Homological stability for automorphism groups of free groups.

Autor(en): **Hatcher, Allen**

Objekttyp: **Article**

Zeitschrift: **Commentarii Mathematici Helvetici**

Band (Jahr): **70 (1995)**

PDF erstellt am: **17.07.2024**

Persistenter Link: https://doi.org/10.5169/seals-52990

Comment. Math. Helvetici 70 (1995) 39-62
0010-2571/95/010039-24\$1.50 + 0.20/0

Homological stability for automorphism groups of free groups

ALLEN HATCHER

Let F_n be a free group on n generators, $\mathrm{Aut}(F_n)$ its group of automorphisms, and $\mathrm{Out}(F_n)$ its outer automorphism group, the quotient of $\mathrm{Aut}(F_n)$ by inner automorphisms. There has been much progress of late in the study of these groups via the one-dimensional model which arises from regarding F_n as the fundamental group of a graph; see e.g., [BH] and [CV]. In this paper we return to a three-dimensional model first used by J. H. C. Whitehead in the 1930's, which involves looking at embedded 2-spheres in a connected sum of $S^1 \times S^2$'s. Refining Whitehead's techniques and applying subsequent results of Laudenbach, we use this three-dimensional model to prove:

THEOREM. (a) *The map* $H_i(\mathrm{Aut}(F_n)) \to H_i(\mathrm{Aut}(F_{n+1}))$ *induced by the natural inclusion* $\mathrm{Aut}(F_n) \subset \mathrm{Aut}(F_{n+1})$ *is an isomorphism for* $n > i^2/4 + 2i - 1$.

(b) *The map* $H_i(\mathrm{Aut}(F_n)) \to H_i(\mathrm{Out}(F_n))$ *induced by the projection* $\mathrm{Aut}(F_n) \to \mathrm{Out}(F_n)$ *is an isomorphism for* $n > i^2/4 + 5i/2$.

Here the inclusion $\mathrm{Aut}(F_n) \subset \mathrm{Aut}(F_{n+1})$ is obtained by extending an automorphism of F_n to an automorphism of F_{n+1} fixing the $(n+1)^{\mathrm{st}}$ basis element.

Homological stability results like the statement in (a) have a long history, going back to the classical Lie groups and continuing more recently with various discrete groups, including symmetric groups, braid groups, many linear groups, and mapping class groups of surfaces; see the References at the end of the paper. In all these earlier cases the function $\varphi(i)$ giving the stable dimension range $n > \varphi(i)$ is linear, so it would be rather surprising if our φ cannot be improved to be linear as well.

Letting Aut_∞ denote the union of the groups $\mathrm{Aut}(F_n)$ under the inclusions $\mathrm{Aut}(F_n) \subset \mathrm{Aut}(F_{n+1}) \subset \cdots$, the theorem implies that the stable groups $H_i(\mathrm{Aut}_\infty) = \varinjlim H_i(\mathrm{Aut}(F_n))$ are finitely generated since $H_i(\mathrm{Out}(F_n))$ is finitely generated for all i and n, by [CV]. (This can also be proved using the techniques of the present paper, as we show in an Appendix.) Other than finite generation, the main thing known about the homology groups of Aut_∞ is that they contain quite a lot of torsion. This comes from the infinite symmetric group $\Sigma_\infty = \bigcup_n \Sigma_n$, which can be viewed as the subgroup of Aut_∞ consisting of automorphisms permuting the basis elements.

PROPOSITION. *The inclusion $\Sigma_\infty \subset \mathrm{Aut}_\infty$ induces an isomorphism of $H_i(\Sigma_\infty)$ onto a direct summand of $H_i(\mathrm{Aut}_\infty)$, for all i.*

This is easily deduced from a theorem of Waldhausen, as follows. Let $\bigvee S^k$ be the wedge sum of an infinite sequence of k-spheres and let $\mathrm{Aut}(\bigvee S^k)$ be the H-space of homotopy equivalences $\bigvee S^k \to \bigvee S^k$ which are the identity on all but finitely many of the spheres. Suspension provides a diagram

$$\begin{array}{ccccc} \mathrm{Aut}(\bigvee S^0) & \longrightarrow & \mathrm{Aut}(\bigvee S^1) & \longrightarrow & \varinjlim_k \mathrm{Aut}(\bigvee S^k) \\ \| & & \downarrow \simeq & & \\ \Sigma_\infty & \longrightarrow & \mathrm{Aut}_\infty & & \end{array}$$

where the vertical map in the middle, taking the induced homomorphism on π_1, is a homotopy equivalence since $\bigvee S^1$ is a $K(F_\infty, 1)$. Now if we apply the functor B^+, B denoting classifying space and "+" denoting the Quillen plus construction, then $\varinjlim_k \mathrm{Aut}(\bigvee S^k)$ becomes the space $A(*)$ of Waldhausen, who showed in [$\mathrm{W}_{1,2}$] that $A(*)$ splits up to homotopy as the product of $B^+\Sigma_\infty$ with another space (essentially the stable smooth pseudo-isotopy space for a disk), the inclusion of the factor $B^+\Sigma_\infty$ being induced by the map across the top of the diagram above. Hence we have splittings $H_i(\mathrm{Aut}_\infty) \approx H_i(\Sigma_\infty) \oplus (?)$.

The complementary summand of $H_i(\Sigma_\infty)$ in $H_i(\mathrm{Aut}_\infty)$ is zero for $i = 1, 2$ (see [G] for the latter case), but whether it is trivial for all i seems to be unknown. In particular, it is not known whether $\tilde{H}_*(\mathrm{Aut}_\infty; \mathbb{Q})$ vanishes.

Our approach to homological stability for $\mathrm{Aut}(F_n)$ and $\mathrm{Out}(F_n)$ is modeled on Harer's proof in [H] for the case of mapping class groups of surfaces, as refined by Ivanov [I]. In our case, we consider the mapping class groups of the 3-manifolds $M_{n,s}$ obtained from the connected sum of n copies of $S^1 \times S^2$ by deleting the interiors of s disjoint 3-balls. Thus $M_{n,s}$ is $M_{n,0}$ with s punctures. When $s = 1$ the puncture provides a basepoint for $M_{n,1}$ and we have a natural homomorphism from the mapping class group $\pi_0 \mathrm{Diff}^+(M_{n,1})$ to $\mathrm{Aut}(\pi_1(M_{n,1})) = \mathrm{Aut}(F_n)$. This is surjective since the classical Nielsen generators of $\mathrm{Aut}(F_n)$ are easily realized by orientation-preserving diffeomorphisms of $M_{n,1}$. And by a theorem of Laudenbach [L_2], the kernel of $\pi_0 \mathrm{Diff}^+(M_{n,1}) \to \mathrm{Aut}(F_n)$ is a finite direct sum of $\mathbb{Z}_2$'s generated by "Dehn twists" – diffeomorphisms supported in a product $S^2 \times I \subset M_{n,1}$, each sphere $S^2 \times \{t\}$ being taken to itself by a rotation through angle $2\pi t$. In the case of the unpunctured manifold $M_{n,0}$ there are corresponding statements with $\mathrm{Aut}(F_n)$ replaced by $\mathrm{Out}(F_n)$.

Associated to $M_{n,s}$ is a simplicial complex $\mathbb{S}(M_{n,s})$ whose k-simplices are isotopy classes of systems of $k+1$ disjointly embedded 2-spheres in $M_{n,s}$, none of which bounds a ball or is isotopic to a boundary sphere of $M_{n,s}$, and no two of which are isotopic. The mapping class group of $M_{n,s}$ acts on $\mathbb{S}(M_{n,s})$, and from Laudenbach's theorem that homotopic 2-spheres in $M_{n,s}$ are isotopic (see [$L_{1,2}$]) it follows easily that twists along 2-spheres act trivially on $\mathbb{S}(M_{n,s})$. Hence there is an induced action of the quotient group $\Gamma_{n,s} = \pi_0\,\mathrm{Diff}^+(M_{n,s})/T$, T being the subgroup generated by twists, a normal subgroup. Note that $\Gamma_{n,1} \approx \mathrm{Aut}(F_n)$ and $\Gamma_{n,0} \approx \mathrm{Out}(F_n)$ by Laudenbach's theorem mentioned in the previous paragraph.

There are two parallel spectral sequence arguments for proving stability of $H_i(\Gamma_{n,s})$ with respect to n and s. The spectral sequences in question are the natural ones associated to the actions of $\Gamma_{n,s}$ on two $\Gamma_{n,s}$-invariant subcomplexes of $\mathbb{S}(M_{n,s})$, consisting of systems with connected complement in one case and systems of spheres which separate off a simply-connected submanifold in the other case. (Technically, in the first case one must consider systems which are ordered and oriented.) The input needed to deduce stability from the spectral sequences is that the subcomplexes and their quotients by the action are highly connected. For the quotients this is easy, and for the subcomplexes it reduces as in [H] to showing high connectivity of the ambient complex $\mathbb{S}(M_{n,s})$.

The main technical result in the paper is that $\mathbb{S}(M_{n,s})$ is contractible if $n > 0$. This is proved by imitating the simple proof in [Hat] of contractibility of the analogous complex of arcs on a punctured surface. However, for this scheme to work one needs the fact that sphere systems can be isotoped into a fairly canonical normal form with respect to a decomposition of $M_{n,s}$ into "pairs of pants," i.e., 3-punctured S^3's. This normal form, which in retrospect seems such an obvious analog of a well-known property of curves on a surface, is the main new idea in the paper. It relies heavily on Laudenbach's homotopy-implies-isotopy theorem for 2-spheres in $M_{n,s}$.

The complex $\mathbb{S}(M_{n,0})$ is closely related to Culler–Vogtmann's "Outer Space" [CV]. As we show in the Appendix, Outer Space can be identified with the dense subspace of $\mathbb{S}(M_{n,0})$ which is the complement of the subcomplex consisting of sphere systems having at least one nonsimply-connected complementary component. The contraction we construct for $\mathbb{S}(M_{n,0})$ restricts to a contraction of Outer Space, giving an alternative proof of the main technical result in [CV]. We also describe in the Appendix how Culler–Vogtmann's calculation of the virtual cohomological dimension of $\mathrm{Aut}(F_n)$ and $\mathrm{Out}(F_n)$ can be rederived via sphere complexes.

The advantage of sphere complexes over Outer Space for the purposes of the present paper is that the punctured manifolds $M_{n,s}$ can be treated just as easily as $M_{n,0}$. In particular, $\mathrm{Aut}(F_n) = \Gamma_{n,1}$ can be handled as well as $\mathrm{Out}(F_n) = \Gamma_{n,0}$. But more importantly from a technical standpoint, punctures are needed in order to identify the simplex-stabilizers for the actions of $\Gamma_{n,0}$ on $\mathbb{S}(M_{n,0})$ and $\Gamma_{n,1}$ on $\mathbb{S}(M_{n,1})$. In view of the connection between Outer Space and actions of F_n on $\mathbb{R}$-trees, it is natural to suspect that there should be a theory of "$\mathbb{R}$-trees with punctures" which would allow homological stability for $\mathrm{Aut}(F_n)$ to be proved without using any 3-manifold topology.

We might remark that contractibility of $\mathbb{S}(M_{n,s})$ contrasts strongly with the fact that the actual space of all smooth 2-spheres in $M_{n,s}$ has a rather complicated homotopy type, with non-finitely generated fundamental group for $n > 1$ for example [M]. In [HM] this complication was an obstruction to proving more than simple-connectivity of sphere complexes without the idea of normal form.

Whitehead's method of studying automorphisms of free groups via 2-spheres in a connected sum of $S^1 \times S^2$'s has also been used profitably in [$\mathrm{GT}_{1,2}$].

It is a pleasure to thank Karen Vogtman for numerous stimulating conversations about this material, without which the paper would probably not have been written.

1. Normal form for sphere systems

Let M be the manifold $M_{n,s}$, the connected sum of n $S^1 \times S^2$'s with s punctures. By a *system of 2-spheres* in M we shall mean a finite collection of disjointly embedded smooth 2-spheres $S_i \subset M$ none of which bounds a ball or is isotopic to a boundary sphere of M, and no two of which are isotopic, i.e., form the boundary of an $S^2 \times I \subset M$.

We shall need Laudenbach's theorem that homotopic systems of 2-spheres in M are isotopic. This is proved in [L_1] for single spheres, and the extension to systems, which is easy, is described in [L_2] (see in particular page 83 and Lemma V.4.2). Actually, in this section and the next we shall need only a rather easy special case of Laudenbach's theorem; this is explained in the Remark at the end of this section.

Let $\Sigma = \bigcup_j \Sigma_j$ be a fixed maximal system. Splitting M along Σ then produces a collection of 3-punctured spheres P_k. A system $S = \bigcup_i S_i$ is said to be in *normal form* with respect to Σ if each S_i either coincides with a sphere Σ_j or meets Σ transversely in a nonempty collection of circles splitting S_i into components called *pieces*, such that the following two conditions hold in each P_k:

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

From (1) it follows that the pieces are either disks, cylinders, or pairs of pants. A cylinder piece connects two components of ∂P_k and a pants piece connects all three components of ∂P_k. A disk piece has boundary on one component of ∂P_k and separates the other two components of ∂P_k, by (2). Hence a P_k cannot contain both disk and pants pieces, and all the disk pieces in a P_k must have their boundaries on the same component of ∂P_k. Each individual cylinder or pants piece in a P_k must be unknotted in P_k since its boundary circles lie on different components of ∂P_k, but a collection of cylinder and pants pieces can be knotted and linked in a complicated fashion. However, since homotopic systems are isotopic, such knotting and linking can always be eliminated by an isotopy of the system in M, though the isotopy will generally have to move outside P_k.

PROPOSITION 1.1. *Every system $S \subset M$ can be isotoped into normal form with respect to Σ.*

Proof. We may assume S has been isotoped so that the maximum number of S_i's coincide with Σ_j's and the remaining S_i's intersect Σ transversely. We shall show that if S is then not in normal form, it can be isotoped to decrease the number of circle components of $S \cap \Sigma$. We may suppose there are no circles bounding trivial disks as in condition (2). Also there can be no S_i's disjoint from Σ since these must be parallel to Σ_j's by the maximality of Σ.

We may perform a sequence of surgeries on all the spheres S_i which meet Σ transversely, surgering along the circles of $S \cap \Sigma$, innermost circles in Σ first, then innermost remaining circles, etc., to produce a collection of disjoint spheres S_{im} in $M - \Sigma$. Inverse to this sequence of surgeries is a sequence of tubing operations, joining the S_{im}'s by disjoint but possibly nested cylinders $S^1 \times I$ crossing the Σ_j's, eventually reconstructing S. If S is not in normal form, then there is a P_k containing a sphere S_{im} from which two tubes run across the same component of ∂P_k. Let T_1 be the first of these tubes to be attached to S_{im}, T_2 the second. Isotope S by dragging the end of T_2 attached to S_{im} first along S_{im} to where T_1 attaches, avoiding other tubes along the way, then dragging this end of T_2 along T_1 across Σ. If there are tubes passing inside T_2, these are to be dragged along too.

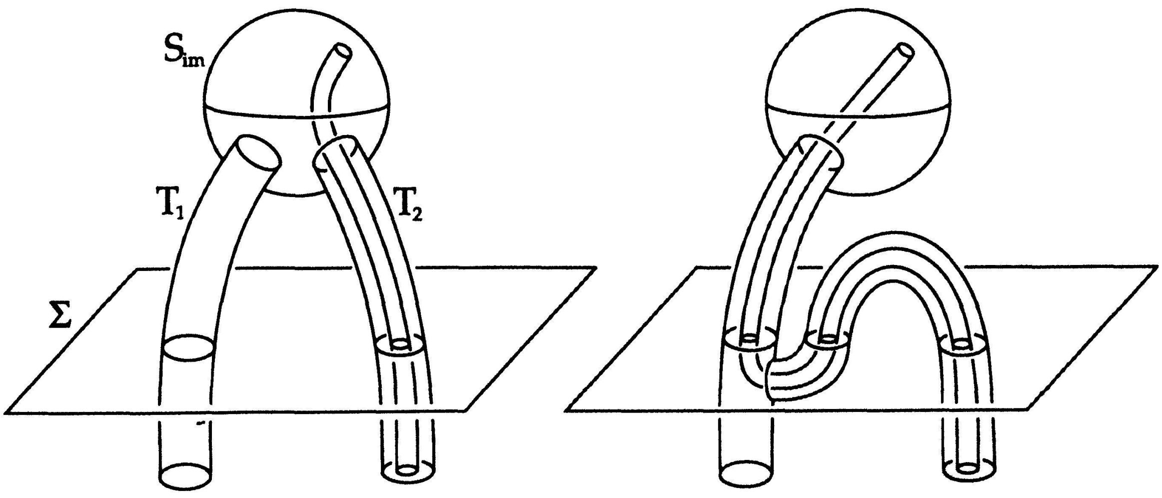

If there are r tubes inside T_2, the net effect of this isotopy is to increase the number of circles of $S \cap \Sigma$ by $2r + 1$: one from T_2 and two for each tube inside T_2, one of these two being inside T_2 and the other near T_1, the result of dragging the tube along with T_2. After this isotopy, the part of T_2 in P_k is a cylinder which can be homotoped into Σ, keeping its ends fixed. Hence by a homotopy of T_2 and the r tubes inside it we can produce a new sphere system S' homotopic and therefore isotopic to S, meeting Σ in $2r + 2$ fewer circles. Comparing S' with the original S, we have decreased by 1 the number of circles of $S \cap \Sigma$ if S is not in normal form. □

Two sphere systems S and S' in normal form with respect to Σ will be called *equivalent* if there is a homotopy $h_t\colon S \to M$ from S to S' with the following properties:

(1) On the S_i's which coincide with Σ_j's, h_t is the identity.
(2) On the other S_i's, h_t remains transverse to Σ for all t, and $h_t(S) \cap \Sigma$ varies only by isotopy in Σ. In particular, the circle components of $h_t(S) \cap \Sigma$ stay disjoint for all t.

Systems which are isotopic through normal form systems are equivalent, but the relation of equivalence is more general than this, allowing cylinder and pants pieces to unknot and unlink in each P_k.

Equivalent sphere systems are isotopic since they are homotopic. Conversely:

PROPOSITION 1.2. *Isotopic sphere systems in normal form are equivalent.*

Proof. The idea is to look in the universal cover $\tilde{M}$, where oriented spheres determine well-defined elements of $\pi_2(\tilde{M}) \approx H_2(\tilde{M})$.

Let $\tilde{\Sigma}$ be the preimage of Σ in $\tilde{M}$. To the pair $(\tilde{M}, \tilde{\Sigma})$ we may associate a graph T by first taking a letter "Y" for each lift $\tilde{P}_k$, the three endpoints of the Y correspnding to the three boundary spheres of $\tilde{P}_k$, then identifying the endpoints of different Y's which correspond to the same sphere of $\tilde{\Sigma}$. Since $\tilde{M}$ is simply-connected, each component of $\tilde{\Sigma}$ separates $\tilde{M}$, hence T is a tree. The endpoints of Y's corresponding to components of $\tilde{\Sigma}$ will be regarded as valence-two vertices of T.

For a sphere S in normal form, choose a lift $\tilde{S}$ to $\tilde{M}$. We observe first that for each $\tilde{P}_k$, $\tilde{S} \cap \tilde{P}_k$ has at most one component. This is trivially true if $\tilde{S}$ is a component of $\tilde{\Sigma}$, and when $\tilde{S}$ is transverse to $\tilde{\Sigma}$ it can be seen by considering the dual tree $T(\tilde{S})$ to $\tilde{S} \cap \tilde{\Sigma}$ in $\tilde{S}$, having a vertex for each component of $\tilde{S} - \tilde{\Sigma}$ and an edge for each circle of $\tilde{S} \cap \tilde{\Sigma}$. The inclusion $\tilde{S} \subset \tilde{M}$ induces a natural map $T(S) \to T$ which is locally injective, hence globally injective, and this implies that $\tilde{S} \cap \tilde{P}_k$ has at most one component, for each $\tilde{P}_k$.

Viewing $T(\tilde{S})$ as a subtree of T, it will be convenient to enlarge $T(\tilde{S})$ slightly so that it is the union of the Y's corresponding to $\tilde{P}_k$'s meeting $\tilde{S}$. Thus a disc piece of $\tilde{S}$ corresponds to an extremal Y of $T(\tilde{S})$, connected at only one of its three ends to the rest of $T(\tilde{S})$; a cylinder piece of $\tilde{S}$ corresponds to a Y in $T(\tilde{S})$ connected to the rest of $T(\tilde{S})$ at two of its ends; and a pants piece of $\tilde{S}$ corresponds to a Y connected to the rest of $T(\tilde{S})$ at all three of its ends. In particular, $T(\tilde{S})$ cannot consist of a single Y. For completeness, when $\tilde{S}$ is a component of $\tilde{\Sigma}$ we define $T(\tilde{S})$ to be the corresponding vertex of T.

Choosing a transverse orientation for $\tilde{S}$, then if $T(\tilde{S})$ is not a point, each end vertex of $T(\tilde{S})$ can be labelled $+$ or $-$ according to whether the corresponding sphere of $\tilde{\Sigma} \cup \partial\tilde{M}$ lies on the $+$ or the $-$ side of $\tilde{S}$ with respect to the chosen transverse orientation. Since $\pi_1(\tilde{M}) = 0$, $\tilde{S}$ separates $\tilde{M}$ and these sides are well-defined. For an extremal Y of $T(\tilde{S})$, connected to the rest of $T(\tilde{S})$ at only one of its three ends, the other two free ends have opposite sign. The non-extremal Y's of $T(\tilde{S})$ have either one free end, labelled $+$ or $-$, or no free ends.

Let $\tilde{S}_+$ be the union of the spheres of $\tilde{\Sigma} \cup \partial\tilde{M}$ corresponding to $+$ ends of $T(\tilde{S})$ and let $\tilde{S}_-$ be the union of the spheres corresponding to $-$ ends. The transverse orientation for $\tilde{S}$ induces transverse orientations for the spheres of $\tilde{S}_\pm$, in such a way that $\tilde{S}$ is homologous in $\tilde{M}$ to $\tilde{S}_+$ and also to $\tilde{S}_-$. Observe that $\tilde{S}_+$ determines $T(\tilde{S})$, namely, $T(\tilde{S})$ is the smallest connected union of Y's in T containing the vertices corresponding to $\tilde{S}_+$. This is because each extremal Y of $T(\tilde{S})$ contains both a $+$ and a $-$ vertex. Since $\tilde{S}_+$ determines $T(\tilde{S})$, it also determines $\tilde{S}_-$.

The homology class of $\tilde{S}$ in $\tilde{M}$ determines $\tilde{S}_+$ uniquely. To see this, suppose we have another sphere $\tilde{S}'$ lifting a normal form sphere S', with $\tilde{S}'$ homologous to $\tilde{S}$. For example, $\tilde{S}'$ might be obtained from $\tilde{S}$ by lifting an isotopy from S to a normal form sphere S'. If $T(\tilde{S}') \neq T(\tilde{S})$, there would be a sphere of $\tilde{\Sigma}$ separating $\tilde{M}$ into two parts, one of which, $\tilde{M}_0$, contained all of $\tilde{S}'_+$, say, and all but exactly one sphere of $\tilde{S}_+$. Then the classes of $\tilde{S}$ and $\tilde{S}'$ in $H_2(\tilde{M})$ would have distinct images in $H_2(\tilde{M}, \tilde{M}_0)$, zero in one case and nonzero in the other. So $T(\tilde{S}') = T(\tilde{S})$. If $\tilde{S}$ and $\tilde{S}'$ produced different choices of signs for the ends of $T(\tilde{S}) = T(\tilde{S}')$ then we could choose $\tilde{M}_0$ as the part of $\tilde{M}$ on one side of a sphere of $\tilde{\Sigma}$ in such a way that there was exactly one different choice of signs outside $\tilde{M}_0$, and again $\tilde{S}$ and $\tilde{S}'$ would give different classes in $H_2(\tilde{M}, \tilde{M}_0)$; there are two subcases here, according to whether the different choice of sign occurs in an extremal Y of $T(\tilde{S})$ or not.

From the preceding we can deduce that $\tilde{S}$ and $\tilde{S}'$ are isotopic staying transverse to $\tilde{\Sigma}$ at all times. Namely, the tree $T(\tilde{S}) = T(\tilde{S}')$ determines the combinatorial pattern of pieces of $\tilde{S}$ and $\tilde{S}'$, so there is an isotopy of the pair $(\tilde{M}, \tilde{\Sigma})$ which deforms $\tilde{S}'$ so that $\tilde{S}' \cap \tilde{\Sigma}$ coincides with $\tilde{S} \cap \tilde{\Sigma}$ as a collection of transversely oriented circles in $\tilde{\Sigma}$. Then $\tilde{S}'$ can be isotoped to $\tilde{S}$, fixing $\tilde{S}' \cap \tilde{\Sigma}$ and without introducing new intersections with $\tilde{\Sigma}$, since the isotopy class rel boundary of a transversely oriented piece in a $\tilde{P}_k$ is uniquely determined by its oriented boundary and, in the case of disk and cylinder pieces, the data of which side the components of $\partial\tilde{P}_k$ disjoint from it lie on, data specified by the signs on $T(\tilde{S})$. In particular there is no problem of knotting or linking since $\tilde{S}$ and $\tilde{S}'$ each contain at most one piece in a $\tilde{P}_k$.

Now we consider a normal form system S having any number of components. Choosing transverse orientations for the spheres of S, it makes sense to say that one piece of S lies on the positive or negative side of another piece in the same P_k. If S' is another normal form system isotopic to S, the isotopy transfers the transverse orientations from S to S'. The considerations of the preceding paragraphs give a bijective correspondence between the pieces of S and S'. We assert:

(*) If s_p and s_q are pieces of S in P_k and s'_p and s'_q are the corresponding pieces of S', then s'_q is on the same side of s'_p as s_q is of s_p.

To see this, consider a lift $\tilde{P}_k$ of P_k in $\tilde{M}$. This contains lifted pieces $\tilde{s}_p, \tilde{s}_q$, etc. Let $\tilde{s}_p$ and $\tilde{s}_q$ be contained in lifts $\tilde{S}_i$ and $\tilde{S}_j$ of components of S, with $\tilde{s}'_p$ and $\tilde{s}'_q$ contained in the corresponding lifts $\tilde{S}'_i$ and $\tilde{S}'_j$. We may isotope $\tilde{S}'_i$ to $\tilde{S}_i$ staying transverse to $\tilde{\Sigma}$ at all times, as noted above. This isotopy extends to an isotopy of the pair $(\tilde{M}, \tilde{\Sigma})$ so we may assume $\tilde{S}_i = \tilde{S}'_i$. If the assertion (*) were false we would have $\tilde{S}_j$ and $\tilde{S}'_j$ lying on opposite sides of $\tilde{S}_i$. This evidently implies that $\tilde{S}_i$ and $\tilde{S}_j$ have the same trees $T(\tilde{S}_i)$ and $T(\tilde{S}_j)$ with the same labels on end vertices, so $\tilde{S}_i$ and $\tilde{S}_j$ would be isotopic. If these two spheres were lifts of different components of S, these components would be homotopic, hence isotopic, contrary to the definition of a sphere system. On the other hand, if $\tilde{S}_i$ and $\tilde{S}_j$ were lifts of the same sphere of S, there would be a deck transformation carrying one to the other, hence preserving the subtree $T(\tilde{S}_i) = T(\tilde{S}_j)$ of T, but this is impossible since a nontrivial deck transformation cannot leave a finite subtree of T invariant.

For the isotopic normal form systems S and S' the tree data gives a bijective correspondence between the transversely-oriented circles of $S \cap \Sigma$ and $S' \cap \Sigma$ in Σ. By (*) the "side" relations for these two systems of circles in Σ also correpsond. It is elementary to see that the isotopy class of a transversely-oriented circle system in a sphere is determined by these side relations, so by an isotopy of (M, Σ) we may deform S' so that corresponding transversely-oriented circles of $S \cap \Sigma$ and $S' \cap \Sigma$ are equal. As noted earlier when we were working in $\tilde{M}$, the individual pieces of S' can then be isotoped rel boundary within each P_k to make $S' = S$, and all these isotopies together give an equivalence of S' with S. □

Remark. The proof of Proposition 1.1 used only a special case of Laudenbach's theorem that homotopic systems are isotopic, namely, the case of a homotopy which consists of "passing one nested family of tubes through another," by which we mean the following. In knot theory there is the familiar operation of changing an overcrossing to an undercrossing by a regular homotopy which passes one strand of the knot through another strand. Imagine this operation as being performed on a tubular neighborhood of the knot. Then on the boundary of the tubular neighborhood the operation amounts to "passing one tube through another tube." More generally, each of these two tubes could contain a number of smaller tubes in its interior, possibly nested in some fashion, but just running straight through the outer tube. Then the crossing change yields the general case of "passing one nested family of tubes through another." In the absence of nesting it is obvious how to realize this homotopy of a sphere system by an isotopy. (This is essentially the "light bulb trick" – unknotting the cord from which a light bulb is suspended by slipping the knot over the bulb.) The general case reduces to this special case by a nice induction argument given on pp. 71–73 of [L_1].

An examination of the proof of Proposition 1.2 shows that the relation of "equivalence" for normal form systems could have been defined more restrictively as isotopy through normal form systems together with operations of passing one nested family of tubes through another in individual P_k's. Using this observation it would not be difficult to use the machinery of the proof of 1.2 to prove Laudenbach's theorem that homotopic systems in $M_{n,s}$ are isotopic, though this proof would probably not be appreciably simpler than the one in [L_2].

2. Contractibility of the full sphere complex

For the manifold $M = M_{n,s}$ define the simplicial complex $\mathbb{S}(M)$ to have as its k-simplices the isotopy classes of systems of $k+1$ spheres in M. The $(k-1)$-dimensional faces of such a simplex are obtained by deleting the various spheres of the system, one at a time. The fact that $\mathbb{S}(M)$ is indeed a simplicial complex, i.e., that simplices are uniquely determined by their vertices, follows immediately from Laudenbach's theorem that homotopic systems are isotopic. The maximal simplices of $\mathbb{S}(M)$ all have the same dimension, namely $3n+s-4$, as one sees by Euler characteristic considerations using the fact that the complementary regions of a maximal system are all 3-punctured spheres.

THEOREM 2.1. *$\mathbb{S}(M_{n,s})$ is contractible if $n>0$.*

In the important cases $s \leq 1$ the contraction will be obtained via a piecewise linear flow on $\mathbb{S}(M)$ which shrinks $\mathbb{S}(M)$ onto a given maximal simplex.

Proof. First we do the cases $s \leq 1$. Consider a system S in normal form with respect to the fixed maximal system Σ. Let (t_i) be the barycentric coordinates of a point in the interior of the simplex of $\mathbb{S}(M)$ defined by S. These can always be normalized to have sum 1. We think of the numbers $t_1 > 0$ as weights on the spheres S_i in S. It is then convenient to replace each S_i by a family $S_i \times [0, t_i]$ of parallel copies of S_i of "thickness" t_i. Then when a weight t_i goes to zero at a face of the simplex, the family $S_i \times [0, t_i]$ shrinks to thickness zero and is deleted. It is convenient also to allow the operation of splitting the family $S_i \times [0, t_i]$ into several parallel families of total thickness t_i, as well as the inverse operation of glueing parallel families together, adding their weights.

From the weighted system S we construct a collection of finite trees T_j, one for each sphere Σ_j which S meets transversely. The vertices of T_j are the components of $\Sigma_j - S$ and the edges are the circles of $S \cap \Sigma_j$. The weights on the S_i's define lengths for these edges, so T_j is a metric tree.

There is a canonical way to shrink each T_j to a point by shortening all extremal edges simultaneously at unit speed. Once an extremal edge has disappeared one continues shrinking all remaining extremal edges. We will show that this shrinking, performed on all the T_j's at once, can be lifted to a path in $\mathbb{S}(M)$ starting with the given weighted system S and ending with a system disjoint from Σ, hence in the simplex determined by Σ.

The extremal vertices v of T_j correspond to disjoint disks $D_v \subset \Sigma_j$ with $D_v \cap S = \partial D_v$. One can use these disks D_v to surger S to a new system in which the circles ∂D_v have been eliminated from $S \cap \Sigma$. Taking the weights into account, one gradually surgers through the appropriate families $S_i \times [0, t_i]$ at unit speed, decreasing the weights of these families while increasing the weights of the new families created by the surgery. The old and new families can be taken to be disjoint, so the surgery can be viewed as simply transferring weights from the old family to the new. It may happen that a family $S_i \times [0, t_i]$ is being surgered by disks on both the $S_i \times \{0\}$ and $S_i + \{t_i\}$ sides simultaneously, which just means that the thickness of this family is decreasing twice as fast. When the thickness of a family $S_i \times [0, t_i]$ has shrunk to zero, this family is deleted and one continues the surgery process on the remaining sphere families.

Let us see what this surgery process does in a single P_k. Surgery on a disk piece produces a sphere isotopic to a component of ∂P_k. Surgery on one end of a cylinder piece produces a disk, and if this disk does not separate two components of ∂P_k then the other end of the cylinder would have to be surgered simultaneously, producing a trivial sphere bounding a ball in P_k. Surgery on one end of a pants piece produces a cylinder joining the other two components of ∂P_k. If an end of this cylinder is being surgered at the same time, we are effectively in the case of a cylinder piece, already considered. We conclude: The surgery process produces spheres which are either in normal form or trivial, bounding a ball or isotopic to a puncture. Moreover, the trivial spheres are disjoint from Σ so are not involved in further surgeries and can be deleted without affecting the surgery process. The hypothesis $s \le 1$ guarantees that not everything is deleted. This is because a normal form sphere transverse to Σ has at least two disk pieces, and a single surgery can make only one of these a trivial sphere since a disk piece in P_k only becomes trivial by becoming isotopic to a component of ∂P_k in ∂M.

Thus the surgery process on a given weighted system S defines a path $S(t)$ in $\mathbb{S}(M)$ starting at $S = S(0)$ and ending in the simplex Σ. This path is obviously continuous, and in fact piecewise linear since the thicknesses of the families of spheres in $S(t)$ vary linearly with the time parameter t.

It might be regarded as obvious that the path $S(t)$ varies continuously and piecewise linearly with the weights $(t_i) \in \Delta^n$ on the various spheres of S. However, to spell this out in some detail, consider the function φ which assigns to each oriented edge of a tree T_j the length of the longest monotone edgepath in T_j which begins with the given oriented edge. Strictly, there are two functions φ_+ and φ_- here, depending on whether the length of the given edge is counted or not. Note that $\varphi_\pm$ depends piecewise linearly on the weights t_i since it is the maximum of a finite collection of linear functions measuring the distances to different endpoints of T_j and the lengths of the edges of T_j are given by weights t_i. Surgery on a family of circles of $(S_i \times [0, t_i]) \cap \Sigma$ begins at the time specified by the smaller of the values of φ_- on the two possible orientations of the edge of a T_j corresponding to the given family of circles. In most cases, the surgery on this family of circles ends at the corresponding value of φ_+, the exception being when the given edge of T_j contains in its interior the "center" of T_j, from which the longest paths in both directions have equal length. In this case the corresponding family of circles is being surgered from both sides simultaneously, so the surgery ends sooner, at the time given by the length of these two longest paths to the midpoint. More generally, we need to single out the t-values when a single sphere in a family $S_i \times [0, t_i]$ is being surgered from both sides at once. This occurs when t is halfway between the starting time for one surgery and the stopping time for the other surgery, i.e., the average of two $\varphi_\pm$'s.

Thus we have a finite number of piecewise linear functions on the weight simplex Δ^n whose graphs divide $\Delta^n \times [0, \infty)$ up into a finite number of closed regions, in each of which the systems $S(t)$ vary linearly within a simplex of $\mathbb{S}(M)$.

As we pass from the simplex containing S to a face by letting some weights go to zero, the paths $S(t)$ become the paths $S(t)$ associated to points in that face. This is clear because all that is happening to the trees T_j is that some edges are shrinking to zero length at the face. So the paths $S(t)$ give a well-defined deformation retraction of $\mathbb{S}(M)$ onto the simplex of Σ. The deformation retraction is actually a piecewise linear flow on $\mathbb{S}(M)$ since the restriction of the path $S(t)$ to $t \ge t_0$ is the path for $S(t_0)$.

The cases $s > 1$ reduce to $s \le 1$ by the next lemma. □

LEMMA 2.2. *If for some* $s \geq 1$ $\mathbb{S}(M_{n,s})$ *is contractible, then so is* $\mathbb{S}(M_{n,s+1})$.

Proof. Let the boundary spheres of $M = M_{n,g+1}$ be denoted $\partial_0, \ldots, \partial_s$. Call a vertex of $\mathbb{S}(M)$ "special" it it is a sphere in M splitting off a 3-punctured S^3 having ∂_0 as one puncture. Let $\mathbb{S}'(M)$ be the subcomplex of $\mathbb{S}(M)$ consisting of simplices having no special vertices. For a fixed special vertex Σ_1, its link $L(\Sigma_1)$ in $\mathbb{S}(M)$ lies in $\mathbb{S}'(M)$ and can be identified with $\mathbb{S}(M_{n,s})$.

A deformation retraction of $\mathbb{S}'(M)$ onto $L(\Sigma_1)$ can be obtained as follows. Enlarge Σ_1 to a maximal system Σ and put other systems S into normal form with respect to Σ. Then systems S not containing Σ_1 will meet the 3-punctured sphere P cut off by Σ_1 in disks separating ∂_0 from the other puncture ∂_k in P. Thus $S \cap \Sigma_1$ consists of parallel circles in Σ_1, and we may surger S along these circles using the disks they bound *on the side of* ∂_0, innermost disks first, as usual. If the given system S lies in $\mathbb{S}'(M)$ then the surgery produces parallel copies of ∂_0, which are to be discarded, together with a new system in $L(\Sigma_1)$. Therefore if we modify the proof of 2.1 so that instead of surgering systems S along all the circles of $S \cap \Sigma$, we surgery only the circles of $S \cap \Sigma_1$, surgering in the way just described, then we obtain a deformation retraction of $\mathbb{S}'(M)$ onto $L(\Sigma_1)$.

$\mathbb{S}(M)$ is the union of $\mathbb{S}'(M)$ with the stars of the special vertices. These stars have disjoint interiors since two distinct special vertices cannot be represented by disjoint spheres. The links of special vertices are contractible, being copies of $\mathbb{S}(M_{n,s})$, so attaching the stars, which are cones on the links, does not affect homotopy type. Thus $\mathbb{S}(M)$ deformation retracts to $\mathbb{S}'(M)$, which is contractible by the preceding paragraph. □

3. Subcomplexes of $\mathbb{S}(M)$

There are two subcomplexes of $\mathbb{S}(M)$ which will be needed for proving homological stability. The first of these, denoted Y, has as its simplices the systems S for which $M - S$ is connected. The other subcomplex is Z, whose simplices are systems S such that for one component M_S of $M - S$ the map $\pi_1(M_S) \to \pi_1(M)$ is an isomorphism, the other components being necessarily punctured S^3's. We shall also need the oriented version $Y^{\pm}$ of Y, whose simplices are systems S defining a simplex of Y together with the additional data of an orientation for each component of S. There will be no need to consider an oriented version of Z since the spheres in a system S defining a simplex in Z have a natural normal orientation determined by which side M_S lies on.

The dimension of Y and $Y^{\pm}$ is $n - 1$, while Z has dimension $s - 2$.

PROPOSITION 3.1. *Y and $Y^{\pm}$ are $(n-2)$-connected, and Z is $(s-3)$-connected. Hence all three complexes are homotopy equivalent to wedges of spheres.*

Proof. Consider first the case of Z. Here we follow the reasoning for Lemma 2.2, with Z in place of $S(M)$. The subcomplex $Z' \subset Z$ consisting of simplices with no special vertices deformation retracts onto the link of the special vertex Σ_1, by surgering systems S in Z' along $S \cap \Sigma_1$ as in the proof of 2.2. The union of Z' with the star of Σ_1 is then a contractible subcomplex $Z'' \subset Z$, and Z is obtained from Z'' by attaching the stars of all the other special vertices along their links. A space homotopy equivalent to Z is produced by collapsing the contractible subcomplex Z'' to a point, and this quotient space is the wedge of the suspensions of the links of all the special vertices other than Σ_1. These links are copies of Z for the manifold M with one fewer puncture, so induction on s gives the result for Z.

For Y, surgery along Σ_1 gives a deformation retraction of $Y(M_{n,s})$ onto $Y(M_{n,s-1})$ if $s > 1$, reducing us to the case $s \le 1$. Then we proceed as in the proof of Theorem 1.1 of [H, pp. 219–220]. A map $f: S^i \to Y$ may be extended to $f: D^{i+1} \to \mathbb{S}(M)$ since $\mathbb{S}(M)$ is contractible. We may assume f is piecewise linear with respect to some triangulation of D^{i+1}. If $f(D^{i+1})$ is not contained in Y, the image of some simplex is a system S with disconnected complement. This means the associated graph G_S, with vertices the components of $M - S$ and edges the components of S, has at least two vertices. The edges of G_S whose endpoints are distinct correspond to a subsystem $S' \subset S$ with the property that each of its component spheres S'_j separates $M - \bigcup_{\ell \neq j} S'_\ell$. Such a system is called purely separating. Let σ be a simplex of D^{i+1} of maximal dimension k such that the system $S = f(\sigma)$ is purely separating. Since σ is not contained in ∂D^{i+1}, the link of σ is a sphere S^{i-k}, and by the maximality assumption, f maps this S^{i-k} into $Y(M')$ where M' is M split open along $f(\sigma)$. In case M' is not connected, then $Y(M')$ is, essentially by definition, the join of the spaces Y for the various components of M'. Since $s \le 1$, each component of M' is a connected sum of fewer than n copies of $S^1 \times S^2$ (with punctures). If the total number of these $S^1 \times S^2$'s in M' is m, then by induction $Y(M')$ is homotopy equivalent to a wedge of S^{m-1}'s. We have $m \ge n - k$ since the purely separating system S has at most $k + 1$ spheres. If $i \le n - 2$ then $i - k \le n - k - 2 \le m - 2$, the connectivity of $Y(M')$, so $f \mid S^{i-k}$ extends to a map $D^{i-k+1} \to Y(M')$. We modify the given f on the interior of the star of σ by taking the join of its values on $\partial\sigma$ with this map $D^{i-k+1} \to Y(M')$. The new f has no simplices in the interior of the star of σ which map to purely separating systems, so by a finite sequence of such modifications we can decrease k.

To prove that $Y^{\pm}$ is $(n-2)$-connected, first choose arbitrarily a positive orientation for each essential sphere in M. The subcomplex Y^+ of positively oriented systems is then a copy of Y, so it will suffice to deform a given piecewise linear map $f: S^i \to Y^{\pm}$ into Y^+ for $i \le n - 2$. Let σ be a simplex of S^i of maximum dimension k such that all the vertices of $S = f(\sigma)$ are negatively oriented. The linking sphere S^{i-k-1} then maps to Y^+ and can be viewed as a map

$S^{i-k-1} \to Y^+(M') = Y(M')$ where M' is M split along S. Since S contains at most $k+1$ spheres, M' has at least $n-k-1$ $S^1 \times S^2$ factors and so $Y(M')$ is $(n-k-3)$-connected. The assumption $i \le n-2$ then implies that $f \mid S^{i-k-1}$ extends to a map $D^{i-k} \to Y^+(M')$. We can modify f on the star of σ by joining with this extension, and the new f is homotopic to the old one by joining. Thus we can step by step deform f to have image in Y^+. □

Let $\Gamma = \Gamma_{n,s}$ be the group of isotopy classes of orientation-preserving diffeomorphisms of $M_{n,s}$ taking each puncture to itself, modulo the normal subgroup generated by twists along 2-spheres, i.e., diffeomorphisms supported in a product $S^2 \times I \subset M$ and taking each sphere $S^2 \times \{t\}$ to itself by a rotation through the angle $2\pi t$ about some chosen axis. Such a twist acts trivially on the sphere complex $\mathbb{S}(M)$ since a sphere system may be isotoped to meet $S^2 \times I$ in tubes transverse to the levels $S^2 \times \{t\}$, and then the effect of the twist is to produce a new system which is obviously homotopic to the old one, hence isotopic to it by Laudenbach's theorem. Thus we get an action of Γ on $\mathbb{S}(M)$. Similarly, Γ acts on the complexes Y, $Y^{\pm}$, and Z.

LEMMA 3.2. *The quotient Z/Γ is contractible.*

Proof. Up to diffeomorphism of M fixing ∂M, a sphere S corresponding to a vertex of Z is determined by the set $P(S)$ of punctures which are not boundary components of M_S. More generally, a system $S = S_0 \cup \cdots \cup S_k$ corresponding to a k-simplex of Z is determined up to diffeomorphism fixing ∂M by the $k+1$ sets $P(S_i)$. There is a distinguished vertex v of Z/Γ corresponding to a sphere S for which $P(S)$ is all the punctures. If a sphere system corresponding to a simplex of Z does not contain such a sphere S, then the system can always be enlarged so that it does contain such an S. This implies that Z/Γ is the star of the distinguished vertex v, hence is contractible. □

One can describe Z/Γ explicitly in terms of the simplicial complex P_s associated to the poset of partitions of $\{1, \ldots, s\}$, partially ordered by refinement. Namely, to each sphere S corresponding to a vertex of Z we may associate the partition consisting of the set of punctures $P(S)$ together with the single-element sets consisting of the punctures not in $P(S)$. This induces an embedding of Z/Γ as a subcomplex of P_s, the link of the partition of $\{1, \ldots, s\}$ into s single-element subsets. The distinguished vertex v corresponds to the partition consisting of the entire set $\{1, \ldots, s\}$.

For the action of Γ on Z, the stabilizer of a simplex $\sigma \subset Z$ is the subgroup $\Gamma_\sigma \subset \Gamma$ represented by diffeomorphisms $f: M \to M$ leaving invariant a sphere system S corresponding to σ. Such an f can be isotoped to be the identity on all the complementary components of S in M except the nonsimply-connected component M_S. Namely, there is at least one simple-connected complementary component of S having all but one of its boundary spheres in ∂M, and f fixes ∂M so it fixes the one other boundary sphere of this component as well. Hence f is isotopic to the identity on this component, and one can proceed inductively to the remaining simply-connected components of $M - S$. Since $M_S = M_{n,t}$ for some $t < s$ we then have a natural surjection $\Gamma_{n,t} \to \Gamma_\sigma$.

LEMMA 3.3. *The natural surjection $\Gamma_{n,t} \to \Gamma_\sigma$ is an isomorphism.*

Proof. Suppose a diffeomorphism $g : M_S \to M_S$ extends to $f: M \to M$ via the identity outside M_S, and f is isotopic to the identity in 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. By looking in the universal cover of M we see that $f(S')$ is homotopic to S' in M_S. Laudenbach's theorem then implies that $f(S')$ is isotopic to S' in M_S, so g may be isotoped to leave S' invariant. Then g must take each component of S' to itself, preserving orientation, otherwise f would act nontrivially on $H_2(M)$. Splitting M_S along S' produces a punctured S^3, and the corresponding splitting of g is isotopic to the identity since it preserves all the punctures. So g itself is isotopic to a product of twists along spheres of S', hence g is zero in $\Gamma_{n,t}$. □

Consider now the question of identifying simplex-stabilizers Γ_σ for the action of $\Gamma_{n,s}$ on the subcomplex $Y \subset \mathbb{S}(M_{n,s})$. If σ corresponds to a system S containing k spheres, there is a natural map $\Gamma_{n-k,s+2k} \to \Gamma_\sigma$ since splitting $M_{n,s}$ along S produces $M_{n-k,s+2k}$. This homomorphism will not be surjective, however, since diffeomorphisms fixing the system S can permute the different spheres in S and reverse their orientations, giving rise to diffeomorphisms of $M_{n-k,s+2k}$ permuting punctures, which is not allowed in $\Gamma_{n-k,s+2k}$. The situation is slightly improved if we take the natural action of $\Gamma_{n,s}$ on the complex $Y^\pm$ of oriented systems in Y. More is needed however: systems which are ordered as well as oriented. An easy way to achieve this is to replace $Y^\pm$ by the complex X whose k-simplices are the simplicial maps $\Delta^k \to Y^\pm$. Simplicial maps take vertices to vertices, but distinct vertices can map to the same vertex. Thus the k-simplices of X are the $(k+1)$-tuples $(S_0, \ldots, S_k)$ of oriented spheres, not necessarily distinct (though repeated spheres have the same orientation) whose union is a simplex of Y. Since spheres may be repeated arbitrarily often, X is infinite-dimensional. There is a natural projection $X \to Y^\pm$ sending a simplicial map $\Delta^k \to Y^\pm$ to its image simplex, and it is a classical fact in algebraic topology that this projection induces an isomorphism on simplicial homology, hence also on singular homology; see e.g. Theorem 4.6.8 in [S]. This uses no special properties of $Y^\pm$, just that it is a simplicial complex.

The action of $\Gamma_{n,s}$ on $Y^{\pm}$ induces an action on X, and for a simplex σ of X whose image in Y is a system of k spheres, so $k \leq \dim \sigma + 1$, there is again a natural map $\Gamma_{n-k,s+2k} \to \Gamma_\sigma$. By the definition of X this is surjective, and in the same way that the preceding lemma was proved one shows:

LEMMA 3.4. *The natural surjection $\Gamma_{n-k,s+2k} \to \Gamma_\sigma$ is an isomorphism.*

The analog of Lemma 3.2 is:

LEMMA 3.5. $X/\Gamma_{n,s}$ is $(n-2)$-connected.

Proof. From basic 3-manifold topology, Γ acts transitively on k-simplices of $Y^{\pm}$, so every k-simplex of $Y^{\pm}$ is equivalent under a diffeomorphism of M to the first $k+1$ spheres of a fixed oriented system $S_1 \cup \cdots \cup S_n$ with connected complement. Further, all permutations of these S_i's can be realized by diffeomorphisms of M. It follows that X/Γ can be identified with the quotient K_n/Σ_n of the simplicial complex K_n of sequences $(n_0, \ldots, n_k)$ of positive integers $n_i \leq n$ by the action of the permutation group Σ_n. The inclusion $K_{n-1} \subset K_n$ induces an inclusion $K_{n-1}/\Sigma_{n-1} \subset K_n/\Sigma_n$, with the $(n-2)$-skeleton of K_n/Σ_n contained in K_{n-1}/Σ_{n-1}, so it will suffice to show that K_{n-1}/Σ_{n-1} is contractible in K_n/Σ_n.

The inclusion $K_{n-1} \subset K_n$ extends to an embedding of the cone CK_{n-1} in K_n taking the cone point to the vertex (n). Factoring out the action of Σ_{n-1}, we get a map $C(K_{n-1}/\Sigma_{n-1}) \to K_n/\Sigma_{n-1} \to K_n/\Sigma_n$ extending the inclusion $K_{n-1}/\Sigma_{n-1} \subset K_n/\Sigma_n$. $\square$

4. Homological stability

We shall be interested in three homomorphisms

$$\alpha : H_i(\Gamma_{n,s}) \to H_i(\Gamma_{n,s+1}), \qquad s \geq 1$$

$$\beta : H_i(\Gamma_{n,s}) \to H_i(\Gamma_{n,s-1}), \qquad s \geq 1$$

$$\gamma : H_i(\Gamma_{n,s}) \to H_i(\Gamma_{n+1,s-2}), \qquad s \geq 2$$

induced by maps $\Gamma_{n,s} \to \Gamma_{m,t}$ which are obtained by enlarging the manifold $M_{n,s}$ to $M_{m,t}$ and extending diffeomorphisms via the identity on the added piece. For α one enlarges $M_{n,s}$ by gluing on a 3-punctured S^3 to a component of $\partial M_{n,s}$. For β one just fills in a puncture with a ball, and for γ one adjoins an $S^2 \times I$ connecting two punctures of $M_{n,s}$, or equivalently one glues together two boundary spheres of $M_{n,s}$.

Since we are factoring out twists along 2-spheres we don't need to worry about whether diffeomorphisms are the identity on each boundary sphere or merely orientation-preserving.

Our main theorem is about $\gamma\alpha^2 : H_i(\Gamma_{n,1}) \to H_i(\Gamma_{n+1,1})$ and $\beta : H_i(\Gamma_{n,1}) \to H_i(\Gamma_{n,0})$. Note that the composition $\beta\alpha$ is the identity since this corresponds to just adding a collar on a boundary component of $M_{n,s}$. So α is injective for $s \geq 1$ and β is surjective for $s \geq 2$. However, the latter fact does not directly apply to the theorem, which involves β for $s = 1$.

Associated to a group Γ acting simplicially on a simplicial complex X there is a well-known spectral sequence, described for example in Chapter 7 of [B] and constructed as follows. Let $E\Gamma$ be a contractible complex on which Γ acts freely with quotient a classifying space $B\Gamma$ for Γ. Then the diagonal action of Γ on $X \times E\Gamma$ is free, and the quotient space $X \times_\Gamma E\Gamma$ has homology the equivariant homology $H^\Gamma_*(X)$ of X, essentially by definition. The cellular chain complex $C_*(X \times_\Gamma E\Gamma)$ can be identified with $C_*(X) \otimes_\Gamma C_*(E\Gamma)$, and is actually a double complex, via the boundary maps in the two factors. Filtering this double complex horizontally and vertically leads to two spectral sequences converging to $H^\Gamma_*(X)$. One filtration has E^1_{pq} equal to the p^{th} homology group of the chain complex $C_*(X) \otimes_\Gamma C_q(E\Gamma)$ with the boundary map from X. Since $C_q(E\Gamma)$ is free with free action of Γ, we therefore have $E^1_{pq} = H_p(X) \otimes_\Gamma C_q(E\Gamma)$. The second filtration has E^1_{pq} the q^{th} homology group of the complex $C_p(X) \otimes_G C_*(E\Gamma)$ with the boundary map from $E\Gamma$, so $E^1_{pq} = H_q(\Gamma; C_p(X))$ in this case, the coefficient group $C_p(X)$ being twisted via the Γ action. Shapiro's lemma implies that $H_q(\Gamma; C_p(X)) \approx \oplus_\sigma H_q(\Gamma_\sigma)$ where $\Gamma_\sigma \subset \Gamma$ is the stabilizer of the p-simplex σ in X and there is one summand $H_q(\Gamma_\sigma)$ for each orbit of the action of Γ on the p-simplices in X. The actions we will consider have the property that if an element $g \in \Gamma$ satisfies $g(\sigma) = \sigma$ then $g \mid \sigma$ is the identity. Then the sum $\oplus_\sigma H_q(\Gamma_\sigma)$ is over the p-simplices in X/Γ, which inherits a natural cell structure from X with simplicial cells. Further, the term E^2_{pq} in this spectral sequence can be interpreted as $H_p(X/\Gamma; \{H_q(\Gamma_\sigma)\})$, the simplicial homology of X/Γ with coefficients in the system of groups $H_q(\Gamma_\sigma)$; see [B]. In our application this system will be constant, so this will be just ordinary homology.

It will be more convenient to use a slight modification of the preceding construction, obtained by replacing $C_*(X)$ by the augmented complex $\tilde{C}_*(X)$ which has an extra $\mathbb{Z}$ in dimension -1. Then the spectral sequence from the first filtration has $E^1_{pq} = \tilde{H}_p(X) \otimes_\Gamma C_q(E\Gamma)$, and for the second filtration we have $E^1_{pq} = \oplus_\sigma H_q(\Gamma_\sigma)$ for $p \geq 0$, with $E^1_{-1,q} = H_q(\Gamma)$ since $\Gamma_\varnothing = \Gamma$, and $E^2_{pq} = \tilde{H}_p(X/\Gamma; \{H_q(\Gamma_\sigma)\})$.

To prove the half of the theorem concerning the stabilization $\gamma\alpha^2 : H_i(\Gamma_{n-1,1}) \to H_i(\Gamma_{n,1})$ we shall construct a function $\varphi : \mathbb{N} \to \mathbb{N}$ such that the following two statements hold:

(4.1) (a_k) $\alpha : H_k(\Gamma_{n-1,s}) \to H_k(\Gamma_{n-1,s+1})$ is an isomorphism for $n > \varphi(k)$ and $s \geq 1$.
(b_k) $\gamma : H_k(\Gamma_{n-1,s+2}) \to H_k(\Gamma_{n,s})$ is an isomorphism for $n > \varphi(k)$ and $s \geq 1$, and a surjection for $n = \varphi(k)$ and $s \geq 1$.

The strategy will be to determine an appropriate value for $\varphi(k)$ by induction on k, so *from now on we shall assume*:

($*$) A function $\varphi : \{0, 1, \ldots, k-1\} \to \mathbb{N}$ has been constructed such that (a_i) and (b_i) hold for $i < k$.

The function φ will also be assumed to be strictly monotonic, so $\varphi(i) < \varphi(j)$ if $i < j$, and to satisfy $\varphi(i) > i$.

A partial result toward condition (a_k) is:

LEMMA 4.2. *$\alpha : H_k(\Gamma_{n,s-1}) \to H_k(\Gamma_{n,s})$ is an isomorphism if $s > k+1$ and $n+1 > \varphi(k-1)$.*

Proof. Consider the action of $\Gamma_{n,s}$ on the subcomplex $Z \subset \mathbb{S}(M_{n,s})$ defined in §3. This action has the good property that simplices which are invariant under an element of Γ are fixed by it. For this action of $\Gamma_{n,s}$ on Z the first spectral sequence has $E^1_{pq} = 0$ for $p < s-2$ by Proposition 3.1, hence $E^\infty_{p,q} = 0$ for $p < s-2$. Since both spectral sequences converge to the same thing, this implies $E^\infty_{p,q} = 0$ for $p+q < s-2$ in the second spectral sequence. In particular the term $E^1_{-1,k} = H_k(\Gamma_{n,s})$ in this spectral sequence must be killed by differentials if $-1+k < s-2$. (The differentials in E^r go r units to the left and $r-1$ units upward.)

In the E^2 array of this spectral sequence we must have all zeros below the k^{th} row by Lemma 3.2 since the coefficient system $\{H_q(\Gamma_\sigma)\}$ is constant below the k^{th} row by the inductive hypothesis ($*$) and the assumption $n+1 > \varphi(k-1)$, all these groups $H_q(\Gamma_\sigma)$ having a canonical isomorphism with $H_q(\Gamma_{n,s})$. Thus the only possibility is that the differential $E^1_{0,k} \to E^1_{-1,k}$ is surjective. The term $E^1_{0,k}$ is $\bigoplus_{t<s} H_k(\Gamma_{n,t})$, the sum over the vertex stabilizers, and the differential is the sum of the stabilizations $H_k(\Gamma_{n,t}) \to H_k(\Gamma_{n,s})$. Each of these factors through $H_k(\Gamma_{n,s-1})$, so $\alpha : H_k(\Gamma_{n,s-1}) \to H_k(\Gamma_{n,s})$ is surjective. (The image of $H_k(\Gamma_{n,t}) \to H_k(\Gamma_{n,s})$ depends only on t, not on the vertex sphere realizing this stabilization since any two such spheres are related by a diffeomorphism of $M_{n,s}$, inducing an inner automorphism of $\Gamma_{n,s}$ hence the identity on $H_k(\Gamma_{n,s})$.)

As noted at the beginning of this section, $\alpha : H_k(\Gamma_{n,s-1}) \to H_k(\Gamma_{n,s})$ is injective if $s > 1$. $\square$

LEMMA 4.3. $\gamma : H_k(\Gamma_{n-1,s+2}) \to H_k(\Gamma_{n,s})$ is surjective if $n \geq \varphi(k-1)+2$.

Proof. Consider the two spectral sequences associated to the action of $\Gamma_{n,s}$ on the complex X defined in §3. The first of these has $E^1_{pq} = 0$ for $p < n-1$ by Proposition 3.1 and the fact the $Y^{\pm}$ and X have isomorphic homology. Hence the second spectral sequence has $E^{\infty}_{pq} = 0$ for $p+q < n-1$. In particular, if $k < n$ the term $E^1_{-1,k} = H_k(\Gamma_{n,s})$ must be killed by differentials originating on the line $p+q = k$. As in the proof of the preceding lemma we shall show that the first differential γ is the only one which can be nonzero, hence γ is surjective.

The term E^1_{pq} is a sum of groups $H_q(\Gamma_{n-j,s+2j})$ for $j \leq p+1$ by Lemma 3.4. So the k^{th} row of the E^1 array, with its differential, is

$$H_k(\Gamma_{n,s}) \xleftarrow{\gamma} H_k(\Gamma_{n-1,s+2}) \xleftarrow{d} H_k(\Gamma_{n-2,s+4}) \oplus H_k(\Gamma_{n-1,s+2}) \longleftarrow \cdots$$

The induction hypothesis $(*)$ implies that $E^2_{pq} = 0$ for $p+q \leq k$ and $q < k$ provided that $n-p-1 \geq \varphi(k-p)$ for $p = 1, 2, \ldots, k$ and that $k < n-1$ so that Lemma 3.5 applies. By monotonicity of φ, the condition $n-p-1 \geq \varphi(k-p)$ reduces to the case $p = 1$, i.e., $n \geq \varphi(k-1)+2$. This implies the earlier condition $k < n-1$ since we assume $\varphi(k-1) > k-1$.

Thus if $n \geq \varphi(k-1)+2$, γ is the only differential which can kill $H_k(\Gamma_{n,s})$, so γ must be surjective. $\square$

Injectivity of γ will require an extra step. What will be shown is that the differential d in the sequence above is zero. This implies injectivity of γ since the term $H_k(\Gamma_{n-1,s+2})$ must disappear by E^{∞} and as shown in the previous proof, there are no terms below the k^{th} row which could kill the kernel of γ. For this argument we need the E^2 terms to vanish for one more unit to the right than in the surjectivity argument, but this is taken care of automatically by the statement (b_k) which has strict inequalities for injectivity and weak inequalities for surjectivity.

The map d is easily described. On the first summand it is the difference between the stabilizations γ obtained by gluing together two different pairs of boundary spheres of $M_{n-2,s+4}$, while on the second summand d is clearly zero – the difference between two coinciding stabilizations. If $\alpha^3 : H_k(\Gamma_{n-2,s+1}) \to H_k(\Gamma_{n-2,s+4})$ were onto, then d would be zero since there is a diffeomorphism of $M_{n-2,s+4}$ supported in $M_{n-2,s+4} - M_{n-2,s+1}$ permuting the two different pairs of punctures being identified.

To show $\alpha^3 : H_k(\Gamma_{n-2,s+1}) \to H_k(\Gamma_{n-2,s+4})$ is onto we would like to use Lemma 4.2, but unfortunately this requires s to be large with respect to k. (This is in contrast to the situation for mapping class groups of surfaces, where the dimension range for stabilization with respect to punctures depends on n instead of s.) To get around this problem, consider the commutative diagram

$$\begin{tikzcd} H_k(\Gamma_{n,s-1}) \arrow[r, "\alpha"] & H_k(\Gamma_{n,s}) \\ H_k(\Gamma_{n-j,s+2j-1}) \arrow[u, "\gamma^j"] \arrow[r, "\alpha"] & H_k(\Gamma_{n-j,s+2j}) \arrow[u, "\gamma^j"] \end{tikzcd}$$

where we need $s \geq 2$ in order for the upper α to be defined. By Lemma 4.3 we know that the γ^j on the right will be onto if $n - j + 1 \geq \varphi(k-1) + 2$. By Lemma 4.2 the lower α in the diagram will be onto if $k + 1 < s + 2j$ and $n - j + 1 > \varphi(k-1)$. The smallest j satisfying $k + 1 < s + 2j$ is $j = [(k - s + 1)/2] + 1$, where $[x]$ denotes the greatest integer in x. Choosing this value of j, we deduce:

(4.4) The upper α in the preceding diagram is surjective if $n > \varphi(k-1) + [(k - s + 1)/2] + 1$.

In particular

(4.5) $\alpha^3 : H_k(\Gamma_{n-2,s+1}) \to H_k(\Gamma_{n-2,s+4})$ is onto, and hence $\gamma : H_k(\Gamma_{n-1,s+2}) \to H_k(\Gamma_{n,s})$ is injective, provided that $n - 2 > \varphi(k-1) + [(k - s - 1)/2] + 1$.

For the application to the statements in (4.1) we may restrict attention to $s \geq 1$ in this inequality, which then holds for $s \geq 1$ if it holds for $s = 1$, i.e., if $n > \varphi(k-1) + [k/2] + 2$. Thus if we define φ recursively by $\varphi(k) = \varphi(k-1) + [k/2] + 2$, starting with $\varphi(0) = 1$, part (b_k) of (4.1) holds. Part (a_k) also holds since by (4.4), the map $\alpha : H_k(\Gamma_{n-1,s}) \to H_k(\Gamma_{n-1,s+1})$ is an isomorphism if $n - 1 > \varphi(k-1) + [(k-1)/2] + 1$.

With the recursive definition $\varphi(k) = \varphi(k-1) + [k/2] + 2$, $\varphi(0) = 1$, it is not hard to check that the inequality $n > \varphi(k)$ is equivalent to $n > k^2/4 + 2k + 1$ for positive integers n. Hence the first half of the main theorem is proved.

To prove the other half of the main theorem asserting that the stabilization $\beta : H_k(\Gamma_{n,1}) \to H_k(\Gamma_{n,0})$ is an isomorphism we use the diagram

$$\begin{tikzcd} H_k(\Gamma_{n-1,2}) \arrow[r, "\gamma"] \arrow[d, shift right, "\alpha"'] & H_k(\Gamma_{n,0}) \\ H_k(\Gamma_{n-1,3}) \arrow[u, shift right, "\beta"'] \arrow[r, "\gamma"] & H_k(\Gamma_{n,1}) \arrow[u, "\beta"'] \end{tikzcd}$$

For the upper γ to be an isomorphism we now need the inequality

$n-2>\varphi(k-1)+[(k-s-1)/2]+1$ in (4.5) for $s=0$ as well as for $s\geq 1$, so we get $n>\varphi(k-1)+[(k-1)/2]+3$. The map α and the lower γ are also isomorphisms if $n>\varphi(k-1)+[(k-1)/2]+3$, hence also the β on the right. For φ defined recursively by $\varphi(k)=\varphi(k-1)+[(k-1)/2]+3$, $\varphi(0)=1$, the inequality $n>\varphi(k)$ translates into $n>k^2/4+5k/2+1$.

Appendix: The Connection with Outer Space

Here we relate sphere complexes to the work of Culler–Vogtmann in [CV].

The points of the rank n Outer Space $\mathbb{O}$ of Culler–Vogtmann are equivalence classes of homotopy equivalences $f: X_0 \to X$ where X_0 is a bouquet of n circles and X is a metric graph which doesn't deformation retract onto any subgraph, the metric being normalized so that the total length of all the edges is 1. The equivalence relation on such "marked metric graphs" $f: X_0 \to X$ is given by homotopy of f and composition with isometries $X \to X'$. Fixing the topological type of X and varying only the lengths of its edges traces out an open simplex in $\mathbb{O}$. Passing to faces of this simplex corresponds to letting the lengths of some edges go to zero. Depending on which edges are collapsing in this way, the face might or might not belong to $\mathbb{O}$.

Let $M=M_{n,0}$ and $\mathbb{S}=\mathbb{S}(M)$, and let $\mathbb{S}_\infty$ be the subcomplex of $\mathbb{S}$ consisting of sphere systems having at least one nonsimply-connected complementary component in M. A sphere system S has a dual graph $G(S)$ having vertices the components of $M-S$ and edges the spheres of S. We may view $G(S)$ as embedded in M by choosing a vertex point in each component of $M-S$ and connecting these vertices by edges crossing the spheres of S, each sphere having a single edge crossing it exactly once. Somewhat more canonically, $G(S)$ is also a quotient of M, obtained by thickening S to a product $S\times[0,1]\subset M$, then collapsing the components of $M-S\times(0,1)$ to points and also the components of $S\times\{t\}$ for each $t\in(0,1)$. If S is in $\mathbb{S}-\mathbb{S}_\infty$ then both maps $G(S)\hookrightarrow M$ and $M\to G(S)$ are isomorphisms on π_1. Fixing a system S_0 with $G(S_0)=X_0$, the composition $G(S_0)\to M\to G(S)$ is then a homotopy equivalence. The barycentric coordinates of a point in the open simplex of $\mathbb{S}$ determined by S give weights on the components of S and hence lengths on the corresponding edges of $G(S)$. In this way we obtain a map $\Phi:\mathbb{S}-\mathbb{S}_\infty\to\mathbb{O}$ sending the weighted system S to $G(S_0)\to G(S)$. On each open simplex of $\mathbb{S}-\mathbb{S}_\infty$ Φ is a linear homeomorphism onto an open simplex of $\mathbb{O}$, and Φ is continuous when we pass to faces of simplices, hence Φ is continuous everywhere. Also, Φ is equivariant with respect to the natural action of $\mathrm{Out}(F_n)$ on $\mathbb{S}-\mathbb{S}_\infty$ and $\mathbb{O}$.

PROPOSITION. $\Phi:\mathbb{S}-\mathbb{S}_\infty\to\mathbb{O}$ is a homeomorphism.

Proof. We construct an inverse map. Let $f: X_0 \to X$ represent a point of $\mathbb{O}$. We may build a manifold $M(X)$ diffeomorphic to M by taking a k-punctured 3-sphere for each valence-k vertex of X and identifying the boundary 2-spheres of these punctured 3-spheres according to the edges of X. These identified boundary 2-spheres, weighted according to the lengths of the corresponding edges of X, then give a weighted sphere system $S(X) \subset M(X)$ whose associated metric graph is isometric to X. If we choose a diffeomorphism $h : M \to M(X)$, then $h^{-1}(S(X))$ is a weighted system $S \subset M$ with $\Phi(S)$ equal to $g : X_0 \to X$ for some homotopy equivalence g. Since all automorphisms of $\pi_1(M)$ are realized by diffeomorphisms of M, we may rechoose h so that $\Phi(S)$ equals the given $f: X_0 \to X$. With this condition on h, the isotopy class of $S = h^{-1}(S(X))$ depends only on the given $f: X_0 \to X$ and not on the choice of h since any other choice h' induces the same isomorphism on π_1 hence by Laudenbach's theorem is isotopic to h modulo twists along 2-spheres, which have no effect on isotopy classes of sphere systems. Any isometry $X \to X'$ can be realized by a diffeomorphism $M(X) \to M(X')$ taking $S(X)$ to $S(X')$, so we have a well-defined map $\mathbb{O} \to \mathbb{S} - \mathbb{S}_\infty$ which is obviously an inverse to Φ. □

The contraction of $\mathbb{S}$ constructed in section 2 restricts to a contraction of $\mathbb{S} - \mathbb{S}_\infty$. This is because each flow line of the contraction represents a sequence of surgeries on a sphere system S, and simple-connectivity of the components of $M - S$ is preserved by the surgery process. Namely, each surgery cuts one complementary component of S along a disk and attaches a 2-handle to another complementary component, and both these operations preserve simple-connectivity of the complementary components. The flow also involves replacing spheres with parallel copies of themselves and throwing away trivial spheres or parallel copies of other spheres, and these operations too preserve simple-connectivity of the complementary components.

The space $\mathbb{O}$ has dimension $3n-4$, and Culler–Vogtmann describe a nice "spine" of $\mathbb{O}$ which is a contractible subcomplex of dimension $2n-3$ on which $\mathrm{Out}(F_n)$ acts with finite stabilizers and finite quotient. Using this they prove that $\mathrm{Out}(F_n)$ has finitely generated homology groups and virtual cohomological dimension $2n-3$. Their argument can be phrased in terms of $\mathbb{S}$ as follows.

For variety let us switch from $M_{n,0}$ to $M_{n,1}$. Let $\mathbb{S}_0$ be the subcomplex of the barycentric subdivision of $\mathbb{S}$ having vertices the systems S with all the components of $M - S$ simply-connected, and with k-simplices the chains $S_0 \subset \cdots \subset S_k$ of such systems. The minimal number of spheres in such a system is n and the maximum number is $3n-2$, so $\mathbb{S}_0$ has dimension $2n-2$. By a standard PL topology argument $\mathbb{S}_0$ is a deformation retract of $\mathbb{S} - \mathbb{S}_\infty$, hence is also contractible. The action of $\Gamma_{n,1} = \mathrm{Aut}(F_n)$ on $\mathbb{S}_0$ has finite stabilizers since an orientation-preserving diffeomorphism of a punctured S^3 is determined up to isotopy by how it permutes the punctures. The quotient $\mathbb{S}_0/\mathrm{Aut}(F_n)$ is finite since there are only finitely many isotopy classes of 2-spheres in a punctured S^3, such a 2-sphere being determined up to isotopy by how it separates the punctures.

At this point we need the fact that $\mathrm{Aut}(F_n)$ has a torsionfree subgroup of finite index. Namely, the composition $\mathrm{Aut}(F_n) \to GL_n(\mathbb{Z}) \to GL_n(\mathbb{Z}_p)$ has torsionfree kernel for $p \geq 3$. For the second of these two maps this is explained on p. 40 of [B], and for the first map it can be deduced as a pleasant exercise from a result of Culler [Cu] that any periodic automorphism of F_n is induced by a periodic homeomorphism of a finite graph having fundamental group F_n. If G denotes the kernel of $\mathrm{Aut}(F_n) \to GL_n(\mathbb{Z}_p)$ then G acts freely on $\mathbb{S}_0$ and so the quotient $\mathbb{S}_0/G$ is a finite CW complex $K(G, 1)$. Applying the Lyndon–Hochschild–Serre spectral sequence to the fibration $K(G, 1) \to K(\mathrm{Aut}(F_n), 1) \to K(\mathrm{Aut}(F_n)/G, 1)$ we can deduce that the homology groups of $\mathrm{Aut}(F_n)$ are finitely generated: The E^2 terms are $H_p(\mathrm{Aut}(F_n)/G; H_q(G))$, $H_q(G)$ is finitely generated since we have a finite $K(G, 1)$, and the homology groups of $\mathrm{Aut}(F_n)/G$ are finitely generated with any finitely generated twisted coefficient system since finite groups have CW $K(\pi, 1)$'s with finite skeleta.

Since $\mathbb{S}_0/G$ has dimension $2n-2$, $\mathrm{Aut}(F_n)$ had v.c.d. at most $2n-2$. On the other hand, the v.c.d. is at least $2n-2$ since $\mathrm{Aut}(F_n)$ contains a free abelian subgroup of rank $2n-2$ generated by the $2n-2$ automorphisms which fix all generators x_i of F_n except for one x_k with $k > 1$, which is sent to x_1x_k or x_kx_1.

REFERENCES

[B] Brown, K. S., *Cohomology of Groups*, Springer-Verlag, 1982.

[BH] Bestvina, M. and Handel, M., *Train tracks and automorphisms of free groups*, Annals of Math. *135* (1992), 1–51.

[Ch$_1$] Charney, R., *Homology stability for GL_n of a Dedekind domain*, Invent. Math. *56* (1980), 1–17.

[Ch$_2$] Charney, R., *A generalization of a theorem of Vogtmann*, J. Pure Appl. Alg. *44* (1987), 107–125.

[C] Cohen, F. R., *Artin's braid groups, classical homotopy theory, and sundry other curiosities*, Contemp. Math. *78* (1988), 167–206.

[Cu] Culler, M., *Finite groups of outer automorphisms of a free group*, Contemp. Math. *33* (1984), 197–207.

[CV] Culler, M. and Vogtmann, K., *Moduli of graphs and automorphisms of free groups*, Invent. Math. *84* (1986), 91–119.

[G] Gersten, S. M., *A presentation for the special automorphism group of a free group*, J. Pure Appl. Alg. *33* (1984), 269–279.

[GT$_1$] Goldstein, R. Z. and Turner, E. C., *Automorphisms of free groups and their fixed points*, Invent. Math. *78* (1984), 1–12.

[GT$_2$] Goldstein, R. Z. and Turner, E. C., *Monomorphisms of finitely generated free groups have finitely generated equalizers*, Invent. Math. *82* (1985), 283–289.

[H] Harer, J. L., *Stability of the homology of the mapping class groups of orientable surfaces*, Annals of Math. *121* (1985), 215–249.

[Hat] Hatcher, A., *On triangulations of surfaces*, Top. and its Appl. *40* (1991), 189–194.

[HM] Hatcher, A. and McCullough, D., *Finite presentation of 3-manifold mapping class groups*, Springer Lecture Notes *1425* (1990), 48–57.

[I] IVANOV, N. V., *Complexes of curves and the Teichmüller modular group*, Russ. Math. Surv. *42: 3* (1987), 55–107.

[L_1] LAUDENBACH, F., *Sur les 2-spheres d'une variété de dimension*, 3, Annals of Math. *97* (1973), 57–81.

[L_2] LAUDENBACH, F., *Topologie de la dimension trois: homotopie et isotopie*, Astérisque *12* (1974), 1–152.

[M] MCCULLOUGH, D., *Homotopy groups of the space of self-homotopy-equivalences*, Trans. A.M.S. *264* (1981), 151–163.

[S] SPANIER, E., *Algebraic Topology*, McGraw-Hill, 1966.

[vK] VAN DER KALLEN, W., *Homology stability for lienar groups*, Invent. Math. *60* (1980), 269–295.

[V] VOGTMANN, K., *Spherical posets and homology stability for $O_{n,n}$*, Topology *20* (1981), 119–132.

[Wa] WAGONER, J. B., *Stability for homology of the general linear group over a local ring*, Topology *15* (1976), 417–423.

[W_1] WALDHAUSEN, F., *Algebraic K-theory of topological spaces II*, Springer Lecture Notes *763* (1979), 356–394.

[W_2] WALDHAUSEN, F., *Algebraic K-theory of spaces, a manifold approach*, Canad. Math. Soc. Conf. Proc. v.2 part 1 (1982), 141–184, A.M.S.

[Wh_1] WHITEHEAD, J. H.C., *On certain sets of elements in a free group*, Proc. London Math. Soc. *41* (1936), 48–56.

[Wh_2] WHITEHEAD, J. H. C., *On equivalent sets of elements in a free group*, Annals of Math. *37* (1936), 782–800.

Department of Mathematics
Cornell University
Ithaca
NY 14853-7901
USA

Received March 6, 1993