CERF THEORY FOR GRAPHS

ALLEN HATCHER AND KAREN VOGTMANN

ABSTRACT

We develop a deformation theory for k-parameter families of pointed marked graphs with fixed fundamental group F_n. Applications include a simple geometric proof of stability of the rational homology of $\mathrm{Aut}(F_n)$, computations of the rational homology in small dimensions, proofs that various natural complexes of free factorizations of F_n are highly connected, and an improvement on the stability range for the integral homology of $\mathrm{Aut}(F_n)$.

1. *Introduction*

It is standard practice to study groups by studying their actions on well-chosen spaces. In the case of the outer automorphism group $\mathrm{Out}(F_n)$ of the free group F_n on n generators, a good space on which this group acts was constructed in [**3**]. In the present paper we study an analogous space on which the automorphism group $\mathrm{Aut}(F_n)$ acts. This space, which we call $\mathbb{A}_n$, is a sort of Teichmüller space of finite basepointed graphs with fundamental group F_n.

As often happens, specifying basepoints has certain technical advantages. In particular, basepoints give rise to a nice filtration $\mathbb{A}_{n,0} \subset \mathbb{A}_{n,1} \subset \ldots$ of $\mathbb{A}_n$ by degree, where we define the degree of a finite basepointed graph to be the number of non-basepoint vertices, counted 'with multiplicity' (see Section 3). The main technical result of the paper, the Degree Theorem, implies that the pair $(\mathbb{A}_n, \mathbb{A}_{n,k})$ is k-connected, so that $\mathbb{A}_{n,k}$ behaves something like the k-skeleton of $\mathbb{A}_n$.

From the Degree Theorem we derive these applications:

(1) We show the natural stabilization map $H_i(\mathrm{Aut}(F_n); \mathbb{Z}) \longrightarrow H_i(\mathrm{Aut}(F_{n+1}); \mathbb{Z})$ is an isomorphism for $n \geqslant 2i+3$, a significant improvement on the quadratic dimension range obtained in [**5**]. For homology with $\mathbb{Q}$ coefficients there is a particularly simple, geometric proof of homology stability which gives an isomorphism for $n \geqslant 3(i+1)/2$.

(2) We give an explicit description of a finite complex which can be used to compute the rational homology $H_i(\mathrm{Aut}(F_n); \mathbb{Q})$. This complex is much smaller than complexes which were previously available, and independent of n for n large with respect to i. As an example, we see immediately that $H_i(\mathrm{Aut}(F_n); \mathbb{Q}) = 0$ for $i = 1, 2$ and all n. In [**6**] we push this approach further to compute $H_i(\mathrm{Aut}(F_n); \mathbb{Q})$ for $i \leqslant 6$ as well, with the following results. The groups $H_3(\mathrm{Aut}(F_n); \mathbb{Q})$, $H_5(\mathrm{Aut}(F_n); \mathbb{Q})$, and $H_6(\mathrm{Aut}(F_n); \mathbb{Q})$ are zero for all n. In dimension four, we have $H_4(\mathrm{Aut}(F_n); \mathbb{Q}) = 0$ for $n \neq 4$, but $H_4(\mathrm{Aut}(F_4); \mathbb{Q}) = \mathbb{Q}$.

(3) Consider the free factorizations $F_n = H_1 * \ldots * H_m$ where the H_i are nontrivial proper subgroups of F_n. The collection of all such factorizations (without regard to the order of the factors H_i) is a partially ordered set with respect to refinement, that is, the further decomposition of the H_i into free products. The simplicial complex associated to this partially ordered set is $(n-2)$-dimensional, and we show it is homotopy equivalent to a wedge of $(n-2)$-dimensional spheres.

Received 20 February 1996; revised 12 November 1996

1991 *Mathematics Subject Classification* 20F32.

J. London Math. Soc. (2) 58 (1998) 633–655

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

The proof of the Degree Theorem is of some interest in its own right, involving a sort of parameterized Morse theory for graphs which has not yet appeared in this subject. This is what the title of the paper refers to, since Cerf theory in differential topology is the study of families of real-valued C^∞ functions on a smooth manifold. Each point in $\mathbb{A}_n$ is represented by a basepointed graph whose edges have assigned lengths, giving a path-metric on the graph. We can then regard the distance to the basepoint as defining a canonical 'Morse function' on the graph. To prove the Degree Theorem, one needs to show that a k-parameter family of such metric graphs can be deformed so as to lower their degrees to at most k for all parameter values. This is done by a process of simplifying the associated Morse functions so as to push as many of the vertices of the graphs down to the basepoint as possible, using general position arguments and a fairly delicate induction.

2. *The space* $\mathbb{A}_n$

2.1. *Points of* $\mathbb{A}_n$

By a *graph* Γ we mean a one-dimensional CW-complex, with *vertices* the 0-cells and *edges* the 1-cells. The *valence* of a vertex v is the number $|v|$ such that a sufficiently small connected neighborhood of v is a cone on $|v|$ points, with v as its cone point. An edge is a *separating edge* if removing a point in its interior disconnects the graph.

A point of $\mathbb{A}_n$ is an equivalence class of triples (Γ, v_0, ϕ) where:

(1) Γ is a finite connected graph with no separating edges.

(2) v_0 is a basepoint vertex of Γ. All other vertices have valence at least 3.

(3) $\phi: F_n \longrightarrow \pi_1(\Gamma, v_0)$ is an isomorphism.

(4) The edges of Γ are assigned positive real lengths in such a way that the sum of the lengths of all the edges is 1. This determines a natural metric on Γ.

(5) (Γ, v_0, ϕ) is equivalent to (Γ', v_0', ϕ') if there is a basepoint-preserving isometry $\rho: \Gamma \longrightarrow \Gamma'$ inducing ρ_* on π_1 with $\phi' = \rho_* \phi$.

2.2. *Topology on* $\mathbb{A}_n$

To define the topology on $\mathbb{A}_n$, consider first the possibility of varying the lengths of the edges of Γ for $(\Gamma, v_0, \phi) \in \mathbb{A}_n$, keeping the length of each edge positive and the sum of all the lengths equal to 1. This defines a collection of points $(\Gamma_s, v_0, \phi) \in \mathbb{A}_n$ parameterized by points s in the interior of a simplex Δ^k, whose barycentric coordinates are the lengths of the edges. These points $(\Gamma_s, v_0, \phi) \in \mathbb{A}_n$ are all distinct since any nontrivial isometry $\Gamma \longrightarrow \Gamma$ acts nontrivially on $H_1(\Gamma)$, and hence changes ϕ.

Each point $(\Gamma, v_0, \phi) \in \mathbb{A}_n$ lies in such an open simplex $\sigma = \sigma(\Gamma, v_0, \phi) \subset \mathbb{A}_n$. Some (but not all) open faces of σ also belong to $\mathbb{A}_n$, as follows. Passing to a face of σ corresponds to shrinking the lengths of certain edges of Γ to zero (while increasing the lengths of the other edges), thereby producing a quotient graph Γ'. If the quotient map $q: \Gamma \longrightarrow \Gamma'$ induces an isomorphism on π_1, then $(\Gamma', q(v_0), q_* \phi)$ is a point in $\mathbb{A}_n$. Varying the lengths of the edges of Γ', which we identify with the edges of Γ which are not shrunk to zero length, we trace out an open simplex τ of $\mathbb{A}_n$ forming a face of σ. The topology of $\mathbb{A}_n$ is defined by these face relations among all such pairs of open simplices σ, $\tau \in \mathbb{A}_n$. More precisely, define a formal closure $\bar{\sigma}$ of each open simplex σ of $\mathbb{A}_n$ by adjoining to an abstract copy of σ those of its open faces which

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

give rise to open simplices of $\mathbb{A}_n$. Then we have a natural map $\bar{\sigma} \longrightarrow \mathbb{A}_n$, and we define a topology on $\mathbb{A}_n$ by saying that a set is open if and only if its inverse images under all such maps $\bar{\sigma} \longrightarrow \mathbb{A}_n$ are open. It follows from Proposition 2.1 below that the projections $\bar{\sigma} \longrightarrow \mathbb{A}_n$ are in fact embeddings, so the formal closure $\bar{\sigma}$ is the same as the actual closure in $\mathbb{A}_n$.

2.3. *Sphere complexes and* $\mathbb{A}_n$

There is another way of looking at $\mathbb{A}_n$ as a subspace of the sphere complex $S(M)$ of [**5**], where M is the compact 3-manifold obtained from the connected sum of n copies of $S^1 \times S^2$ by deleting an open ball. Vertices of $S(M)$ are isotopy classes of embedded 2-spheres in M, excluding those which bound a ball or are isotopic to the boundary sphere of M. A collection of $k+1$ vertices of $S(M)$ spans a k-simplex of $S(M)$ if the corresponding $k+1$ isotopy classes of 2-spheres can be represented by a system of $k+1$ disjoint 2-spheres. Let $\mathbb{A}_n^*$ be the subcomplex of $S(M)$ corresponding to systems of nonseparating spheres, that is, spheres $S \subset M$ such that $M-S$ is connected. Let $\partial\mathbb{A}_n^*$ be the subcomplex of $\mathbb{A}_n^*$ consisting of sphere systems $S_0 \cup \ldots \cup S_k$ such that at least one component of $M-(S_0 \cup \ldots \cup S_k)$ is not simply-connected. There is a natural map from $\mathbb{A}_n^* - \partial\mathbb{A}_n^*$ to $\mathbb{A}_n$, obtained in the following way. Each sphere system $S_0 \cup \ldots \cup S_k$ has a dual graph Γ, whose vertices are the components of $M-(S_0 \cup \ldots \cup S_k)$ and whose edges are the spheres S_i. There is a natural quotient map $M \longrightarrow \Gamma$, inducing a homomorphism $\phi : \pi_1(M, v_0') \cong F_n \longrightarrow \pi_1(\Gamma, v_0)$ where v_0' is a basepoint in the boundary of M and v_0 is its image in Γ. The barycentric coordinates of a point in the open simplex of $S(M)$ spanned by $S_0 \cup \ldots \cup S_k$ determine lengths on the edges of Γ. Thus to each point in $S(M)$ we have an associated triple (Γ, v_0, ϕ) consisting of a metric graph Γ, a basepoint vertex v_0, and a homomorphism $\phi : F_n \longrightarrow \pi_1(\Gamma, v_0)$. The points in the subspace $\mathbb{A}_n^* - \partial\mathbb{A}_n^*$ of $S(M)$ are precisely those for which (Γ, v_0, ϕ) is a point in $\mathbb{A}_n$. Namely, restricting to nonseparating spheres corresponds to Γ having no separating edges, and ϕ being an isomorphism corresponds to the components of $M-(S_0 \cup \ldots \cup S_k)$ being simply-connected.

PROPOSITION 2.1. *The map* $\mathbb{A}_n^* - \partial\mathbb{A}_n^* \longrightarrow \mathbb{A}_n$ *constructed above is a homeomorphism. Hence* $\mathbb{A}_n$ *is contractible.*

This is proved in the Appendix of [**5**] for the analogous unbasepointed situation. Adding basepoints is completely straightforward and we shall leave the proof as an exercise, with the following remarks. Separating spheres were allowed in [**5**]. This makes no difference for proving the first statement of the proposition. For the second statement one can argue as follows. The contraction of the sphere complex $S(M)$ constructed in the proof of Theorem 2.1 of [**5**] restricts to a contraction of the subspace consisting of sphere systems having simply-connected complementary regions. This subspace retracts in turn onto its subspace consisting of sphere systems without separating spheres, by simply deleting any separating spheres. Since a retract of a contractible space is contractible, this finishes the argument.

REMARK 2.2. The space $\mathbb{A}_n^*$ can also be identified with the space of minimal actions of the free group F_n on rooted simplicial $\mathbb{R}$-trees with trivial edge stabilizers, where two actions are considered to be the same if they differ only by a constant scaling of the tree. The subspace $\mathbb{A}_n$ corresponds to actions whose vertex stabilizers

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

are trivial as well, that is, free actions. Each such action determines a Lyndon length function $L\colon F_n \longrightarrow \mathbb{R}$, where the length of an element is the distance it moves the root in the tree. The map sending an action to its Lyndon length function gives an embedding from $\mathbb{A}_n^*$ to an infinite-dimensional real projective space. The topology on $\mathbb{A}_n$ is equivalent to the subspace topology, but the subspace topology on $\mathbb{A}_n^*$ is not the same as the simplicial topology.

2.4. *Spine of* $\mathbb{A}_n$

Our description of $\mathbb{A}_n$ as a union of open simplices does not give it the structure of a simplicial complex since not all faces of its open simplices are in $\mathbb{A}_n$. However, the geometric realization of the partially ordered set of open simplices is a simplicial complex, denoted $S\mathbb{A}_n$. The vertices of $S\mathbb{A}_n$ are the open simplices of $\mathbb{A}_n$, and a k-simplex is a chain of $k+1$ open simplices, each of which is a face of the next. There is a natural embedding of $S\mathbb{A}_n$ in $\mathbb{A}_n$ which sends each vertex to the barycenter of the corresponding open simplex, and each k-simplex to the convex hull of the corresponding barycenters. There is a natural deformation retraction of $\mathbb{A}_n$ onto $S\mathbb{A}_n$, by pushing within each open simplex of $\mathbb{A}_n$ away from its missing faces, that is, the faces of $\partial\mathbb{A}_n^*$. For this reason, $S\mathbb{A}_n$ is called the *spine* of $\mathbb{A}_n$.

The dimension of $\mathbb{A}_n$ is $3n-3$ since the maximum number of edges which a graph satisfying (1)–(3) can have is $3n-2$. The dimension of the spine is $2n-2$.

3. *The Degree Theorem*

3.1. *Degree of a graph*

DEFINITION 3.1. The *degree* of a finite connected graph Γ with basepoint v_0 is

$$\deg(\Gamma) = \sum_{v \neq v_0} |v|-2.$$

For each vertex v of Γ we can think of the number $|v|-2$ as the 'multiplicity' of v. If two vertices v_1 and v_2 are joined by an edge, then collapsing this edge produces a single vertex v in the quotient graph with $|v|-2 = |v_1|-2+|v_2|-2$, so the multiplicity of v is the sum of the multiplicities of v_1 and v_2. By collapsing the edges in a maximal tree, one by one, we see that the sum $\sum_v(|v|-2)$ over all vertices of Γ is equal to $2n-2$ if $\pi_1(\Gamma) \cong F_n$, since this is its value when Γ has a single vertex. It follows that the degree of Γ can also be computed as $2n-|v_0|$, and is at most $2n-2$.

EXAMPLE 3.2. The only graphs of degree 0 are *roses* (graphs with exactly one vertex, the basepoint). Modulo edges with both endpoints attached to the basepoint, the only graph of degree 1 is the *theta graph* (the graph with two vertices and three edges, each with distinct endpoints), and the only graphs of degree 2 are the five indicated in Figure 1.

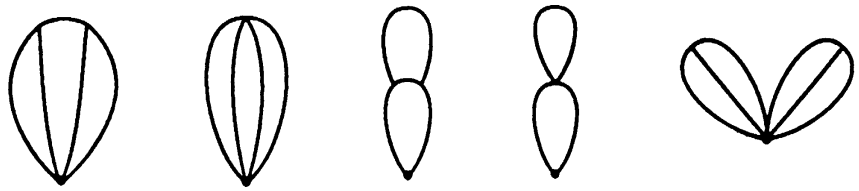


FIG. 1.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

3.2. *Statement of the Degree Theorem*

Let $\mathbb{A}_{n,k} \subset \mathbb{A}_n$ be the subspace consisting of triples (Γ, v_0, ϕ) for which Γ has degree at most k. Since the degree of Γ depends only on the topological type of (Γ, v_0), $\mathbb{A}_{n,k}$ is a union of open simplices of $\mathbb{A}_n$. Also, $\mathbb{A}_{n,k}$ is a closed subspace of $\mathbb{A}_n$ since passing to a face of the open simplex of $\mathbb{A}_n$ containing (Γ, v_0, ϕ) collapses certain edges of Γ, which can only change the valence of the basepoint by increasing it, thus lowering the degree of Γ. Since graphs in $\mathbb{A}_n$ have degree at most $2n-2$, we have $\mathbb{A}_{n,k} = \mathbb{A}_n$ for $k \geqslant 2n-2$.

Our main technical result is the following.

THEOREM 3.3. *A piecewise linear map $f_0 : D^k \longrightarrow \mathbb{A}_n$ is homotopic to a map $f_1 : D^k \longrightarrow \mathbb{A}_{n,k}$ by a homotopy f_t during which degree decreases monotonically, that is, if $t_1 < t_2$ then $\deg(f_{t_1}(s)) \geqslant \deg(f_{t_2}(s))$ for all $s \in D^k$.*

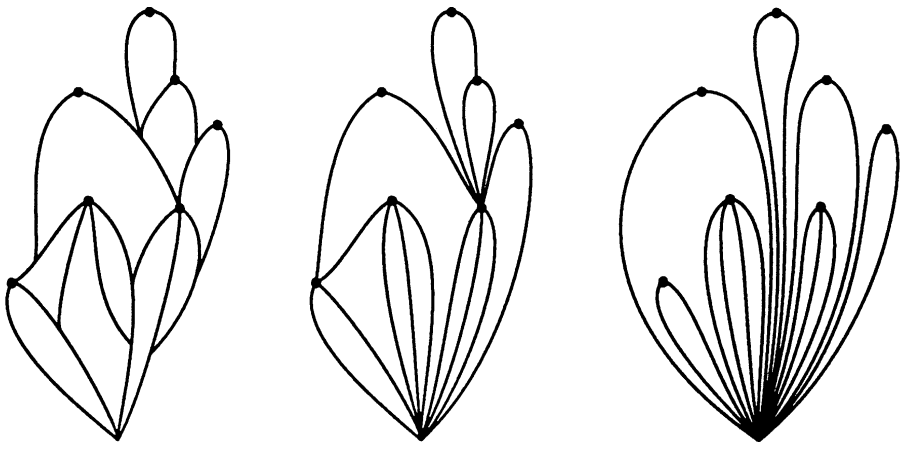
COROLLARY 3.4. *The pair $(\mathbb{A}_n, \mathbb{A}_{n,k})$ is k-connected. Equivalently, since $\mathbb{A}_n$ is contractible, $\mathbb{A}_{n,k}$ is $(k-1)$-connected.*

The proof of the theorem will be given in the next section. In the rest of this section we make various definitions we will need and explain something of the idea of the proof.

3.3. *Height and critical points*

Given a point $(\Gamma, v_0, \phi) \in \mathbb{A}_n$, let $h : \Gamma \longrightarrow \mathbb{R}$ be the function measuring distance to v_0, using the given metric on Γ. We think of h as a *height function* on Γ, as in Morse theory. A point $x \in \Gamma$ is called a *critical point* if near x there are at least two different paths leading down from x, that is, paths along which h decreases. The critical points of Γ are isolated, clearly. As motivation for this terminology, note that the homotopy type of the subspace $h^{-1}([0, r])$ of Γ changes with varying r only when r passes through a critical value of h, just as for Morse functions on manifolds.

An example is shown in Figure 2(a). In this figure we think of h literally as the height function, so the lengths of edges are determined by their vertical heights rather



(a) (b) (c)

FIG. 2.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

than their lengths as arcs in the plane. The height function has seven critical points, indicated by the dots. Four of them are vertices and the other three occur at interior points of edges.

The function h measuring distance to the basepoint has one special property not satisfied by an arbitrary piecewise linear function $\Gamma \longrightarrow \mathbb{R}$, namely, it has only one local minimum, the basepoint. Conversely, it is not hard to see that any piecewise linear function $\Gamma \longrightarrow \mathbb{R}$ with a unique local minimum is the distance function for some metric on Γ, with the minimum as the basepoint. Thus by restricting our 'Morse functions' to distance functions we are not straying too far from a general 'Morse theory' of graphs, and we avoid quite a bit of unwanted noise which unrestricted local minima would generate.

3.4. *Cones and canonical splitting*

To give an idea of how the proof of the Degree Theorem proceeds, let us look at the unparameterized case, when $k = 0$. Here one is given a single point $(\Gamma, v_0, \phi) \in \mathbb{A}_n$ which one wants to deform so as to reduce its degree monotonically to 0, eliminating all vertices except v_0. The obvious way to do this would be simply to shrink the lengths of all edges in a maximal tree to zero, but this approach does not seem to work very well in the parameterized case.

For a critical point $x \in \Gamma$ which is a vertex, the *cone* C_x is the union of x with all open edges of Γ incident to x and lying below x. In case x is a critical point which is not a vertex, define C_x to be the open edge containing x. In either case C_x consists of two or more *branches* leading down from x, which are arcs along which h decreases. If x is not a vertex, then there are just two branches. If x is a vertex there can be two or more branches, and there may also be edges coming into x from above which are not part of C_x. The *ϵ-cone* C_x^ϵ consists of the points of C_x within distance ϵ of x.

If one begins travelling down a branch of a cone, it is possible to continue moving monotonically downward toward the basepoint without ambiguity until one first meets another critical point. If no critical points are encountered, one can travel all the way to the basepoint. Each such *extended branch*, connecting x with a lower critical point or the basepoint, is embedded in Γ, but different extended branches can merge at noncritical vertices. There is a canonical way of deforming Γ to eliminate these junctions, by splitting the extended branches apart, splitting downward to the next critical point, or all the way to the basepoint in favorable cases. One can start at the top of Γ and split downward at each noncritical vertex as one comes to it, like pulling a comb through matted hair. We refer to this process as *canonically splitting* Γ. For example, canonically splitting the graph in Figure 2(a) produces the graph in Figure 2(b).

Canonical splitting does not increase degree, and decreases it if some noncritical vertices split all the way down to the basepoint. Thus in the example of Figure 2 the degree drops from 16 to 11.

A canonically split Γ is the disjoint union of the basepoint and the cones C_x. We can regard Γ as being built in stages, starting with the basepoint and attaching the cones in sequence, according to the heights of the critical points.

3.5. *Sliding in the ϵ-cones*

Having canonically split Γ, we may perturb the attaching points of the branches of the cones C_x by sliding them downwards, so that they remain in the ϵ-cones which

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

contain them, but no longer lie at critical points. This perturbation can be viewed as putting the attaching points in 'general position' with respect to the critical points. Such a perturbation of attaching points has no effect on degree, but canonical splitting of the perturbed Γ will reduce its degree further. In fact, this canonical splitting will deform Γ so that all the cones C_x attach directly to the basepoint, as in Figure 2(c).

The preceding deformations of Γ do not change its ϵ-cones C_x^ϵ. If the ϵ-cones of the original Γ each had just two branches, we would now have succeeded in deforming Γ to a graph of degree 0. So it would have sufficed to have done a preliminary perturbation of Γ to make all its critical points interior points of edges. Such a preliminary perturbation is easy to do since if a vertex x is a critical point, then by small changes in the lengths of the edges of C_x we can insure that x is no longer a critical point.

4. *Proof of the Degree Theorem*

The proof proceeds in two main stages. In the first, a k-parameter family of graphs is perturbed to make the critical points as simple as possible for a generic k-parameter family. This involves only perturbations of the lengths of the edges of the graphs, and does not affect their topological types, so there is no reduction of degree. In the second stage, deformations are made which do not change the local structure of the critical points, that is, the ϵ-cones, but the degree is monotonically reduced to be at most k throughout the family.

4.1. *Codimension*

DEFINITION 4.1. The *codimension* of a critical point x is defined to be 0 if x is an interior point of an edge, and $b-1$ if x is a vertex whose ϵ-cone C_x^ϵ has b branches. The codimension of a graph Γ is the sum of the codimensions of its critical points.

If a critical point x is an interior point of an edge, small perturbations of Γ clearly produce a critical point near x in the same edge, which is why we say such critical points have codimension 0. On the other hand, if the critical point x is a vertex, the b branches of C_x^ϵ are the initial segments of b edgepaths from x to v_0 of equal length, and the condition that these paths have equal length is given by $b-1$ linear equations among the lengths of the edges of Γ. These equations are independent since the corresponding edgepaths all begin with different edges. Hence the solution set of these equations is a codimension $b-1$ linear subspace of the open simplex σ containing (Γ, v_0, ϕ) in $\mathbb{A}_n$.

EXAMPLE 4.2. Let Γ be the graph with two vertices v_0 and v_1 joined by $n+1 \geqslant 3$ edges of length $1/(n+1)$. Then the nonbasepoint vertex v_1 is a critical point of codimension n. The open simplex σ containing Γ is the interior of an n-simplex with barycentric coordinates the lengths of the edges of Γ, and Γ itself lies at the barycenter of this simplex. Any small perturbation of Γ produces critical points of smaller codimension. The case $n = 2$ is illustrated in Figure 3. Here v_1 becomes a critical point of codimension 1 along the three line segments from the barycenter to the vertices of σ, and elsewhere in σ there are two critical points of codimension 0.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

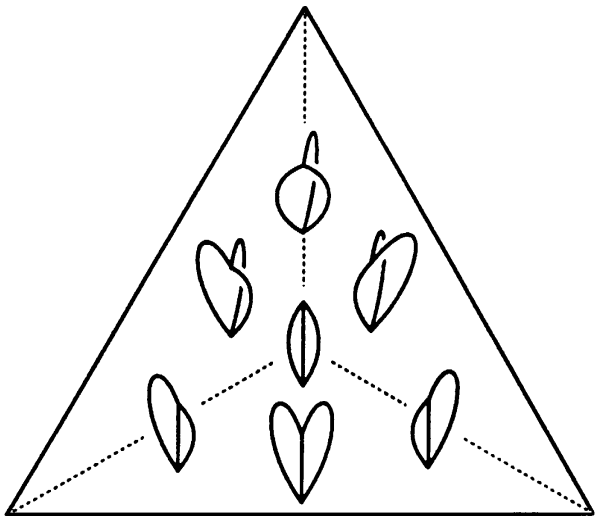


FIG. 3.

4.2. *Critical planes and a stratification of* $\mathbb{A}_n$

Consider a point $(\Gamma, v_0, \phi) \in \mathbb{A}_n$ lying in an open simplex σ obtained by varying the lengths of the edges of Γ. Using these lengths as coordinates, we identify σ with the open simplex in $\mathbb{R}^m$ consisting of points with positive coordinates whose sum is 1, and (Γ, v_0, ϕ) with a point s in this simplex.

The minimum distance from a vertex x of Γ to v_0 is always realized by a path which passes through each edge of Γ at most once. We call such oriented edgepaths in Γ *preferred*. The length l_γ of a preferred path γ is a linear function on $\mathbb{R}^m$, the sum of certain of the coordinates of $\mathbb{R}^m$. By a *critical hyperplane* we mean the codimension-one linear subspace of $\mathbb{R}^m$ defined by equating two length functions l_γ corresponding to distinct preferred paths γ from the same vertex x to v_0. A *critical plane* is an intersection of critical hyperplanes.

The critical hyperplanes determine a stratification of $\mathbb{R}^m$, hence also a stratification of σ. The strata are open convex polyhedra of various dimensions, each defined by a set of equalities and strict inequalities among pairs of length functions l_γ associated to preferred paths γ.

As we pass from σ to a codimension-one open face τ in $\mathbb{A}_n$, an edge of Γ shrinks to zero length, forming a quotient graph Γ'. Preferred paths γ in Γ map to preferred paths γ' in Γ' under the quotient map $\Gamma \longrightarrow \Gamma'$, so each critical hyperplane in σ meets τ in a critical hyperplane in τ. Thus we have a stratification of all of $\mathbb{A}_n$.

As Γ varies over σ, its topological type does not change. The following lemma says that if we restrict to individual strata within σ, the critical points of Γ do not essentially change either.

LEMMA 4.3. *If we let a graph* $(\Gamma, v_0, \phi) \in \sigma$ *vary over a stratum in* σ, *then*:

(i) *Each vertex of* Γ *which is a critical point remains a critical point throughout the stratum, and the homeomorphism type of its* ϵ*-cone does not change.*

(ii) *A critical point which is not a vertex varies continuously within the edge containing it.*

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

(iii) *The codimension of each critical point is constant over the stratum, hence the codimension of* Γ *remains constant. If the codimension of* Γ *is* i*, then the stratum is contained in a critical plane of codimension* i.

Proof. For a given (Γ, v_0, ϕ), the branches of an ϵ-cone C_x^ϵ at a critical vertex x are contained in initial edges of a certain collection of preferred paths γ from x to v_0, all having the same length which is shorter than the length of any other preferred path from x. As we move through the stratum containing (Γ, v_0, ϕ) these same γ remain the shortest paths from x to v_0, and all other preferred paths from x remain longer than these, so C_x^ϵ does not change over the stratum.

For a critical point x interior to an edge, there are two shortest paths from x to the basepoint passing through the two adjacent vertices v_1 and v_2. These two paths contain subpaths γ_1 and γ_2 which are shortest paths from v_1 and v_2 to the basepoint. As we vary over the stratum, γ_1 and γ_2 remain shortest paths. The critical point x can only vary within the open edge containing it while the parameter s remains within the stratum, since if x moved to one of its endpoints v_i, this would create a new shortest path from v_i to the basepoint.

The last statement of the lemma follows from the observation that the length functions for a collection of preferred paths with distinct initial edges are independent.

4.3. *Stage* 1: *Simplifying the critical points*

Let us relabel the given map f_0 as $f\colon D^k \longrightarrow \mathbb{A}_n$, and let $f(s) = (\Gamma_s, v_s, \phi_s)$. We will deform f by deforming the family Γ_s in such a way that it will be obvious how to carry along the basepoint v_s and the isomorphism $\phi_s\colon F_n \longrightarrow \pi_1(\Gamma_s, v_s)$. Therefore no confusion should ensue if we are somewhat free with our notation, sometimes writing Γ_s when we really mean (Γ_s, v_s) or (Γ_s, v_s, ϕ_s).

Let $\Sigma_i \subset \mathbb{A}_n$ be the subspace consisting of triples (Γ, v_0, ϕ) of codimension at least i, so we have a decreasing filtration $\mathbb{A}_n = \Sigma_0 \supset \Sigma_1 \supset \Sigma_2 \supset \ldots$. By Lemma 4.3, Σ_i is a union of strata of $\mathbb{A}_n$ having codimension at least i in each open simplex of $\mathbb{A}_n$. Moreover, Σ_i is a closed subset of $\mathbb{A}_n$ since the codimension of a critical point cannot decrease as we pass to the boundary of a stratum.

LEMMA 4.4. *The map* $f\colon D^k \longrightarrow \mathbb{A}_n$ *can be homotoped, preserving degree, to a piecewise linear map such that, for all* i, $f^{-1}(\Sigma_i)$ *has codimension at least* i *in* D^k.

Proof. Since the given f is piecewise linear, there is a triangulation of D^k for which f maps each simplex linearly to the closure of an open simplex of $\mathbb{A}_n$. If we barycentrically subdivide this triangulation, we will have the additional property that the image of each simplex meets at most one codimension-one face of each open simplex σ of $\mathbb{A}_n$. We shall homotope f by perturbing the images of the vertices within the open simplices of $\mathbb{A}_n$ which contain them. Thus the homotopy will consist only of perturbing the lengths of edges of the graphs Γ_s, so degree will be preserved.

By a preliminary perturbation of f on vertices we may assume each image $f(\Delta^l)$ of a simplex Δ^l of D^k has maximal dimension, either l or the dimension of the open simplex of $\mathbb{A}_n$ containing it. Subsequent small perturbations of f on vertices will not affect this property.

Consider an open simplex σ of $\mathbb{A}_n$. As before, we regard σ as lying in the Euclidean space $\mathbb{R}^m$ whose coordinates are the lengths of edges of graphs Γ for $(\Gamma, v_0, \phi) \in \sigma$. For a simplex Δ^l of D^k mapped by f to the closure of σ, our aim is to put $f(\Delta^l)$ in general

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

position with respect to certain critical planes H in $\mathbb{R}^m$, in the sense that if L is the linear subspace of $\mathbb{R}^m$ spanned by $f(\Delta^l)$, then either $L \cap H = 0$ or L and H span $\mathbb{R}^m$. Note that this is an open condition, preserved by small perturbations of f on vertices of Δ^l.

We call a collection $\mathscr{C} = \{\gamma_j\}$ of preferred paths in Γ *compatible* if no edge of Γ is traversed in different directions by different γ_j. The critical planes H we are interested in are those defined by equating lengths of some set of pairs of paths γ_j in a compatible collection $\mathscr{C}$. Note that $\Sigma_i \cap \sigma$ is contained in a union of such critical planes H of codimension i. If we arrange that $f(\Delta^l)$ is in general position with respect to such planes H, for all simplices Δ^l in D^k, then the lemma will follow since the codimension of $f(\Delta^l) \cap H$ in $f(\Delta^l)$ is at least as great as the codimension of H in $\mathbb{R}^m$, and codimension is preserved under pullbacks via linear maps.

If $f(\Delta^l)$ is contained entirely in σ, not just in the closure of σ, then the desired general position for $f(\Delta^l)$ is easy to achieve just by small perturbations of f on the vertices of Δ^l. The difficulty in the general case is that certain vertices of Δ^l may map to faces of σ, and f on these vertices can only be perturbed while staying within the respective faces. We need to show that general position can still be achieved with these constraints.

Let τ be an open simplex of $\mathbb{A}_n$ which is a face of σ. The simplex τ spans a linear subspace $\mathbb{R}^p$ of $\mathbb{R}^m$ obtained by letting the lengths of some edges of Γ go to zero, producing a quotient graph Γ'. We shall show first that the codimension of $H \cap \mathbb{R}^p$ in $\mathbb{R}^p$ equals the codimension of H in $\mathbb{R}^m$.

The critical plane H is the intersection of hyperplanes of the form $l_{\gamma_i} = l_{\gamma_j}$ for a compatible collection of paths γ_i. Since the γ_i are compatible, we can orient Γ by assigning each edge which is in some γ_i the orientation coming from γ_i and orienting edges not in any γ_i arbitrarily. To each hyperplane $l_{\gamma_i} = l_{\gamma_j}$ we then associate the cellular 1-cycle $\gamma_i - \gamma_j$. Such 1-cycles $\gamma_i - \gamma_j$ generate a subspace of $H_1(\Gamma; \mathbb{R})$ whose dimension is the codimension of H since there are no cellular 2-chains. The statement that the codimension of $H \cap \mathbb{R}^p$ in $\mathbb{R}^p$ equals the codimension of H in $\mathbb{R}^m$ now follows: the quotient map $q\colon \Gamma \longrightarrow \Gamma'$ takes the given collection of compatible preferred paths γ_i to another such collection γ_i' in Γ'; a hyperplane $l_{\gamma_i} = l_{\gamma_j}$ meets $\mathbb{R}^p$ in the hyperplane $l_{\gamma_i'} = l_{\gamma_j'}$, so the critical hyperplanes in $\mathbb{R}^m$ whose intersection is H meet $\mathbb{R}^p$ in critical hyperplanes whose intersection is $H \cap R^p$; and finally, the induced map $q_*\colon H_1(\Gamma; \mathbb{R}) \longrightarrow H_1(\Gamma'; \mathbb{R})$ is an isomorphism since q is a homotopy equivalence, so the subspace generated by the cycles $\gamma_i - \gamma_j$ is taken isomorphically to the subspace generated by the cycles $\gamma_i' - \gamma_j'$.

Now we can put $f(\Delta^l)$ into general position with respect to H by an inductive procedure. By our preliminary barycentric subdivision of D^k, $f(\Delta^l)$ is contained in the union of a chain of open simplices $\sigma_1, \ldots, \sigma_q = \sigma$ of $\mathbb{A}_n$, each a proper face of the next. Let V_i be the subspace of $\mathbb{R}^m$ spanned by σ_i and let $H_i = H \cap V_i$. Also, let Δ^{l_i} be the subsimplex of Δ^l spanned by the vertices mapping to V_i and let L_i be the subspace of $\mathbb{R}^m$ spanned by $f(\Delta^{l_i})$. To start the induction there is no difficulty in perturbing f on the vertices of Δ^{l_1} so that L_1 is in general position with respect to H_1 in V_1, as we noted before. By the preceding paragraph the codimension of H_i in V_i is independent of i, so L_1 is then automatically in general position with respect to H_2 in V_2. There is then no difficulty in perturbing f on vertices of $\Delta^{l_2} - \Delta^{l_1}$ so that L_2 is in general position with respect to H_2 in V_2. The process is then repeated, putting L_i into general position with respect to H_i in V_i for each i. In the end $f(\Delta^l)$ is in general position with respect to H in $\mathbb{R}^m$, as desired.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

Then one repeats the procedure for other critical planes H and other simplices Δ^l. Since general position is an open condition, if subsequent perturbations are sufficiently small they will not destroy a general position already created.

4.4. *Stage* 2: *Overview*

Having made the deformation of f in Lemma 4.4, all subsequent deformations of f will leave sufficiently small ϵ-cones unchanged.

The filtration of $\mathbb{A}_n$ defined by the subspaces Σ_i consisting of graphs of codimension at least i pulls back via f to a filtration of D^k by subpolyhedra $f^{-1}(\Sigma_i)$ of codimension at least i in D^k. Let $\mathscr{S}^i = f^{-1}(\Sigma_{k-i})$, so $\dim(\mathscr{S}^i) \leqslant i$ and $\mathscr{S}^0 \subset \mathscr{S}^1 \subset \ldots \subset \mathscr{S}^k = D^k$. We assume by induction on i that f has already been deformed so that the degree of Γ_s is at most k for s in a neighborhood of $\mathscr{S}^{i-1}$.

Let S be a component of $\mathscr{S}^i - \mathscr{S}^{i-1}$. For convenience, delete from S the points lying in a small open neighborhood of $\mathscr{S}^{i-1}$ so that the new S is a compact polyhedron with the degree of Γ_s at most k in a neighborhood of its frontier ∂S in $\mathscr{S}^i$. As s varies over S, the ϵ-cones of Γ_s do not change, since the codimension of critical points of Γ_s does not change.

A downward path in Γ_s from one critical point to another is called a *connecting path*. A connecting path may contain critical points in its interior. The number of connecting paths in Γ_s is its *complexity* c_s. We shall also be interested in the number e_s of connecting paths in Γ_s which are extended branches, containing no critical points in their interior. If $e_s > i$ for some $s \in S$, we will describe a procedure for deforming the family Γ_s by a homotopy supported in a neighborhood of S so that:

(1) The ϵ-cones of Γ_s do not change during the homotopy, and the degree of Γ_s is not increased.

(2) The maximum complexity of Γ_s over S is decreased.

Since the maximum complexity cannot decrease infinitely often, finitely many repetitions of the process will reduce us to the case that $e_s \leqslant i$ for all $s \in S$. We will then construct a further deformation splitting Γ_s to reduce its degree to at most k near S, thereby completing the induction step.

4.5. *Canonical splitting and extension to a neighborhood*

The set K of points $s \in S$ for which the complexity c_s is maximal is a closed subset, and in fact a subpolyhedron, since complexity changes only when lower endpoints of extended branches move off critical points, decreasing the complexity. Let K_0 be a connected component of K. Since complexity is constant over K_0, the extended branches of Γ_s are independent of s over K_0. Hence we can define a homotopy Γ_{st} for $s \in K_0$ which performs the canonical splitting of Γ_s. As s passes from K_0 to nearby points in S, the lower endpoints of extended branches of Γ_s can move off critical points, but sufficiently near K_0 they at least stay in the associated ϵ-cones. So we can extend the homotopy Γ_{st} over a neighborhood of K_0 in S by splitting the extended branches down to these ϵ-cones. As we move off S into D^k, each critical point x of Γ_s over S can bifurcate into several critical points. Let $H_x \subset \Gamma_s$ be the convex hull of these critical points, and H_x^ϵ be the downward ϵ-neighborhood of this hull, consisting of H_x together with points of Γ_s closer to the basepoint within distance ϵ of H_x. The extended branches of critical points of Γ_s for $s \in K_0$ can be viewed as attaching at their

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

lower endpoints to the H_x^ϵ for nearby s in D^k, so we can extend the homotopy Γ_{st} over a neighborhood of K_0 in D^k by splitting extended branches down to the H_x^ϵ where they attach.

This homotopy defined in a neighborhood of K_0 can be damped down near the boundary of the neighborhood so as to give a homotopy Γ_{st} defined throughout D^k and supported near K_0. For simplicity we rename the result of this homotopy again Γ_s.

4.6. *Reducing complexity by sliding in the ϵ-cones*

After the preceding homotopy, the attaching point α_j of a branch β_j of Γ_s lies either in an ϵ-cone C_x^ϵ or at the basepoint, for s in a neighborhood N_0 of K_0 in S. We shall construct a homotopy of Γ_s for $s \in N_0$ by perturbing the attaching points α_j which lie in ϵ-cones C_x^ϵ. Each such attaching point can be regarded as a map $\alpha_j : N_0 \longrightarrow C_x^\epsilon$, $s \mapsto \alpha_j(s)$, with $\alpha_j(s) = x$ for $s \in K_0$. Taking the product of these maps for all such branches gives a map $\alpha : N_0 \longrightarrow \Pi_j C_x^\epsilon$, the latter product consisting of copies of the various ϵ-cones C_x^ϵ, one copy of C_x^ϵ for each branch β_j attaching to C_x^ϵ. Inside this product is the point p whose coordinates are the critical points $x \in C_x^\epsilon$. For $s \in K_0$ we have $\alpha(s) = p$, and $\alpha(s) \neq p$ for $s \in N_0 - K_0$ since the complexity of Γ_s decreases once we leave K_0. We would like to perturb α to be in general position with respect to the point p, and argue that this implies that $\alpha(s) \neq p$ for all $s \in N_0$, lowering the complexity of Γ_s over K_0.

We reduce the problem to one of putting a map into general position with respect to a subcomplex of a manifold, then apply standard general position theory for manifolds. Choose an embedding of each ϵ-cone C_x^ϵ into a 2-dimensional disk D_x, such that C_x^ϵ meets ∂D_x in just the endpoints of the branches of C_x^ϵ. Since C_x^ϵ does not change over S, one such embedding works over all of S. Under the natural retraction $D_x \longrightarrow C_x^\epsilon$ the preimage of x is a star-shaped graph Y_x homeomorphic to C_x^ϵ itself (see Figure 4).

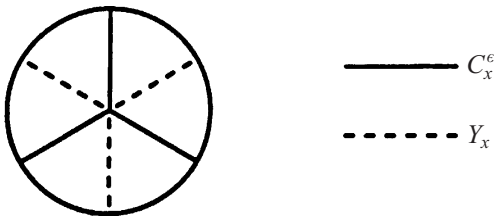


FIG. 4.

We regard α as a map $N_0 \longrightarrow \Pi_j D_x$, and perturb α to be in general position with respect to the subcomplex $Y = \Pi_j Y_x$. After this perturbation, the codimension of $\alpha^{-1}(Y)$ in N_0 is at least as large as the codimension of Y in $\Pi_j D_x$, which equals the number of subscripts j, that is, the number of branches attaching to ϵ-cones C_x^ϵ. This is exactly the number e_s since Γ_s has been canonically split over K_0. Thus if $e_s > i$, a perturbation of α makes its image disjoint from Y. Composing with the retractions of the disks D_x onto the ϵ-cones C_x^ϵ, we have then deformed attaching points of branches so that at no point in N_0 are all the attaching points α_j at critical points simultaneously.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

Any perturbation of the attaching points α_j so that not all of them are at the critical points reduces complexity. This is because there is a natural quotient map of the perturbed Γ_s onto the original Γ_s with all α_j at critical points, so the collection of connecting paths in the perturbed Γ_s can be regarded as a subcollection of the connecting paths of the original Γ_s. This will be a proper subcollection if at least one α_j is not at a critical point.

Thus we have succeeded in homotoping Γ_s over N_0 to decrease the maximum complexity there. We wish to extend this homotopy Γ_{st} over a neighborhood of K_0 in D^k. There is a canonical map $r_x: C_x^\epsilon \longrightarrow H_x^\epsilon$ taking endpoints of C_x^ϵ to the corresponding endpoints of H_x^ϵ, the point x to the barycenter of H_x^ϵ, and each segment of C_x^ϵ joining an endpoint to x to the segment in H_x^ϵ joining the corresponding endpoint to the barycenter. By composition with r_x, the homotopy of attaching points in the preceding paragraph for $s \in N_0$ gives a homotopy of attaching points for nearby s in D^k. For an attaching point $\alpha_j(s) \in H_x^\epsilon$ this homotopy starts at $r_x(\alpha_j(s))$ rather than at $\alpha_j(s)$ itself, but we can precede the homotopy by the homotopy from $\alpha_j(s)$ to $r_x(\alpha_j(s))$ along the canonical path in H_x^ϵ. If $s \in S$, this preliminary homotopy is constant.

Again we damp this homotopy Γ_{st} defined in a neighborhood of K_0 off near the boundary of the neighborhood to get a homotopy defined over all of D^k and supported near K_0. Doing this for all components K_0, the result is a new family Γ_s, whose maximum complexity on S has decreased. So by iteration, as described earlier, we eventually deform Γ_s to make $e_s \leqslant i$ over S.

4.7. *Reducing the degree*

It remains to reduce the degree to at most k near S. Define the *split-degree* $\mathrm{sdeg}(\Gamma_s)$ to be the degree of the canonical splitting of Γ_s. For a critical point x of Γ_s, let a_x be the number of extended branches which terminate at x, and b_x the number of downward directions from x. Then

$$\mathrm{sdeg}(\Gamma_s) = \sum_x (a_x + b_x - 2) \leqslant \sum_x a_x + \sum_{\mathrm{codim}(x) \geqslant 1} (b_x - 1) = e_s + \mathrm{codim}(\Gamma_s).$$

Our operations thus far have reduced e_s to at most i on S. Since the dimension of S is i, the codimension of graphs in the image of S is equal to $k-i$, by Lemma 4.4. Therefore the split-degree of Γ_s is at most k over S.

To go from split-degree at most k to degree at most k, choose a triangulation of S such that split-degree is constant over open simplices. Consider a simplex σ of this triangulation, and assume inductively that we have already reduced the degree to at most k near $\partial\sigma$. The discrepancy between degree and split-degree arises from extended branches which can be canonically split all the way down to the basepoint. The collection of extended branches having their lower endpoint at the basepoint does not change over the interior of σ since split-degree is constant there. We can apply the canonical splitting procedure from earlier in the proof just to these extended branches ending at the basepoint, damping this splitting off near $\partial\sigma$ where the degree is already at most k. In a neighborhood of σ in S, split-degree can decrease because more extended branches go down to the basepoint, but we ignore these and just split the extended branches which we split over σ. Similarly, as we pass to points near σ in D^k we continue to split the same collection of extended branches, regarding these as extending downward from hulls, as earlier in the proof. The resulting homotopy is then damped off as we move away from σ in D^k, finishing the induction step.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

5. *Rational homology*

5.1. *Definition of* $S\mathbb{A}_{n,k}$ *and* $Q_{n,k}$

In this section we will study the rational homology of $\mathrm{Aut}(F_n)$ by studying the action of $\mathrm{Aut}(F_n)$ on subcomplexes of the spine $\mathrm{S}\mathbb{A}_n$ of $\mathbb{A}_n$. The spine has computational advantages over the whole space since it is of lower dimension and its quotient by $\mathrm{Aut}(F_n)$ is compact.

The subcomplexes we use correspond to the subspaces $\mathbb{A}_{n,k}$ consisting of graphs of degree at most k. The deformation retraction of $\mathbb{A}_n$ onto $S\mathbb{A}_n$ restricts to a deformation retraction of $\mathbb{A}_{n,k}$ onto the full subcomplex $S\mathbb{A}_{n,k}$ of $S\mathbb{A}_n$ spanned by barycenters of open simplices in $\mathbb{A}_{n,k}$. It follows from Corollary 3.4 that $S\mathbb{A}_{n,k}$ is $(k-1)$-connected.

The group $\mathrm{Aut}(F_n)$ acts on $\mathbb{A}_n$, taking open simplices to open simplices, by

$$(\Gamma, v_0, \phi)\psi = (\Gamma, v_0, \phi\psi).$$

This action restricts to simplicial actions on the spines $S\mathbb{A}_n$ and $S\mathbb{A}_{n,k}$. Let Q_n be the quotient $S\mathbb{A}_n/\mathrm{Aut}(F_n)$ and $Q_{n,k}$ the quotient $S\mathbb{A}_{n,k}/\mathrm{Aut}(F_n)$, so that Q_n is filtered by the subspaces $Q_{n,k}$. Since marked graphs which differ by basepoint-preserving homeomorphisms of the graph are equivalent in $S\mathbb{A}_n$, and since the action of $\mathrm{Aut}(F_n)$ changes the markings ϕ arbitrarily, the vertices of the quotient Q_n are simply the homeomorphism types of basepointed graphs. An r-simplex of $S\mathbb{A}_n$ is given by a sequence of r forest collapses, where a *forest* is a disjoint union of trees. This sequence can be regarded as an increasing sequence of forests $\Phi_1 \subset \Phi_2 \subset \ldots \subset \Phi_r$ in the initial pointed marked graph Γ. Another sequence $\Phi_1' \subset \Phi_2' \subset \ldots \subset \Phi_r'$ in Γ is equivalent under the action of $\mathrm{Aut}(F_n)$ if and only if there is a basepoint-preserving homeomorphism of Γ taking Φ_i to Φ_i' for all i. Since all vertices of the simplex are graphs with distinct homeomorphism types, the action does not identify any points within a simplex, but only identifies different simplices together. Thus Q_n has a CW structure with each cell a simplex; however, this is not a simplicial complex structure since different simplices can have the same set of vertices, as we shall see in examples below. Since there are only finitely many topological types of basepointed graphs in $\mathbb{A}_n$, Q_n is a finite complex, as are its subcomplexes $Q_{n,k}$.

Every simplex of Q_n or $Q_{n,k}$ is a subsimplex of a simplex obtained by the following procedure. Choose a basepointed graph Γ in which all nonbasepoint vertices have valence 3, and choose a maximal tree T in Γ. If T has r edges, then we get an r-simplex by collapsing these edges one by one, in some order. Different orders may give the same or different r-simplices of Q_n, depending on whether they are equivalent under basepoint-preserving symmetries of Γ. In the case of $Q_{n,k}$, Γ will have k vertices other than the basepoint, so T will have k edges joining the $k+1$ vertices of Γ, and $Q_{n,k}$ will be k-dimensional.

As an example let us determine $Q_{n,2}$. Modulo wedging on circles at the basepoint, there are only five homeomorphism types of pairs (Γ, T) as in the previous paragraph, with the degree of Γ equal to 2. These are shown in Figure 5. In each case T has two edges. In graphs 1, 3 and 5 there is a symmetry switching the two edges of T, so only one 2-simplex is produced in these cases. For graphs 2 and 4 there are two 2-simplices, depending on the order of collapsing the two edges of T. Thus $Q_{n,2}$ has a total of seven 2-simplices if $n \geqslant 4$, and these are attached together as shown. Note that the 2-simplices 1 and 2b have two edges in common, so form a 'pocket', and similarly for

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

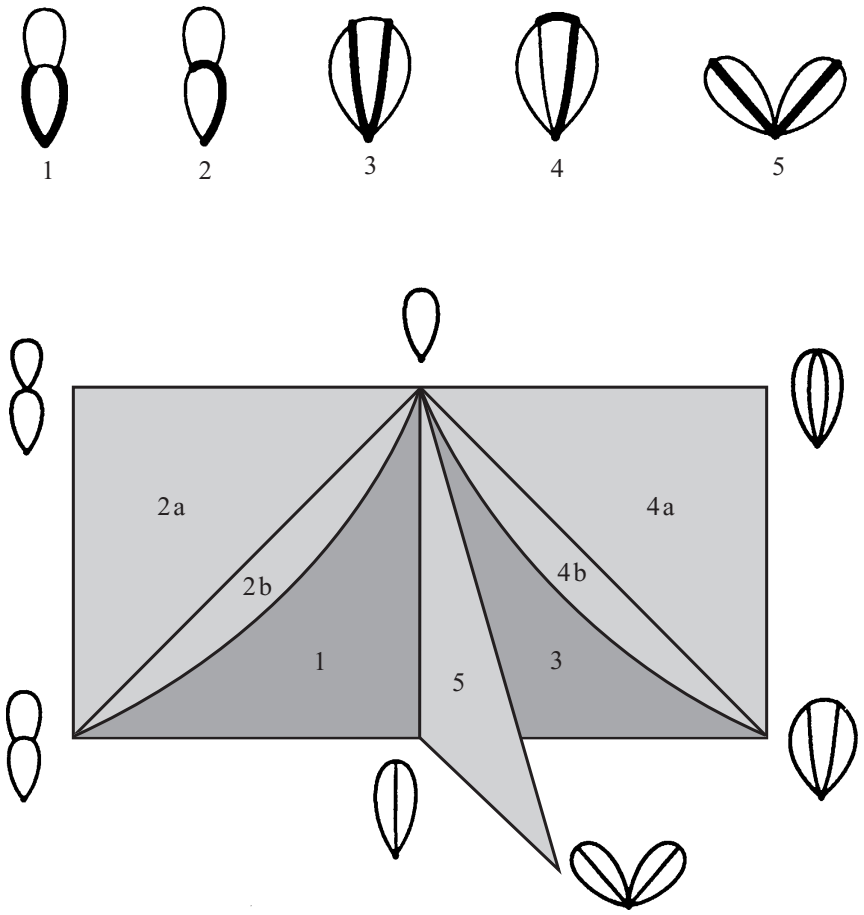


FIG. 5.

the 2-simplices 3 and 4b. When $n = 3$ the simplex labelled 5 is omitted, and when $n = 2$ the simplices 3 and 4ab are also omitted. Note that $Q_{n,2}$ is contractible for all n.

5.2. *Stability*

The stabilizer of a simplex of $S\mathbb{A}_n$ under the action of $\mathrm{Aut}(F_n)$ is contained in the symmetry group of a finite basepointed graph; in particular it is a finite group. From this finite-stabilizer property and the fact that $S\mathbb{A}_{n,k}$ is $(k-1)$-connected, it follows that $Q_{n,k}$ has the same rational homology as $\mathrm{Aut}(F_n)$ in dimensions less than k (see [**2**, Exercise 2, p. 174]). More precisely, we have the following lemma.

LEMMA 5.1. *$H_i(Q_{n,k}; \mathbb{Q})$ is isomorphic to $H_i(\mathrm{Aut}(F_n); \mathbb{Q})$ for $k > i$, and $H_k(Q_{n,k}; \mathbb{Q})$ maps onto $H_k(\mathrm{Aut}(F_n); \mathbb{Q})$.*

For example, since $Q_{n,2}$ is contractible for all n we deduce that $H_i(\mathrm{Aut}(F_n); \mathbb{Q}) = 0$ for $i = 1, 2$. This is consistent with the classical fact that $H_1(\mathrm{Aut}(F_n); \mathbb{Z}) = \mathbb{Z}_2$ and the calculations of $H_2(\mathrm{Aut}(F_n); \mathbb{Z})$ in [**4**, **7–9**].

Stability for the rational homology of $\mathrm{Aut}(F_n)$ will follow from the fact that the homotopy type of the complexes $Q_{n,k}$ is independent of n, for n sufficiently large with

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

respect to k. To prove this, we need the following observations on the relation between the rank, degree and homeomorphism type of a graph. Recall that a vertex of a graph is called a *cut vertex* if deleting it disconnects the graph.

LEMMA 5.2. *Let Γ be a graph of degree k and rank n.*

(i) *If $k < n/2$, then Γ has a loop at the basepoint.*

(ii) *If $k < 2n/3$ then Γ has either a loop at the basepoint or a theta graph wedge summand.*

(iii) *If $k < n-1$, then the basepoint is a cut vertex.*

Proof. Without loss of generality, we may assume that all vertices of Γ other than the basepoint are trivalent. Let Γ_1 be the full subgraph of Γ spanned by all nonbasepoint vertices. Then Γ_1 has $V_1 = k$ vertices and E_1 edges.

If Γ has no loops at the basepoint, there are $|v_0| = 2n-k$ edges not in E_1, and we compute $\chi(\Gamma) = 1-n = (1+V_1)-(E_1+|v_0|) = 1+k-E_1-(2n-k)$, so that $n = 2k-E_1$. In particular, $n \leqslant 2k$ and statement (i) follows. If $k < 2n/3$, then $E_1 = 2k-n < k/2$, so that Γ_1 has an isolated vertex, and statement (ii) follows. If $k < n-1$, then $E_1 = 2k-n < k-1$, so $\chi(\Gamma_1) = V_1-E_1 > k-(k-1) = 1$, and Γ_1 is disconnected, proving (iii).

We can embed $S\mathbb{A}_{n,k}$ in $S\mathbb{A}_{n+1,k}$ by sending (Γ, v_0, ϕ) to $(\Gamma \vee S^1, v_0, \phi')$, where ϕ' is ϕ extended to F_{n+1} by mapping the last generator to the simple loop S^1. There are two possibilities for ϕ', depending on the orientation chosen for S^1, but these differ by an isometry of $\Gamma \vee S^1$, and hence give the same point in $S\mathbb{A}_{n+1,k}$. This embedding of $S\mathbb{A}_{n,k}$ in $S\mathbb{A}_{n+1,k}$ induces an embedding $Q_{n,k} \longrightarrow Q_{n+1,k}$, sending a chain of forests in Γ corresponding to a cell in $Q_{n,k}$ to the same chain of forests in $\Gamma \vee S^1$. This embedding is onto if $n \geqslant 2k$ by statement (i) of Lemma 5.2, so we have:

PROPOSITION 5.3. *The map $Q_{n,k} \longrightarrow Q_{n+1,k}$ is a homeomorphism for $n \geqslant 2k$.*

For n between $3k/2$ and $2k$, the inclusions are no longer homeomorphisms, but we have:

PROPOSITION 5.4. *For $n \geqslant 3k/2$, the inclusion $Q_{n,k} \longrightarrow Q_{n+1,k}$ is a homotopy equivalence.*

Proof. A graph Γ which is in $Q_{n+1,k}$ but not in the image of $Q_{n,k}$ has degree at most k, rank $n+1 > 3k/2$ and no loops at the basepoint. Therefore, by statement (ii) of Lemma 5.2, contains $r \geqslant 1$ theta graph wedge summands. Let Θ be the union of all the theta graph summands of Γ, so that $\Gamma = \Theta \vee \Gamma'$, with $\text{rank}(\Gamma') = n+1-2r$. Let $\text{St}_{m,i}(\Gamma)$ denote the union of all the simplicial cells in $Q_{m,i}$ which contain Γ. Then $\text{St}_{n+1,k}(\Gamma)$ is the product of $\text{St}_{2r,r}(\Theta)$ with $\text{St}_{n+1-2r,k-r}(\Gamma')$. Since Θ has only one maximal chain of forests up to automorphisms of Θ, $\text{St}_{2r,r}(\Theta)$ consists of exactly one simplex, all of whose vertices other than Θ have lower degree. Shrinking one edge of Θ collapses this simplex, giving a deformation retraction of $\text{St}_{n+1,k}(\Gamma)$ into the image of $Q_{n-1,k-1}$ in $Q_{n+1,k}$.

The embedding $Q_{n,k} \longrightarrow Q_{n+1,k}$ is natural with respect to the inclusion $\text{Aut}(F_n) \longrightarrow \text{Aut}(F_{n+1})$. Therefore Propositions 5.3 and 5.4, and Lemma 5.1 imply the following theorem.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

THEOREM 5.5. *The map* $H_i(\mathrm{Aut}(F_n);\mathbb{Q}) \longrightarrow H_i(\mathrm{Aut}(F_{n+1});\mathbb{Q})$ *induced by inclusion is an isomorphism for* $n \geqslant 3(i+1)/2$.

6. *Free factorizations and split factorizations*

In this section we consider two simplicial complexes with natural $\mathrm{Aut}(F_n)$ actions, and show that they are highly connected. Both are geometric realizations of partially ordered sets (posets) of factorizations of F_n.

6.1. *Free factorizations and cut vertices*

We first consider factorizations $F_n = H_1 * \ldots * H_m$ where the H_i are nontrivial proper subgroups of F_n and two such factorizations which differ only by a permutation of the factors H_i are regarded as the same. The set of all such factorizations is partially ordered by refinement: $H \geqslant K$ if K can be obtained from H by splitting some of the factors of H. Define FF_n to be the geometric realization of this poset. Since the number of factors can range from 2 to n, FF_n has dimension $n-2$. We will prove:

THEOREM 6.1. FF_n *is* $(n-3)$*-connected, hence is homotopy equivalent to a wedge of* $(n-2)$*-spheres*.

Free factorizations of F_n are mirrored geometrically by pointed marked graphs whose basepoint is a cut vertex. With this in mind, we define a subcomplex C_n of $S\mathbb{A}_n$ by $C_n = \{(\Gamma, v_0, \phi)\colon \Gamma - v_0 \text{ is disconnected}\}$.

We first prove the geometric version of Theorem 6.1.

PROPOSITION 6.2. C_n *is* $(n-3)$*-connected.*

Proof. Consider a map $S^i \longrightarrow C_n$ where $i \leqslant n-3$. Since $S\mathbb{A}_n$ is contractible, this extends to a map $D^{i+1} \longrightarrow S\mathbb{A}_n$. By the Degree Theorem, this extension can be deformed into $S\mathbb{A}_{n,n-2}$. By part (iii) of Lemma 5.2, the base vertex of any pointed graph of rank n and degree at most $n-2$ is a cut vertex, so that $S\mathbb{A}_{n,n-2} \subset C_n$, proving the proposition.

To go from complexes of graphs to complexes of splittings of F_n, we use Quillen's Theorem A [**10**]. If X and Y are posets, a map $f\colon X \longrightarrow Y$ is called a *poset map* if $f(x_1) \leqslant f(x_2)$ whenever $x_1 \leqslant x_2$. For $y \in Y$, the *fiber* $f_{\leqslant y}$ is the set of $x \in X$ with $f(x) \leqslant y$.

THEOREM 6.3 (Quillen's Theorem A). *A poset map* $f\colon X \longrightarrow Y$ *is a homotopy equivalence if and only if the fibers* $f_{\leqslant y}$ *are contractible for all* $y \in Y$.

Proof of Theorem 6.1. For (Γ, v_0, ϕ) a vertex of C_n, let $\Gamma_1, \ldots, \Gamma_k$ be the closures of the connected components of $\Gamma - v_0$. The map $f\colon C_n \longrightarrow FF_n$ sending (Γ, v_0, ϕ) to the factorization $\phi^{-1}(\pi_1(\Gamma_1, v_0)) * \ldots * \phi^{-1}(\pi_1(\Gamma_k, v_0))$ of F_n is a poset map, since collapsing a forest of Γ can only refine the associated factorization.

Let $H_1 * H_2 * \ldots * H_k$ be a factorization of F_n. The fiber $f_{\leqslant H_1 * H_2 * \ldots * H_k}$ is the join of the contractible complexes $S\mathbb{A}_{n_i}$, where n_i is the rank of H_i, so is contractible. We now apply Quillen's Theorem A to prove the theorem.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

6.2. *Loops at the basepoint and split factorizations*

To prove integral homology stability, it is convenient to consider the simplicial complex SF_n whose k-simplices are the (unordered) factorizations $\mathbb{Z}_0 * \ldots * \mathbb{Z}_k * H$ of F_n where each $\mathbb{Z}_i$ is a copy of $\mathbb{Z}$. The various faces of dimension $k-1$ of such a k-simplex are obtained by amalgamating one $\mathbb{Z}_i$ with H. The maximal simplices of SF_n are the factorizations $F_n = \mathbb{Z}_0 * \ldots * \mathbb{Z}_{n-1}$, so SF_n is $(n-1)$-dimensional.

PROPOSITION 6.4. *SF_n is $[(n-3)/2]$-connected.*

Proof. We first compare SF_n with the subcomplex SF'_n of FF_n defined by the poset of free factorizations of F_n having at most one noncyclic factor. A vertex $\mathbb{Z}_0 * \ldots * \mathbb{Z}_k * H$ of SF'_n is uniquely determined by the adjacent vertices $\mathbb{Z}_0 * H_0, \ldots, \mathbb{Z}_k * H_k$, where H_i is the free product of H with all of the cyclic factors except $\mathbb{Z}_i$. This means that SF'_n is a subcomplex of the barycentric subdivision of SF_n. Since cyclic factors H are permitted in free factorizations $\mathbb{Z}_0 * \ldots * \mathbb{Z}_k * H$ for SF_n but not for SF'_n, SF'_n intersects each $(n-1)$-simplex of SF_n in the cone on the $(n-3)$-skeleton of the $(n-1)$-simplex. The simplex deformation retracts onto this cone, so SF_n deformation retracts onto SF'_n since each $(n-2)$-simplex of SF_n is a face of only one $(n-1)$-simplex.

Next we compare SF'_n with the subcomplex L_n of $S\mathbb{A}_n$ consisting of pointed marked graphs with at least one loop at the basepoint. Let (Γ, v_0, ϕ) be a point in L_n, so that Γ is the wedge sum of circles $c_1, \ldots, c_k$ and a graph Γ_0 with no loops at the basepoint. A map $f: L_n \longrightarrow SF'_n$ is defined by sending (Γ, v_0, ϕ) to $\phi^{-1}(\pi_1(c_1, v_0)) * \ldots * \phi^{-1}(\pi_1(c_k, v_0)) * \phi^{-1}(\pi_1(\Gamma_0, v_0))$. This is a poset map, since a forest collapse can only split off more $\mathbb{Z}$-factors from $\pi_1(\Gamma_0, v_0)$. The fibers $f_{\leqslant H}$ are isomorphic to $S\mathbb{A}_{n_0}$, where n_0 is the rank of $\pi_1(\Gamma_0)$. Since these fibers are contractible, f is a homotopy equivalence by Quillen's Theorem A.

It remains to show that L_n is $[(n-3)/2]$-connected. By Lemma 5.2, if (Γ, v_0) is a pointed graph of rank n and degree $k < n/2$, then Γ has at least one loop at the basepoint. By the Degree Theorem, a map $D^k \longrightarrow S\mathbb{A}_n$ extending $S^{(k-1)} \longrightarrow L_n$ can be deformed into L_n as long as $k < n/2$, that is, L_n is $([n-1)/2]-1) = [(n-3)/2]$-connected.

We do not know whether SF_n is more highly connected than the proposition states, but even if it is, this would not improve the homological stability range we obtain in the next section.

7. *Integral homology stability*

In [**5**], it is shown that the inclusion $\mathrm{Aut}(F_n) \longrightarrow \mathrm{Aut}(F_{n+1})$ induces an isomorphism on homology in dimension i if $n > i^2/4 + 2i - 1$. The quadratic dependence of n on i was surprising in light of similar bounds for homology stability of various linear groups and mapping class groups, which are linear in i. In this section we use a spectral sequence argument to show that the stability range can be improved as follows.

THEOREM 7.1. *The map $H_i(\mathrm{Aut}(F_n); \mathbb{Z}) \longrightarrow H_i(\mathrm{Aut}(F_{n+1}); \mathbb{Z})$ induced by inclusion is surjective for $n \geqslant 2i+1$ and an isomorphism for $n \geqslant 2i+3$.*

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

REMARK 7.2. In [**5**] it is also shown that the natural map $\mathrm{Aut}(F_n) \longrightarrow \mathrm{Out}(F_n)$ induces an isomorphism on homology in dimension i if $n > i^2/4 + 5i/2$. The above result does not improve the homology stability range for $\mathrm{Out}(F_n)$.

7.1. *Construction of the spectral sequence*

To streamline the notation in what follows, we will denote $\mathrm{Aut}(F_n)$ by A_n.

A_n acts on the complex SF_n of 'split factorizations' of F_n defined at the end of the last section. This action gives rise to a spectral sequence as follows.

Let $C_{n-1} \longrightarrow C_{n-2} \longrightarrow \ldots \longrightarrow C_0 \longrightarrow \mathbb{Z}$ be the augmented simplicial chain complex of SF_n. Since SF_n is $m = [(n-3)/2]$-connected, the truncated chain complex

$$0 \longrightarrow Z_m \longrightarrow C_m \longrightarrow \ldots \longrightarrow C_0 \longrightarrow \mathbb{Z} \longrightarrow 0$$

is exact, where Z_m is the subgroup of cycles in dimension m. We denote this truncated complex by C_*, with $C_{-1} = \mathbb{Z}$ and $C_{m+1} = Z_m$. Let

$$\ldots \longrightarrow E_2 A_n \longrightarrow E_1 A_n \longrightarrow E_0 A_n \longrightarrow \mathbb{Z}$$

be a free $\mathbb{Z}A_n$-resolution of $\mathbb{Z}$. Tensoring C_* over $\mathbb{Z}A_n$ with $E_* A_n$ gives a double complex:

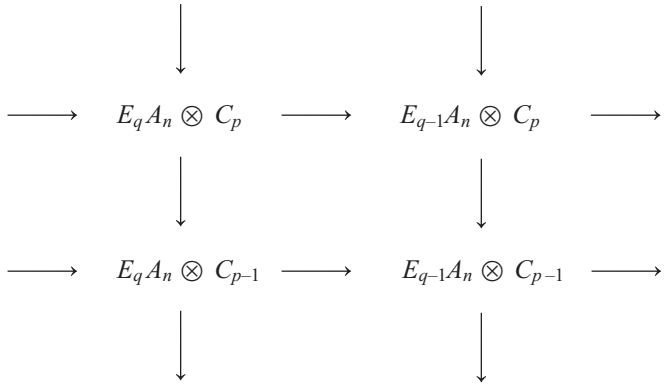


FIG. 6.

with $q \geqslant 0$ and $-1 \leqslant p \leqslant m+1$. The horizontal arrows are $d \otimes 1$, where d is the differential in $E_* A_n$, and the vertical arrows are $1 \otimes (-1)^q \partial$, where ∂ is the differential in C_*. The columns are exact since C_* is exact and $E_q A_n$ is a free $\mathbb{Z}A_n$-module. Therefore the spectral sequence associated to the vertical filtration of this complex converges to zero. Taking the spectral sequence associated to the horizontal filtration gives another spectral sequence converging to zero, whose $E^1_{p,q}$-term is $H_q(A_n; C_p)$.

For each n, fix an ordered basis $x_1, \ldots, x_n$ for F_n, so that F_n appears naturally as the subgroup of F_{n+1} generated by the first n basis elements. A_n acts transitively on p-simplices of SF_n, so that

$$C_p \cong \mathbb{Z}A_n \otimes_{\mathbb{Z}S_p} \mathbb{Z}_\sigma \text{ for } 0 \leqslant p \leqslant m$$

where S_p is the stabilizer of the standard p-cell $\langle x_1 \rangle * \langle x_2 \rangle * \ldots * \langle x_{p+1} \rangle * \langle x_{p+2}, \ldots, x_n \rangle$, and $\mathbb{Z}_\sigma$ is $\mathbb{Z}$ with the action of an element g of S_p being plus or minus the identity depending on whether g preserves or reverses the orientation of the standard p-cell (see [**2**, Example 5.5(b), p. 68]). By Shapiro's lemma, the inclusion of S_p into A_n induces an isomorphism $H_q(S_p; \mathbb{Z}_\sigma) \cong H_q(A_n; C_p)$ [**2**, p. 73].

The stabilizer S_p is equal to $\Sigma_{p+1} \times A^+_{n-p-1}$, where Σ_{p+1} is generated by

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

automorphisms which permute and invert the first $p+1$ basis elements of F_n, and A^+_{n-p-1} is the group of automorphisms of the subgroup of F_n generated by the last $n-p-1$ basis elements. We can now identify the $E^1_{p,q}$ terms of the horizontal spectral sequence as

$$E^1_{p,q} = H_q(\Sigma_{p+1} \times A^+_{n-p-1}; \mathbb{Z}_\sigma).$$

Note that elements of A^+_{n-p-1} act trivially on the module $\mathbb{Z}_\sigma$, as do inversions of the first $p+1$ basis elements. We can apply the Künneth formula [**2**, Exercise, p. 106] to compute

$$H_q(\Sigma_p \times A^+_{n-p-1}; \mathbb{Z}_\sigma) \cong \bigoplus_{s+t=q} H_s(\Sigma_{p+1}; \mathbb{Z}_\sigma) \otimes H_t(A^+_{n-p-1}; \mathbb{Z})$$
$$\oplus \bigoplus_{s+t=q-1} \mathrm{Tor}(H_s(\Sigma_{p+1}; \mathbb{Z}_\sigma), H_t(A^+_{n-p-1}; \mathbb{Z})).$$

The inclusion $F_n \longrightarrow F_{n+1}$ induces a map of the corresponding complexes $SF_n \longrightarrow SF_{n+1}$, mapping a p-cell $\langle w_1 \rangle * \ldots * \langle w_{p+1} \rangle * H$ to $\langle w_1 \rangle * \ldots * \langle w_{p+1} \rangle * (H * \langle x_{n+1} \rangle)$. We denote by C'_* the truncated chain complex of SF_{n+1}. The inclusion $A_n \longrightarrow A_{n+1}$ sends $\Sigma_{p+1} \times A^+_{n-p-1}$ to $\Sigma_{p+1} \times A^+_{n-p}$ by the identity on the first factor and the standard inclusion on the second factor. We form a relative spectral sequence from the double complex (Figure 7).

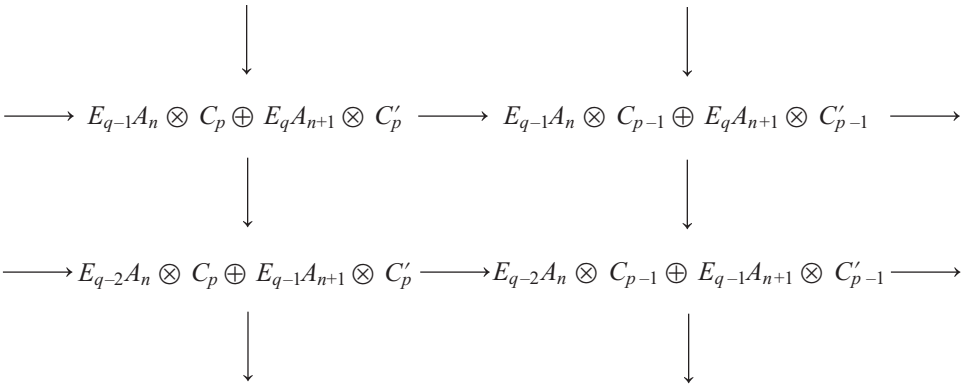


FIG. 7.

The columns are the mapping cones of the chain map $E_*A_n \otimes C_p \longrightarrow E_*A_{n+1} \otimes C_p$. The $E^1_{p,q}$-terms of the spectral sequence are

$$E^1_{p,q} = H_q(\Sigma_{p+1} \times A^+_{n-p}, \Sigma_{p+1} \times A^+_{n-p-1}; \mathbb{Z}_\sigma)$$
$$= H_q((\Sigma_{p+1}, \varnothing) \times (A^+_{n-p}, A^+_{n-p-1}); \mathbb{Z}_\sigma).$$

By the Künneth formula (see [**11**, pp. 234–235]) this becomes

$$E^1_{p,q} \cong \bigoplus_{s+t=q} H_s(\Sigma_{p+1}; \mathbb{Z}_\sigma) \otimes H_t(A^+_{n-p}, A^+_{n-p-1}; \mathbb{Z})$$
$$\oplus \bigoplus_{s+t=q-1} \mathrm{Tor}(H_s(\Sigma_{p+1}; \mathbb{Z}_\sigma), H_t(A^+_{n-p}, A^+_{n-p-1}; \mathbb{Z})).$$

7.2. *Stability with slope* 3

We first show that $H_i(A_n; \mathbb{Z}) \longrightarrow H_i(A_{n+1}; \mathbb{Z})$ is surjective for $n \geqslant 3i+1$ and an isomorphism for $n \geqslant 3i+2$.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

We prove this by induction on i. Specifically, we assume that for $k < i$ the following statements hold:

$(a)_k$ For $n \geqslant 3k$, $H_k(A_n) \longrightarrow H_k(A_{n+1})$ is surjective.

$(b)_k$ For $n \geqslant 3k+1$, $H_k(A_n) \longrightarrow H_k(A_{n+1})$ is an isomorphism.

In particular, $(a)_k$ and $(b)_k$ imply that $H_k(A_{n+1}, A_n) = 0$ for $n \geqslant 3k$. In the spectral sequence constructed above, the map $d^1\colon E^1_{0,i} \longrightarrow E^1_{-1,i}$ is

$$H_i(A_n^+, A_{n-1}^+; \mathbb{Z}) \longrightarrow H_i(A_{n+1}, A_n; \mathbb{Z})$$

since $H_0(\Sigma_{p+1}; \mathbb{Z}_\sigma) = \mathbb{Z}$ for $p \leqslant 0$, the coefficients $\mathbb{Z}_\sigma$ being untwisted in this case. By induction the terms $E^1_{s,t}$ with $t \leqslant i-1$ and $s+t = i$ are all zero. Since the only differentials that can ever hit $E^k_{-1,i}$ originate in $E^k_{s,t}$ with $s+t = i$, and since the spectral sequence converges to zero, this d^1 map must be onto.

The subgroup A_n^+ of A_{n+1} is conjugate to the standard A_n by an automorphism of F_{n+1} (sending x_i to x_{i+1} for $i \leqslant n$ and x_{n+1} to x_1). Therefore the composition $A_n \longrightarrow A_n^+ \longrightarrow A_{n+1}$ is equal to the standard inclusion $A_n \longrightarrow A_{n+1}$ on the level of homology.

A diagram chase of Figure 8 now shows $H_i(A_{n+1}, A_n) = 0$, which implies $(a)_i$.

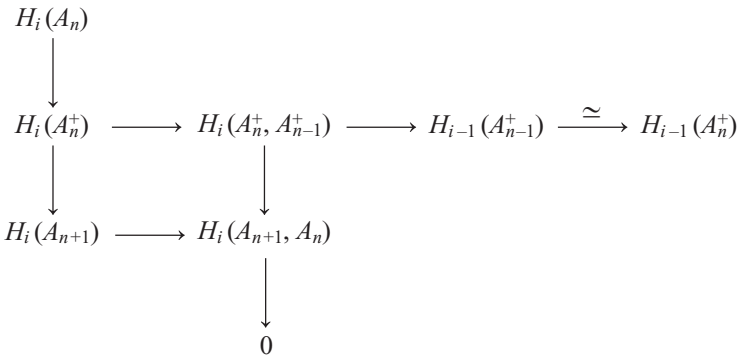


FIG. 8.

To prove $(b)_i$, we consider the same spectral sequence, assuming $n \geqslant 3i+1$. By $(a)_i$, the terms $E^1_{s,t}$ with $t \leqslant i$ and $s+t = i+1$ are all zero, so that the differential $d^1\colon E^1_{0,i+1} \longrightarrow E^1_{-1,i+1}$ is onto.

We now chase Figure 9 to prove $(b)_i$:

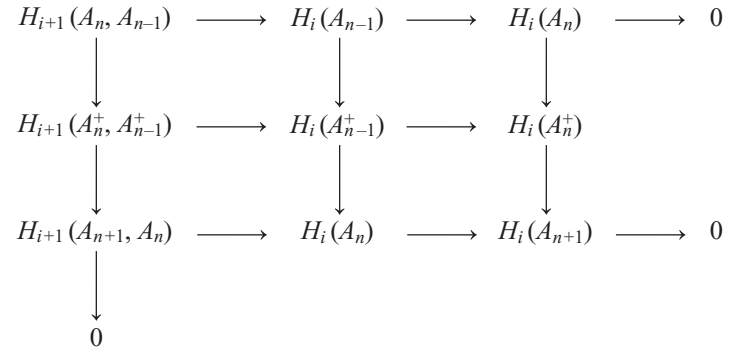


FIG. 9.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

7.3. *Stability with slope* 2

Finally, we improve the stability range to prove the theorem. Specifically, we show that $H_i(A_{n+1}, A_n) = 0$ for $n \geqslant 2i+1$, so that the map $H_i(A_n; \mathbb{Z}) \longrightarrow H_i(A_{n+1}; \mathbb{Z})$ induced by inclusion is surjective for $n \geqslant 2i+1$ and an isomorphism for $n \geqslant 2i+3$.

Assume by induction that $H_j(A_k, A_{k-1}) = 0$ for $j < i$ and $k \geqslant 2j+1$. We know that $H_i(A_n, A_{n-1})$ is zero for n sufficiently large with respect to i, either by the slope 3 argument above or by the theorem of [**5**]. Therefore it suffices to show the map $H_i(A_n^+, A_{n-1}^+) \longrightarrow H_i(A_{n+1}, A_n)$ induced by inclusion is injective for $n \geqslant 2i+1$.

This map is the differential $d^1\colon E^1_{0,i} \longrightarrow E^1_{-1,i}$ in the relative spectral sequence for the complex SF_{n+1} constructed in the previous section. By induction, the terms $E^1_{s,t}$ with $t < i$ and $s+t = i+1$ are zero, so the only differential with a chance of killing $E^1_{0,i}$ is $d^1\colon E^1_{1,t} \longrightarrow E^1_{0,i}$. Since $E^\infty_{0,i} = 0$, injectivity of $d^1\colon E^1_{0,i} \longrightarrow E^1_{-1,i}$ will follow if we show that $d = d^1\colon E^1_{1,i} \longrightarrow E^1_{0,i}$ is the zero map.

Recall that $E^1_{1,i} = H_i(A_{n+1}, A_n; C_1)$ and $E^1_{0,i} = H_i(A_{n+1}, A_n; C_0)$. The differential d is given on the chain level by $1 \otimes \partial$, where $\partial\colon C_1 \longrightarrow C_0$ is the boundary map on the chains of SF_{n+1}. In particular, if e_0 is the standard 1-cell $\langle x_1 \rangle * \langle x_2 \rangle * \langle x_3, \ldots, x_{n+1} \rangle$ and v_0 is the standard 0-cell $\langle x_1 \rangle * \langle x_2, \ldots, x_{n+1} \rangle$, then $\partial e_0 = v_0 - \alpha v_0$, where α is the automorphism interchanging x_1 and x_2 and fixing the other generators.

Recall also that $S_1 = \Sigma_2 \times A_{n-1}^+$ is the stabilizer of e_0 in A_{n+1}. By Shapiro's lemma, the map

$$S\colon H_i(\Sigma_2 \times (A_{n-1}^+, A_{n-2}^+); \mathbb{Z}_\sigma) \longrightarrow H_i(A_{n+1}, A_n; C_1)$$

induced by inclusion is an isomorphism.

By induction, all of the terms in the Künneth expansion of $H_i(\Sigma_2 \times (A_{n-1}^+, A_{n-2}^+); \mathbb{Z}_\sigma)$ vanish except one, and we are left with an isomorphism

$$K\colon H_0(\Sigma_2; \mathbb{Z}_\sigma) \otimes H_i(A_{n-1}^+, A_{n-2}^+) \longrightarrow H_i(\Sigma_2 \times (A_{n-1}^+, A_{n-2}^+); \mathbb{Z}_\sigma)$$

given by the homology cross product.

We will compute the maps K, S and d on the chain level, and show that the composition $K \circ S \circ d$ is zero on homology; since K and S are isomorphisms, this will show that d is zero.

Let F be a free resolution of $\mathbb{Z}$ over $\mathbb{Z}S_1$. Then F is naturally a free resolution over $\mathbb{Z}\Sigma_2$ and over $\mathbb{Z}A_{n-1}^+$ by restriction, so that $F \otimes F$ is a free resolution of $\mathbb{Z}$ over $\mathbb{Z}(\Sigma_2 \times A_{n-1}^+)$, where the action is given by $(x \otimes y)(\sigma, g) = x\sigma \otimes yg$. On the chain level, the map

$$K\colon (F_0 \otimes_{\Sigma_2} \mathbb{Z}) \otimes (F_i \otimes_{A_{n-1}^+} \mathbb{Z}) \longrightarrow (F_0 \otimes F_i) \otimes_{(\Sigma_2 \times A_{n-1}^+)} \mathbb{Z}$$

is given by $(x \otimes s) \otimes (y \otimes t) \mapsto (x \otimes y) \otimes st$.

From the resolution $F \otimes F$ for S_1, we get a free resolution for A_{n+1} by inducing up: $F' = (F \otimes F) \otimes_{S_1} \mathbb{Z}A_{n+1}$. On the chain level, the map

$$S\colon F \otimes F \otimes_{S_1} \mathbb{Z} \longrightarrow F' \otimes_{A_{n+1}} C_1 = (F \otimes F \otimes_{S_1} \mathbb{Z}A_{n+1}) \otimes_{A_{n+1}} (\mathbb{Z}A_{n+1} \otimes_{S_1} \mathbb{Z})$$

is given by $(x \otimes y) \otimes_{S_1} st \mapsto (x \otimes y \otimes_{S_1} 1) \otimes_{A_{n+1}} (1 \otimes_{S_1} st)$.

Finally, we compute d on the chain level. For any free resolution F' of $\mathbb{Z}$ over $\mathbb{Z}A_{n+1}$, d is given on the chain level $F'_i \otimes C_1 \longrightarrow F'_i \otimes C_0$ by $x \otimes e \mapsto x \otimes \partial e$, where ∂ is the boundary map $C_1 \longrightarrow C_0$. Each edge e is equal to ge_0 for some $g \in A_{n+1}$, so that

$$x \otimes e = x \otimes ge_0 \mapsto x \otimes (gv_0 - g\alpha v_0) = (xg - xg\alpha) \otimes v_0$$

that is, the map $1 \otimes \partial$ is the same as the map $(1 - \alpha) \otimes 1$.

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press

Note that $\alpha \in \Sigma_2 \leqslant S_1$. Thus when we apply d to $(x \otimes y \otimes_{S_1} 1) \otimes_A (1 \otimes_{S_1} st)$ we get

$$\begin{aligned} &((x \otimes y \otimes_{S_1} 1) - (x \otimes y \otimes_{S_1} \alpha)) \otimes_{A_{n+1}} (1 \otimes_{S_1} st) \\ &= ((x \otimes y \otimes_{S_1} 1) - (x\alpha \otimes y \otimes_{S_1} 1)) \otimes_{A_{n+1}} (1 \otimes_{S_1} st) \\ &= ((x - x\alpha) \otimes y \otimes_{S_1} 1) \otimes_{A_{n+1}} (1 \otimes_{S_1} st). \end{aligned}$$

But $x - x\alpha = 0$ in $H_0(\Sigma_2; Z) \cong \mathbb{Z}/2\mathbb{Z}$ since $\alpha \in \Sigma_2$ acts nontrivially on $\mathbb{Z}$; thus the last expression is zero on homology.

References

1. A. Borel, 'Stable real cohomology of arithmetic groups', *Ann. Sci. École Norm. Sup.* 7 (1974) 235–272.
2. K. S. Brown, *Cohomology of groups* (Springer, New York, 1982).
3. M. Culler and K. Vogtmann, 'Moduli of graphs and automorphisms of free groups', *Invent. Math.* 84 (1986) 91–119.
4. S. M. Gersten, 'A presentation for the special automorphism group of a free group', *J. Pure Appl. Algebra* 33 (1984) 269–279.
5. A. Hatcher, 'Homology stability for groups of automorphisms of free groups', *Comment. Math. Helv.* 70 (1995) 39–62.
6. A. Hatcher and K. Vogtmann, 'Rational homology of Aut(F_n)', *Math. Res. Lett.* (to appear).
7. J. Kiralis, 'The second homology groups of $\mathrm{Aut}_+(F_3)$ and $\mathrm{Aut}_+(F_4)$', *Proc. Amer. Math. Soc.* (to appear).
8. E. Ya. Markushina, 'Universal central extensions of groups of proper automorphisms of free groups', *Math. Notes* 47 (1990) 342–348.
9. E. Ya. Markushina, 'Schur multipliers of the automorphism groups of free groups', *Algebra and Logic* 29 (1990) 370–376.
10. D. Quillen, 'Homotopy properties of the poset of p-subgroups of a finite group', *Adv. Math.* 28 (1978) 129–137.
11. E. H. Spanier, *Algebraic topology* (McGraw–Hill, New York, 1966).

Department of Mathematics
Cornell University
Ithaca
NY 14853-7901
USA

https://doi.org/10.1112/S0024610798006644 Published online by Cambridge University Press