

Contractibility of Outer space: reprise

Karen Vogtmann

Abstract.

This note contains a newly streamlined version of the original proof that Outer space is contractible.

§1. Introduction

In a series of lectures in August 2014 at the Seventh Seasonal Institute of the Mathematical Society of Japan I began by recalling the construction of Outer space for a free group F_n . This is a finite-dimensional contractible space with a proper action of the group $Out(F_n)$ of outer automorphisms of F_n [4]. I then discussed the idea of developing an analogous outer space for the outer automorphism group of a general right-angled Artin group (RAAG) A_Γ . Such a space was introduced in [2] for the case that A_Γ has no twist automorphisms. We also understand the case that A_Γ is generated entirely by twist automorphisms and signed permutations, where the relevant space is a contractible subspace of the symmetric space for $SL(n, \mathbb{R})$. Our candidate for an Outer space for a general RAAG is a hybrid of these two spaces.

The key result of [2] is that the outer space constructed there is contractible. Although there are now several proofs that the original Outer space is contractible, the combinatorial techniques used in the original proof [4] are ultimately what worked for us in the more general RAAG setting. Thus the proof in [2] follows the original proof but also incorporates simplifications, including some which were already introduced in [5]. Unfortunately, new complications also arise due to the fact that the outer automorphism group of a general RAAG is more complicated than the outer automorphism group of a free group. In this note I will avoid these complications by just giving the complete simplified argument in the case of a free group. One motivation for doing this is to

Received February 2, 2015.

Revised May 8, 2015.




Fig. 1. A marked graph (G, g)

clarify the original proof, another is to make it easier for those who want to understand the general RAAG case to follow the argument.

§2. Outer space and its spine

In this section we very briefly recall the definition of Outer space for a free group F_n and its spine K_n . For a more detailed introduction to these spaces, see [8].

A *rose* is a graph with one vertex and n edges (so the edges are loops at the vertex). We begin by fixing a specific rose R_n whose fundamental group we identify with F_n . A point in Outer space is then an equivalence class of *marked metric graphs* (G, g) (see Figure 1), i.e.

- (1) G is a finite metric graph with all vertices of valence at least 3.
- (2) The volume of G (i.e. the sum of the lengths of its edges) is 1.
- (3) $g: R_n \rightarrow G$ is a homotopy equivalence.
- (4) (G, g) is equivalent to (G', g') if there is an isometry $h: G \rightarrow G'$ with $h \circ g \simeq g'$.

In the rest of this note we will not be careful about distinguishing between a marked graph and its equivalence class.

Remark 2.1. Requiring that the volume be equal to one is a means of normalizing the projective class of a metric graph. It is also sometimes convenient to consider other normalizations, or even to consider the unprojectivized version of Outer space, where the edges of G are allowed to have any positive lengths.

Outer space is a union of open simplices, where the simplex containing (G, g) consists of all marked graphs one can obtain by varying the (positive) edge-lengths of G while keeping the volume equal to one. Passing to a face of the simplex corresponds to shrinking some edges to points. Some faces of each simplex are missing, since if an entire loop is

shrunk to a point the fundamental group is no longer F_n and the induced marking is no longer a homotopy equivalence. Formally including these missing faces gives a simplicial complex, called the *simplicial closure* of Outer space; the simplices which are not in Outer space are said to be *at infinity*.

Remark 2.2. The simplicial closure of Outer space is also called the *free splitting complex* or the *sphere complex*; these terminologies arise from different (equivalent) descriptions of Outer space, the first as a space of actions of F_n on metric simplicial trees and the second as a space of weighted sphere systems in a doubled handlebody.

The set of all open simplices of Outer space is partially ordered by the face relation, and the geometric realization of this partially ordered set (poset) is called the *spine* of Outer space. Thus a simplex in the spine K_n is a chain of open simplices $\sigma_0 \subset \dots \subset \sigma_k$ with σ_i a proper face of σ_{i+1} . In other words, K_n is a subcomplex of the barycentric subdivision of the simplicial closure. There is a natural equivariant deformation retract of all of Outer space onto K_n , performed by pushing linearly from the (missing) simplices at infinity onto the spine.

Thus to prove that Outer space is contractible it suffices to show that the spine K_n is contractible.

§3. Structure of the spine K_n and plan of attack

Since vertices of K_n sit at the barycenters of simplices of Outer space all edges have the same length and we may think of these as purely combinatorial (as opposed to metric) objects or, equivalently, assume all edges have length one. We take this point of view for the rest of the paper.

To describe the simplices of K_n , recall that a *forest* in a graph G is a subgraph which contains no loops, i.e. a forest is a disjoint union of trees. Collapsing each tree of a forest F to a point gives a new graph $G//F$ and the collapsing map $c_F: G \rightarrow G//F$ is a homotopy equivalence, so the composition $c_F \circ g$ is a marking of $G//F$. Forest collapse gives the vertices of K_n the structure of a partially ordered set (poset). The entire complex K_n is the *geometric realisation* (also called the *order complex*) of this poset. In other words, there is an edge in K_n from (G, g) to (G', g') whenever (G', g') can be obtained from (G, g) by a forest collapse, and there is a k -simplex for every chain of k forest collapses

$$(G_0, g_0) \rightarrow (G_1, g_1) \rightarrow \dots \rightarrow (G_k, g_k).$$

Every vertex in K_n is connected by an edge to at least one marked rose (R, r) , obtained from (G, g) by collapsing a maximal tree. Thus K_n is the union of the simplicial stars of its marked roses. In order to prove that K_n is contractible the idea is to build K_n by starting with the star of a single marked rose (which is contractible), then attach the rest of the stars in some order and prove that at each stage we are attaching along something contractible. In order to carry out this plan we need to

- (1) Define a *norm* on marked roses and prove that this norm well-orders the marked roses.
- (2) Identify which marked graphs in the star of a marked rose are *reductive*, i.e. are adjacent to marked roses of smaller norm.
- (3) Prove that the subcomplex of reductive marked graphs in a star is contractible.

We actually perform a bit of sleight-of-hand because it is easy to show that K_n is connected. We then show that *if* the subcomplex of reductive marked graphs is non-empty, *then* it is contractible. This shows that K_n is a union of contractible components, but since it is connected it is actually contractible.

§4. Connectivity of K_n via Stallings folds

If $h: R_n \rightarrow R$ is a homeomorphism then the marked rose $\rho_0 = (R, h)$ is called the *standard rose*. Since every vertex of K_n is connected to a marked rose, to prove that K_n is connected it suffices to connect any marked rose to ρ_0 . The fact that you can do this follows easily from Nielsen's theorem that $Out(F_n)$ is generated by signed permutations and transvections (i.e. automorphisms which multiply one generator by another). But there is also a very slick, completely elementary way to see this, due to Stallings (which also reproves Nielsen's theorem).

A map $g: G \rightarrow H$ between two graphs is called a *graph morphism* if it sends vertices to vertices and edges either to single edges or to vertices. If we are allowed to add bivalent vertices to G then any continuous map is homotopic to a graph morphism, so we will assume all of our maps are graph morphisms.

If a graph morphism $g: G \rightarrow H$ is not locally injective then either some edge of G must be mapped to a point in H or two edges of G emanating from the same vertex must map to the same edge in H . In either case, g factors through a morphism $G \rightarrow G_1$, which either collapses an edge (in the first case) or folds two edges together (in the second case). In the second case this morphism is called a *Stallings fold*. An example is illustrated by the top arrow in Figure 2.




Fig. 2. A fold is a blowup followed by a forest collapse

Proposition 4.1. K_n is connected.

Proof. As remarked above, it suffices to connect any marked rose $\rho = (R, r)$ to the standard rose ρ_0 . We begin by representing a homotopy inverse to r by a graph morphism $s: R \rightarrow R_n$. To do this we need to subdivide the edges of R suitably. Technically this is not allowed in K_n since it introduces bivalent vertices, but we can recover the point of K_n by simply ignoring the bivalent vertices. If s is not locally injective, then either some edge collapses or you can fold two edges which start at the same vertex. Note that these edges must have distinct terminal vertices, since otherwise they would form a loop with null-homotopic image, which can't happen because s is a homotopy equivalence.

Recall that collapsing a forest in a marked graph gives an edge in K_n (unless the forest contains only proper subsets of subdivided edges, in which case collapsing does not change the point of K_n). The reverse of a single edge collapse is called a *blowup*. A fold corresponds to a blowup followed by a forest collapse when the folded edges have distinct terminal vertices; thus a Stallings fold gives a path in the 1-skeleton of K_n . See Figure 2 for the case that neither edge is a loop. There is a similar picture if one edge is a loop.

If s is not locally injective, perform a fold. (If a univalent vertex is produced, also collapse the adjacent edge; such a fold and collapse does not change the point of K_n .) If the induced map is not locally injective then fold or collapse again, thus producing a path in K_n . This process has to stop because each time the total number of edges in the graph decreases. When it stops, the induced map $f: G \rightarrow R_n$ is locally injective.

We now claim that a locally injective map f is actually a homeomorphism, so the path in K_n has arrived at $(f^{-1}, G) = \rho_0$. To see this, let x_i be the i -th petal of R_n . Since f is a homotopy equivalence, there is some loop ℓ_i in G with $f(\ell_i) \simeq x_i$. Since f is locally injective, ℓ_i is a simple loop in G . If $i \neq j$, then the loops ℓ_i and ℓ_j can intersect in at most a point, since that is true of their images x_i and x_j , so any overlap would have to collapse to a point. The union of the ℓ_i must be all of G , since otherwise the complement would be a forest which must collapse. Finally the ℓ_i must all intersect in the same point, forming a rose. Q.E.D.

§5. The norm of a rose

The next task is to find a ‘‘Morse function’’ which totally orders the roses. The idea is that more complicated markings should come later in the ordering.

For any loop γ in a graph G , let $\ell_G(\gamma)$ denote the length of the shortest loop in the (free) homotopy class of γ , where we think of each edge of G as having length one. Note that γ is shortest in its homotopy class if and only if it is locally injective, in which case we call it a *tight* loop.

Fix a basis x_1, \dots, x_n for F_n and list the conjugacy classes in order of increasing (cyclically reduced) word-length:

$$\mathcal{W} = (w_1, w_2, \dots) = (x_1, x_2, \dots, x_n, x_1^2, \dots, x_1x_2, \dots, x_1x_2^{-1}, \dots, x_1^3, \dots).$$

Note that it is redundant to include both w and w^{-1} in \mathcal{W} , so we won't. Let $\mathbb{Z}^{\mathcal{W}}$ be the associated ordered abelian group, with the lexicographical order. For any rose $\rho = (r, R)$ and any $w \in \mathcal{W}$, define $\|\rho\|_w$ to be equal to $\ell_R(r(w))$. The norm of ρ is then defined by

$$\|\rho\| = \|\rho\|_{\mathcal{W}} = (\|\rho\|_{w_1}, \|\rho\|_{w_2}, \dots) \in \mathbb{Z}^{\mathcal{W}}.$$

This norm totally orders the roses, according to the following basic theorem, proved independently by Alperin and Bass and by Culler and Morgan:

Theorem 5.1. [1, 3] *A free action of F_n on a simplicial tree is determined by its translation length function.*

We will show that our norm has the following stronger property, which we will need to do induction:

Proposition 5.2. *The set of roses is well-ordered in this norm.*

Lemma 5.3. *Any free minimal action of F_n on a simplicial tree is determined by the translation lengths of finitely many conjugacy classes.*

Proof. Assume first that the quotient by the action is a rose. We claim that in this case the action is determined by the translation lengths of the $2n^2 + 2n$ classes which have length at most 2. These are represented by oriented loops of length at most 2 in the quotient rose. The petals of the rose correspond to some basis w_1, \dots, w_n for F_n .

Suppose there was another free minimal action with quotient a rose in which these classes have the same lengths. This action is the first action twisted by an (outer) automorphism $\hat{\phi}$. Any automorphism ϕ representing $\hat{\phi}$ must permute the conjugacy classes of $w_1, w_1^{-1}, \dots, w_n, w_n^{-1}$ since these are all of the classes of length 1. In fact, we may assume ϕ sends each w_i to a conjugate of itself, since permuting and inverting the w_i can be realized by an isometry of the rose, which lifts to an equivariant isometry of the tree.

Take a representative ϕ for $\hat{\phi}$ with $\phi(w_1) = w_1$, and suppose $\phi(w_2) = uw_2u^{-1}$. Then $\phi(w_1w_2) = w_1uw_2u^{-1}$. Since this is conjugate to an element of length 2, u must be a power of w_1 . Therefore, after composing ϕ with conjugation by u^{-1} we may assume $\phi(w_1) = w_1$ and $\phi(w_2) = w_2$.

Now consider $\phi(w_i) = vw_iv^{-1}$ for $i > 2$. The argument above shows that v must be a power of w_1 and a power of w_2 , so in fact $v = 1$ and ϕ is the identity.

If the quotient by the action is a marked graph which is not a rose, choose a maximal tree and collapse it to get a rose. We can distinguish this rose from any other rose by the lengths of finitely many conjugacy classes. Our original marked graph is obtained from this rose by blowing up the vertex into a tree. There are only finitely many ways to do this, which by Theorem 5.1 can be distinguished by finitely many lengths. Thus the entire action is determined by the lengths of finitely many conjugacy classes. Q.E.D.

Remark 5.4. The finite set of conjugacy classes found in Lemma 5.3 depends on the action. It can be shown that for any fixed finite set of conjugacy classes there are two roses in which those conjugacy classes have the same translation length. (In fact, you can find an arbitrary (finite) number of roses in which those conjugacy classes have the same translation length, see [7]).

Proof of Proposition 5.2. We have to show that any subset U of roses has a least element. Set $U = U_0$ and define a decreasing chain

$$U = U_0 \supset U_1 \supset U_2 \dots,$$

where U_i is the set of elements in U_{i-1} for which $\|\rho\|_{w_i}$ is minimal possible, say $\|\rho\|_{w_i} = \ell_i$. Note that each element of U_i is \leq each element of U_{i-1} in the ordering.

The function $f: F_n \rightarrow \mathbb{Z}$ defined by $f(w_i) = \ell_i$ satisfies the axioms for a translation length function, so corresponds to an action on a tree. This action is free since $f(w_i) = \ell_i \neq 0$ for all i . Therefore it corresponds to a marked graph γ (with all edges of length 1). By Lemma 5.3 γ is determined by the lengths of finitely many elements of F_n . So for N sufficiently large U_N has only one element, a rose which must be equal to γ , which is smaller than any other element of U . Q.E.D.

§6. Separating edges

It is convenient to prove contractibility just for the subcomplex L_n of K_n spanned by graphs with no separating edges. This is justified by the following observation.

Proposition 6.1. *K_n deformation retracts to the subcomplex L_n spanned by graphs with no separating edges.*

The deformation retraction is easy to see: one just uniformly shrinks all separating edges to zero. Since K_n is the geometric realisation of a poset, one can give a formal proof using Quillen's Poset Lemma, which will also come in handy later.

Lemma 6.2. [*Quillen's Poset Lemma [6]*] *Let P be a poset and $f: P \rightarrow P$ a poset map (i.e. $x \leq y$ implies $f(x) \leq f(y)$). If in addition $f(x) \leq x$ for all x , then $|P|$ deformation retracts to $|f(P)|$, where vertical bars denote geometric realization.*

Note that by using the opposite poset we can draw the same conclusion if $f(x) \geq x$ for all x . The proof of the Poset Lemma is a straightforward application of the prism operator, and is left to the reader.

Proof of Proposition 6.1. The map $f: K_n \rightarrow K_n$ which contracts each separating edge is a poset map with image L_n . Q.E.D.

We reiterate our plan of attack, as an excuse to introduce some notation. All roses are in L_n , and we view L_n as the union of the simplicial stars of its roses:

$$L_n = \bigcup_{\text{roses } \rho} st(\rho).$$

We construct L_n by starting with the star of the (unique) rose of minimal norm and adding stars of roses in the order dictated by the norm, i.e.

for each rose ρ , define

$$L_{<\rho} = \bigcup_{\|\rho'\| < \|\rho\|} st(\rho').$$

We will prove that if $st(\rho) \cap L_{<\rho}$ is non-empty, then it is contractible. This will show that L_n is a union of contractible components. But we already know L_n is connected, so it is contractible.

§7. Reductive graphs and the Factorization lemma

We call a marked graph *reductive* if it is in $st(\rho) \cap L_{<\rho}$; thus (G, g) is reductive if and only if G contains maximal trees Φ and F such that collapsing Φ gives ρ and collapsing F gives a different rose ρ' with $\|\rho'\| < \|\rho\|$.

For each edge $e \in G$ of a marked graph (G, g) , let $|e| \in \mathbb{Z}^{\mathcal{W}}$ be the element whose coefficient $|e|_w$ is the number of times a tight representative of $g(w)$ crosses e , in either direction. Since collapsing a forest sends tight paths to tight paths, we get

$$\|\rho'\| = \|\rho\| + \sum_{\alpha \in \Phi} |\alpha| - \sum_{e \in F} |e|.$$

Proposition 7.1. *If $\Phi = \{\alpha_1, \dots, \alpha_k\}$ and $F = \{e_1, \dots, e_k\}$ are maximal trees in a graph G , then there is a permutation σ of $\{1, \dots, n\}$ such that $e_{\sigma(i)}$ connects the two components of $\Phi - \alpha_i$ for each i .*

Proof. First of all we want σ to be the identity on common edges of Φ and F . We can then contract these edges to reduce the problem to the case that Φ and F have no edges in common.

It's easy to move from Φ to F by replacing one edge at a time, but we want to do something a little more subtle than that. Here's one proof you can do it:

$\Phi \cup F$ is a graph of rank k , and the edges of Φ or of F each give a basis for $H_1(\Phi \cup F)$. The change of basis matrix B has entry $b_{ij} = \pm 1$ if the unique path in F connecting the endpoints of α_j crosses e_i ; otherwise $b_{ij} = 0$. Since this matrix is non-singular, some term

$$sign(\sigma) \prod_{i=1}^n b_{i, \sigma(i)}$$

in the expression for the determinant of B is non-zero. Then $e_{\sigma(i)}$ joins the two components of $\Phi - \alpha_i$. Q.E.D.

Corollary 7.2 (Factorization lemma). *In $st(\rho)$, every reductive (G, g) is adjacent to a 2-vertex reductive graph.*

Proof. With Φ , F and σ as in Lemma 7.1, we have

$$\|\rho'\| = \|\rho\| + \sum_i |\alpha_i| - \sum_i |e_{\sigma(i)}| = \|\rho\| + \sum_i (|\alpha_i| - |e_{\sigma(i)}|).$$

Since $\|\rho'\| < \|\rho\|$, we must have $|\alpha_i| - |e_{\sigma(i)}| < 0$ for some i . Then the two-vertex graph obtained by collapsing all edges of $\Phi - \alpha_i$ is reductive (and is connected to (G, g)).

Q.E.D.

Thus we may view $st(\rho) \cap L_{<} \rho$ as a union of stars of 2-vertex graphs. If we are lucky there is only one reductive 2-vertex graph so $st(\rho) \cap L_{<} \rho$ is contractible. We are seldom so lucky, however, and need to work harder. In order to do this we will reinterpret reductive graphs and the norm using a neat combinatorial model originally introduced by Whitehead in the context of sphere complexes in doubled handlebodies. This model (translated into the language of partitions and graphs instead of sphere complexes) is explained in the next section.

§8. Ideal edges

We now fix a rose $\rho = (r, R)$ and re-interpret graphs in $st(\rho)$ in terms of partitions of the set H of half-edges of R . We denote the natural involution on H by $e \mapsto \bar{e}$. This section explains the translation from graphs to partitions.

A marked graph (G, g) is in $st(\rho)$ if and only if G has a maximal tree Φ so that the composition of the collapsing map c_Φ with g is homotopic to r . In particular, the edges of $G - \Phi$ are mapped homeomorphically onto the edges of R , so if you snip each edge of $G - \Phi$ you obtain a tree whose leaves are labelled by the elements of H (see Figure 3). Each edge α of Φ gives a partition of H into two subsets, called the *sides* of α , and different edges α and β give *compatible* partitions, in the sense that $A \cap B = \emptyset$ for some choice of sides A of α and B of β . Conversely, given any set of compatible partitions of H we can reconstruct the tree Φ with leaves labeled by H , and recover (G, g) by reconnecting the paired elements of H .

Definition 8.1. A partition of the set H of half-edges of R is an *ideal edge* if it separates some pair $\{e, \bar{e}\}$. An ideal edge is *trivial* if one side is a singleton. An *ideal tree* is a set of non-trivial ideal edges which are pairwise compatible.




Fig. 3. A graph in $st(\rho)$ is equivalent to a partition of the half-edges of ρ

Note that a trivial ideal edge corresponds to a graph with a bivalent vertex, and a partition which is not an ideal edge corresponds to a graph with a separating edge. Since none of our graphs have bivalent vertices or separating edges, the vertices of $st(\rho)$ correspond to ideal trees. Collapsing an edge of a tree corresponds to removing a partition from the associated ideal tree. Thus the simplicial complex $st(\rho)$ is the geometric realization of the poset of ideal trees, ordered by inclusion.

We say an ideal tree or edge is *reductive* if the corresponding graph is reductive. The Factorization Lemma, reinterpreted in this language, says that if Φ is a reductive ideal tree, then it contains a reductive ideal edge.

If α is a reductive ideal edge, then there is some pair $\{e, \bar{e}\}$ in H separated by α with $|\alpha| - |e| < 0$. If A is the side of α containing e , we call (A, e) a *reductive pair* for α .

§9. The star graph and the norm

We have reinterpreted graphs in $st(\rho)$ as partitions of the half-edges of ρ . We now need to interpret the norm of a rose in this model. In order to understand which partitions are in the reductive subcomplex of $st(\rho)$, we also need to interpret $|e|$ and $|\alpha|$ in this model.

Fix a rose $\rho = (r, R)$ and let H be the half-edges of R . For any conjugacy class w we associate a graph Γ_w called the *star graph* of w . The vertices of Γ_w are the elements of H . To define the edges take a tight loop representing $r(w)$ in R , then snip the edges of R to make a tree with $2n$ leaves. This cuts $r(w)$ into segments joining the cuts; these are the edges of Γ_w (Figure 4). Since $\|\rho\|_w$ is the length of $r(w)$, the sum of the valences of Γ_w is twice $\|\rho\|_w$.

Let (G, g) be the two-vertex graph represented by a single ideal edge α . Recall that $|e|_w$ denotes the number of times a tight loop representing




Fig. 4. Star graph of the loop $w = e_2e_4^{-1}e_3^2$

$g(w)$ passes over e . If $e \neq \alpha$, then we could also measure $|e|_w$ in ρ where $|e|_w$ is equal to the valence of e (or of \bar{e}) in the star graph Γ_w . If $e = \alpha$, then $|\alpha|_w$ is equal to the number of edges of Γ_w with one vertex on each side of α .

§10. Contractibility of $st(\rho) \cap L_{<\rho}$

Since we have reinterpreted $st(\rho)$ as the geometric realization of the poset of ideal trees, we will make free and frequent use of Quillen’s Poset Lemma (Lemma 6.2).

The complex we are interested in, $st(\rho) \cap L_{<\rho}$, is the poset of *reductive* ideal trees; call it P . As we remarked in section 9, the Factorization Lemma says that each $\Phi \in P$ contains at least one reductive ideal edge. Therefore the map $P \rightarrow P$ which throws out all of the *non-reductive* ideal edges is a well-defined map. It satisfies the hypotheses of the Poset Lemma, so $|P|$ deformation retracts to its image, which is the realization of the subposet Q of *strictly reductive* ideal trees, i.e. ideal trees all of whose edges are reductive.

Now choose a maximally reductive ideal edge μ , i.e. there is a reductive pair (M, m) for μ with $|m| - |\mu|$ maximal among all reductive pairs (A, a) . If every other reductive ideal edge is compatible with μ , then μ can be added to every strictly reductive ideal tree Φ , and the poset maps $\Phi \mapsto \Phi \cup \{\mu\} \mapsto \{\mu\}$ retract $|Q|$ to the single point μ .

If there are edges α which cross μ , we have to work harder. Here is the Key Lemma (see Figure 5):

Lemma 10.1 (Key Lemma). *Let μ be a maximally reductive ideal edge, with maximal pair (M, m) , let α be a reductive edge that crosses*




Fig. 5. Key Lemma: One of the dotted ideal edges is reductive

μ , and let A be the side of α that contains m . Then $A \cup M$ or $\bar{A} \cap M$ is one side of a reductive ideal edge γ .

Assuming the Key Lemma, we proceed as follows:

Choose α , with side A containing m , such that

- If β is compatible with α and the side B of β containing m also contains A , then $B \supset M$.

By the Key Lemma, one of $M \cap \bar{A}$ or $M \cup A$ determines a reductive ideal edge γ . We now observe that γ is compatible with α , μ and with every β compatible with α . Therefore

$$\Phi \mapsto \begin{cases} \Phi \cup \gamma & \text{if } \alpha \in \Phi \\ \Phi & \text{if } \alpha \notin \Phi \end{cases}$$

is a poset map. It satisfies the condition of the Poset Lemma, so retracts $|Q|$ to its image.

In the image, everything that contains α also contains γ . Then the map throwing α out of every Φ that contains it is also a poset map. The final effect is to replace the edge α which crosses μ by the edge γ which is compatible with μ , i.e. the image is now all reductive ideal trees which do not contain α .

We repeat this procedure until we have eliminated every ideal edge which crosses μ . Then we can retract to μ as before.

§11. Proof of the Key Lemma

In this section we prove the Key Lemma needed in the proof of contractibility. Recall we are working with a fixed rose $\rho = (r, R)$ and partitions of the set of half edges H of R . We first define the *dot product* $A.B$ of disjoint subsets A and B of H as the element of $\mathbb{Z}^{\mathcal{W}}$ with coordinate $(A.B)_w$ equal to the number of edges in the star graph Γ_w with one vertex in A and one vertex in B .




Fig. 6. Diagram for Lemma 11.2

For $A \subseteq H$, set $\bar{A} = H \setminus A$ and $|A| = A.\bar{A}$. As noted in section 9, for $e \in H$ $|e|_w$ is just the valence of e in the star graph $\Gamma(w)$, and if A is either side of an ideal edge α , then $|A|_w = |\alpha|_w$.

We use “+” to denote disjoint union of sets, as well as addition in \mathbb{Z}^W , resulting in the following pleasing formulas.

Lemma 11.1. *If A, B and C are disjoint subsets of H , then*

- (1) $A.B = B.A$
- (2) $A.(B + C) = A.B + A.C$

Proof. Straightforward.

Q.E.D.

Lemma 11.2. *Let A and B be subsets of H . Then $|A \cap B| + |A \cup B| \leq |A| + |B|$.*

Proof. A and B together partition H into disjoint subsets $Z = A \cap B$, $W = \bar{A} \cap \bar{B}$, $X = A \cap \bar{B}$ and $Y = B \cap \bar{A}$ (see Figure 6). We compute

$$|A| = (X + Z).(Y + W) = X.Y + X.W + Z.Y + Z.W$$

$$|B| = (Y + Z).(X + W) = Y.X + Y.W + Z.X + Z.W$$

and

$$|A \cap B| = Z.(X + Y + W) = Z.X + Z.Y + Z.W$$

$$|A \cup B| = (X + Z + Y).W = X.W + Z.W + Y.W$$

So altogether we have $|A| + |B| = |A \cap B| + |A \cup B| + 2(X.Y)$; in particular

$$|A \cap B| + |A \cup B| \leq |A| + |B|.$$

Q.E.D.

Lemma 11.3 (Key Lemma). *Let μ be a maximally reductive ideal edge, with maximal pair (M, m) , let α be a reductive edge that crosses μ , and let A be the side of α that contains m . Then $A \cup M$ or $\overline{A} \cap M$ is one side of a reductive ideal edge γ .*

Proof of Key Lemma. Since α and μ cross, together they partition H into four disjoint subsets, which we will call *sectors*. Since α is reductive, there is $a \in A$ with $|a| - |\alpha| > 0$. The proof falls into cases depending on the locations of a, \overline{a} and \overline{m} .

If each sector contains one of $a, \overline{a}, m, \overline{m}$, then the only possibility is $a \in A \cap \overline{M}, \overline{a} \in M \cap \overline{A}$ and $\overline{m} \in \overline{A} \cap \overline{M}$ (since we already have $m \in A \cap M$). In this case Lemma 11.2 gives us

$$|A \cap M| + |A \cup M| + |\overline{A} \cap M| + |\overline{A} \cup M| \leq 2|A| + 2|M|$$

or

$$\begin{aligned} (|m| - |A \cap M|) + (|m| - |A \cup M|) + (|a| - |\overline{A} \cap M|) + (|a| - |\overline{A} \cup M|) \\ \geq 2(|a| - |A|) + 2(|m| - |M|) \end{aligned}$$

so

$$\begin{aligned} (|m| - |A \cup M|) + (|a| - |\overline{A} \cap M|) \geq \\ 2(|a| - |A|) + [(|m| - |M|) - (|m| - |A \cap M|)] + [(|m| - |M|) - (|a| - |\overline{A} \cup M|)] \end{aligned}$$

Since (A, a) is a reductive pair and (M, m) is maximally reductive, each of the three terms on the bottom line is positive, so the sum on the next line up is positive, which implies that at least one of $(A \cup M, m)$ or $(\overline{A} \cap M, a)$ is a reductive pair, as required.

We may now assume some sector contains none of $a, \overline{a}, m, \overline{m}$. Since a and \overline{a} (resp. m and \overline{m}) can't be in the same sector, some sector has a or \overline{a} and m or \overline{m} . Replacing (M, m) by $(\overline{M}, \overline{m})$ if necessary, we may assume $a, m \in A \cap M$. The rest of the proof breaks into cases depending on the positions of \overline{a} and \overline{m} .

If $\overline{m} \in \overline{A} \cap \overline{M}$, then the inequality $|A \cap M| + |A \cup M| \leq |A| + |B|$ from Lemma 11.2 gives

$$(|a| - |A \cap M|) + (|m| - |A \cup M|) \geq (|a| - |A|) + (|m| - |M|)$$

so

$$|m| - |A \cup M| \geq [|a| - |A|] + [(|m| - |M|) - (|a| - |A \cap M|)].$$

Since (M, m) is maximally reductive, both terms on the right hand side are positive, showing that $(A \cup M, m)$ is a reductive pair.

If $\bar{m} \in A$ and $\bar{a} \in \bar{A} \cap \bar{M}$, the same proof with the roles of a and m switched shows $(A \cup M, a)$ is a reductive pair. The only remaining case is $\bar{m} \in A, \bar{a} \in M$; in this case we use Lemma 11.2 with the sets \bar{A} and M to get

$$|\bar{A} \cap M| + |\bar{A} \cup M| \leq |\bar{A}| + |M| = |A| + |M|$$

so

$$\begin{aligned} |m| - |\bar{A} \cap M| + |a| - |\bar{A} \cup M| &\geq |a| - |\bar{A}| + |m| - |M| \\ |m| - |\bar{A} \cap M| &\geq (|a| - |\bar{A}|) + [(|m| - |M|) - (|a| - |\bar{A} \cup M|)], \end{aligned}$$

and again both terms on the right are positive since (M, m) is maximally reductive, showing $(\bar{A} \cap M, m)$ is a reductive pair.

Q.E.D.

References

- [1] R. ALPERIN AND H. BASS, *Length functions of group actions on Λ -trees*, Combinatorial group theory and topology (Alta, Utah, 1984), 265–378, Ann. of Math. Stud., **111**, Princeton Univ. Press, Princeton, NJ, 1987.
- [2] R. CHARNEY, N. STAMBAUGH AND K. VOGTMANN, *Outer space for untwisted automorphisms of right-angled Artin groups*, Geom. Topol. **21** (2017), no. 2, 1131–1178.
- [3] M. CULLER AND J. W. MORGAN, *Group actions on R -trees*, Proc. London Math. Soc. (3) **55** (1987), no. 3, 571–604.
- [4] M. CULLER AND K. VOGTMANN, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), no. 1, 91–119.
- [5] S. KRSTIC AND K. VOGTMANN, *Equivariant outer space and automorphisms of free-by-finite groups*, Comment. Math. Helv. **68** (1993), no. 2, 216–262.
- [6] D. QUILLEN, *Homotopy properties of the poset of nontrivial p -subgroups of a group*, Adv. Math. **28** (1978), 101–128.
- [7] J. SMILLIE AND K. VOGTMANN, *Length functions and outer space*, Michigan Math J. **39** (1992) 485–493.
- [8] K. VOGTMANN, *Automorphisms of free groups and Outer space*, Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000). Geom. Dedicata **94** (2002), 1–31.

Cornell University in Ithaca, NY, U.S.A
The University of Warwick in Coventry, U.K.
E-mail address: K.Vogtmann@warwick.ac.uk