedure can be repeated for the subcomplex L' , and that after $c(TS(v))$ such reductions $TS(v)$ has been retracted to a

subcomplex satisfying the hypotheses of the claim. This completes the proof that $TS(v)$ is contractible. \square

Lemma 6.2.3. *Let I be a complete set of zero-reductive ideal edges of ρ , and $v = \rho^I$. Then $S(v)$ is contractible.*

Proof. The map $S(I) \rightarrow TS(I)$ which sends any element J of $S(I)$ to the subset of J consisting of reductive ideal edges is a poset map, and hence induces a deformation retraction $S(v) \rightarrow TS(v)$ by the Poset Lemma. \square

We are now set up to prove the main lemma of the theorem.

Main Lemma 6.2.4. *Let I be a complete set of zero-reductive ideal edges of ρ , corresponding to the vertex v of $st(\rho)$. Let L be a subcomplex of $R(v)$ such that*

$$\text{(i) } L \text{ contains } S(v),$$

$$\text{and (ii) } L = \bigcup_{w \in L \cap TR(v)} st_+(w).$$

Then L is contractible.

Proof. Define

$$d(L) = \min_{w \in (L - S(v)) \cap TR(v)} \{\dim(st_+(w))\}.$$

The proof proceeds by induction on $d(L)$.

To begin the induction, we will show that if $d(L) = 0$, then $L = S(v)$, and hence L is contractible by the preceding lemma. By hypothesis (i) we have $S(v) \subset L$. To show the opposite inclusion, let w be any vertex of L . By (ii), $w \in st_+(w')$ for some totally reductive vertex $w' = \rho^J$ of L . If each ideal edge of J is zero-reductive, then the Existence Theorem says that we can find an ideal edge k which is compatible with each ideal edge in J and is strictly reductive. But then $J \cup \{k, \bar{k}\}$ is a complete set of ideal edges representing a vertex of $lk_+(w')$; thus $\dim st_+(w') \geq 1$, contradicting the assumption that $d(L) = 0$. Therefore w' must be strictly reductive, which implies that $w \in S(v)$.

Assume now that $d(L) = k > 0$. Let L' be the union of $S(v)$ and all of the upper stars of vertices $w \in TR(v)$ such that $\dim st_+(w) < d(L)$. Then L' clearly satisfies the hypotheses of the Main Lemma, and $d(L') < d(L)$. Thus L' is contractible by induction.

If $w \in L - L'$, then $\dim st_+(w) = d(L)$ and the star of w in L is equal to $st_+(w)$. Note that if w and w' are any two vertices of $L - L'$, then $st_+(w) \cap st_+(w') \subset L'$. Thus to see that L is contractible, we need only show that $st_+(w) \cap L'$ is contractible for each $w \in L - L'$.

We now observe that if $w \in L - L'$ is reductive, then w must be zero-reductive since $S(v) \subset L'$. In addition, $S(w) \subset st_+(w) \cap L' \subset R(w)$, and $st_+(w) \cap L'$ equals the union of the upper stars of all vertices u of $st_+(w) \cap L'$ which are contained in $TR(v)$ and hence in $TR(w)$. Since $d(st_+(w) \cap L') < d(L)$, $st_+(w) \cap L'$ is contractible by induction. This completes the proof of the Main Lemma. \square

To finish the proof of the theorem, we show that $L = st(\rho) \cap K_{<\rho}$ satisfies the hypotheses of the Main Lemma (6.2.4), and is hence contractible.

It is clear that L satisfies condition (i).

To verify (ii), let $v = \rho^J$ be any vertex in L . We must find a vertex $v' = \rho^{J'}$ of $L \cap TR(v)$ with $v \in st_+(v')$, i.e. $J' \subset J$. Since $v \in st(\rho) \cap st(\rho')$, where $\|\rho\| \leq \|\rho'\|$, the Factorization Lemma implies that some ideal edge in J must have non-negative reductivity. Thus we can take J' to be the set of all reductive edges in J . \square

To prove that K is contractible, let $\{x_i\}$ be the standard basis of F_n corresponding to the petals of R_0 , and set

$$W_0 = \{x_i\}_{i=1}^n \cup \{x_i x_j\}_{i < j} \cup \{x_i \bar{x}_j\}_{i < j}.$$

Then the star graph $S_{W_0}(\rho_0)$ is the complete graph on $2n$ nodes, i.e. there is exactly one unoriented edge between every pair of nodes.

Proposition 6.2.5. *Let ρ be a rose. Then $\|\rho\|_{W_0} \leq n(2n-1)$, with equality if and only if $\rho = \rho_0$.*

Proof. Consider the Whitehead diagram for an ideal edge i in the star graph $S_{W_0}(\rho_0)$. If i contains s edges, the simple closed curve representing i intersects the star graph in $s(2n-s)$ points. Thus the reductivity of the Whitehead move (i, e) for any edge e of ρ_0 is

$$\text{red}(i, e) = |e| - s(2n-s) = (2n-1) - s(2n-s) = (s-1)(s-(2n-1)).$$

Since i is an ideal edge, we have $1 < s < 2n-1$, so $\text{red}(i, e) < 0$.

We have just shown that ρ_0 has no reductive ideal edges. Since ρ_0 has no strictly reductive edges, Corollary 4.2.3 implies that $\|\rho_0\|_{W_0} = n(2n-1)$ must be minimal. Since ρ_0 has no zero-reductive edges, Proposition 4.2.2 implies that ρ_0 is the unique rose with minimal norm. \square

This proposition says that for $W = W_0$, K_{\min} consists of the star of the single rose ρ_0 , and is hence contractible. Thus we have proved Corollary 6.1.2, i.e. the complex K is contractible. Q.E.D.

References

1. Serre, J.-P., Bass, H.: Arbres, Amalgames, SL_2 . *Asterisque* **46** (1977)
2. Culler, M., Morgan, J.: Group actions on \mathbb{R} -trees. (Preprint)
3. Fathi, A., Laudenbach, F., Poenaru, V.: Travaux de Thurston sur les surfaces. *Asterisque*, **66-67** (1979)
4. Gersten, S.M.: On fixed points of certain automorphisms of free groups. *Proc. Lond. Math. Soc.* **48**, (3) 72-90 (1984)
5. Hoare, A.H.M.: Coinitial graphs and Whitehead automorphisms. *Can. J. Math.* **30**, 112-123 (1979)
6. Higgins, P.J., Lyndon, R.C.: Equivalence of elements under automorphisms of a free group. *J. Lond. Math. Soc.* **8**, 254-258 (1974)
7. McCool, J.: A presentation for the automorphism group of a free group of finite rank. *J. Lond. Math. Soc.* **8**, 259-266 (1974)

8. McCool, J.: Some finitely presented subgroups of the automorphism group of a free group. *J. Algebra* **35**, 205–213 (1975)
9. Morgan, J., Shalen, P.B.: Valuations, trees, and degeneration of hyperbolic structures. *Ann. of Math.* **120**, 401–476 (1984)
10. Quillen, D.: Higher Algebraic K-Theory, I. *Lect. Notes Math.* **341**, 85–147 (1974)
11. Thurston, W.: On the topology and dynamics of diffeomorphisms of surfaces. (Preprint)
12. Whitehead, J.H.C.: On equivalent sets of elements in a free group. *Ann. Math.* **37**, 782–800 (1936)

Oblatum 30-VII-1984 & 18-VI-1985