THE COMPLEX OF FREE FACTORS OF A FREE GROUP

ALLEN HATCHER[1] AND KAREN VOGTMANN[1]

Abstract. We show that the geometric realization of the partially ordered set of proper free factors in a finitely generated free group of rank n is homotopy equivalent to a wedge of spheres of dimension $n-2$.

The original version of this paper, published in 1998 in the Oxford Quarterly, contained an error, and the purpose of the present version is to provide a correction. The main results of the paper are unchanged. The error was in Lemma 2.3 which was used in the proof of Theorem 2.5. Unfortunately, the statement of Lemma 2.3 was false; it is replaced here with a new lemma which suffices to prove Theorem 2.5 after a few small changes in the proof. Section 2 has been reorganized to accommodate the changes, and we have taken the occasion to make some further clarifications in this section.

1. Introduction

An important tool in the study of the group $GL(n,\mathbb{Z})$ is provided by the geometric realization of the partially ordered set (poset) of proper direct summands of $\mathbb{Z}^n$. The natural inclusion $\mathbb{Z}^n \to \mathbb{Q}^n$ gives a one-to-one correspondence between proper direct summands of $\mathbb{Z}^n$ and proper subspaces of $\mathbb{Q}^n$, so that this poset is isomorphic to the spherical building X_n for $GL(n,\mathbb{Q})$. The term "spherical" comes from the Solomon-Tits theorem [10], which says that X_n has the homotopy type of a bouquet of spheres:

Theorem (Solomon-Tits Theorem)**.** *The geometric realization of the poset of proper subspaces of an n-dimensional vector space has the homotopy type of a bouquet of spheres of dimension $n-2$.*

The building X_n encodes the structure of parabolic subgroups of $GL(n,\mathbb{Q})$: they are the stabilizers of simplices. X_n also parametrizes the Borel-Serre boundary of the homogeneous space for $GL(n,\mathbb{R})$. The top-dimensional homology $H_{n-2}(X_n)$ is the Steinberg module I_n for $GL(n,\mathbb{Q})$, and is a dualizing module for the homology of $GL(n,\mathbb{Z})$, i.e. for all coefficient modules M there are isomorphisms

$$H^i(GL(n,\mathbb{Z});M) \to H_{d-i}(GL(n,\mathbb{Z});M \otimes I_n),$$

where $d = n(n-1)/2$ is the virtual cohomological dimension of $GL(n,\mathbb{Z})$.

If one replaces $GL(n,\mathbb{Z})$ by the group $Aut(F_n)$ of automorphisms of the free group of rank n, the natural analog FC_n of X_n is the geometric realization of the poset of proper free factors of F_n. The abelianization map $F_n \to \mathbb{Z}^n$ induces a map from FC_n to the poset of summands of $\mathbb{Z}^n$. In this paper we prove the analog of the Solomon-Tits theorem for FC_n:

Theorem 1.1. *The geometric realization of the poset of proper free factors of F_n has the homotopy type of a bouquet of spheres of dimension $n-2$.*

[1]Partially supported by NSF grant DMS-9307313.

By analogy, we call the top homology $H_{n-2}(FC_n)$ the *Steinberg module* for $Aut(F_n)$. This leaves open some intriguing questions. It has recently been shown that $Aut(F_n)$ is a virtual duality group [1]; does the Steinberg module act as a dualizing module? [This is answered in the negative for $n = 5$ in the preprint [5] posted in 2022.] There is an analog, called Autre space, of the homogeneous space for $GL(n, \mathbb{Z})$ and the Borel-Serre boundary; what is the relation between this and the "building" of free factors?

In [9], Quillen developed tools for studying the homotopy type of the geometric realization $|X|$ of a poset X. Given an order-preserving map $f\colon X \to Y$ (a "poset map"), there is a spectral sequence relating the homology of $|X|$, the homology of $|Y|$, and the homology of the "fibers" $|f/y|$, where

$$f/y = \{x \in X | f(x) \leq y\}$$

with the induced poset structure.

To understand FC_n then, one might try to apply Quillen's theory using the poset map $FC_n \to X_n$. However, it seems to be difficult to understand the fibers of this map. Instead, we proceed by modeling the poset of free factors topologically, as the poset B_n of simplices of a certain subcomplex of the "sphere complex" $S(M)$ studied in [2]. There is a natural poset map from B_n to FC_n; we compute the homotopy type of B_n and of the Quillen fibers of the poset map, and apply Quillen's spectral sequence to obtain the result.

2. Sphere systems

Let M be the compact 3-manifold obtained by taking a connected sum of n copies of $S^1 \times S^2$ and removing the interior of a closed ball. A *sphere system* in M is a non-empty finite set of disjointly embedded 2-spheres in the interior of M, no two of which are isotopic, and none of which bounds a ball or is isotopic to the boundary sphere of M. The complex $S(M)$ of sphere systems in M is defined to be the simplicial complex whose k-simplices are isotopy classes of sphere systems with $k + 1$ spheres.

Fix a basepoint p on ∂M. The fundamental group $\pi_1(M, p)$ is isomorphic to F_n. Any automorphism of F_n can be realized by a homeomorphism of M fixing ∂M. A theorem of Laudenbach [6] implies that such a homeomorphism inducing the identity on $\pi_1(M, p)$ acts trivially on isotopy classes of sphere systems, so that in fact $Aut(F_n)$ acts on $S(M)$.

For H a subset of $\pi_1(M, p)$, define S_H to be the subcomplex of $S(M)$ consisting of isotopy classes of sphere systems S such that $\pi_1(M - S, p) \supseteq H$. When H is trivial, S_H is $S(M)$, and in this case the following result was proved in [2].

Theorem 2.1. *The complex S_H is contractible for each $n \geq 1$.*

The proof will be a variant of the proof in [2], using the following fact.

Lemma 2.2. *Any two simplices in S_H can be represented by sphere systems Σ and S such that every element of H is representable by a loop disjoint from both Σ and S.*

Proof. Enlarge Σ to a maximal sphere system Σ', so the components of $M - \Sigma'$ are three-punctured spheres. By Proposition 1.1 of [2] we may isotope S to be in normal form with respect to Σ'. This means that S intersects each component of $M - \Sigma'$ in a collection of surfaces, each having at most one boundary circle on each of the three punctures; and if one of these surfaces is a disk then it separates the two punctures not containing its boundary.

We can represent a given element of H by a loop γ_0 based at p, such that γ_0 is disjoint from S and transverse to Σ'. The points of intersection of γ_0 with Σ' divide γ_0 into a finite set of arcs, each entirely contained in one component of $M - \Sigma'$. Suppose one of these arcs α, in a component P of $M - \Sigma'$, has both endpoints on the same boundary sphere σ of P. Since the map $\pi_0(\sigma - (S \cap \sigma)) \to \pi_0(P - (S \cap P))$ is injective (an easy consequence of normal form), there is an arc α' in $\sigma - (S \cap \sigma)$ with $\partial\alpha' = \partial\alpha$. Since P is simply-connected, α is homotopic to α' fixing endpoints. This homotopy gives a homotopy of γ_0 eliminating the two points of $\partial\alpha$ from $\gamma_0 \cap \Sigma'$, without introducing any intersection points with S. After repeating this operation a finite number of times, we may assume there are no remaining arcs of $\gamma_0 - (\gamma_0 \cap \Sigma')$ of the specified sort.

Now consider a homotopy $F\colon I \times I \to M$ of γ_0 to a loop γ_1 disjoint from Σ. Make F transverse to Σ' and look at $F^{-1}(\Sigma')$. This consists of a collection of disjoint arcs and circles. These do not meet the left and right edges of $I \times I$ since these edges map to the basepoint p.

We claim that every arc component of $F^{-1}(\Sigma')$ with one endpoint on $I \times \{0\}$ must have its other endpoint on $I \times \{1\}$. If not, choose an "edgemost" arc with both endpoints on $I \times \{0\}$, i.e. an arc such that the interval of $I \times \{0\}$ bounded by the endpoints contains no other point of $F^{-1}(\Sigma')$. Then γ_0 maps this interval to an arc α in $M - S$ which is entirely contained in one component P of $M - \Sigma'$ and has both endpoints on the same boundary sphere of P, contradicting our assumption that all such arcs have been eliminated.

Since the loop γ_1 is disjoint from Σ, it follows that γ_0 must be disjoint from Σ, and by construction γ_0 was disjoint from S. $\square$

Proof of Theorem 2.1. Following the method in [2], a contraction of S_H can be constructed by performing a sequence of surgeries on an arbitrary system S in S_H to eliminate its intersections with a fixed system Σ in S_H, after first putting S into normal form with respect to a maximal system Σ' containing Σ. In [2] the system Σ itself was maximal but for the present proof Σ must be in S_H so it cannot be maximal if H is nontrivial. We can choose Σ to be a single sphere defining a vertex of S_H for example. Once S has been surgered to be disjoint from Σ it will lie in the star of S which is contractible so this will finish the proof. (Alternatively, we could use the simpler contraction technique of [3], which reverses the roles of S and Σ.)

Each surgery on S is obtained by taking a circle of intersection of S and Σ which is innermost on Σ among the remaining circles of $S \cap \Sigma$, bounding a disk D in Σ with $D \cap S = \partial D$, then taking the two spheres obtained by attaching parallel copies of D to parallel copies of the two disks of $S - \partial D$. By the lemma, elements of H are representable by loops disjoint from Σ and S, so these loops remain disjoint from sphere systems obtained by surgering S along Σ because such surgery produces spheres lying in a neighborhood of $S \cup \Sigma$.

In order to ensure a continuous retraction the surgery process is made canonical by performing surgery on all innermost spheres at once. As a result, surgery can produce trivial spheres bounding balls in M or balls punctured by the sphere ∂M. At the end of the surgery process we discard all trivial spheres in the resulting sphere system, and it must be checked that at least one nontrivial sphere remains. To check this we note that the end result could be achieved by doing a single surgery at a time and then renormalizing, so it suffices to show that a single surgery cannot produce two trivial spheres.

Suppose, to the contrary, that a nontrivial sphere s is surgered to produced two trivial spheres s' and s''. The spheres $s \cup s' \cup s''$ form the boundary of a three-punctured sphere P in M. If s' or s'', say s', bounds a ball or punctured ball on the same side of s' as P then P will be contained in this ball or punctured ball, hence so will s, contradicting the nontriviality of s. Thus s' and s'' both bound balls or punctured balls on the opposite side from P. They cannot both bound punctured balls since M has only one puncture, so one of them bounds a ball. This forces the other one to be isotopic to s, but this is not possible since s is nontrivial while s' and s'' are trivial. $\square$

For an inductive argument in the next theorem we will need a generalization of the preceding theorem to manifolds with more than one boundary sphere. Let M_k be the manifold obtained from the connected sum of n copies of $S^1 \times S^2$ by deleting the interiors of k disjoint closed balls rather than just a single ball. Choose the basepoint p on one of the spheres in ∂M. For $H \subseteq \pi_1(M_k, p)$, define $S_H(M_k)$ to be the complex of isotopy classes of sphere systems S in M_k no two of which are isotopic and none of which bounds a ball or is isotopic to a sphere of ∂M_k, and such that $\pi_1(M - S, p) \supseteq H$.

Lemma 2.3. *For $n \geq 1$ and $k \geq 1$ the complex $S_H(M_{k+1})$ deformation retracts onto a subcomplex isomorphic to $S_H(M_k)$.*

Proof. When H is trivial this is Lemma 2.2 in [2] and the following proof extends the proof there to the general case after one small refinement to take H into account. Let the spheres of ∂M_{k+1} be $\partial_0, \cdots, \partial_k$ with $p \in \partial_0$. Call a vertex of $S_H(M_{k+1})$ *special* if it splits off a three-punctured sphere from M_{k+1} having ∂_k as one of its boundary components. Let $S'_H(M_{k+1})$ be the subcomplex of $S_H(M_{k+1})$ consisting of simplices with no special vertices. Then $S_H(M_{k+1})$ is obtained from $S'_H(M_{k+1})$ by attaching the stars of the special vertices to $S'_H(M_{k+1})$ along the links of the special vertices. These links can be identified with copies of $S_H(M_k)$ so they are contractible by induction on k. The interiors of the stars are disjoint since there are no edges joining special vertices. Hence $S_H(M_{k+1})$ deformation retracts to $S'_H(M_{k+1})$ since it is obtained by attaching contractible complexes along contractible subcomplexes.

The proof will be completed by showing that $S'_H(M_{k+1})$ deformation retracts onto a subcomplex isomorphic to $S_H(M_k)$. Let Σ be a sphere in M_{k+1} splitting off a three-punctured sphere P having ∂_0 and ∂_k as its other two boundary components. After putting a system S in $S'_H(M_{k+1})$ into normal form with respect to a maximal system containing Σ, S will intersect P in a set of parallel disks separating ∂_0 from ∂_k. We can eliminate these disks from $S \cap P$ by pushing them across ∂_k one by one and then outside P. This can also be described as surgering S along the circles of $S \cap \Sigma$ using the disks they bound in Σ on the side of ∂_k. Such a surgery on a sphere s of S produces a pair of sphere s' and s'' with s' the sphere we have pushed across ∂_k and s'' a trivial sphere parallel to ∂_k which is then discarded.

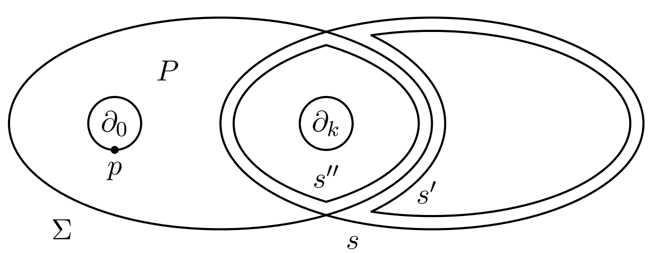

We claim that s' is neither trivial nor special. If it were either of these, it would separate M_{k+1} into two components, hence s would also separate. The effect of replacing s by s' is to move the puncture ∂_k from one side of s to the other side. It is then not hard to check that s being neither trivial nor special implies the same is true for s'.

The surgery defines a path in $S(M_{k+1})$ from S to $S \cup s'$ and then to $(S \cup s') - s$. The fundamental group of the component of $M_{k+1} - S$ containing p is unchanged during this process so the path lies in $S'_H(M_{k+1})$. By eliminating all the disks of $S \cap P$ by surgeries in this way we see that $S'_H(M_{k+1})$ deformation retracts onto its subcomplex of the systems disjoint from Σ. This subcomplex can be identified with $S_H(M_k)$ by identifying $M_{k+1} - P$ with M_k and choosing a new basepoint p in Σ. $\square$

For the proof of the main theorem in the paper we will work with certain subcomplexes of $S(M_k)$ and $S_H(M_k)$, the subcomplexes $Y(M_k) \subset S(M_k)$ and $Y_H(M_k) \subset S_H(M_k)$ consisting of sphere systems S with $M_k - S$ connected. Eventually only the case $k = 1$ will be needed, but to prove the key property that $Y(M_1)$ and $Y_H(M_1)$ are highly connected we will need to consider larger values of k. It will be convenient to extend the definition of $Y(M_k)$ and $Y_H(M_k)$ to allow $n = 0$, with M_k the sphere S^3 with k punctures. In this case $Y(M_k)$ and $Y_H(M_k)$ are empty since all spheres in M_k are separating.

Definition 2.4. *A simplicial complex K is m-spherical if it is m-dimensional and $(m-1)$-connected. A complex is spherical if it is m-spherical for some m.*

Theorem 2.5. *Let H be a free factor of $F_n = \pi_1(M_k, p)$. Then $Y_H(M_k)$ is $(n - rk(H) - 1)$-spherical, where $rk(H)$ is the rank of H.*

In particular, when H is trivial $Y(M_k)$ is $(n-1)$-spherical. This special case is part of Proposition 3.1 of [2] whose proof contained the same error that is corrected below. A corrected proof of the special case already appeared in Proposition 3.2 of [4].

Proof. $Y_H(M_k)$ has dimension $n - rk(H) - 1$ since a maximal simplex of $Y_H(M_k)$ has $n - rk(H)$ spheres. This follows because all free factors of F_n of the same rank are equivalent under automorphisms of F_n, and all sphere systems with connected complement and the same number of spheres are equivalent under orientation-preserving homeomorphisms of M_k. Thus it suffices to prove $Y_H(M_k)$ is $(n - rk(H) - 2)$-connected.

Let $i \leq n - rk(H) - 2$. Any map $g\colon S^i \to Y_H(M_k)$ can be extended to a map $\widehat{g}\colon D^{i+1} \to S_H(M_k)$ since $S_H(M_k)$ is contractible. We can assume $\widehat{g}$ is a simplicial map with respect to some triangulation of D^{i+1} compatible with its standard piecewise linear structure. We will repeatedly redefine $\widehat{g}$ on the stars of certain simplices in the interior of D^{i+1} until eventually the image of $\widehat{g}$ lies in $Y_H(M_k)$.

To each sphere system S we associate a dual graph $\Gamma(S)$, with one vertex for each component of $M - S$ and one edge for each sphere in S. The endpoints of the edge corresponding to $s \in S$ are the vertices corresponding to the component or components adjacent to s. We say a sphere system S is *purely separating* if $\Gamma(S)$ has no edges which begin and end at the same vertex. Each sphere system S has a *purely separating core*, consisting of those spheres in S which correspond to the core of $\Gamma(S)$, i.e., the subgraph spanned by edges with distinct vertices. The purely separating core of $S \in S_H(M_k)$ is empty if and only if S is in $Y_H(M_k)$.

Let σ be a simplex of D^{i+1} of maximal dimension among the simplices τ with $\widehat{g}(\tau)$ purely separating. Note that all such simplices τ lie in the interior of D^{i+1} since the boundary of D^{i+1} maps to Y_H. Let $S = \widehat{g}(\sigma)$, and let $N_0, \cdots, N_r$ $(r \geq 1)$ be the connected components of $M - S$, with $p \in N_0$. A simplex τ in the link $lk(\sigma)$ maps to a system T in the link of S, so that each $T_j = T \cap N_j$ is a sphere system in N_j and $H \leq \pi_1(N_0 - T_0, p)$. Furthermore $N_j - T_j$ must be connected for all j since otherwise the core of $\Gamma(S \cup T)$ would have more edges than $\Gamma(S)$, contradicting the maximality of σ. Thus $\widehat{g}$ maps $lk(\sigma)$ into a subcomplex of $S_H(M_k)$ which can be identified with $Y_H(N_0) * Y(N_1) * \cdots * Y(N_r)$. Some of the factors $Y_H(N_0)$ and $Y(N_j)$ can be empty if $rk(H) = rk(\pi_1(N_0))$ or $rk(\pi_1(N_j)) = 0$. Such factors are (-1)-spherical and contribute nothing to the join.

Since σ is a simplex in the interior of D^{i+1}, $lk(\sigma)$ is a sphere of dimension $i - dim(\sigma)$. Each N_j has fundamental group of rank $n_j \leq n$ with equality only if N_j has fewer than k boundary components, so by induction on the lexicographically ordered pair (n, k), $Y_H(N_0)$ is $(n_0 - rk(H) - 1)$-spherical and, for $j \geq 1$, $Y(N_j)$ is $(n_j - 1)$-spherical. The induction can start with the cases $(n, k) = (0, k)$ when the theorem is obvious. For the join $Y_H(N_0) * Y(N_1) * \cdots * Y(N_r)$ it then follows that this is spherical of dimension $(\sum_{i=0}^{r} n_j) - rk(H) - 1$.

Now $n = (\sum_j n_j) + rk(\pi_1(\Gamma(S))) = (\sum_j n_j) + m - r$ where m is the number of spheres in S, i.e., edges in $\Gamma(S)$. Since a simplicial map cannot increase dimension, we have $\dim(\sigma) \geq m - 1$. Therefore

$$\begin{aligned} i - \dim(\sigma) &\leq n - rk(H) - 2 - \dim(\sigma) \\ &\leq n - rk(H) - m - 1 \\ &= \Big(\sum_j n_j\Big) - rk(H) - 1 - r \\ &< \Big(\sum_j n_j\Big) - rk(H) - 1. \end{aligned}$$

Hence the map $\widehat{g}\colon lk(\sigma) \to Y_H(N_0) * Y(N_1) * \cdots * Y(N_r)$ can be extended to a map of a disk D^k into $Y_H(N_0) * Y(N_1) * \cdots * Y(N_r)$, where $k = i + 1 - dim(\sigma)$. The system S is compatible with every system in the image of D^k, so this map can be extended to a map $\sigma * D^k \to S_H$. We replace the star of σ in D^{i+1} by the disk $\partial(\sigma) * D^k$, and define $\widehat{g}$ on $\partial(\sigma) * D^k$ using this map.

What have we improved? The new simplices in the disk $\partial(\sigma) * D^k$ are of the form $\sigma' * \tau$, where σ' is a face of σ and $\widehat{g}(\tau) \subset Y_H(N_0) * Y(N_1) * \cdots * Y(N_r)$. The image of such a simplex $\sigma' * \tau$ is a system $S' \cup T$ such that in $\Gamma(S' \cup T)$ the edges corresponding to T are all loops. Therefore any simplex in the disk $\partial(\sigma) * D^k$ with purely separating image must lie in the boundary of this disk, where we have not modified $\widehat{g}$.

We continue this process, eliminating purely separating simplices until there are none in the image of $\widehat{g}$. Since every system in $S_H(M_k) - Y_H(M_k)$ has a non-trivial purely separating core, in fact the whole disk maps into $Y_H(M_k)$, and we are done. □

3. Free factors

We now turn to the poset FC_n of proper free factors of the free group F_n, partially ordered by inclusion. A k-simplex in the geometric realization $|FC_n|$ is a flag $H_0 < H_1 < \cdots < H_k$ of proper free factors of F_n, each properly included in the next. Each H_i is also a free factor of H_{i+1} (see [8], p. 117]), so that a maximal simplex of $|FC_n|$ has dimension $n - 2$.

We want to model free factors of F_n by sphere systems in $Y = Y(M)$, by taking the fundamental group of the (connected) complement. Here M is the manifold obtained from the connected sum of n copies of $S^1 \times S^2$ by removing the interior of a closed ball. A sphere system with n spheres and connected complement, corresponding to an $(n-1)$-dimensional simplex of Y, in fact has simply-connected complement. But we only want to consider proper free factors, so instead we consider the $(n-2)$-skeleton $Y^{(n-2)}$. Since Y is $(n-2)$-connected by Theorem 2.5, $Y^{(n-2)}$ is $(n-2)$-spherical.

In order to relate $Y^{(n-2)}$ to FC_n, we take the barycentric subdivision B_n of $Y^{(n-2)}$. Then B_n is the geometric realization of a poset of isotopy classes of sphere systems, partially ordered by inclusion. If $S \subseteq S'$ are sphere systems, we have $\pi_1(M-S,p) \geq \pi_1(M-S',p)$, reversing the partial ordering. Taking fundamental group of the complement thus gives a poset map $f\colon B_n \to (FC_n)^{op}$, where $(FC_n)^{op}$ denotes FC_n with the opposite partial ordering.

Proposition 3.1. *$f\colon B_n \to (FC_n)^{op}$ is surjective.*

Proof. Every simplex of FC_n is contained in a simplex of dimension $n-2$ so it suffices to show f maps onto all $(n-2)$-simplices. The group $Aut(F_n)$ acts transitively on $(n-2)$-simplices of FC_n, and all elements of $Aut(F_n)$ are realized by homeomorphisms of M, so f will be surjective if its image contains a single $(n-2)$-simplex, which it obviously does. □

Corollary 3.2. *FC_n is connected if $n \geq 3$.*

Proof. Theorem 2.5 implies that B_n is connected for $n \geq 3$. So, given any two vertices of FC_n, lift them to vertices of B_n by Proposition 3.1, connect the lifted vertices by a path, then project the path back down to FC_n. □

For any proper free factor H, let $B_{\geq H}$ denote the fiber f/H, consisting of isotopy classes of sphere systems S in B_n with $\pi_1(M-S,p) \geq H$.

Proposition 3.3. *Let H be a proper free factor of F_n. Then $B_{\geq H}$ is $(n - rk(H) - 1)$-spherical. If $rk(H) = n-1$ then $B_{\geq H}$ is a single point.*

Proof. The fiber $B_{\geq H}$ is the barycentric subdivision of Y_H, so is $(n - rk(H) - 1)$-spherical by Theorem 2.5.

Suppose $rk(H) = n-1$, so that $\pi_1(M,p) = H * \langle x \rangle$ for some x. An element of $B_{\geq H}$ is represented by a sphere system containing exactly one sphere s, which is non-separating with $\pi_1(M-s,p) = H$. Suppose s and s' are two such spheres. Since s and s' are both non-separating, there is a homeomorphism h of M taking s to s'. Since automorphisms of H can be realized by homeomorphisms of M fixing s, we may assume that the induced map on $\pi_1(M,p)$ is the identity on H.

Claim. *The induced map $h_*\colon \pi_1(M,p) \to \pi_1(M,p)$ must send x to an element of the form $Ux^{\pm 1}V$, with $U, V \in H$.*

Proof. Let $\{x_1, \cdots, x_{n-1}\}$ be a basis for H, and let W be the reduced word representing $h_*(x)$ in the basis $\{x_1, \cdots, x_{n-1}, x\}$ for $\pi_1(M,p)$. By looking at the map induced by h on homology, we see that the exponent sum of x in W must be ± 1. Since h_* fixes H, $\{x_1, \cdots, x_{n-1}, W\}$ is a basis for $\pi_1(M,p)$. If W contained both x and x^{-1}, we could apply Nielsen automorphisms to the set $\{x_1, \cdots, x_{n-1}, W\}$ until W was of the form $x^{\pm 1}W_0x^{\pm 1}$. But $\{x_1, \cdots, x_{n-1}, x^{\pm 1}W_0x^{\pm 1}\}$ is not a basis, since it is Nielsen reduced and not of the form $\{x_1^{\pm 1}, \cdots, x_{n-1}^{\pm 1}, x^{\pm 1}\}$ (see [7], Prop. 2.8)). □

The automorphism fixing H and sending $x \mapsto Ux^{\pm 1}V$ can be realized by a homeomorphism h' of M which takes s to itself (see [6], Lemme 4.3.1). The composition $h'h^{-1}$ sends s' to s and induces the identity on π_1, hence acts trivially on the sphere complex. Thus s and s' are isotopic. □

Corollary 3.4. *FC_n is simply connected for $n \geq 4$.*

Proof. Let $e_0, e_1, \cdots, e_k$ be the edges of an edge-path loop in FC_n, and choose lifts $\tilde{e}_i$ of these edges to B_n. Let $e_{i-1}e_i$ or e_ke_0 be two adjacent edges of the path, meeting at the vertex H. The lifts $\tilde{e}_{i-1}$ and $\tilde{e}_i$ may not be connected, i.e. $\tilde{e}_{i-1}$ may terminate at a sphere system S' and $\tilde{e}_i$ may begin at a different sphere system S. However, both S and S' are in the fiber $B_{\geq H}$, which is connected by Proposition 3.3, so we may connect S and S' by a path in $B_{\geq H}$. Connecting the endpoints of each lifted edge in this way, we obtain a loop in B_n, which may be filled in by a disk if $n \geq 4$, by Proposition 3.3. The projection of this loop to FC_n is homotopic to the original loop, since each extra edge-path segment we added projects to a loop in the star of some vertex H, which is contractible. Therefore the projection of the disk kills our original loop in the fundamental group. □

Remark 3.5. It is possible to describe the complex Y_H purely in terms of F_n. Suppose first that H is trivial. Define a simplicial complex Z to have vertices the rank $n-1$ free factors of F_n, with a set of k such factors spanning a simplex in Z if there is an automorphism of F_n taking them to the k factors obtained by deleting the standard basis elements $x_1, \cdots, x_k$ of F_n one at a time. There is a simplicial map $f\colon Y \to Z$ sending a system of k spheres to the set of k fundamental groups of the complements of these spheres. These fundamental groups are equivalent to the standard set of k rank $n-1$ factors under an automorphism of F_n since the homeomorphism group of M acts transitively on simplices of Y of a given dimension, and the standard k factors are the fundamental groups of the complements of the spheres in a standard system in Y. The last statement of Proposition 3.3 says that f is a bijection on vertices, so f embeds Y as a subcomplex of Z. The maximal simplices in Y and Z have dimension $n-1$ and the groups $Homeo(M)$ and $Aut(F_n)$ act transitively on these simplices, so f must be surjective, hence an isomorphism. When H is nontrivial, f restricts to an isomorphism from Y_H onto the subcomplex Z_H spanned by the vertices which are free factors containing H.

We are now ready to apply Quillen's spectral sequence to compute the homology of B_n and thus prove the main theorem.

Theorem 3.6. *FC_n is $(n-2)$-spherical.*

Proof. We prove the theorem by induction on n. If $n \leq 4$, the theorem follows from Corollaries 3.2 and 3.4.

Quillen's spectral sequence ([9], 7.7) applied to $f\colon B_n \to (FC_n)^{op}$ becomes

$$E^2_{p,q} = H_p(FC_n; H \mapsto H_q(B_{\geq H})) \Rightarrow H_{p+q}(B_n),$$

where the E^2-term is computed using homology with coefficients in the functor $H \mapsto H_q(B_{\geq H})$.

For $q = 0$, Corollary 3.2 gives $H_0(B_{\geq H}) = \mathbb{Z}$ for all H, so $E^2_{p,0} = H_p(FC_n, \mathbb{Z})$.

For $q > 0$, we have $E^2_{p,q} = H_p(FC_n; H \mapsto \widetilde{H}_q(B_{\geq H}))$, and we follow Quillen ([9], proof of Theorem 9.1) to compute this.

For a subposet A of FC_n, let L_A denote the functor sending H to a fixed abelian group L if $H \in A$ and to 0 otherwise. Set $U = FC_{\leq H} = \{H' \in FC_n | H' \leq H\}$ and $V = FC_{<H} = \{H' \in FC_n | H' < H\}$. Then

$$H_i(FC_n; L_V) = H_i(V; L), \text{ and}$$

$$H_i(FC_n; L_U) = H_i(U; L) = \begin{cases} L & \text{if i=0 and} \\ 0 & \text{otherwise,} \end{cases}$$

since $|U|$ is contractible. The short exact sequence of functors

$$1 \to L_V \to L_U \to L_{\{H\}} \to 1$$

gives a long exact homology sequence, from which we compute

$$H_i(FC_n; L_{\{H\}}) = \widetilde{H}_{i-1}(FC_{<H}; L).$$

Now, $H \mapsto \widetilde{H}_q(B_{\geq H})$ is equal to the functor

$$\bigoplus_{rk(H)=n-q-1} \widetilde{H}_q(B_{\geq H})_{\{H\}}$$

since $B_{\geq H}$ is $(n - rk(H) - 1)$-spherical by Proposition 3.3. Thus

$$\begin{aligned} E^2_{p,q} &= H_p(FC_n; H \mapsto \widetilde{H}_q(B_{\geq H})) \\ &= \bigoplus_{rk(H)=n-q-1} H_p(FC_n; \widetilde{H}_q(B_{\geq H})_{\{H\}}) \\ &= \bigoplus_{rk(H)=n-q-1} \widetilde{H}_{p-1}(FC_{<H}; \widetilde{H}_q(B_{\geq H})) \end{aligned}$$

Free factors of F_n contained in H are also free factors of H. Since H has rank $< n$. $FC_{<H}$ is $(rk(H) - 2)$-spherical by induction. Therefore $E^2_{p,q} = 0$ unless $p - 1 = (n - q - 1) - 2$, i.e. $p + q = n - 2$. Since all terms in the E^2-term of the spectral sequence are zero except the bottom row for $p \leq n - 2$ and the diagonal $p + q = n - 2$, all differentials are zero and we have $E^2 = E^\infty$ as in the following diagram:

$$\begin{array}{c|ccccccc}
n-2 & E^2_{0,n-2} & 0 & 0 & 0 & \cdots & & \\
n-3 & 0 & E^2_{1,n-3} & 0 & 0 & \cdots & & \\
n-4 & 0 & 0 & E^2_{2,n-4} & 0 & \cdots & & \\
\vdots & \vdots & \vdots & \vdots & \ddots & 0 & & \\
1 & 0 & 0 & 0 & \cdots & E^2_{n-3,1} & 0 & \\
0 & H_0(FC_n) & H_1(FC_n) & H_2(FC_n) & \cdots & H_{n-3}(FC_n) & H_{n-2}(FC_n) & 0 \\
\hline
 & 0 & 1 & 2 & \cdots & n-3 & n-2 &
\end{array}$$

Since FC_n is connected and the spectral sequence converges to $H_*(B_n)$, which is $(n-3)$-connected, we must have $\tilde{H}_i(FC_n) = 0$ for $i \neq n - 2$. Since FC_n is simply-connected by Corollary 3.4, this implies that FC_n is $(n - 3)$-connected by the Hurewicz theorem. The theorem follows since FC_n is $(n - 2)$-dimensional. □

4. The Cohen-Macaulay Property

In a PL triangulation of an n-dimensional sphere, the link of every k-simplex is an $(n-k-1)$-sphere. A poset is said to be *Cohen-Macaulay* of dimension n if its geometric realization is n-spherical and the link of every k-simplex is $(n-k-1)$-spherical (see [9]). Spherical buildings are Cohen-Macaulay, and we remark that FC_n also has this nice local property.

To see this, let $\sigma = \{H_0 < H_1 < \cdots < H_k\}$ be a k-simplex of FC_n. The link of σ is the join of subcomplexes $FC_{H_i,H_{i+1}}$ of FC_n spanned by free factors H with $H_i < H < H_{i+1}$ $(-1 \leq i \leq k$, with the conventions $H_{-1} = 1$ and $H_{k+1} = F_n$). Counting dimensions, we see that it suffices to show that for each r and s with $0 \leq s < r$, the poset $FC_{r,s}$ of proper free factors H of F_r which properly contain F_s is $(r-s-2)$-spherical. The proof of this is identical to the proof that FC_n is $(n-2)$-spherical, after setting $n = r$ and replacing the complex Y by Y_{F_s}.

5. The map to the building

As mentioned in the introduction, the abelianization map $F_n \to \mathbb{Z}^n$ induces a map from the free factor complex FC_n to the building X_n, since summands of $\mathbb{Z}^n$ correspond to subspaces of $\mathbb{Q}^n$. Since the map $Aut(F_n) \to GL(n, \mathbb{Z})$ is surjective, every basis for $\mathbb{Z}^n$ lifts to a basis for F_n, and hence every flag of summands of $\mathbb{Z}^n$ lifts to a flag of free factors of F_n, i.e. the map $FC_n \to X_n$ is surjective.

Given a basis $\{v_1, \cdots, v_n\}$ for $\mathbb{Q}^n$, consider the subcomplex of X_n consisting of all flags of subspaces of the form $\langle v_{i_1} \rangle \subset \langle v_{i_1}, v_{i_2} \rangle \subset \cdots \subset \langle v_{i_1}, \cdots, v_{i_{n-1}} \rangle$. This subcomplex can be identified with the barycentric subdivision of the boundary of an $(n-1)$-dimensional simplex, so forms an $(n-2)$-dimensional sphere in X_n, called an *apartment*.

The construction above applied to a basis of F_n instead of $\mathbb{Q}^n$ yields an $(n-2)$-dimensional sphere in FC_n. In particular, $H_{n-2}(FC_n)$ is non-trivial. This sphere maps to an apartment in X_n showing that the induced map $H_{n-2}(FC_n) \to H_{n-2}(X_n)$ is also non-trivial.

The property of buildings which is missing in FC_n is that given any two maximal simplices there is an apartment which contains both of them. For example, for $n = 3$ and F_3 free on $\{x, y, z\}$ there is no "apartment" which contains both the one-cells corresponding to $\langle x \rangle \subset \langle x, y \rangle$ and $\langle yxy^{-1} \rangle \subset \langle x, y \rangle$, since x and yxy^{-1} do not form part of a basis of F_3.

References

[1] M. Bestvina and M. Feighn, The topology at infinity of $Out(F_n)$, *Invent. Math.* 140 (2000), 651–692.

[2] A. Hatcher, Homological stability for automorphism groups of free groups, *Comment. Math. Helv.* 70 (1995), 39–62.

[3] A. Hatcher and K. Vogtmann, Isoperimetric inequalities for automorphisms of free groups, *Pacific Journal* 173 (1996), 425–442.

[4] A. Hatcher and N. Wahl, Stabilization for the automorphisms of free groups with boundaries, *Geometry and Topology* 9 (2005), 1295–1336.

[5] Z. Himes, J. Miller, S. Nariman, and A. Putman, The free factor complex and the dualizing module for the automorphism group of a free group, arXiv: 2203.03102.

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

[7] R. Lyndon and P. Schupp, *Combinatorial Group Theory*, Springer-Verlag, Berlin, 1977.

[8] W. Magnus, A. Karrass, and D. Solitar, *Combinatorial Group Theory*, Dover, New York, 1976.

[9] D. Quillen, Homotopy properties of the poset of p-subgroups of a group, *Advances in Math* 28 (1978), 101–128.

[10] L. Solomon, The Steinberg character of a finite group with BN-pair, *Theory of Finite Groups* (Brauer and Sah, eds.), Benjamin, 1969.