

Contents

1	Introduction and Preliminaries	2
1.1	Group Automorphisms	2
1.2	Group Cohomology	3
1.3	Graphs	5
1.4	Cutting and Pasting	6
1.5	Surgery Theory	7
1.6	Homotopy Theory and Spectral Sequences	8
1.6.1	Homotopy Theory	8
1.6.2	Spectral Sequences	9
2	Culler-Vogtmann Outer Space and the Sphere Complex	10
2.1	Culler-Vogtmann Outer Space	10
2.2	The Sphere Complex	12
3	Normal Form	13
3.1	Definition and Existence of Normal Forms	13
3.2	Equivalence of Normal Form Systems	14
4	Contractibility of Sphere Complex	16
5	Relation to Culler-Vogtmann Outer Space	17
6	Outline of Homological Stability	19
7	Connectivity of Subcomplexes Part I	20
8	Connectivity of Subcomplexes Part II	22
8.1	The Degree Theorem	24
8.2	Proof of Connectivity	26
9	Connectivity of Subcomplexes Part III	26
10	Spectral Sequence Arguments	27
10.1	Stabilisation of α	27
10.2	Stabilisation of μ	28
10.3	Stabilisation of β for Small s	29
11	Whitehead's Algorithm	29
11.1	Transformations and Sphere Bases	30
11.2	Recognising and Modifying Sphere Bases	31
11.3	Normal Forms of Sphere Bases	32
12	Conclusion	34

1 Introduction and Preliminaries

In this essay, we will introduce and describe a nice topological space on which $\text{Out}(F_n)$, the outer automorphism group of the free group, acts. We will define the space in two different ways, prove that it is contractible, describe a nice subspace on which $\text{Out}(F_n)$ acts with finite stabilisers and finite quotient, and use this to deduce some cohomological information about $\text{Out}(F_n)$. We will largely follow [7] and [4] in this part. After this, we will proceed to follow [7, 9, 10] in proving that the i th homology groups of $\text{Out}(F_n)$ are isomorphic after a certain value of n , streamlining many duplicated arguments across the 3 papers, and uniting the many errata of [7] in one place. We will then use some of the 3-manifold topology developed to discuss Whitehead's algorithm for detecting when two elements of a free group can be mapped to each other under some automorphism, following [16, 6]. At the end of the essay, we will indicate some other views that may be taken on outer space that we did not have space to mention in this essay.

1.1 Group Automorphisms

We begin by defining $\text{Out}(F_n)$. For completeness, we give the following definitions first.

Definition 1.1. *The **automorphism group** of a group G , denoted $\text{Aut}(G)$, is the group of automorphisms of G under composition. An automorphism is called **inner** if it is of the form $\phi_g: x \mapsto gxg^{-1}$ for some $g \in G$. We denote the set of inner automorphisms by $\text{Inn}(G)$.*

Remark 1.1. The set $\text{Inn}(G)$ is a normal subgroup of $\text{Aut}(G)$. We have that $\phi_h\phi_g(x) = h(gxg^{-1})h^{-1} = (hg)x(hg)^{-1} = \phi_{hg}(x)$, $\phi_g^{-1} = \phi_{g^{-1}}$ and $\psi\phi_g\psi^{-1}(x) = \phi_{\psi(g)}(x)$. Therefore we further see that the map $g \mapsto \phi_g$ has image $\text{Inn}(G)$ and kernel $Z(G)$, so $\text{Inn}(G) \cong G/Z(G)$.

From the above remark if we understand G well, then we understand the inner automorphisms well. It is therefore useful, if we wish to study the automorphisms of a group, to study the following group:

Definition 1.2. *The quotient $\text{Aut}(G)/\text{Inn}(G)$ is called the **outer automorphism group**, and is denoted $\text{Out}(G)$.*

Example 1. It is well known that

$$\text{Aut}(S_n) = \begin{cases} e & n = 2 \\ S_6 \cdot 2 & n = 6 \\ S_n & n \neq 2, 6 \end{cases}$$

where $S_6 \cdot 2$ is a certain group which contains S_6 as an index 2 subgroup. It follows (as $Z(S_n) = e$, unless $n = 2$), that $\text{Inn}(S_n) = S_n$ (unless $n = 2$) and so

$$\text{Out}(S_n) = \begin{cases} C_2 & n = 6 \\ e & n \neq 6. \end{cases}$$

We make one final definition which will be needed for our discussion on cohomology.

Definition 1.3. *A group is **virtually P** for a property P if it has a finite index subgroup with property P .*

1.2 Group Cohomology

There will be no proofs in this section, as the purpose of this section is to collect those results we need from [2] to prove the facts mentioned in the introduction about cohomology. We note that we may dualise everything to get homology as usual. We make no further comment on this, and expect the reader to understand how to do so.

We begin by defining group cohomology. As usual, we will try to define it as the derived functor of the hom-functor.

We fix a group G and a projective resolution

$$\cdots \rightarrow P_2 \rightarrow P_1 \rightarrow P_0 \rightarrow \mathbb{Z} \rightarrow 0$$

of the trivial $\mathbb{Z}G$ -module $\mathbb{Z}$ as a $\mathbb{Z}G$ -module.

Definition 1.4. We define the **group cohomology** of G with coefficients in M to be the cohomology of the cochain complex

$$0 \rightarrow \text{Hom}_{\mathbb{Z}G}(P_0, M) \rightarrow \text{Hom}_{\mathbb{Z}G}(P_1, M) \rightarrow \text{Hom}_{\mathbb{Z}G}(P_2, M) \rightarrow \cdots$$

Remark 1.2. By standard abstract nonsense, this is independent of choice of projective resolution.

Definition 1.5. We define the **cohomological dimension** of a group G to be

$$\text{cd}(G) = \min\{n \mid H^k(G, M) = 0 \text{ for all } k > n \text{ and all } M\}$$

Theorem 1.1. For every exact sequence of $\mathbb{Z}G$ -modules $0 \rightarrow M' \rightarrow M \rightarrow M''$ there is a natural exact Mayer-Vietoris sequence

$$0 \rightarrow H^0(G, M') \rightarrow H^0(G, M) \rightarrow H^0(G, M'') \rightarrow H^1(G, M') \rightarrow \cdots$$

Proof. Brown, Proposition III.6.1(ii'). □

Corollary 1.1. The cohomological dimension of G is exactly the minimal length of a projective resolution

$$\cdots \rightarrow P_2 \rightarrow P_1 \rightarrow P_0 \rightarrow \mathbb{Z} \rightarrow 0.$$

Proof. Brown, Proposition VIII.2.1. □

Corollary 1.2. If the cohomological dimension of G is finite, it is exactly the minimal n such that $H^k(G, \mathbb{Z}G) = 0$ for all $k > n$.

Proof. Brown, Proposition III.6.2. □

Proposition 1.1. If $H \leq G$, then $\text{cd } H \leq \text{cd } G$.

Proof. Brown, Proposition VIII.2.4. □

Proposition 1.2. If the cohomological dimension of G is finite, $H \leq G$ and $[G : H] < \infty$ then $\text{cd } G = \text{cd } H$.

Proof. Brown, Proposition VIII.2.4. □

Proposition 1.3. $\text{cd } \mathbb{Z}/n\mathbb{Z} = \infty$ for every $n \neq 0$.

Proof. Brown, III.1 Example 2. □

We may use the above three propositions to note an interesting phenomenon. Suppose G is virtually torsion-free, but is not itself torsion-free. Then by Propositions 1.1 and 1.3, $\text{cd } G = \infty$. However if $H \leq G$ is finite index and torsion free, then $\text{cd } H$ may be finite, and by Proposition 1.2, this number does not depend on the choice of torsion-free finite index subgroup. We therefore make the following definition:

Definition 1.6. The *virtual cohomological dimension* of a virtually torsion-free group G is the (common) cohomological dimension of its torsion free finite index subgroups.

We now relate our above definitions to the topology of certain complexes related to the group G .

Definition 1.7. A $K(G, 1)$ for a group G is a cell complex X such that $\pi_1(X) = G$ and the universal cover of X is contractible.

Remark 1.3. It is a fact of homotopy theory that the condition that the universal cover be contractible is equivalent to all higher homotopy groups being trivial, and (separately) to all higher homology groups being trivial.

Example 2. The torus is a $K(\mathbb{Z}^2, 1)$, as its universal cover is $\mathbb{R}^2$.

This matters in the context of cohomology for the following reason:

Proposition 1.4. Let X be a contractible cell complex on which G acts freely. Then $C_*(X)$ is a free resolution of $\mathbb{Z}$ over $\mathbb{Z}G$.

Proof. Brown, Proposition III.4.1. □

Corollary 1.3. The cohomological dimension of G is bounded above by the minimal dimension of a $K(G, 1)$.

Proof. Brown, Proposition VIII.2.2. □

Remark 1.4. In fact, this is known to be an equality if the cohomological dimension is not equal to 2, as stated in Chapter VIII of [2]. It is also known that there is a $K(G, 1)$ of dimension 3 for G of cohomological dimension 2, and there is no case where the above corollary is known to not be an equality.

We now wish to translate the above discussion to a topological interpretation of virtual cohomological dimension for virtually torsion-free groups.

Definition 1.8. An action $G \curvearrowright X$ of a group G on a cell complex X is *proper* if the stabiliser of any cell is finite.

Remark 1.5. This agrees with the more general definition of a proper action for general topological space X , but we shall not need that definition.

Proposition 1.5. Suppose we have a contractible complex X of dimension n on which a group G acts properly. Then $\text{vcd}(G) \leq n$.

Proof. Brown, Proposition VIII.2.2. □

There is one further finiteness property we wish to discuss.

Definition 1.9. A group G is of type *FL* if there is a finite length free resolution of $\mathbb{Z}$ as a $\mathbb{Z}G$ -module such that every module in the resolution is of finite rank.

Definition 1.10. A group G is of type VFL if there is a finite index subgroup of type FL .

Corollary 1.4. In the situation of Proposition 1.5, if X/G has finitely many cells, then G is of type VFL .

Proof. Brown, Proposition VIII.11.1. □

We now finally state the precise theorem we will prove about $\text{Out}(F_n)$.

Theorem 1.2. $\text{Out}(F_n)$ is of type VFL and $\text{vcd}(\text{Out}(F_n)) = 2n - 3$.

Remark 1.6. It is true that $\text{Out}(F_n)$ and $\text{Aut}(F_n)$ have torsion-free groups of finite index, due to Culler et al. in [5], however the proof is beyond the scope of this essay. We will also prove a similar result for $\text{Aut}(F_n)$, as it requires little more effort than the $\text{Out}(F_n)$ case.

Proof. We will prove that $\text{vcd}(\text{Out}(F_n)) \geq 2n - 3$ by presenting a subgroup of cohomological dimension $2n - 3$, following [4]. The remainder of the proof is deferred to later. We note that $(S^1)^{2n-3}$ is a $K(\mathbb{Z}^{2n-3}, 1)$, so $\text{cd}(\mathbb{Z}^{2n-3}) = 2n - 3$. Now there is a subgroup of $\text{Out}(F_n)$ which is free abelian of rank $2n - 3$, generated by α_i given as $x_i \mapsto x_1 x_i$ (and the identity on all other generators) and β_i given as $x_i \mapsto x_i x_1$ (and the identity on all other generators) for $i > 1$, so by Proposition 1.1 we are done. □

To conclude this subsection, we state a theorem about homology relating homology of submodules to that of the group which will be needed for homological stability. We assume knowledge of induced modules.

Proposition 1.6. (*Shapiro's Lemma*) If $H \leq G$ and M is a H module, then $H_*(G, \text{Ind}_H^G M) \cong H_*(H, M)$.

Proof. Brown, Proposition III.6.2. □

1.3 Graphs

By a **graph**, we mean a connected 1-dimensional cell complex. In particular, we allow edges from a node to itself, and we allow multiple edges between two vertices.


A **metric graph** is a graph with a non-zero length assigned to each edge. For finite graphs, this naturally forms a metric space with $d(x, y)$ being the infimum of the lengths of paths between x and y (for infinite graphs, a nice necessary, but not sufficient, condition for this pseudometric to be a metric is that the lengths be bounded below by a uniform positive constant).

For a finite metric graph X , we define the **volume** of X to be the sum of all the lengths of the edges. This will be used once later to provide a normalisation for Culler and Vogtmann's construction of outer space.

A **separating edge** of a graph $X = (V, E)$ is some $e \in E$ such that $X \setminus \{e\}$ is disconnected.

We provide an example of the latter two definitions. The label of each edge denotes its length.

Example 3. The following is a metric graph:



This has volume 17. Note that the distance between two vertices need not be the length of the edge between them (eg. $d(a, b) = 2$) and the only separating edge is the one between c and d . Note that graphs need not have a separating edge, for example the triangle.

1.4 Cutting and Pasting

In 3-manifold topology, as in much of maths, a lot can be proven by ‘cutting’ our spaces into smaller and simpler components. In this section we will explore this process, mostly via analogy to the 2-dimensional case.

We first define the operation under which we may join our spaces back together, as this will provide motivation for some of our later definitions. Much of the material in this section can be found in [13].

Definition 1.11. Given two oriented n -manifolds with $n \leq 3$, M and N we define the **connected sum** $N \# M$ to be the manifold obtained when we remove n -balls B, B' from M and N respectively, and then glue along the boundaries via an orientation reversing diffeomorphism $\partial B \rightarrow \partial B'$.

Remark 1.7. This definition is well-defined, and the proof can be found in [13].

Remark 1.8. This definition can be made in the unoriented case, and for $n \geq 4$, however in those cases the map $S^n \rightarrow S^n$ which we paste along does matter, and therefore we restrict to the cases of interest, in which the choice of map does not matter as they produce manifolds diffeomorphic to each other.

Example 4. If S_g denotes the genus g surface, then $S_g \# S_h = S_{g+h}$.


Example 5. If M is an n -manifold and D is the n -dimensional ball, then $M \# D \cong M$.

Proposition 1.7. If $n \geq 3$, then $\pi_1(M \# N) = \pi_1(M) * \pi_1(N)$.

Proof. Seifert van Kampen. □

We see that cutting along the S^{n-1} along which we glued, and then attaching a ball on both sides of the cut returns us to our original manifolds. We therefore wish to study which S^{n-1} we may cut along for this process. Before moving to a formal definition, we examine this process for surfaces. For the sake of brevity, we will consider S_2 .

Example 6. We cut S_2 along the indicated red line



After some consideration, it is clear we get two manifolds homeomorphic to the tori with holes on the right. If we further cut these tori, say along their top sides, we get a decomposition of S_2 into ‘irreducible’ surfaces, and it is clear how this generalises to any S_g .

This motivates the following definition:

Definition 1.12. The 3-manifold M is **irreducible** if every sphere in its interior bounds a ball.

It is clear that every irreducible manifold is not the connect sum of two manifolds. The converse only has one counterexample: $S^2 \times S^1$, the proof of which we refer to [13].

We will later find it convenient to not need to re-glue a ball back in after cutting, and we will therefore call a manifold with boundary irreducible if it is irreducible after capping off any spherical boundary components.

Example 7. An important case of the latter comment: the 3-sphere with k balls removed is irreducible.

We make one final definition in this section, which will later allow us to define the action of $\text{Out}(F_n)$ on a set of spheres realising part of a decomposition into irreducible manifolds.

Definition 1.13. A **Dehn twist** along a 2-sphere $S \subseteq M$ inside a 3-manifold is given by choosing an axis l cutting out a tubular neighbourhood of S which we denote $S \times [0, 1]$, sending $S \times \{t\}$ to the rotation of S by $2\pi t$ radians around l , and then gluing the neighbourhood back in.

1.5 Surgery Theory

We will describe the process of surgery, and demonstrate that it may be used ‘within’ a manifold to modify submanifolds. This process will be a key part of our definition of the deformation retraction exhibiting the contractibility of outer space, as it will allow us to modify our sphere systems until they are a subset of a fixed sphere system. Much of the detail in this section may also be found in [13].


We make the observation that for a manifold $X \times Y$ the boundary is $\partial(X \times Y) = \partial X \times Y \cup X \times \partial Y$. In particular for $X = D^p, Y = D^q$ we have $\partial(D^p \times D^q) = S^{p-1} \times D^q \cup D^p \times S^{q-1}$. We use this observation to modify a given manifold M .

Definition 1.14. If we have a manifold M of dimension $n = p + q$ with $p < n$ and an embedding $\phi: S^p \times D^q \rightarrow M$, then we may remove the interior of $\text{im}(\phi)$ and glue $D^{p+1} \times D^q$ to M along the first part of its boundary to get a manifold (after smoothing out the corners created) $M' = (M \setminus \text{Int}(\text{Im } \phi)) \cup_{\phi|_{S^p \times S^{q-1}}} D^{p+1} \times S^{q-1}$. We call this process *surgery*.

Remark 1.9. In future sections, we will want to use surgery to modify submanifolds S of a fixed manifold M . In this setting, the process is the same, except that our $D^{p+1} \cup D^q$ is a tubular neighbourhood of an embedding of D^{p+1} with boundary $\phi(S^p \times \{t\})$. In the case that $p+1 = q$, the resulting subspace is (after smoothing corners) a new submanifold of M , and so we have performed the desired surgery within M .

We provide some examples of surgery.

Example 8. If we perform surgery on a S^1 , then we must have $p = 0$ (as $p < 1$), and the result is two copies of S^1 , see the diagram below, where the blue lines are our $S_0 \times D_1$, and the red lines are the other part of the boundary of a $D_1 \times D_1$:




Example 9. We will primarily perform surgery on copies of S^2 in what follows, so we examine what happens in this case. We must have $p < 2$.

We first examine the case of $p = 0$. Then the embedding of $S^0 \times D^2$ is two disjoint disks in S^2 , so the result of removing them and replacing them with a $S^1 \times D^1$ is a torus.

If $p = 1$, then $S^1 \times D^1$ is a cylinder, which we take to be around the equator. Removing this and replacing with a $S^0 \times D^2$, we get two copies of S^2 . It is this latter surgery that will primarily interest us, as it turns spheres into different spheres, disjoint from the equator of the original sphere.

We now provide an example on the torus of performing surgery on submanifolds to obtain new submanifolds:

Example 10. Consider the submanifold of T^2 which is isomorphic to S^1 given by the blue line. We conduct surgery on the S^0 given by the intersections of the red and blue lines. We embed a D^1 between them via the red line, and take the tubular neighbourhood within the green lines. The resulting submanifolds are given on the left, and, after smoothing, are truly submanifolds



1.6 Homotopy Theory and Spectral Sequences

1.6.1 Homotopy Theory

During our proofs of homological stability, we shall need to prove that certain homology groups of spaces are trivial. The way we shall do so is given by homotopy theory, and we shall give the basic

definitions and theorems we shall need here. All results we give can be found in [1]. We assume all spaces under consideration are CW-complexes (in fact, we will only need these results for unions of open simplices of simplicial complexes)

Definition 1.15. The *n th homotopy group* of a space X with basepoint x_0 , denoted $\pi_n(X, x_0)$, is the set of homotopy classes of maps $D_n \rightarrow X$ such that ∂D_n maps to x_0 , with all homotopies relative to ∂D_n . We say a path-connected space is *n -connected* if $\pi_i(X, x_0) = e$ for $0 < i \leq n$.

Example 11. A contractible space is n -connected for every $n \geq 1$. As with fundamental groups, if X is path connected, $\pi_n(X, x_0) \cong \pi_n(X, x_1)$ for all $x_0, x_1 \in X$, but the isomorphism is not canonical.

Theorem 1.3. If X is n -connected, then $H_i(X) = 0$ for $0 < i \leq n$.

Proof. See [1], Proposition 2.4.19. □

Definition 1.16. The *join* of topological spaces X and Y is

$$X * Y = X \times Y \times I / \{(x, y, 0) \sim (x, y', 0), (x, y, 1) \sim (x', y, 1)\}.$$

We will need the following two theorems later to prove connectivity/contractibility results for certain complexes of interest.

Theorem 1.4. If X is n -connected and Y is m -connected, then $X * Y$ is $(p + q + 2)$ -connected.

Proof. See [1], Proposition 6.2.11. □

Theorem 1.5. If a CW complex X is n -connected for every $n \geq 1$, then it is contractible.

Proof. See [1], Proposition 2.4.7. □

Finally, for one of our connectivity arguments, we will need to appeal to a certain long exact sequence in homotopy, for which we give the relevant definition and theorem statement.

Definition 1.17. A map of based spaces $p: E \rightarrow B$ is a **fibration** if for every map of based spaces $f: Z \rightarrow E$ and every homotopy $g: Z \times I \rightarrow B$ such that $g(\cdot, 0) = p \circ f(\cdot)$, there is a homotopy $h: Z \times I \rightarrow E$ such that $p \circ h = g$, and $h(\cdot, 0) = f(\cdot)$. The **fibre** of this fibration is the space $F = p^{-1}(x_0)$, where x_0 is the basepoint of B .

Theorem 1.6. Given a fibration $E \rightarrow B$ with fibre F , there is a long exact sequence of homotopy groups

$$\pi_1(F) \rightarrow \pi_1(E) \rightarrow \pi_1(B) \rightarrow \pi_0(F) \rightarrow \pi_0(E) \rightarrow \pi_0(B).$$

(This sequence continues to the left, but this is all we shall need).

Proof. See [1], Proposition 4.2.19. □

1.6.2 Spectral Sequences

The arguments for homological stability in [7, 9, 10] are largely based on spectral sequences. We give here the bare minimum to understand these arguments, however we strongly recommend that the interested reader read Chapter VII of [2], from which most of the content in this subsection derives.

Definition 1.18. A (*homological*) **spectral sequence** is a sequence $E_{\bullet, \bullet}^r$ of bigraded abelian groups for $r \geq 1$ with maps $d_{p, q}^r: E_{p, q}^r \rightarrow E_{p-r, q+r-1}^r$ such that $d^r \circ d^r = 0$ and $E_{p, q}^{r+1} = \ker(d_{p, q}^r) / \text{im}(d_{p+r, q-r+1}^r)$.

We care about these for the following reason. Suppose we have a double complex $C_{\bullet,\bullet}$ (ie, a chain complex of chain complexes such that the vertical and horizontal differentials commute). We can obtain from this the **total chain complex** $T(C)_{\bullet}$ with $T(C)_n = \bigoplus_{p+q=n} C_{p,q}$ with differential given by $d^0 + (-1)^p d^1$ where d^0 is the differential of $C_{p,q}$ in the first coordinate, and d^1 in the second.

We may filter $T(C)$ in two ways, with $F_p(T(C)_n) = \bigoplus_{i \leq p} C_{i,n-i}$ and with $G_p(T(C)_n) = \bigoplus_{i \leq p} C_{n-i,i}$. In these case we get a spectral sequence by setting $E_{p,q}^1 = H_q(C_{p,\bullet})$ (resp. $E_{p,q}^1 = H_q(C_{\bullet,p})$). The differential of $E_{\bullet,\bullet}^1$ in these cases are the maps induced by the differential $C_{\bullet,p} \rightarrow C_{\bullet,p-1}$ (resp. $C_{p,\bullet} \rightarrow C_{p-1,\bullet}$), and so $E_{\bullet,\bullet}^2$ can be described as taking the homology in the first coordinate, and then the second (resp. second coordinate, and then the first).

It is a fact (see [2], Chapter VII) that $E_{p,q}^r$ ‘converges to’ (in some sense that will be clear in the cases we care about, as in those cases we will have that d^r is eventually 0, and hence $E_{p,q}^r = E_{p,q}^{r+k}$ for all $k \geq 0$). We define this limit to be $E_{p,q}^\infty$ $F_p H_{p+q}(TC)$ (resp. $G_p H_{p+q}(TC)$) and hence $\bigoplus_{p+q=n} E_{p,q}^\infty$ is the same in both cases. This is the main fact about spectral sequences that we will use.

In the case where our double complex is the tensor product of a $\mathbb{Z}G$ resolution of $\mathbb{Z}$ and the chain complex related to a simplicial complex X , we may compute via this technique (see [2], Proposition VII.5.6) that $E_{p,q}^2 = H_p(G, H_q(X))$, and that $E_{p,q}^\infty = F_p H_{p,q}(X/G)$. Most of our arguments (phrased in the original sources as proving the local coefficient system is trivial, or untwisted) consist of proving that we may apply this proposition. Intuitively, this consists of proving that all the induced maps $C_p(X) \rightarrow C_{p-1}(X)$ are ‘compatible’ isomorphisms.

2 Culler-Vogtmann Outer Space and the Sphere Complex

2.1 Culler-Vogtmann Outer Space

We fix an n , and following [4] we wish to construct a nice space on which $\text{Out}(F_n)$ acts. We will do this by constructing a space on which $\text{Aut}(F_n)$ acts, where $\text{Inn}(F_n)$ acts trivially.

Definition 2.1. A *(n -petal) rose* is the topological space $R_n = \bigvee_{i=1}^n S^1$.


Remark 2.1. The n -petal rose will be used as our ‘prototypical’ graph with fundamental group F_n to allow us to ‘label’ the edges of any other such graph.

Definition 2.2. A **marking** on a metric graph G with fundamental group $\pi_1(G) = F_n$ is a homotopy equivalence $R_n \rightarrow G$. A metric graph G with a marking $g: R_n \rightarrow G$ is called a **marked graph**, and two marked graphs (G_1, g_1) and (G_2, g_2) are **equivalent** if there is an isometry $i: G_1 \rightarrow G_2$ such that

$$\begin{array}{ccc} & G_1 & \\ g_1 \nearrow & & \downarrow i \\ R_n & & G_2 \\ g_2 \searrow & & \end{array}$$

commutes.


Example 12. Consider the following graph, where every edge has length 1:




By contracting the middle edge to a point, we see this is homotopy equivalent to R_2 , where we label the generators of $\pi_1(R_2)$ as shown:



We choose a marking given by sending the a loop around the middle then left edges, and the b loop around the middle then right edges. We may represent this marking by labelling the non-collapsed edges with where they are mapped to under the homotopy inverse to the marking given by contracting the middle edge:



Note that we could choose a different homotopy inverse (for example, contracting the right edge) and we would get a different (but equivalent) labelling:



Remark 2.2. Similarly to the above example, we can get infinitely many different (but equivalent) graphs. An important case is that we can conjugate the labelling by $w \in F_n$ (by choosing the homotopy inverse sending an edge labelled w' to $ww'w^{-1}$).

We are now in a position to define our outer space:

Definition 2.3. We define the **Culler-Vogtmann outer space** for F_n , denoted X_n , to be the space of all equivalence classes of connected marked graphs with fundamental group F_n , volume 1 and no 2 or 1-valent vertices.

We wish to topologise the space above. As we have a metric on our marked graphs, and our equivalence relation is defined using isometries between them, a natural choice of topology is induced on X_n by considering the lengths of words in F_n defined as follows:

Definition 2.4. For a marked graph (G, g) we define the associated length function $l_G: F_n \rightarrow \mathbb{R}$ by sending $w \in F_n$ to the length of the associated path in G .

Remark 2.3. If marked graphs (G, g) and (H, h) are equivalent via an isometry $i: G \rightarrow H$, then $l_G = l_H$. This is because if g' is a homotopy inverse to g , then $i \circ g \simeq h$, so $i \simeq h \circ g'$, but both h and g' send the representative of the word w in R_0 and G to a representative of the word in H and R_0 respectively, and hence $l_G(w) = l_H(w)$. We therefore get a map $l: X_n \rightarrow \mathbb{P}(F_n)$ (the projective space on functions from F_n to $\mathbb{R}$) via $(G, g) \mapsto (l(w))_{w \in F_n}$, and we may pull back the topology, allowing us to topologise X_n .

Remark 2.4. X_n topologised as such is a disjoint union of open simplices, where each simplex contains points which are isomorphic as graphs (but not necessarily as metric graphs), where we just vary the metric while retaining the same graph structure. Each face is given by setting the length of a given set of edges equal to 0 (and hence removing it). The resulting space is not a simplicial complex as upon removing edges, we may find that the graph is either not connected, or no longer has fundamental group F_n , and is hence not in X_n .

We define an action of $\text{Aut}(F_n)$ on X_n as follows. The automorphism $\phi: \text{Aut}(F_n) \rightarrow \text{Aut}(F_n)$ is the map induced on the fundamental group by the map $\psi: R_n \rightarrow R_n$ sending the edge e of R_n labelled by a generator x_i of F_n to the path labelled by $\phi(x_i)$. Given $\phi \in \text{Aut}(F_n)$, we can therefore send the marked graph (G, g) to a different marked graph $(G, g \circ \psi)$. This respects equivalence of marked graphs as that is defined by post-composition, and our action is defined by precomposition, so they commute. This descends to an action of $\text{Out}(F_n)$ on X_n , as by Remark 2.2, conjugation by an element of F_n sends every marked graph to an equivalent marked graph, and hence $\text{Inn}(F_n)$ acts trivially on F_n .

2.2 The Sphere Complex

We now describe a different construction of a related topological space, first defined in [7], which we will use to prove that X_n is contractible. We wish to use theory from low dimensional topology to prove this, and so we choose a different model of space with fundamental group F_n .

We define the manifold $M_{n,s}$ to be the n -fold connected sum of $S^2 \times S^1$ with n punctures. If we consider the mapping class group of $M_{n,1}$, we get a map from $\text{Mod}(M_{n,1}) \rightarrow \text{Aut}(F_n)$ via the action on the fundamental group of $M_{n,0}$ based at the puncture (as the puncture must be fixed by any representative of an element of the mapping class group, and hence the basepoint is fixed). As the basepoint need not be fixed for $\text{Mod}(M_{n,0})$, by choosing fixed paths from some basepoint $x_0 \in M_{s,0}$ to x for all $x \in M_{s,0}$, we get an action on the fundamental group which is defined up to conjugacy, and hence a map $\text{Mod}(M_{n,0}) \rightarrow \text{Out}(F_n)$.

The space we wish to define is motivated by the cut and paste constructions of Section 1.4:

Definition 2.5. We define a **system of 2-spheres** in $M_{n,s}$ to be a collection of disjointly embedded smooth 2-spheres in $M_{n,s}$ such that no sphere bounds a 3-ball or is isotopic to a boundary component of $M_{n,s}$, and no two spheres are isotopic to each other. We define the **sphere complex**, denoted $\mathbb{S}(M_{n,s})$, to be the simplicial complex with k -simplices given by isotopy classes of systems of 2-spheres of size $k + 1$, with faces given by taking subsystems (up to isotopy).

Remark 2.5. A system of spheres is exactly the data of a partial decomposition of $M_{n,s}$ into manifolds, and example 6 can be easily turned into (the equivalent one dimension down of) a system of spheres.

Remark 2.6. To show this is a simplicial complex, we need to know that systems are determined by the isotopy classes of their individual spheres. This is a theorem of Laudenbach [12] that states that homotopic systems of spheres are isotopic, as we can easily find a homotopy between any two sphere systems which are sphere-wise isotopic, namely by isotoping isotopic spheres to each other, not worrying about intersections of different spheres.

Remark 2.7. Both of the above described maps $\text{Mod}(M_{n,1}) \rightarrow \text{Aut}(F_n)$ and $\text{Mod}(M_{n,0}) \rightarrow \text{Out}(F_n)$ are surjective as all of the Nielsen generators of $\text{Aut}(F_n)$ can be realised as orientation preserving diffeomorphisms of $M_{n,0}$ and of $M_{n,1}$, and by Laudenbach [12], the kernels are generated by Dehn twists. Dehn twists act trivially on the sphere complex, as proven in 3.3. Hence we get a left action of $\text{Aut}(F_n)$ (resp. $\text{Out}(F_n)$) on $\mathbb{S}(M_{n,1})$ (resp. $\mathbb{S}(M_{n,0})$) as the mapping class group gives a well-defined map between sphere systems. For consistency with the action on X_n , we will use the corresponding right action (where $S \cdot \phi = \phi^{-1}(S)$) for the rest of this essay.

Remark 2.8. The maximal simplices of the sphere complex all have dimension $3n + s - 4$. We may see this in the following way. The graph with vertices corresponding to $M_{n,s} \setminus S$ and one for each puncture, and edges between adjacent components has fundamental group F_n (as it is a deformation retract of $M_{n,0}$ via capping off the punctures), is trivalent (except at boundary components) as the components of $M_{n,s} \setminus S$ are 3-punctured spheres, and so if we have $v + s$ vertices, then we have $n + v + s - 1$ edges. As

$3v + s = 2n + 2v + 2s - 2$ (because each of the non-boundary vertices are trivalent), we have $3n + 2s - 3$ edges and hence we have $3n + s - 3$ separating spheres, as the edges to boundary components don't contribute a sphere. Therefore the dimension of maximal simplices is $3n + s - 4$.

We will prove a link, originally proven in [7], between the sphere complex for $M_{n,0}$ and Culler and Vogtmann's construction of outer space in section 5.

3 Normal Form

3.1 Definition and Existence of Normal Forms


We now consider a maximal sphere system, that is, a sphere system which is contained in no other such system. Throughout this section we such fix a system which we denote by $\Sigma = \bigcup_j \Sigma_j$, where Σ_j are the disjoint spheres of the system. By cutting along this system, we produce a collection of 3-punctured spheres which we call P_k .

Definition 3.1. A sphere system $S = \bigcup_i S_i$ is said to be in **normal form** with respect to Σ if for each S_j , either S_i is isotopic to some Σ_j , or Σ meets S_j transversely in a (necessarily non-empty, by maximality) collection of circles splitting S_i into **pieces** such that the following hold for each punctured sphere P_k :

1. Each piece of S_j in P_k meets each boundary component of P_k in at most one circle;
2. No piece in P_k is a disk isotopic, fixing its boundary, to a disk in ∂P_k .

Remark 3.1. As P_k is a 3-punctured sphere, it follows from (1) that each piece is a sphere, cut in at most 3 disjoint circles, and hence is a disk, cylinder or pair of pants (corresponding respectively to having a boundary on one, two or three components of the boundary of P_k). In the first case, the disk must separate the other two boundary components, by (2).

We illustrate the first two cases of this remark with the situation in 2 dimensions (with the last not being possible in dimension 2). If we call the boundary components of a 3-punctured 2-sphere B_1, B_2, B_3 , then by cutting along disjoint paths from B_1 to B_2 and B_1 to B_3 , we get that the 3-punctured 2-sphere is a decagon with identifications shown below. The two labelled paths are examples of the cases described.



The coloured edges denote those edges which respectively make up B_1, B_2, B_3 .

Condition 1 of this normal form allows us to perform surgery on Σ with respect to S , which will later be vital in our construction of a contraction of $\mathbb{S}(M_{n,s})$. Condition 2 states that every intersection is ‘fundamental’, in that it provides meaningful information about the topology of our embedding.

The following proofs are due to Hatcher in [7].

Proposition 3.1. *Every system $S \subseteq M$ can be isotoped into normal form with respect to Σ .*

Proof. We may assume that S has been isotoped so that a maximal number of S_j correspond with some Σ_i , and so that every intersection between S and Σ is transverse.

We may ensure that condition (2) is satisfied by isotoping any such pieces (starting with the inner-most) through the Σ_j which they intersect.

If S is now not in normal form, we aim to decrease the number of circle components of this intersection. We may conduct surgery along spheres S_i which meet Σ along the circles where they meet (starting with the inner-most ones) to produce a collection of disjoint spheres S_{im} which do not intersect Σ . The inverse operation of this involves cutting holes in these S_{im} and attaching tubes $S_1 \times I$, so now all our intersections with Σ are contained within these tubes.

If S is not in normal form, some S_{im} has two tubes leading from it which pass through the same component of ∂P_k . Denote these by T_1 and T_2 , and we may assume T_1 is not contained in any larger tube. We pull T_1 (and any tubes contained within it) along the surface of S_{im} until it intersects at the same circle as T_2 . We then pull T_1 (and any tubes inside it) down along T_2 until it passes through Σ again. We now have two intersections of T_1 with Σ (its original intersection and the extra intersection made to go through Σ so as to create the intersection at T_2) and 3 for any tube contained within it (the same two as T_1 , and the intersection that lies within T_2). The first and last of these, however, can be isotoped through Σ as they form a cylinder with both boundaries on the same component of ∂P_k , and so in total we have reduced the number of intersections by 1 (as T_1 and T_2 share an intersection).

By induction, we can therefore put S into normal form. $\square$

3.2 Equivalence of Normal Form Systems

Definition 3.2. *Two sphere systems S, S' in normal form with respect to Σ are said to be **equivalent** if there is a homotopy $h_t: S \rightarrow M$ from S to S' such that:*

1. h_t is the identity on S_j if $S_j = \Sigma_i$ for some i ;
2. On all other S_j , $h_t(S_j)$ is always transverse to Σ , with the intersection varying via isotopy, and in particular remain disjoint.

Remark 3.2. Systems isotopic through normal form systems are equivalent, however equivalence also allows us to unknot pieces within each P_k . It is clear that equivalent systems are isotopic, as they are homotopic. We prove that if the systems are in normal form, this is in fact an equivalence.

Proposition 3.2. *Isotopic sphere systems in normal form are equivalent.*

Proof. We look at the universal cover, where oriented spheres give well-defined elements of $\pi_2(\tilde{M})$, which is isomorphic to $H_2(\tilde{M})$ by Hurewicz (as $\pi_1(\tilde{M}) = 0$).

We take the preimage $\tilde{\Sigma}$ of Σ in $\tilde{M}$, and associate to this a graph, T , with vertices corresponding to lifts of P_k as k ranges over every punctured sphere, and to the lifts of the boundary spheres, with edges between two vertices if one is a boundary sphere of the other.

Since $\tilde{M}$ is simply connected, each component σ of $\tilde{\Sigma}$ separates M (as otherwise we could form a loop intersecting σ exactly once, and the intersection form yields that such a loop is non-trivial in homology, and hence in the fundamental group), and hence T is a tree.

Now suppose S is a sphere in normal form with Σ and choose a lift $\tilde{S}$ to $\tilde{M}$. For each $\tilde{P}_k$, $\tilde{S} \cap \tilde{P}_k$ has at most one component. This is trivial if $\tilde{S}$ is a component of $\tilde{\Sigma}$. Otherwise, we can form a graph similar to T , but using $\tilde{S}$ instead of all of $\tilde{M}$, which we call $T(S)$. The inclusion $\tilde{S} \rightarrow \tilde{M}$ induces a map $T(S) \rightarrow T$, which is locally injective by normal form, and hence globally injective. Therefore $T(S)$ is a tree and so $\tilde{S} \cap \tilde{P}_k$ has at most one component.

We may now view $T(S)$ as a subtree of T , and it is convenient to enlarge it to contain all the vertices corresponding to any $\tilde{P}_k$ (and its boundary spheres) which meets $\tilde{S}$. A disk piece of $\tilde{S}$ corresponds to a part of the tree of the following form



where dots correspond to the attachment to the rest of the tree. A cylinder corresponds to a



and a pair of pants component to



In particular, $T(S)$ can't consist of only a single instance of one of these. For completeness, if $\tilde{S}$ corresponds to a component of $\tilde{\Sigma}$, we use $T(S)$ to denote the corresponding vertex.

We choose a transverse orientation for $\tilde{S}$. If $T(S)$ is not a point, each end vertex of $T(S)$ can be labelled as to whether it lies on the positive or negative side of $\tilde{S}$. By the same argument as before, $\tilde{S}$ separates $\tilde{M}$, so these sides are well-defined. For terminal vertices in a component of form 3.1, by condition (2) for normal form, the two vertices have opposite sign. The sign in 3.2 can be either sign.

Let $\tilde{S}_+$ be the union of spheres with positive sign, and $\tilde{S}_-$ those with negative sign. We get induced orientations of these so that $\tilde{S}$ is homologous to each of them separately. As each extremal Y is of the form 3.1 having both a $+$ and a $-$ sign, $\tilde{S}_+$ determines $T(S)$ (and identically so does $\tilde{S}_-$), by taking the smallest connected subset of T containing all the S_+ , such that if the vertex corresponding to $\tilde{P}_k$ is included, so is all of its boundary spheres. In this way, $\tilde{S}_+$ also determines $\tilde{S}_-$.

The homology class of $\tilde{S}$ determines $\tilde{S}_+$ uniquely. Suppose that $\tilde{S}$ is homologous to $\tilde{S}'$ (as happens when they are isotopic through normal forms, by lifting the isotopy). If $T(S) \neq T(S')$, then there is a component of $\tilde{\Sigma}$ separating $\tilde{M}$ into M_0 , which (without loss of generality) contains all but a single vertex of $T(S)$, and contains all of $T(S')$. Then $\tilde{S}$ and $\tilde{S}'$ are different in $H_2(\tilde{M}, \tilde{M}_0)$, contradiction so $T(S) = T(S')$.

If the choice of sign S and S' induce on some component of $\tilde{\Sigma}$ are different, again by separating off such a component, we get a contradiction in relative homology.

We may now deduce that $\tilde{S}, \tilde{S}'$ are isotopic transverse to $\tilde{\Sigma}$. As the trees $T(S), T(S')$ are equal, the combinatorial data of the pattern of pieces are the same for each, so there is an isotopy of $(\tilde{M}, \tilde{\Sigma})$ mapping the intersections with $\tilde{S}, \tilde{S}'$ onto each other, and we can hence isotope $\tilde{S}, \tilde{S}'$ onto each other, as there is one piece in each P_k (and hence no knotting issues), and the signs of the boundaries (along with which boundaries $\tilde{S}, \tilde{S}'$ intersect) uniquely identify them up to isotopy.

We now consider isotopic normal form systems S, S' , and we choose transverse orientations on them that are equal under the isotopy. We have a bijective correspondence between pieces of S and of S' by the preceding arguments, and we assert that if s_p, s_q are pieces of S in P_k and s'_p, s'_q the corresponding pieces of S' , then s'_p is on the same side relative to s'_q as s_p is relative to s_q . To see this, we consider a lift $\tilde{P}_k$ of P_k to $\tilde{M}$. Denoting the lifts to this piece of everything we have considered using tildes, we may isotope $\tilde{S}_i$ to $\tilde{S}'_i$ containing $\tilde{s}_p, \tilde{s}'_p$ as we have argued for the single sphere case, and hence, as the isotopy is an isotopy of the pair $(\tilde{M}, \tilde{\Sigma})$, we may assume they are equal. If the assertion were false, then $\tilde{S}_j, \tilde{S}'_j$ containing $\tilde{s}_q, \tilde{s}'_q$ would be on opposite sides of $\tilde{S}_i$. Now as we have argued before, the trees of S_j, S'_j are equal and lie on opposite sides of $\tilde{S}_i$, and hence $T(S_i) = T(S_j) = T(S'_j)$, as S_i lies in between S_j and S'_j , and therefore must pass through the same components in the same manner, so $S_i = S_j$. We therefore have a deck transformation taking $\tilde{S}_i$ to $\tilde{S}_j$ which preserves $T(S_i)$, but this is impossible as a deck transformation can't fix a finite subtree by considering its barycentre.

We therefore have an orientation preserving bijection between circles of $\Sigma \cap S$ and those of $\Sigma \cap S'$, so we may isotope these to each other, and then as before we may isotope within P_k to get the desired equivalence homotopy. $\square$

Remark 3.3. We prove Dehn twists act trivially on the sphere complex. Take a sphere system S . If the sphere Σ we are twisting around is in S , then trivially the Dehn twist preserve S . Otherwise, by the above proposition, we may assume our sphere system intersects Σ transversely in tubes. It is then clear that S is homotopic to its Dehn twist, as the induced knotting does not matter for homotopy, and so it is isotopic to its Dehn twist by Laudenbach [12].

4 Contractibility of Sphere Complex

We now wish to prove that the sphere complex is contractible, so we can get cohomological information about $\text{Out}(F_n)$ from it, as in Section 1.2. The following proof is due to [7].

Theorem 4.1. $\mathbb{S}(M_{n,s})$ is contractible if $n > 0$.

Proof. We prove only the case of primary interest to us, namely $s \leq 1$. The result is true for $s > 1$, however the proof is similar, and not any more enlightening. In the case $s \leq 1$, we will construct a piecewise linear deformation retraction onto the maximal simplex corresponding to a maximal system Σ . We consider a sphere system S , and by Proposition 3.1, we may assume that S is in normal form with respect to Σ . We take a point with coordinates (t_i) in the simplex corresponding to S , and we will construct the path from this point to the simplex corresponding to Σ .

We think of the number t_i as the ‘weight’ corresponding to S_i , and it will be convenient to thicken S_i (while keeping all the S_i disjoint), and think of the width of this slice as being t_i . If the weight is 0, we remove the sphere from the simplex. Of course, we may also split this t_i into different slices of positive width summing to t_i , and we still consider this system as representing the same point.

For each sphere Σ_i of Σ which S meets transversely, we form a metric tree T_i similarly to the proof of Proposition 3.2, with vertices the components of $\Sigma_i \setminus S$ and edges between neighbouring vertices. We metrize the edge $e \subseteq \Sigma_i \cap S_j$ by giving it length t_j . This is a finite tree, so there is a canonical deformation retraction to a point by iteratively shrinking the extremal edges to 0 length at constant

speed. We aim to surger S so that the deformation induced on the trees T_i is this deformation. We will therefore obtain a sphere system disjoint from Σ , which must be in the simplex given by Σ as Σ is maximal.

We consider the extremal vertices in T_j . These correspond to the disks $D_v \subseteq \Sigma_j$ such that $D_v \cap S = \partial D_v$ (ie, those disks bounded by intersections of Σ_j with S which contain no other intersection). We can therefore use these disks to perform surgery on S , surgering through the relevant widened S_i (of width t_j) at constant speed, moving the weight to the new families constructed by the surgery. These families may be taken to be disjoint to S , so we may view the surgery in this way. We delete families when they disappear as described above. There is no problem if surgery is being conducted from both sides of S_j simultaneously - the width just decreases twice as fast.

We consider a single component P_k of $M_{n,s} \setminus \Sigma$ and consider what this surgery does in P_k . Surgery on a disk piece produces a sphere isotopic to a component of ∂P_k (and hence to some Σ_i or a puncture by definition). Surgery on one side of a cylinder C produces a disk piece D . Any sphere of S intersecting the other end of C must be contained in $P_k \setminus C'$, where C' is C surgered along D . Therefore C' either separates two boundary components, or there is no sphere on one side of C' , and hence we are performing surgery on both sides, and therefore get a sphere contained in P_k , hence bounding a ball. If we form a pair of pants, surgery on one end produces a cylinder, and we are in the case described above. We can remove any trivial spheres bounding a ball or puncture, as they are not involved in further surgery, and therefore the result of surgery is still normal with respect to Σ . We have at least two disk pieces (ie, vertices which are extremal in the graph with vertices the components of $S_j \setminus \Sigma$), and as we only have one puncture, the result is a non-trivial system as only one of these two pieces can be homotopic to a puncture, and hence be removed.

As the paths $S(t)$ corresponding to S will always conduct the same surgeries depending only on the spheres in existence at any given time, and will conduct those surgeries for a time given as a piecewise linear function of the weights, the paths are continuous, and even piecewise linear, within a simplex (including at the faces), and so we are done for $s \leq 1$. $\square$

5 Relation to Culler-Vogtmann Outer Space

We earlier provided two constructions of an outer space. We now prove that these constructions are closely related following [7], before moving on next section to prove some cohomological finiteness facts about $\text{Out}(F_n)$ which follow from the existence of these spaces.

It is clear, from the fact that $\mathbb{S} = \mathbb{S}(M_{n,0})$ is a simplicial complex and X_n is not, that we will need to remove some points from the former. It turns out that our deformation retraction constructed in the previous section avoids these points, as long as we don't start at them, and therefore we will have provided a proof that X_n is contractible.

The proof of this will rely heavily on a dual graph, $G(S)$, constructed similarly to Proposition 3.2. More explicitly, $G(S)$ has vertices the components of $M_{n,0} \setminus S$ with edges between neighbouring components. $G(S)$ will be the point of outer space corresponding to S , and we will need to determine which $G(S)$ do not lie in X_n . We may embed $G(S)$ into $M_{n,0}$ by choosing a vertex in each component, and sending edges to paths between the corresponding vertices. This can also be seen as a quotient of $M_{n,0}$, by thickening every sphere of S to $S \times [0, 1]$, then projecting onto the second component, and sending every point of $M_{n,0} \setminus (S \times [0, 1])$ to the vertex in the same component.

We know that the fundamental group of $M_{n,0}$ is F_n , and that for $G(S)$ to be in X_n , its fundamental group must be F_n . Further, by Seifert van Kampen, the fundamental group of $M_{n,0}$ is the fundamental group of the graph of groups corresponding to cutting along S . This graph of groups has $G_e = 1$ for every edge e , and hence by [15] I.5.1 Example 1, the fundamental group of $G(S)$ is a free factor of $M_{n,0}$.

We hence need the other factor, which is given by the free product of the fundamental groups of the components of $M_{n,0} \setminus S$, to be trivial. Hence the points we wish to remove are those with some non-simply-connected complementary component. We denote this set $\mathbb{S}_\infty$. We now need to define a marking on $G(S)$. If $S \in \mathbb{S} \setminus \mathbb{S}_\infty$, then we know $G(S) \rightarrow M_{n,0}$ is an isomorphism on fundamental groups, and so is $M_{n,0} \rightarrow G(S)$ as we defined above. We therefore fix $S_0 \in \mathbb{S} \setminus \mathbb{S}_\infty$. Then $G(S_0) \rightarrow M_{n,0} \rightarrow G(S)$ is a homotopy equivalence, and hence is a marking. The map $\phi: \mathbb{S} \setminus \mathbb{S}_\infty \rightarrow X_n$ given by $S \mapsto (G(S_0) \rightarrow G(S))$ is linear on open simplices and continuous on faces, and hence continuous everywhere. It is also $\text{Out}(F_n)$ -equivariant (noting that $\text{Out}(F_n)$ acts on $\mathbb{S}$ via pullback). We wish to prove the following theorem.

Theorem 5.1. $\phi: \mathbb{S} \setminus \mathbb{S}_\infty \rightarrow X_n$ is a homeomorphism.

Proof. We construct an inverse. Let $f: R_n \rightarrow X$ be a point of X_n . We build a manifold $M(X) \cong M_{n,0}$ by taking a k -punctured sphere for each k -valence vertex of X , relating each puncture to an edge, and gluing the two punctures which correspond to a given edge.

We have a natural sphere system $S(X)$ given by the spheres along which we glued, weighted by the lengths of the corresponding edges.

The corresponding graph is easily seen to be isometric to X , and if we choose a diffeomorphism $M_{n,0} \cong M(X)$, then the pullback of $S(X)$ is a weighted system S with $\Phi(S) = (X, g)$ for some homotopy equivalence g .

Since every homotopy equivalence $R_0 \rightarrow R_0$ is realised by a diffeomorphism of $M_{s,0}$, we may rechoose our diffeomorphism so $g = f$.

The isotopy class of S doesn't depend on the choice of diffeomorphism, as every choice induces the same isomorphism on fundamental groups, and hence by [12] is a composition of Dehn twists, which act trivially on $\mathbb{S}$.

Any isometry $X \rightarrow X'$ gives a diffeomorphism $M(X) \rightarrow M(X')$ which takes $S(X)$ to $S(X')$, and hence our inverse is well-defined, and clearly an inverse to ϕ . $\square$

Remark 5.1. The deformation retraction of $\mathbb{S}$ that we constructed maps $\mathbb{S} \setminus \mathbb{S}_\infty$ to itself, as each surgery removes a disk from one complementary component of $S \in \mathbb{S} \setminus \mathbb{S}_\infty$ and adds a 2-handle to another, and the only other operations we need to do are adding or merging parallel spheres, and throwing away trivial ones. All of these operations preserve simple connectivity, and so we are done.

We now describe a nice subcomplex, both in $\mathbb{S}$ and in X_n , which has dimension $2n - 3$, and on which $\text{Out}(F_n)$ acts with finite stabilisers and finite quotient. This will complete our proof of Theorem 1.2 as promised. The description in X_n comes from [4] and is as follows.

Proof. We define the subcomplex Y_n of X_n to be the subcomplex of graphs with no separating edges. We denote the simplicial complex containing X_n (ie, the complex of all connected marked graphs with no 1 or 2 valent vertices with fundamental group of rank at most n) by X_n^* , and we denote the subcomplex containing Y_n by Y_n^* .

We take the barycentric subdivision of Y_n^* and denote it Y_n' , and we let K_n be the full subcomplex of Y_n' which is disjoint from ∂Y_n , so that K_n is the full subcomplex of Y_n' with a vertex for every simplex of Y_n . We can define an $\text{Out}(F_n)$ -equivariant deformation retraction of Y_n' onto K_n by sending each face of Y_n' onto the face it has in K_n (note that K_n is full so this is well-defined).

A vertex of K_n is a face of Y_n , so to determine the dimension of K_n it suffices to determine the minimal dimension of a face of Y_n . This is easily seen to be $n - 1$, so that the dimension of K is $3n - 4 - (n - 1) = 2n - 3$.

As K_n is homotopy equivalent to Y_n , and Y_n is homotopy equivalent to X_n (as separating edges can be uniformly reduced in length to 0, providing a deformation retraction), we see that K_n is contractible.

We note that the stabiliser of any point of K is a subgroup of the isometry group of any graph represented by that point, and hence is finite, and that the quotient is finite (as any two bases of F_n can be mapped to each other by an automorphism, and so any two points with the same underlying combinatorial graph are in the same $\text{Out}(F_n)$ orbit). We have therefore completed the proof of Theorem 1.2 in the X_n construction.

We wish to translate the above discussion into the sphere complex now. Our replacement for K_n shall be $\mathbb{S}_0$ with vertices the systems such that $M \setminus S$ is simply connected, and with k -simplices those k -simplices of the barycentric subdivision of $\mathbb{S}_\infty$.

It is immediately clear from our discussions above that this is the subcomplex of $\mathbb{S}$ corresponding to K_n , and the same argument gives the same dimension. The graph theoretic argument that we should have finite stabilisers and finite quotient becomes the geometric topology fact that our sphere complex is determined by how many punctures each complementary sphere has (which corresponds to the valence of a vertex in the graph), and how those punctures are attached, and any automorphisms of the system are (up to isotopy) given by choosing a different gluing in this manner. $\square$

Remark 5.2. The above argument, in the sphere complex case, also works for $M_{n,1}$, where we can deduce, in an identical manner, that $\text{vcd}(\text{Aut}_n(F_n)) = 2n - 2$.

6 Outline of Homological Stability

In [7], after proving the contractibility of the sphere complex (and hence of outer space), the main aim of the paper is to gain more information about the relationships of certain groups' homology, including as previously outlined, $\text{Out}(F_n)$ and $\text{Aut}(F_n)$. We shall explain the broad proof strategy, as it is not clear, being split across 4 papers, before giving greater details in the next few sections.

We state all theorems for general $M_{n,s}$ as $s = 2$ is needed in the proof, making it simpler to state for all s , and indicate where the corrections from [9, 10, 11] are needed as appropriate. For the main theorem, [9] gives a stronger lower bound than [7], and so we shall state it as given there. We first define several relevant maps, as defined in [7, 9, 10] (we use the naming system of [9, 10], which it is important to note is not the same as [7!]). We denote $\Gamma_{n,s} = \text{Mod}(M_{n,s})/T$, where T is the normal subgroup generated by Dehn twists. As we have discussed, for $s = 0, 1$ this is just $\text{Out}(F_n)$ (resp. $\text{Aut}(F_n)$).

Definition 6.1.

- $\beta: \Gamma_{n,s+2} \rightarrow \Gamma_{n+1,s}$ where $s \geq 0$ is defined as the map induced by gluing a tube between two punctures of $M_{n,s}$.
- $\mu: \Gamma_{n,s} \rightarrow \Gamma_{n,s+1}$ where $s \geq 1$ is defined as the map induced by gluing a 3-punctured 3-sphere to one of the punctures of $M_{n,s}$.
- $\alpha: \Gamma_{n,s} \rightarrow \Gamma_{n+1,s}$ where $s \geq 1$ is defined by $\beta\mu^2$.
- $\gamma: \Gamma_{n,1} \rightarrow \Gamma_{n,0}$ is given by capping off the boundary component.
- $\eta: \Gamma_{n,s} \rightarrow \Gamma_{n+1,s-1}$ where $s \geq 2$ is defined as the map induced by gluing a 3-punctured 3-sphere to two of the punctures of $M_{n,s}$.

Theorem 6.1. α, β, μ are isomorphisms from $H_i(\Gamma_{n,s})$ to their respective images for $n \geq 2i + 2$ if $s > 0$. For β , if $s = 0$ then it is an isomorphism for $n \geq 2i + 3$. γ is an isomorphism for $n \geq 2i + 4$.

We briefly indicate the path that the proof of this theorem takes, as with it left across 4 different papers, it is hard to trace the critical path, and which theorems are not proven correctly, in the literature. The proofs and definitions given in Sections 2, 3 and 4 of this essay cover the first two sections of [7], and so we shall pick up from there.

The next steps involve defining a further nice subcomplex of $\mathbb{S}(M_{n,s})$, called Y . Y is closely related to the concept of a sphere basis, consisting of those systems which are either sphere bases, or may be extended to a sphere basis.

It is then proven that Y is $n - 2$ connected which is a generalisation of simply connected to higher homotopy groups. We may thus deduce the homotopy type of Y . The proof of this given in [7] is wrong, and a correct proof is given as Proposition 3.2 of [11].

An ordered version of Y is then defined, and it is then proven that $X/\Gamma_{n,s}$ is contractible, and that the stabilisers of the action are $\Gamma_{n-k,s+2k}$ for some $t < s$. This brings us to the end of section 3 of [7] (having omitted the information about the subcomplex Z defined in [7], as it is not used in any correct argument for the theorem).

Section 4 of [7] is replaced in its entirety by [9], and so we shall now follow that paper.

We define another complex, $W_{n,s}$ with k simplices being the systems of size $k + 1$, together with an enveloping sphere separating them from the rest of the sphere. The same sorts of connectivity and stabilisation information about W is proven, and then it is proven that α induces an isomorphism on homology through an argument based on a spectral sequence related to W .

There is then an error in proving β is an isomorphism on homology for $s \geq 1$, which is corrected in [10], so we will temporarily defer the discussion of this proof. It is then proven that β is an isomorphism on homology for $s = 0$. This is proven in a very similar way to the theorem for α , but this time using X , however in absence of the enveloping spheres that we have in W , we must choose them arbitrarily, and then prove the choice does not matter.

We now return to the proof for $s \geq 1$. Now following [10], we wish to prove that μ induces an isomorphism. It is easy to show that μ is injective, and has the same stable range as α , so it remains to show that μ is eventually surjective on homology. This is the purpose for which [10] introduces η .

As before, we introduce another complex $U_{n,s}$ and wish to prove properties about its connectivity and stabilisers, finishing with a spectral sequence argument.

We have $\alpha = \mu\eta$, and the proof proceeds by proving that $\eta\mu$ is an isomorphism which commutes with α , so we can pass to the direct limit and consider $U_{\infty,s} = \varinjlim U_{n,s}$. The next step of the proof is to consider the action of $\Gamma_{\infty,s} = \varinjlim \Gamma_{n,s}$ on $U_{\infty,s}$. From the results on $W_{n,s}$, we may deduce the quotient is contractible, and another spectral sequence argument using $U_{\infty,s}$ proves $\mu\eta$, and hence μ induces an isomorphism on homology.

As μ and α are isomorphisms, we deduce that β is too, and hence that γ is, from $\gamma\alpha = \beta\mu$, and so we are done.

7 Connectivity of Subcomplexes Part I

Having achieved the first goal of the essay, we now look towards proving a homological stability result for the group homology of $\text{Mod}(M_{n,s})$. We define the 5 subcomplexes of $\mathbb{S}(M_{n,s})$ which we called $Y_{n,s}, X_{n,s}, W_{n,s}, U_{n,s}$ and $Z_{n,s}$, and prove the relevant results about their connectivity and point stabilisers. We will then proceed with the spectral sequence arguments described in the previous section.

We begin with $Y_{n,s}$.

Definition 7.1. *We define $Y_{n,s}$ to be the subcomplex of $\mathbb{S}(M_{n,s})$ consisting of all those simplices where $M_{n,s} \setminus S$ is connected. We define $Y_{n,s}^{\pm}$ to be the complex with simplices consisting of those sphere systems such that $M \setminus S$ is connected, together with an orientation of each component.*

Remark 7.1. We have $\dim Y_{n,s} = \dim Y_{n,s}^\pm = n - 1$.

The following theorem is stated, but proven incorrectly in [7] (Proposition 3.1). The proof that follows is from [11], Proposition 3.2, however we have written it in the notation of [7] as we shall use that notation for Y in this essay.

Theorem 7.1. $Y_{n,s}$ and $Y_{n,s}^\pm$ are $(n - 2)$ -connected.

Proof. The result is vacuous for $n = 0$, so we assume $n > 0$.

We consider a map $f: S^m \rightarrow Y_{n,s}$ where $m \leq n - 2$. Since we know $\mathbb{S}(M_{n,s})$ is contractible, this extends to a map $F: D^m \rightarrow \mathbb{S}(M_{n,s})$.

Let σ be a p -dimensional simplex in D^m whose image is a q -dimensional simplex τ of $\mathbb{S}(M_{n,s})$ such that every sphere in the system representing τ separates the complement of all other spheres in the system.

Let the system representing τ be $S_0 \cup \dots \cup S_q$. Then let L_σ be the subcomplex of D^m consisting of the simplices γ such that $\gamma \cup \sigma$ is a simplex, and $\gamma \cap \sigma = \emptyset$. We call this the link of σ .

This is mapped to the link of γ , but this is just the join of all $Y(M_1), \dots, Y(M_d)$, where M_i are the components of the complement of $S_0 \cup \dots \cup S_q$.

By induction on (n, s) , $Y(M_i)$ is $(n_i - 2)$ -connected, so the join of all $Y(M_1), \dots, Y(M_d)$ is $((\sum n_i) - 2)$ -connected. Now this is larger than $n - q - 2 \geq m - p$ as $n - 2 \geq m$, so we can extend the map $S^{m-p} = \text{Link}(\sigma) \rightarrow \text{Link}(\tau)$ to the disc D^{m-p+1} . We use this to change F on the inside of $\text{Link}(\sigma)$, and after finitely many such modifications, we have a map into $Y_{n,s}$, so $Y_{n,s}$ is $(n - 2)$ -connected.

For $Y_{n,s}^\pm$, we may perform an identical procedure, but this time with a simplex mapping into a maximal number of negatively oriented components (with respect to a fixed positive orientation on all spheres), to obtain the same result. $\square$

Now it is not trivial to obtain a nice description of stabilisers for $Y_{n,s}$, as given a simplex σ represented by a sphere system S , we would like to say the stabiliser of σ is just $\Gamma(M_{n,s} \setminus S) = \Gamma_{n-k, s+2k}$. This however is not true, as any of the last $2k$ punctures can be permuted while still obtaining an element of the stabiliser. To remedy this issue, we must work with a version of $Y_{n,s}^\pm$ with ordered vertices.

Definition 7.2. We define $X_{n,s}$ to be the complex with k -simplices the simplicial maps $\Delta^k \rightarrow Y_{n,s}^\pm$ where Δ^k is the standard k -simplex.

Remark 7.2. This complex is infinite dimensional, as spheres may be repeated, but it is a classical fact in algebraic topology that it has the same simplicial homology as $Y_{n,s}^\pm$, and hence the same singular homology.

The action of $\Gamma_{n,s}$ on $Y_{n,s}^\pm$ gives an action on $X_{n,s}$. If the image of the simplex $\sigma \in X_{n,s}$ is k -dimensional in $Y_{n,s}^\pm$, then we may expect that the stabiliser of σ , denoted Γ_σ , is now $\Gamma_{n-k, s+2k}$, and this is indeed the case due to the following lemma from [7].

Lemma 7.1. The natural surjection $\Gamma_{n-k, s+2k} \rightarrow \Gamma_\sigma$ given by acting on $M_{n,s} \setminus \text{Im}(\sigma)$, and then gluing M' back up into $M_{n,s}$ is an isomorphism.

Proof. Let $M_S = M_{n,s} \setminus S$ and suppose $g: M_S \rightarrow M_S$ is a diffeomorphism which is isotopic to the identity as a map $M \rightarrow M$.

Let S' be a sphere system in M_S whose complement is connected and simply connected. Then $f(S')$ is isotopic to S' in M , and hence by passing to the universal cover of M , homotopic in M_S . By Laudenbach [12], S' and $f(S')$ are isotopic in M_S , and so g may be isotoped to leave S' invariant.

We must preserve orientation as otherwise f would act non-trivially on $H_2(M)$. Splitting M_S along S' gives a punctured S^3 , and g is isotopic to the identity on this S^3 as it preserves all the punctures.

Therefore g is isotopic to a product of twists around S' in $\text{Mod}(M_S)$, and so is 0 in $\Gamma(M_S)$ and so we are done. $\square$

We now wish to prove something about the connectivity of the quotient $X/\Gamma_{n,s}$. The following lemma (as every theorem in this section so far) is from [7].

Lemma 7.2. *$X/\Gamma_{n,s}$ is $(n-2)$ -connected.*

Proof. By a simple argument via cutting along the relevant sphere systems, $\Gamma_{n,s}$ acts transitively on $Y_{n,s}^\pm$, so every k -simplex is equivalent to a subsimplex of a fixed oriented system of size n which we denote $S = S_1 \cup \dots \cup S_n$. All permutations of S can be realised via diffeomorphisms of $M_{n,s}$, and so X/Γ can be identified with the quotient of K_n , the complex with k -simplices the lists of numbers below n (possibly with repeats), under the action of the symmetric group S_n .

Now $K_{n-1} \subseteq K_n$ descends to an inclusion $K_{n-1}/S_{n-1} \subseteq K_n/S_n$ with the $(n-2)$ -skeleton completely contained in the former, so it suffices to show that K_{n-1}/S_{n-1} is contractible in K_n/S_n . This is immediate, as adding n to the end of every k -simplex in the former gives us a cone over K_{n-1}/S_{n-1} in K_n/S_n . $\square$

8 Connectivity of Subcomplexes Part II

We now define $W_{n,s}$ and give a result on its connectivity. This is by far the most involved of the 5 complexes to give this result for, and so we give it its own section. The main observation is a generalisation of the results in section 5 allowing us to relate $W_{n,s}$ to similar complexes of graphs, and deduce the results for those graphs using the degree theorem of [8].

We give a formal definition of $W_{n,s}$.

Definition 8.1. *The complex $W_{n,s}$ is defined by having k -simplices be ordered, oriented collections of sphere systems of size $k+1$ with connected complement, together with an enveloping sphere S which separates those spheres, along with a fixed boundary component ∂_0 , from the rest of the manifold. By this, we mean that the component of $M_{n,s} \setminus S$ which contains ∂_0 has boundary consisting only of S and ∂_0 , and is simply connected after removing the sphere system.*

Remark 8.1. The enveloping sphere is unique up to isotopy if the sphere system has size n , but is genuinely meaningful extra data if that is not the case.

Remark 8.2. In a similar way to before, this has a natural action by $\Gamma_{n,s}$. We also define $V_{n,s}$ (resp. $V_{n,s}^\pm$) to be the version of $W_{n,s}$ without ordering or orientation (resp. without ordering). Since $\Gamma_{n,s}$ acts transitively on $V_{n,s}^\pm$ and all reorderings of vertices are realisable by $\Gamma_{n,s}$, we see that $W_{n,s}/\Gamma_{n,s} = K_n/S_n$, and hence is $(n-1)$ -connected.


To motivate our following discussion, we look back at Section 5, and in particular the discussion for $\text{Aut}(F_n)$. We note that we needed to add in a seemingly ad-hoc exception for one vertex of valence 1 in order to account for the boundary component. This was the $s=1$ case, and, following [9], we will define thorned roses as an analogue of the n -petaled roses of Culler-Vogtmann outer space to account for the existence of boundary components. We will then deduce the connectivity of $W_{n,s}$ in a similar way to the quotient maps of Section 5.

Definition 8.2. *An s -thorned metric graph is a based metric graph of volume 1 with s distinguished points $v_0, \dots, v_{s-1}$ (which may be equal), where v_0 is the basepoint. We fix a positive length x and attach edges of length x to each of the $v_1, \dots, v_{s-1}$, and these are called the thorns of the graph. We call the other endpoints of the thorns w_i .*

Definition 8.3. The thorned rose $R_{n,s}$ is the graph with s vertices $v_0, \dots, v_{s-1}$, with edges (v_0, v_i) for $1 \leq i \leq s-1$, and n edges from v_0 to itself.

Remark 8.3. The thorned roses are our analogue of R_n from Culler-Vogtmann outer space in this more general situation, in that they are what will be used to define a ‘marking’ on s -thorned graphs with fundamental group F_n .

Example 13. The thorned rose $R_{3,4}$ is depicted below:



Definition 8.4. A marking of an s -thorned graph Γ is a homotopy equivalence $g: R_{n,s} \rightarrow \Gamma$ such that $g(v_i) = w_i$ for $1 \leq i \leq s-1$. Markings $g: R_{n,s} \rightarrow \Gamma$ and $g': R_{n,s} \rightarrow \Gamma'$ are equivalent if there is an isometry $h: \Gamma \rightarrow \Gamma'$ sending w_i to w'_i which commutes with the markings in the same sense as the $s = 0$ case of Section 2.

In an identical manner to how we defined X_n in Section 2, we define $A_{n,s}$ for this definition of marking, and it is immediate that $X_n = A_{n,0}$.

We note that in the situation of Section 5, we had a bijection between X_n and those simplices of $\mathbb{S}(M_{n,0})$ with every complement simply connected. This motivates the following definition:

Definition 8.5. A sphere system S is simple if every component of $M_{n,s} \setminus S$ is simply connected.

Remark 8.4. We denote the subspace of simple sphere systems as $SS(M_{n,s})$, which is a union of open simplices, but not a subcomplex of $\mathbb{S}(M_{n,s})$.

We define $\Phi: SS(M_{n,s}) \rightarrow A_{n,s}$ in the same way as in Section 5, except we also take a neighbourhood of the boundary spheres except for ∂_0 (in addition to those elements of the sphere system), which gives us our thorns. We get a marking in an identical way to in Section 5, by fixing a sphere system S_0 with graph isomorphic (as thorned graphs) to $R_{n,s}$, and marking using the graph corresponding to it. The requirement for S_0 to have such a graph is that the complement be connected, and clearly such a system does exist.

The result we desire (that $SS(M_{n,s})$ is homeomorphic to $A_{n,s}$) will immediately follow just like Theorem 5.1 if we can prove that $G_{n,s}$ (the group of homotopy self-equivalences of $R_{n,s}$ which fix the marked points) is isomorphic to $\Gamma_{n,s}$ (ie, when we prove a higher dimensional equivalent of $\text{Out}(F_n) \cong \Gamma_{n,0}$ from [12]).

We define the homomorphism $\phi: \Gamma_{n,s} \rightarrow G_{n,s}$ which will be the isomorphism we desire in the only way possible: $\text{Mod}(M_{n,s})$ acts on $M_{n,s}$, and so we induce a map $R_{n,s}$ to itself by including it into $M_{n,s}$, applying $f \in \Gamma_{n,s}$ to $M_{n,s}$, and then applying the quotient. Since Dehn twists have no effect on the class of $f \circ i$, this is well defined, and for it to be a homomorphism we need the following two maps to be homotopic, where $q: M_{n,s} \rightarrow R_{n,s}$ is the quotient map analogous to that of Section 5.

$$\begin{aligned} R &\xrightarrow{i} M \xrightarrow{f} M \xrightarrow{q} R \xrightarrow{i} M \xrightarrow{g} M \xrightarrow{q} R \\ R &\xrightarrow{i} M \xrightarrow{gf} M \xrightarrow{q} R \end{aligned}$$

Now $i \circ r$ is not necessarily homotopic to the identity, but $f \circ i$ can be homotoped to have image in $i(R)$ so that $i \circ r$ is the identity on its image, as $i(R)$ has the same fundamental group as $M_{n,s}$.

Proposition 8.1. *The above homomorphism is an isomorphism.*

Proof. As remarked before, for $s \leq 1$ this is a theorem in [12]. We induct on s for the general case.

We use the group of diffeomorphisms of $M_{n,0}$ fixing s distinguished points $v_0, \dots, v_{s-1}$ and we denote this group by $\text{Diff}_{n,s}$. We use this in place of $\text{Diff}(M_{n,s})$, noting that $\pi_0 \text{Diff}_{n,s} = \text{Mod}(M_{n,s})/G$ where G is the subgroup generated by Dehn twists around the boundary components. We do this as we wish to be able to evaluate at the boundaries.

Evaluation at the s th distinguished point gives a fibration $\text{Diff}_{n,s} \rightarrow \text{Diff}_{n,s-1} \rightarrow M_{n,0} \setminus \{v_0, \dots, v_{s-2}\}$. Similarly, we get a fibration for the homotopy equivalences of $R_{n,s}$ and the homotopy equivalences where we view the metric graph $R_{n,s}$ as an $(s-1)$ -thorned graph by forgetting the last distinguished point. The LES for these fibrations yield the top and bottom rows of the following respectively:

$$\begin{array}{ccccccc} F_n & \longrightarrow & \pi_0 \text{Diff}_{n,s} & \longrightarrow & \pi_0 \text{Diff}_{n,s} & \longrightarrow & 1 \\ \parallel & & \downarrow \psi_n & & \downarrow \psi_{n-1} & & \\ F_n & \longrightarrow & G_{n,s} & \longrightarrow & G_{n,s-1} & \longrightarrow & 1 \end{array}$$

where ψ_i is the composition $\pi_0 \text{Diff}_{n,i} \rightarrow \Gamma_{n,i} \rightarrow G_{n,i}$. The second row extends to a short exact sequence, as if two maps $R_{n,s} \rightarrow R_{n,s}$ are homotopic, they must send the s th thorn to homotopic loops in $R_{n,s}$, and $a \in F_n$ is sent to a map which sends the s th thorn to a . Therefore as the diagram commutes, the top row extends to a short exact sequence. The snake lemma shows that the kernels of ψ_n and ψ_{n-1} are isomorphic, and so by induction both are generated by all the Dehn twists. The induced map $\Gamma_{n,s} \rightarrow G_{n,s}$ is therefore an isomorphism. $\square$

In view of the fact (cf Section 5) that $SS(M)$ is contractible, the preceeding proposition yields that $A_{n,s}$ is contractible (at least for $n \geq 1$, which is the case we are interested in).

8.1 The Degree Theorem

In the discussion of the connectivity properties of $Y_{n,s}$, the key insight was to notice that we could inductively reduce how ‘far’ a homotopy lying in our total space (which was contractible) was from lying in the subspace we were interested in. We wish to acquire a similar tool for $A_{n,s}$, however this shall turn out to be more technical than it was for $Y_{n,s}$. The primary tool needed will be a variation of the degree theorem of [8], as stated in [9]. In the former paper, the degree of a basepointed graph was defined to be $2n - |v_0|$ where $|v|$ is the valence of v . The degree theorem then stated that a k -parameter family of graphs in $A_{n,1}$ (such as a D^k -indexed family of graphs) can be deformed into the subspace of graphs of degree at most k via a deformation during which the degree is non-increasing. This then immediately implies that the subspace of degree k graphs in $A_{n,1}$ is $(k-1)$ -connected. We make the following definition, with the aim of generalising this theorem to thorned graphs.

Definition 8.6. *The degree of an s -thorned graph Γ is $2n + s - 1 - |v_0|$. We denote the set of those graphs of degree at most k by $A_{n,s}^k$.*

Remark 8.5. The degree of $R_{n,s}$ is 0, and the degree of a graph in $A_{n,1}$ agrees with that as defined in [8].


The version of the degree theorem we shall need is that stated in [9]. As the proof of this would take us too far a field into Morse theory, the author shall content themselves by directing the reader to [8] (for the full details) and [9] (for the minor modification needed in the case of thorned graphs).

Theorem 8.1. *A map $f_0: D^k \rightarrow A_{n,s}$ is homotopic to a map $f_1: D^k \rightarrow A_{n,s}^k$ by a homotopy during which the degree decreases monotonically.*

The following lemma from [9] shall be needed to prove that a homotopy we shall use later will be well defined:

Lemma 8.1. *If $k < n/2$ then a degree k thorned graph of rank n has at least one loop at the basepoint (which we require has no thorns in its interior)*

Proof. Suppose Γ has rank n and degree k with basepoint v_0 . By definition $k = 2n + s - 1 - |v_0|$. By performing the operation below (called *blowing up*, see [4] for more details - this is essentially the operation of moving from a face to an open simplex containing it. The operation turns the vertex labelled 1 into those labelled 2 and 3), we may assume every vertex other than v_0 is trivalent or univalent.



For such a graph, the degree equals the number of trivalent vertices, as collapsing such a vertex back down into v_0 reduces the degree by 1, and eventually ends up at $R_{n,s}$. Let Γ_1 be the subgraph spanned by all vertices other than the basepoint. If Γ has no loops at the basepoint, then $E(\Gamma) = E(\Gamma_1) + |v_0|$ where E is the number of edges. We therefore calculate the Euler characteristic as

$$\begin{aligned}
 1 - n &= \chi(\Gamma) \\
 &= V(\Gamma) - E(\Gamma) \\
 &= (1 + V(\Gamma_1)) - (E(\Gamma_1) + |v_0|) \\
 &= (1 + k + s - 1) - (E(\Gamma_1) + (2n + s - 1 - k))
 \end{aligned}$$

So $n = 2k - E(\Gamma_1)$, and so $n \leq 2k$. This was all under the assumption that we have no loops at the basepoint, so we are done. $\square$

We finish our combinatorial diversion by proving the connectivity of the subcomplex of $A_{n,s}$ which we will compare $V_{n,s}$ to. This subspace is $L_{n,s}$, the subspace of all graphs with at least one loop at the basepoint, and no separating edges other than thorns. Note that this is not a subcomplex, only a union of open simplices.

Theorem 8.2. *$L_{n,s}$ is $(n - 3)/2$ -connected.*

Proof. As in [4], the condition on separating edges is trivial to work around, by uniformly shrinking their lengths to 0, and so we shall ignore it.

Now suppose $k < n/2$ and take a map of S^{k-1} into $L_{n,s}$. As before, we get an extension $f_0: D^k \rightarrow A_{n,s}$, which, by the degree theorem, is homotopic to a map $f_1: D^k \rightarrow A_{n,s}^k$ via a homotopy monotonically reducing the degree. By the preceding lemma, every graph in $A_{n,s}^k$ is in $L_{n,s}$, so f_1 maps into $L_{n,s}$.

The proof of the degree theorem yields a homotopy which can only ever increase the number of loops at the basepoint, and so the image of S^{k-1} stays in $L_{n,s}$ for all time, and is filled in at $t = 1$, and hence the image of $S^{k-1} \times [0, 1] \cup D^k \times \{1\} \cong D^k$ fills in $S^{k-1} \times \{0\}$, and lies in $L_{n,s}$, so we are done. $\square$

8.2 Proof of Connectivity

Once again taking a degree of inspiration from Section 5, we wish to produce a nice subspace of $L_{n,s}$ which is both a simplicial complex (after barycentric subdivision) and a deformation retraction of $L_{n,s}$. We construct $SL_{n,s}$, called the *spine* of $L_{n,s}$ from $L_{n,s}$ in exactly the same manner as we constructed K_n from Y_n in Section 5.

We wish to translate this into the language of sphere systems, as some structure is lost when we take the quotient down to graphs. It is fairly immediate to see that $L_{n,s}$ corresponds to those sphere systems such that:

- The system is simple;
- No sphere in the system is separating;
- The component of the complement containing ∂_0 also contains both sides of at least one sphere.

Theorem 8.3. $V_{n,s}$ is $(n-3)/2$ -connected for $s \geq 0$.

Proof. We define a map $f: SL_{n,s} \rightarrow V_{n,s}$ as follows: for a sphere system S we define M_0 to be the component of $M_{n,s} \setminus S$ containing ∂_0 , and we send the sphere system S to the subsystem T consisting of those spheres with both sides contained in M_0 , with enveloping sphere the unique such sphere contained in M_0 .

We claim this is an order reversing poset map, for the posets defined by inclusion of faces. Suppose S' is a face of S in $SL_{n,s}$. Then it must contain all the spheres with both sides in M_0 , as otherwise the system would not be simple. Therefore M'_0 , the component of $M_{n,s} \setminus S'$ containing ∂_0 , is (not necessarily strictly) larger than M_0 , so the enveloping sphere can be taken to be either the same as, or disjoint from, that of T .

To prove that this is a homotopy equivalence, we will invoke Quillen's Poset Lemma [14]. To do so, we need to prove that for any $\tau = (T, T_e) \in V_{n,s}$, the fibre $f_{\geq \tau} = \{S \in SL_{n,s} \mid f(S) \geq \tau\}$ is contractible.

We let k be the number of spheres in τ . If $k = n$, then $f_{\geq \tau}$ is a single point since all spheres in $SL_{n,s}$ are non-separating. Now suppose $k < n$. Then $f_{\leq \tau}$ is the set of all simple systems of non-separating spheres which contain T and are disjoint from T_e . Now T_e splits $M_{n,s}$ into two parts, one of which is homeomorphic to $M_{n-k,s}$ (and the other of which is already simply connected, so contains no sphere of any element of $f_{\geq \tau}$ except for those already in T). This is the space $A_{n-k,s}$, after contracting separating edges as described before, and so is contractible. Hence we may apply Quillen's Poset Lemma, and so f is a homotopy equivalence and we are done. $\square$

Remark 8.6. In an identical manner to how we proved that $Y_{n,s}^\pm$ has the same connectivity as $Y_{n,s}$ in the previous section, $V_{n,s}^\pm$ is also $(n-3)/2$ -connected.

9 Connectivity of Subcomplexes Part III

We now aim to define $U_{n,s}$ from [10], and prove that the limit $U_{\infty,s} = \varinjlim U_{n,s}$ has nice connectivity and stabiliser properties.

Definition 9.1. We define the complex $Z_{n,s}$ for $s \geq 2$ as follows. The vertices are equivalence classes of pairs (S, a) where S is a non-separating sphere and a is an arc from the first boundary component of $M_{n,s}$ to the last such that $|S \cap a| = 1$ with transverse intersection, up to isotopy. A set of $k+1$ vertices $(S_1, a_1), \dots, (S_{k+1}, a_{k+1})$ give a k -simplex if all the sets $S_i \cup a_i$ are pairwise disjoint and the spheres S_i form a sphere system in $Y_{n,s}$.

As we wish to take a limit of these, we need define an inclusion $Z_{n,s} \rightarrow Z_{n+1,s}$. We define this via an inclusion of $M_{n,s} \rightarrow M_{n+1,s}$ such that every curve ending on the final boundary component of the former can be isotoped to end on the final boundary component of the latter. We do this by gluing half the boundary of an $M_{1,1}$ to half of the final boundary component of $M_{n,s}$. It is immediate that this gives a map as desired. This map is injective by uniqueness of normal forms (see Section 2), and we denote the limit by $Z_{\infty,s}$.

Theorem 9.1. $Z_{\infty,s}$ is contractible.

Proof. As it is a CW-complex, it suffices to prove that every homotopy group is trivial. We therefore start with a map $g: S^{k-1} \rightarrow Z_{\infty,s}$, which we may assume is simplicial, and wish to extend it to D^k . As the triangulation of S^{k-1} has finitely many simplices, its image lands in some $Z_{n,s}$. We consider $M_n \subseteq M_{n+1} \subseteq \dots \subseteq M_{n+k+1}$, the inclusions described above which induce α , and choose non-separating spheres T_i in $M_{n+i+1} \setminus M_{n+i}$.

We triangulate D^k by coning off the triangulation of S^{k-1} to a new vertex v , and define $g(v)$ to be (T_0, b_v) , where b_v is any valid arc.

We extend this over the other simplices of D^k by induction on the dimension of the simplex as follows: For a simplex σ of dimension i not contained in S^{k-1} , σ is the coning off of a $(i-1)$ -dimensional simplex which is mapped to $\{(S_0, a_0), \dots, (S_{i-1}, a_{i-1})\}$ in Z_n . The rest of the boundary of σ is formed by lower dimensional simplices in D^k not contained in S^{k-1} , which are mapped to a subcomplex with vertices $\{(S_0, a_0), \dots, (S_{i-1}, a_{i-1}), (T_j, b_\tau)\}$ where $0 \leq j < i$, and $\tau < \sigma$. We send the barycentre of σ to a valid (T_i, b_σ) , which exists as the complement of the already chosen spheres is connected, and we extend g linearly. As we have now extended g , we are done. $\square$

We note that the following diagram commutes, and so $\Gamma_{\infty,s} = \varinjlim \Gamma_{n,s}$ acts on $Z_{\infty,s}$

$$\begin{array}{ccc} Z_{n,s} \times \Gamma_{n,s} & \xrightarrow{(z,\gamma) \mapsto \gamma z} & Z_{n,s} \\ (i,\alpha) \downarrow & & \downarrow i \\ Z_{n+1,s} \times \Gamma_{n+1,s} & \xrightarrow{(z,\gamma) \mapsto \gamma z} & Z_{n+1,s} \end{array}$$

This action is transitive on k -simplices for any k . As usual we consider the ordered version of $Z_{\infty,s}$, which we call $U_{\infty,s}$, and we have the following:

1. $U_{\infty,s}$ has trivial homology, as $Z_{\infty,s}$ does;
2. $U_{\infty,s}/\Gamma_{\infty,s} = \varinjlim U_{n,s}/\Gamma_{n,s} \cong \varinjlim W_{n,s}/\Gamma_{n,s}$, so it is contractible as it has arbitrarily high connectivity;
3. The vertex stabilisers in U_n are copies of $\Gamma_{n-1,s}$ called $\Gamma'_{n-1,s}$, with inclusion given by $\mu\eta$. This commutes with α , so the vertex stabilisers in $\Gamma_{\infty,s}$ are $\Gamma'_{\infty,s}$ (and the inclusions of higher dimensional stabilisers are $(\mu\eta)^n$).

10 Spectral Sequence Arguments

10.1 Stabilisation of α

We will use the spectral sequence associated to the action of $\Gamma_{n+1,s}$ on $W_{n+1,s}$ to prove the following result from [9]:

Theorem 10.1. *For $s \geq 1$, the map $\alpha_*: H_*(\Gamma_{n,s}) \rightarrow H_*(\Gamma_{n+1,s})$ is an isomorphism for $n \geq 2i + 2$ and a surjection for $n \geq 2i + 1$.*

Proof. Let C_* be the chain complex associated to W , augmented with a $\mathbb{Z}$ in degree -1 , and let $E\Gamma_*$ be a free resolution of $\mathbb{Z}$. We have a double complex $W_* \otimes E\Gamma_*$, and so we get two spectral sequences associated to the horizontal and vertical filtrations, which converge to the same thing. First taking homology with respect to the vertical maps (the $E\Gamma_*$ factor), we get

$$E_{p,q}^1 = H_q(\Gamma; C_p) \cong \bigoplus_{\tau} H_q(\Gamma_{\tau}; \mathbb{Z})$$

where the sum is taken over the Γ -orbits of p -simplices, and the isomorphism is by Shapiro's lemma. The coefficient module $\mathbb{Z}$ is trivial as every stabiliser fixes the simplex pointwise, so we will drop it from here on. We note that like before, $\Gamma_{n+1-k,s} \cong \Gamma_{\tau}$.

If we take homology in the horizontal direction first, we get that for this spectral sequence, $E_{p,q}^1 = 0$ for $p \leq m$. We therefore get that in the vertical sequence, $E_{p,q}^{\infty} = 0$ for $p + q \leq m$, as the sum of all these terms is 0 for the horizontal sequence, and the sums for the horizontal and vertical sequences are the same. We proceed by induction on i .

We first show that α_* is a surjection in the desired range. Now this α_* is the differential $d^1: E_{0,i}^1 \rightarrow E_{-1,i}^1$, so it suffices to show that $E_{-1,i}^{\infty} = 0$ and $E_{p,i-p}^2 = 0$ for all $p > 0$. The former ensures that every element of $E_{-1,i}^1$ is the image of a differential in some page, and the latter ensures that all differentials into $E_{-1,i}^k$ for $k > 1$ are coming from 0.

Now if $i - 1 \leq (n - 2)/2$, we have already proven above that $E_{-1,i}^{\infty} = 0$, as $(n - 2)/2$ is a lower bound on the connectivity of W . For the second part, in the q th row, using the description of d^1 from [2] VII.3, we see that these differentials compute the homology of W/Γ with coefficients in $H_q(\Gamma_{\tau})$. If $\tau' \leq \tau$ and $q + p, q \leq i$, then by induction we know $H_q(\Gamma_{\tau}) \rightarrow H_q(\Gamma_{\tau'})$ is an isomorphism, and the maps from the $p + q = i + 1$ terms to the $p + q = i$ terms are surjective, so for $p + q \leq i$, we are taking homology with locally constant coefficients. As permuting faces is induced by a diffeomorphism which acts trivially on the part of $M_{n,s}$ which is cut away by the enveloping sphere, and Γ_{τ} only acts on the part of $M_{n,s}$ which is cut away, these are in fact constant coefficients. As W/Γ is $(n - 1)$ -connected, the desired result holds for $i \leq n - 1$.

To show that d^1 is injective, we need to show that $E_{0,i}^{\infty} = 0$ and $E_{p+1,i-p}^2 = 0$ for all $p > 0$, and also that $d^1: E_{1,i}^1 \rightarrow E_{0,i}^1$ is 0. The former two are identical to the previous argument, with connectivity condition $n \geq 2i + 2$. For the final part, there are two orbits of 1-cells, one of which is degenerate and the other of which isn't. The coefficient system is trivially untwisted for the former, and is untwisted for the latter by the above argument. There is only one zero cell in W/Γ , so it immediately follows that the differential is 0. $\square$

We note that $\eta\mu = \alpha$, so from the above result, μ is injective on homology.

10.2 Stabilisation of μ

We wish to prove that μ is an isomorphism on homology for $n \geq 2i + 2$. We note that it now suffices to prove that μ is surjective on homology. To do this, we will prove $(\mu\eta)_*$ is surjective. As this commutes with α , it suffices to prove that it is surjective on $\varinjlim H_i(\Gamma_{n,s}) = H_i(\Gamma_{\infty,s})$, as we can then use that α is an isomorphism above $2i + 2$ to see that $H_i(\Gamma_{\infty,s}) \cong H_i(\Gamma_{2i+2,s})$. The previous discussion and following theorem are due to [10].

Theorem 10.2. *For $s \geq 1$, the map $(\mu\eta)_*: H_*(\Gamma'_{\infty,s}) \rightarrow H_*(\Gamma_{\infty,s})$ is an isomorphism for $s \geq 2$.*

Proof. We only need to argue it is surjective, by the comments above. To do this, we follow the same proof as for α . The only part which does not immediately translate to our new situation is the proof that the local coefficient system $H_j(\text{Stab}(\sigma_p))$ is trivial. We know that each face of σ_p is of the form $h\sigma_{p-1}$ for some $h \in \Gamma_{\infty,s}$, and therefore that the differential d^1 is induced by conjugation by h . By induction on i , we may assume that the vertical maps in

$$\begin{array}{ccc} H_j(\text{Stab}(\sigma_p)) & \xrightarrow{h_*} & H_j(\text{Stab}(\sigma_{p-1})) \\ \cong \downarrow & & \downarrow \cong \\ H_j(\Gamma_{\infty,s}) & \xrightarrow{h_*} & H_j(\Gamma_{\infty,s}) \end{array}$$

are isomorphisms. We know the lower h_* is the identity, as it is an inner automorphism, so the upper h_* is also trivial and we are done. $\square$

10.3 Stabilisation of β for Small s

For $s > 2$, we know the result on the stabilisation of β_* as $\beta = \alpha\mu^{-2}$, but this identity does not work for $s = 2$, and hence we need a different argument from [10]. We give it here.

Theorem 10.3. $\beta: \Gamma_{n,2} \rightarrow \Gamma_{n+1,0}$ is an isomorphism on homology for $n \geq 2i + 3$ and a surjection for $n = 2i + 2$.

Proof. Once again, we only concern ourselves with the local triviality of the coefficient system.

The problem, compared to the case of α , is the lack of enveloping spheres. We solve this by choosing them arbitrarily, and we see that the choice does not affect the map on homology. After cutting $M_{n+1,0}$ into two along the enveloping sphere, we see that the coefficients will be untwisted if the map $H_q(\Gamma_{n+1-k,1}) \rightarrow H_q(\Gamma_{n+1-k,2k})$ given by μ^{2k-1} is surjective, so that we can take any element of the latter (which is $H_q(\Gamma_\tau)$) to be supported on the far side of the enveloping sphere (ie, in the former). In the proof of surjectivity for α , the highest value $n + 1 - k$ will take is $n - 2$ (as the differential has domain $E_{2,i-1}^1$, whence the relevant simplices contain 3 spheres), so we have proven surjectivity for $n - 2 \geq 2(i - 1) + 2$, that is, $n \geq 2i + 2$. For injectivity, the worst case is for the differential $E_{1,i}^1 \rightarrow E_{0,i}^1$, giving $n \geq 2i + 3$. $\square$

This immediately gives the following corollary:

Corollary 10.1. The map $\text{Aut}(F_n) \rightarrow \text{Out}(F_n)$ is an isomorphism on H_i for $n \geq 2i + 4$ and a surjection for $n = 2i + 3$.

11 Whitehead's Algorithm

We now change focus to look at another application of the 3-manifold topology we have developed, using it to determine when two finite lists of elements $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{b} = (b_1, \dots, b_n)$ of F_n are equivalent under $\text{Aut}(F_n)$, ie when there is $\phi \in \text{Aut}(F_n)$ such that $\phi(a_i) = b_i$ for all $1 \leq i \leq n$. Our exposition will largely follow [6].

We will treat these words as strings in a_i, a_i^{-1} , reserving the term ‘element’ for the strings seen as elements of F_n , and ‘cyclic word’ when we do not care about the start or end point of the words, and read them in a cycle.

We define several types of transformation of words which will help us reduce $\mathbf{a}$ into a normal form that we may compare with that of $\mathbf{b}$. We will assume that F_n is generated by some fixed generators $x_1, \dots, x_n$.

11.1 Transformations and Sphere Bases

Definition 11.1. An *elementary transformation* of a word a is the insertion or removal of subwords of the form $x_i x_i^{-1}$ or $x_i^{-1} x_i$. Two words which may be taken to each other by such transformations are said to be **equivalent**. An element with no subwords of the forms $x_i x_i^{-1}$ or $x_i^{-1} x_i$ is said to be **(elementary) reduced**.

Definition 11.2. Two sets of words $\mathbf{a}, \mathbf{b}$ are **equivalent** if there is an automorphism $\phi \in \text{Aut}(F_n)$ such that $\phi(\mathbf{a}) = \mathbf{b}$.

We will ordinarily work with elements.

Definition 11.3. For a set of reduced words $\mathbf{a}$, we denote by $L(\mathbf{a})$ the sum of their lengths.

Definition 11.4. A *simple transformation* of a set of words $\mathbf{a}$ is any one of the following:

1. A transformation which fixes $x_k^{\pm 1}$ for $k \neq i$ and sends $x_i^{\pm 1} \mapsto (x_i x_j^\varepsilon)^{\pm 1}$;
2. A transformation which fixes $x_k^{\pm 1}$ for $k \neq i$ and sends $x_i^{\pm 1} \mapsto (x_j^\varepsilon x_i)^{\pm 1}$;
3. A transformation which fixes $x_k^{\pm 1}$ for $k \neq i$ and sends $x_i^{\pm 1} \mapsto x_i^{\mp 1}$;
4. A transformation that swaps two elements of the basis,

where ε is a choice of sign.

Definition 11.5. A *T-transformation* is a composition of simple transformations of type 1 and 2, each performed on a different x_i , where j and ε are constant for all of the simple transformations performed (note that we may perform both a type 1 and 2 transformation on the same x_i , whence we get $x_i \mapsto x_j^{-\varepsilon} x_i x_j^\varepsilon$). If every word in $\mathbf{a}$ is elementary reduced, and $L(T\mathbf{a}) < L(\mathbf{a})$ for some T -transformation T , we say T is a **T-reduction**, and if $L(T\mathbf{a}) = L(\mathbf{a})$ we say T is a **level transformation**. We also call simple transformation of type 3 **level transformations**. If no T -reductions can be performed to $\mathbf{a}$, we say it is **minimal**.

Remark 11.1. There are a finite number of T -transformations that can be applied to any set of words, and we can determine which are reductions or level transformations, and hence we can algorithmically reduce a set of words to a minimal set. We aim to prove that any two minimal sets for $\mathbf{a}$ can be taken to each other via level transformations, and this will provide an algorithm to determine if $\mathbf{a}$ and $\mathbf{b}$ are equivalent as described, as there are finitely many lists of words of a given length, and so we may enumerate all of those equivalent to a minimal form of $\mathbf{a}$.

The next few propositions and lemmas (until otherwise stated) are due to [6]. We begin by relating certain sphere systems to bases of F_n . For each of the copies of $S^2 \times S^1$ which make up $M_n = M_{n,0}$, we choose a sphere $S_i = S^2 \times \{t\}$ which is disjoint from any disks removed to form the connect sum. We know that $\pi_1(M_n) = F_n$. If we remove these spheres we get an n -punctured S^3 , and a version of Seifert van Kampen for gluing a space to itself proves:

Proposition 11.1. $\pi_1(M_n) = F(x_1, \dots, x_n)$ where, if the spheres (in order, counted with repetition) that a representative γ of the element $[\gamma]$ passes through are $S_{i_1}, \dots, S_{i_n}$, then $[\gamma] \mapsto x_{i_1} \dots x_{i_n}$.

An automorphism of F_n may therefore be seen as a change of choice of the spheres S_i .

Definition 11.6. We call any image of the spheres S_i under a homeomorphism $h: M_n \rightarrow M_n$ a **sphere basis**.

Definition 11.7. Two sphere bases Σ and Σ^0 are **equivalent** if they represent the same basis of F_n .

Proposition 11.2. For any automorphism ϕ of F_n , there is a sphere basis Σ such that $\phi([\gamma]) = \gamma_\Sigma$.

Proof. We let $\Sigma = h(S)$ where h^{-1} induces ϕ on π_1 . The result then follows from the fact that every automorphism is a product of Nielsen generators, and every Nielsen generator is induced by a diffeomorphism. $\square$

11.2 Recognising and Modifying Sphere Bases

While we have largely worked with M_n up to this point, we will find it useful to work with the space obtained after cutting along all the S_i . This is a sphere S^3 with $2n$ punctures, where we call the two removed open balls corresponding to S_i , D_i and D'_i . We denote this space by R_n . We have a natural quotient map $\pi: R_n \rightarrow M_n$ which identifies each pair D_i and D'_i .

We will also call the preimage of any sphere basis under π a sphere basis.

We now imagine another set of spheres $\Sigma_1, \dots, \Sigma_n$. A natural question arises as to when this set of spheres is also a sphere basis. We note that, when viewed as a subset of S^3 , one side of each of the spheres $\partial D_i, \partial D'_i$ contains no other sphere. We call this **canonical position**, and attempt to put the Σ_i into this position. This process will also prove our assertion earlier that every Nielsen automorphism can be realised by a diffeomorphism of M_n , as we may then take Σ_i containing S_i and Σ_j containing S_i and S_j , and every other $\Sigma_k = S_k$. These will satisfy the conditions of our below theorem, so we will then have the result.

The following result is the key observation that allows us to modify sphere bases:

Lemma 11.1. If S is a sphere basis in M_n in canonical position and Σ_1 is a sphere disjoint from any sphere of $\pi^{-1}(S)$ such that the two components of $\pi^{-1}(S_1)$ are on opposite sides of Σ_1 , then $\{\pi(\Sigma_1), S_2, \dots, S_n\}$ is a sphere basis in M_n .

Proof. We remove the component of R inside Σ_1 containing the positive side of S_1 . This component is a k -punctured sphere for some $k \geq 2$ (as we have punctures corresponding to Σ_1 and S_1 inside the sphere). We can therefore glue the puncture corresponding to S_1 to the puncture outside Σ_1 which corresponds to S_1 . This gives us a new $2n$ -punctured sphere where the spheres are in canonical position. $\square$

Proposition 11.3. Let Σ be a family of n disjoint 2-spheres in M_n . Then the following are equivalent:

1. No finite subset of Σ bounds a submanifold of M_n ;
2. Σ is a sphere basis;
3. $M_n \setminus \Sigma \cong R$;
4. $M_n \setminus \Sigma$ is connected.

Proof. (2) implies (3) implies (4) implies (1) are all easy. As we did in section 3, we wish to use the above lemma with Σ to modify S until they are the same, while preserving the property of being a sphere basis. We divide our proof into 3 cases.

1. If S and Σ differ only in orientation, we are done.
2. Suppose every Σ_i is either in S , or disjoint from S . Reordering and reorienting, we may assume $\Sigma_1, \dots, \Sigma_{k-1}$ are the first $k-1$ elements of S , and that all remaining spheres in Σ are disjoint from S . If Σ_k does not separate some S_j and S'_j , then Σ_j and all the spheres of Σ inside it (except for those corresponding to elements of S) bound a submanifold, contradiction, so by the lemma we may inductively replace S with Σ to get a sphere basis.

3. If Σ and S are not disjoint, then without loss of generality, $S_1 \cap \Sigma_1 \neq \emptyset$. We take an inner-most intersection C which bounds D on Σ_1 and E on S chosen so that $D \cup E$ separates S_1 and S'_1 after a minor perturbation. Then $\{D \cup E, S_2, \dots, S_n\}$ is a sphere basis with one less intersection with Σ so we are done by induction.

□

11.3 Normal Forms of Sphere Bases

The following proposition is original, although inspired by [6], and is needed so that we may strengthen 3.1 to say that the resulting normal form system is equivalent (in the sense of 11.7).

Lemma 11.2. *If sphere bases S and S' are isotopic, then they are equivalent in the sense of 11.7.*

Proof. Translating our problem into R , we see that the set of spheres corresponding to S is isotopic to the set of spheres corresponding to S' . We may extend this to an isotopy H_t on all of R as it is clear we may do so for some finite interval of time, and then compactness yields the isotopy desired. If we have a path γ representing the element $\gamma_S \in F_n$, then $H_t(\gamma)$ yields an isotopy between γ , and a path γ' which must therefore represent the same element in F_n . But then $\gamma_S = \gamma'_{S'} = \gamma_{S'}$, with the first equality holding as $\gamma' = H_1(\gamma)$ and $S = H_1(S')$, and the second as $\gamma \simeq \gamma'$. □

We therefore deduce the following proposition from Goldstein and Turner:

Proposition 11.4. *For a pair of sphere bases S, Σ there are equivalent bases S', Σ' which are in normal form relative to each other.*

Proof. 3.1, combined with the above lemma. □

We have one final theorem from [6] before we prove Whitehead's algorithm, which is useful in that it tells us that we may reduce a word with respect to two bases simultaneously.

Theorem 11.1. *Suppose that S and Σ are sphere bases in normal form.*

1. *If γ is a path on S_i between points not in $S \cap \Sigma$, then γ is homotopic relative to its endpoints to a path γ' such that γ'_Σ is reduced.*
2. *If γ is a path starting and ending in $M \setminus (S \cup \Sigma)$ then γ is homotopic relative to its endpoints to a path γ' such that γ'_S and γ'_Σ are both reduced.*

Proof.

1. Suppose γ_Σ is not reduced, so that $\gamma_\Sigma = \gamma_1 x_i x_i^{-1} \gamma_2$. The normal form ensures that the points where γ meets Σ_i lie on the same circle of intersection in $S \cap \Sigma$, so then we can homotope γ back through this circle, so that it no longer meets Σ_i .
2. After a small homotopy, we may assume that γ does not meet $S \cap \Sigma$. Assume $\gamma_\Sigma = \gamma_1 x_i x_i^{-1} \gamma_3$ and let γ_2 be defined to be the part of the path in the middle. We may homotope γ_2 to lie completely in Σ_i , and then by (1) we may reduce the resulting path so that we get some $\gamma' \simeq \gamma_2$ with γ'_S reduced. We then may homotope γ' slightly below Σ_i , which preserves intersections with S by transversality, and we have removed $x_i x_i^{-1}$ while not increasing the number of such pairs in γ_S , so we are done by induction.

□

We are finally ready to prove the main theorem we need to show that Whitehead's algorithm works. The proof largely mirrors that of Proposition 11.3, and is originally due to [16], although we have taken inspiration from [6] in our exposition of the proof.

Theorem 11.2. *Any two equivalent minimal sets $\mathbf{a}, \mathbf{b}$ of elements can be taken to each other via level transformations.*

Proof. As the inverse of a level transformation is a level transformation, it suffices to show one direction.

By Theorem 11.1 and Proposition 11.1, we may assume that $\mathbf{a}, \mathbf{b}$ are reduced and represented by the same set of paths w relative to sphere bases S, Σ respectively, and we assume that every path in w passes through a fixed common basepoint.

We first consider the case where S and Σ only meet in complete spheres. We assume that $\Sigma_1, \dots, \Sigma_k$ are disjoint from S , and $\Sigma_{k+1} = S_{k+1}, \dots, \Sigma_n = S_n$.

If $k = 0$ then we are done, as S and Σ differ only in orientation, and the required transformation is a level transformation.

In $H_2(M_n; \mathbb{Z}/2\mathbb{Z})$, we have bases given by Σ, S and hence $\Sigma_i = \sum_{j=1}^n r_{ij} S_j$ in $H_2(R; \mathbb{Z}/2\mathbb{Z})$ for some $r_{ij} \in \mathbb{Z}/2\mathbb{Z}$ where $1 \leq i, j \leq n$.

As the elements of S are independent in $H_2(M_n; \mathbb{Z}/2\mathbb{Z})$ by Proposition 11.3, we must have that $\det(r_{ij})_{1 \leq i, j \leq k} = 1$, and hence (after reordering S) we may assume that $r_{ii} = 1$ for all $1 \leq i \leq k$. It therefore follows that $\pi^{-1}(\Sigma_i)$ separates the two components of $\pi^{-1}(S_i)$ by considering the map $H_2(\pi; \mathbb{Z}/2\mathbb{Z}): H_2(R; \mathbb{Z}/2\mathbb{Z}) \rightarrow H_2(M_n; \mathbb{Z}/2\mathbb{Z})$.

It therefore follows from Lemma 11.1 that we may replace S_i with Σ_i in S and get a new sphere basis. By minimality, it follows that this can only increase or maintain the total length of the words that $\mathbf{a}$ represents.

After performing this replacement for all $1 \leq i \leq k$ we arrive at $\mathbf{b}$, and so none of the steps increased the length of the words w represents.

We now note that each substitution is equivalent to performing a T -transformation on the elements represented by w as, up to isotopy, we may take Σ_i to be a small shell around the spheres S_j which it contains, along with thin tubes not intersecting w between them, and this representative of Σ_i clearly has the effect of a T -transformation.

The result is therefore proven in this case, as the transformation is a level transformation since we have already remarked that it does not increase the length of w .

We now consider the general case, where S and Σ are in normal form (which we may do as the process of normalising only decreases the number of components of intersection) but may intersect in some curves. We induct on the number of components in the intersection as in Proposition 11.3.

We work in $R = M_n \setminus S$, defining π as before, and fix a Σ_i . If this has no intersection with S (or is contained within S), then we are done with this Σ_i . Else we take an inner-most intersection curve c of Σ_i with S , say on the component S_j .

As the intersection is inner-most, one of the two parts into which c divides Σ_i has no intersection with S . We label this part A . Now label the two parts that c divides S_j into as E and E' . Either $A \cup E$ or $A \cup E'$ separates one preimage of S_j from the other preimage of S_j under π . Without loss of generality, it is $A \cup E$. By Lemma 11.1, we may therefore use a T -transformation to replace S_j with $E \cup A$. If this is a level transformation, we are done by induction.

Suppose it is not a level transformation, so that w has strictly more intersections with $E \cup A$ than with S_j . Now let $\tilde{\Sigma}_i$ be the preimage under $M_n \rightarrow M_n \setminus \Sigma$ incident with E . We wish to prove more than half of the intersections of w with Σ_i lie in A . If so, then the same process with an inner-most curve on the other side of c must give a level transformation, so we are done by induction. As $S_j = E \cup E'$, that w has strictly more intersections with $E \cup A$ than with S_j is equivalent to saying it has more

intersections with A than with E' , and therefore by choice of E , has more intersections with A than A^c , so we are done. $\square$

Remark 11.2. We will not provide an example of the algorithm here, and will instead give an idea of the time complexity to indicate why.

We assume that computation of the lengths of words takes negligible time compared to generating and performing the T -transformations, as indeed is true if one performs it by hand. Under this assumption, there are $n! \cdot 2^n \cdot 4n(n-1)$ different T transformations. The $n!$ comes from reordering the basis, the 2^n from the possible choices of orientation of the spheres, and the $4n(n-1)$ for the number of transformations of ways to apply simple transformations of types 1 and 2.

Even for $n = 2$, we see that one step would require computing 64 different T -transformations. Using this, we immediately see that if it takes k steps to show that $\mathbf{a}$ and $\mathbf{b}$ are equivalent, then this requires the application of $(n! \cdot 2^n \cdot 4n(n-1))^k$ different transformations.

It is immediately clear that the presentation of any non-trivial examples in this essay would therefore involve an unreasonable amount of computing power and space.

12 Conclusion

In this conclusion, we will briefly note other possible views on Outer Space, other than that seen through 3-manifold theory. In [4], a much more combinatorial view is taken, following arguments closely related to those touched on in our discussion of the degree theorem, and in Section 5. In [3] an alternate, more geometric view on $\text{Out}(F_n)$ largely obtained from taking the universal cover of the graphs in X_n is taken, and this generalises more easily to other groups, by considering trees on which those groups act. The spaces constructed are not in general nice to work with, however when they are, they may allow us to generalise some of our results from this essay.

References

- [1] M. Arkowitz. *Introduction to homotopy theory / Martin Arkowitz*. Universitext. New York ; Springer, 2011.
- [2] Kenneth S. Brown. *Cohomology of Groups / by Kenneth S. Brown*. 1st ed. 1982. Graduate Texts in Mathematics, 87. New York, NY: Springer New York, 1982.
- [3] Matt Clay. “Contractibility of deformation spaces of G-trees”. *Algebraic & geometric topology* 5.4 (2005), pp. 1481–1503.
- [4] Marc Culler and Karen Vogtmann. “Moduli of graphs and automorphisms of free groups”. *Inventiones mathematicae* 84.1 (1986), pp. 91–119.
- [5] Marc Culler et al. “Finite groups of outer automorphisms of a free group”. *Contributions to Group Theory*. Vol. 33. Contemporary Mathematics. Providence, Rhode Island: American Mathematical Society, 1984, pp. 197–207.
- [6] Richard Z. Goldstein and Edward C. Turner. “Automorphisms of free groups and their fixed points”. *Inventiones mathematicae* 78.1 (1984), pp. 1–12.
- [7] Allen Hatcher. “Homological stability for automorphism groups of free groups”. *Commentarii mathematici Helvetici* 70.1 (1995), pp. 39–62.
- [8] Allen Hatcher and Karen Vogtmann. “Cerf theory for graphs”. *Journal of the London Mathematical Society* 58.3 (1998), pp. 633–655.
- [9] Allen Hatcher and Karen Vogtmann. “Homology stability for outer automorphism groups of free groups”. *Algebraic & geometric topology* 4.2 (2004), pp. 1253–1272.
- [10] Allen Hatcher, Karen Vogtmann, and Nathalie Wahl. “Erratum to: Homology stability for outer automorphism groups of free groups”. *Algebraic & geometric topology* 6.2 (2006), pp. 573–579.
- [11] Allen Hatcher and Nathalie Wahl. “Stabilization for the automorphisms of free groups with boundaries”. *Geometry & topology* 9.3 (2005), pp. 1295–1336.
- [12] François. Laudenbach. *Topologie de la dimension trois : homotopie et isotopie / François Laudenbach*. Astérisque ; 12. Paris: Société mathématique de France, 1974.
- [13] Bruno Martelli. *An Introduction to Geometric Topology*. 2016.
- [14] Daniel Quillen. “Homotopy properties of the poset of nontrivial p-subgroups of a group”. *Advances in mathematics (New York. 1965)* 28.2 (1978), pp. 101–128.
- [15] Jean Pierre Serre. *Trees*. Berlin ; New York: Springer-Verlag, 1980.
- [16] J. H. C. Whitehead. “On Equivalent Sets of Elements in a Free Group”. *Annals of mathematics* 37.4 (1936), pp. 782–800.