

THE FREE FACTOR COMPLEX AND THE DUALIZING MODULE FOR THE AUTOMORPHISM GROUP OF A FREE GROUP

ZACHARY HIMES, JEREMY MILLER, SAM NARIMAN, AND ANDREW PUTMAN

ABSTRACT. Answering a question of Hatcher–Vogtmann, we prove that the top homology group of the free factor complex is not the dualizing module for $\text{Aut}(F_n)$, at least for $n = 5$.

1. INTRODUCTION

A group Γ is a *rational duality group* of dimension d if there exists a $\mathbb{Q}[\Gamma]$ -module \mathbb{D} called the *dualizing module* such that for all $\mathbb{Q}[\Gamma]$ -modules M , we have

$$H^{d-i}(G; M) \cong H_i(G; M \otimes_{\mathbb{Q}} \mathbb{D}) \quad \text{for all } i.$$

The rational cohomological dimension of such a G is d , and the dualizing module \mathbb{D} is unique:

$$H^d(G; \mathbb{Q}[G]) \cong H_0(G; \mathbb{Q}[G] \otimes_{\mathbb{Q}} \mathbb{D}) \cong \mathbb{D}.$$

Many geometrically important groups are rational duality groups,¹ e.g., lattices in semisimple Lie groups [2], mapping class groups [17], and outer automorphism groups of free groups [1].

1.1. Identifying the dualizing module. The dualizing module of a rational duality group often has a simple geometric description. Here are two examples:

Example 1.1. For $\text{SL}_n(\mathbb{Z})$, the associated *Tits building* is the geometric realization $\text{Tits}(\mathbb{Q}^n)$ of the poset of nontrivial proper subspaces of \mathbb{Q}^n . The Solomon–Tits theorem [34] says that $\text{Tits}(\mathbb{Q}^n)$ is homotopy equivalent to a wedge of $(n - 2)$ -spheres. The *Steinberg representation* of $\text{SL}_n(\mathbb{Q})$, denoted $\text{St}(\mathbb{Q}^n)$, is its unique nonzero reduced homology group:

$$\text{St}(\mathbb{Q}^n) = \widetilde{H}_{n-2}(\text{Tits}(\mathbb{Q}^n); \mathbb{Q}).$$

Borel–Serre [2] proved $\text{SL}_n(\mathbb{Z})$ is a rational duality group with dualizing module $\text{St}(\mathbb{Q}^n)$. \square

Example 1.2. For the mapping class group $\text{Mod}(\Sigma_g)$ of a compact oriented genus- g surface Σ_g , the *curve complex* is the simplicial complex \mathcal{C}_g whose k -simplices are sets $\{\gamma_0, \dots, \gamma_k\}$ of distinct isotopy classes of nontrivial simple closed curves on Σ_g that can be realized disjointly. Harer [17] proved that for $g \geq 2$ the curve complex \mathcal{C}_g is homotopy equivalent to a wedge of $(2g - 2)$ -spheres. The *Steinberg representation* of $\text{Mod}(\Sigma_g)$, denoted $\text{St}(\Sigma_g)$, is its unique nonzero reduced homology group:

$$\text{St}(\Sigma_g) = \widetilde{H}_{2g-2}(\mathcal{C}_g; \mathbb{Q}).$$

Harer [17] proved that $\text{Mod}(\Sigma_g)$ is a rational duality group with dualizing module $\text{St}(\Sigma_g)$. \square

Date: November 7, 2022.

JM was supported in part by a Simons Foundation Collaboration Grant.

SN was supported in part by NSF grant DMS-2113828 and a Simons Foundation Collaboration Grant.

AP was supported in part by NSF grant DMS-1811210.

¹What these references actually prove is that these groups are virtual duality groups. This is a stronger condition: all virtual duality groups are rational duality groups, but the converse is false. For instance, Deligne [11] constructed a central extension $1 \rightarrow \mathbb{Z}/2 \rightarrow \widetilde{\text{Sp}}_{2g}(\mathbb{Z}) \rightarrow \text{Sp}_{2g}(\mathbb{Z}) \rightarrow 1$ such that $\widetilde{\text{Sp}}_{2g}(\mathbb{Z})$ has no torsion-free subgroup of finite index. This implies that $\widetilde{\text{Sp}}_{2g}(\mathbb{Z})$ is not a virtual duality group; however, since $\text{Sp}_{2g}(\mathbb{Z})$ and $\mathbb{Z}/2$ are rational duality groups, the group $\widetilde{\text{Sp}}_{2g}(\mathbb{Z})$ is as well.

These descriptions are useful since they allow calculations of the high-dimensional rational cohomology of these groups and their finite-index subgroups; see, e.g., [7, 8, 9, 13, 28, 29, 30].

1.2. Steinberg module for automorphism group of free group. Let F_n be a free group of rank n . Bestvina–Feighn [1] proved² that $\text{Aut}(F_n)$ is a rational duality group of dimension $2n - 2$. However, their proof does not give an explicit model for its dualizing module. Identifying the dualizing module of $\text{Aut}(F_n)$ remains a basic open problem.

Hatcher–Vogtmann [19, 20] suggested studying the following. A *free factor* of F_n is a subgroup $A < F_n$ such that there exists another subgroup $B < F_n$ with $F_n = A * B$. The nontrivial proper free factors of F_n form a poset called the *free factor complex*. In analogy with the ordinary Tits building, we will denote its geometric realization by $\text{Tits}(F_n)$, though it is not actually a building. Just like for $\text{Tits}(\mathbb{Q}^n)$, Hatcher–Vogtmann [19, 20] proved that $\text{Tits}(F_n)$ is homotopy equivalent to a wedge of $(n - 2)$ -dimensional spheres. The *Steinberg module* for $\text{Aut}(F_n)$, denoted $\text{St}(F_n)$, is its unique nonzero reduced homology group:

$$\text{St}(F_n) = \widetilde{H}_{n-2}(\text{Tits}(F_n); \mathbb{Q}).$$

On [19, 20, p. 1], Hatcher–Vogtmann asked the following question:

Question 1.3 (Hatcher–Vogtmann). *Is $\text{St}(F_n)$ the dualizing module for $\text{Aut}(F_n)$?*

A consequence of our main theorem is that this question has a negative answer in general.

1.3. Main theorem. Our main theorem is as follows:

Theorem A. $H_i(\text{Aut}(F_n); \text{St}(F_n)) = 0$ for $n \geq 2$ and $i = 0$ or 1 .

If $\text{St}(F_n)$ were the dualizing module for $\text{Aut}(F_n)$, this would imply that

$$H^{2n-2}(\text{Aut}(F_n); \mathbb{Q}) \cong H_0(\text{Aut}(F_n); \text{St}(F_n)) = 0 \quad \text{for } n \geq 2$$

and

$$H^{2n-3}(\text{Aut}(F_n); \mathbb{Q}) \cong H_1(\text{Aut}(F_n); \text{St}(F_n)) = 0 \quad \text{for } n \geq 3.$$

However, Gerlits [15] used a computer³ to prove that $H^7(\text{Aut}(F_5); \mathbb{Q}) \cong \mathbb{Q}$, contradicting the second assertion. We deduce the following:

Corollary 1.4. *The Steinberg module $\text{St}(F_5)$ is not the dualizing module for $\text{Aut}(F_5)$.*

Remark 1.5. We do not know all the rational cohomology of $\text{Aut}(F_n)$ for any $n \geq 6$, so it is unclear whether or not our theorem implies that $\text{St}(F_n)$ is not the dualizing module for $\text{Aut}(F_n)$ for $n \geq 6$. In any case, we conjecture that it is never the dualizing module except possibly in some low-rank degenerate cases. \square

Remark 1.6. In [5], Brück–Gupta study an $\text{Out}(F_n)$ -variant of $\text{Tits}(F_n)$. It would be interesting to adapt our techniques to study this complex. \square

1.4. Presentation. To prove Theorem A, we construct an explicit presentation for the $\mathbb{Q}[\text{Aut}(F_n)]$ -module $\text{St}(F_n)$. Our inspiration is a beautiful presentation for $\text{St}(\mathbb{Q}^n)$ constructed by Bykovskii [6]. Church–Putman [9] gave an alternate topological proof of Bykovskii’s theorem, and we adapt their approach to $\text{Aut}(F_n)$. The key is to find a highly connected simplicial complex that forms an “integral model” for the free factor complex $\text{Tits}(F_n)$. The simplicial complex we use is a variant of one used by Hatcher–Vogtmann’s [19, 20] to prove that $\text{Tits}(F_n)$ is homotopy equivalent to a wedge of $(n - 2)$ -spheres.

²What they actually proved is that $\text{Out}(F_n)$ is a rational duality group. This fits into a short exact sequence $1 \rightarrow F_n \rightarrow \text{Aut}(F_n) \rightarrow \text{Out}(F_n) \rightarrow 1$, and since both $\text{Out}(F_n)$ and F_n are rational duality groups it follows that $\text{Aut}(F_n)$ is as well.

³See [16] for the code.

1.5. **Sphere complex.** To describe this simplicial complex, we need to make two definitions.

Definition 1.7. Let $M_{n,b}$ be the connect sum of n copies of $S^2 \times S^1$ with b disjoint open balls removed. Our convention is that the connect sum of 0 copies of $S^2 \times S^1$ is the unit S^3 for the connect sum operation, so $M_{0,b}$ is S^3 with b disjoint open balls removed. \square

Definition 1.8. A 2-sphere embedded in $M_{n,b}$ is *essential* if it is not homotopic to $\partial M_{n,b}$ or a point. A *rank- k sphere system* in $M_{n,b}$ is a set $\{S_0, \dots, S_k\}$ of distinct isotopy classes of essential 2-spheres embedded in $M_{n,b}$ that can be realized disjointly. \square

We will systematically conflate 2-spheres in $M_{n,b}$ with their isotopy classes. We can now define the key simplicial complex, which was originally introduced by Hatcher [18].

Definition 1.9. The *sphere complex* for $M_{n,b}$, denoted $\mathcal{S}(M_{n,b})$, is the simplicial complex whose k -simplices are rank- k sphere systems $\{S_0, \dots, S_k\}$ in $M_{n,b}$. \square

To connect this to $\text{Aut}(F_n)$, fix a basepoint $*$ in $\partial M_{n,1}$. The mapping class group

$$\text{Mod}(M_{n,1}) = \pi_0(\text{Diff}^+(M_{n,1}, \partial M_{n,1}))$$

acts on $\pi_1(M_{n,1}, *) \cong F_n$, giving a homomorphism $\text{Mod}(M_{n,1}) \rightarrow \text{Aut}(F_n)$. Laudenbach [26, 27] proved that this homomorphism is surjective and that its kernel is a rank- n abelian 2-group generated by “sphere twists”:

$$1 \longrightarrow (\mathbb{Z}/2)^n \longrightarrow \text{Mod}(M_{n,1}) \longrightarrow \text{Aut}(F_n) \longrightarrow 1.$$

See [4] for an alternate proof that shows that this exact sequence splits. Laudenbach also proved that all elements of the kernel fix the isotopy classes of all 2-spheres in $M_{n,1}$. It follows that the action of $\text{Mod}(M_{n,1})$ on $\mathcal{S}(M_{n,1})$ factors through an action of $\text{Aut}(F_n)$ on $\mathcal{S}(M_{n,1})$.

1.6. **Hatcher–Vogtmann’s proof.** The starting point of Hatcher–Vogtmann’s proof that $\text{Tits}(F_n)$ is homotopy equivalent to a wedge of $(n-2)$ -spheres is a theorem of Hatcher [18] saying that $\mathcal{S}(M_{n,1})$ is contractible. Hatcher–Vogtmann consider the following subcomplex:

Definition 1.10. The *nonseparating sphere complex*⁴ of $M_{n,b}$, denoted $\mathcal{N}(M_{n,b})$, is the subcomplex of $\mathcal{S}(M_{n,b})$ whose k -simplices are rank- k sphere systems $\{S_0, \dots, S_k\}$ such that the union of the S_i does not separate $M_{n,b}$. \square

Using the fact that $\mathcal{S}(M_{n,1})$ is contractible, Hatcher–Vogtmann prove that $\mathcal{N}(M_{n,1})$ is homotopy equivalent to a wedge of $(n-1)$ -spheres. For each simplex $\{S_0, \dots, S_k\}$ of $\mathcal{N}(M_{n,1})$, the fundamental group of the complement

$$\pi_1(M_{n,1} \setminus \bigcup_{i=0}^k S_i, *)$$

is a proper free factor of F_n . It is a nontrivial free factor precisely when the simplex is not a maximal simplex. This provides a bridge between $\mathcal{N}(M_{n,1})$ and $\text{Tits}(F_n)$, which Hatcher–Vogtmann use to prove that $\text{Tits}(F_n)$ is homotopy equivalent to a wedge of $(n-2)$ -spheres.

⁴This is different from the complex of nonseparating spheres, which we discuss in Definition 9.1. In the complex of nonseparating spheres, the vertices are required to be nonseparating spheres, but there is no condition on the higher-dimensional simplices. The nonseparating sphere complex is also sometimes called the *purely* nonseparating sphere complex.

1.7. Our approach. It turns out that the fact that $\mathcal{N}(M_{n,1})$ is homotopy equivalent to a wedge of $(n-1)$ -spheres is sufficient to construct generators for $\text{St}(F_n)$, which are a sort of free group version of the spherical apartments in a Tits building. Following Church–Putman [9], to find the relations between these generators we need to add simplices to $\mathcal{N}(M_{n,1})$ to increase its connectivity from $(n-1)$ to n . The result is what we call the “augmented nonseparating sphere complex”, and most of this paper is a detailed study of its topology.

Remark 1.11. Though he did not state it explicitly and it takes some work to extract it from his paper, in [10] Sadofschi Costa proved results that are equivalent to our generating set for $\text{St}(F_n)$. \square

1.8. Outline. There are three main parts to this paper. The first (§2 – §5) is several sections of topological preliminaries. The second (§6 – §11) is a detailed study of the sphere complex and its subcomplexes. Finally, the third (§12 – §13) contains the proof of Theorem A.

1.9. Acknowledgments. We thank Benjamin Brück and Peter Patzt for helpful conversations and Richard Wade for comments on a previous version of this paper. We also would like to thank the referee for their careful reading of our paper.

2. TOPOLOGICAL PRELIMINARIES

We begin with some generalities about the topology of simplicial complexes.

2.1. Connectivity conventions. For any $d \in \mathbb{Z}$, we say that a space X is d -connected if for $k \leq d$, all maps $\mathbb{S}^k \rightarrow X$ from the k -sphere \mathbb{S}^k to X extend over the $(k+1)$ -disc \mathbb{D}^{k+1} . There are two important edge cases to this convention:

- For $k \leq -2$, we have $\mathbb{S}^k = \mathbb{D}^{k+1} = \emptyset$, so all spaces are d -connected for $d \leq -2$.
- We have $\mathbb{S}^{-1} = \emptyset$ and $\mathbb{D}^0 = \{\text{pt}\}$, so a space X is (-1) -connected precisely when $X \neq \emptyset$.

2.2. Links. For a simplex σ of a simplicial complex X , the *link* of σ , denoted $\text{Link}_X(\sigma)$, is the subcomplex of X consisting of all simplices σ' such that the join $\sigma * \sigma'$ is a simplex of X . Notice that this implies that σ' does not share any vertices with σ . The join $\sigma * \text{Link}_X(\sigma)$ is a subcomplex of X called the *star* of σ .

2.3. Combinatorial triangulations. A *combinatorial triangulation* of an n -manifold with boundary is defined inductively as follows. First, any 0-dimensional simplicial complex is a combinatorial triangulation of a 0-manifold. Next, for $n \geq 1$ a combinatorial triangulation of an n -manifold with boundary is a simplicial complex M^n that is a topological n -manifold with boundary such that for all simplices σ of M^n , the complex $\text{Link}_{M^n}(\sigma)$ is as follows:

- If σ does not lie in ∂M^n , then $\text{Link}_{M^n}(\sigma)$ is a combinatorial triangulation of an $(n - \dim(\sigma) - 1)$ -sphere.
- If σ lies in ∂M^n , then $\text{Link}_{M^n}(\sigma)$ is a combinatorial triangulation of an $(n - \dim(\sigma) - 1)$ -disc.

Any subdivision of a combinatorial triangulation of a manifold with boundary is also a combinatorial triangulation.

3. LOCALLY INJECTIVE MAPS, SPHERES, AND DISCS

We now turn to some more technical aspects of simplicial complexes.

3.1. Local injectivity. If $f: Y \rightarrow X$ is a simplicial map between simplicial complexes, then for all simplices σ of Y we have $\dim(\sigma) \geq \dim(f(\sigma))$. For technical reasons, it will be important for us to make this an equality whenever possible:

Definition 3.1. A simplicial map $f: Y \rightarrow X$ between simplicial complexes is *locally injective* if for all simplices σ of Y , we have $\dim(\sigma) = \dim(f(\sigma))$. \square

3.2. Sphere and disc local injectivity properties. To keep our homotopies from getting out of control, it will be helpful to represent homotopy classes in simplicial complexes by locally injective maps. We therefore make the following two definitions.

Definition 3.2. We say that a simplicial complex X has the *sphere local injectivity property* up to dimension d if the following holds for all $k \leq d$. Let $f: \mathbb{S}^k \rightarrow X$ be a continuous map. Then f can be homotoped to a map $f: \mathbb{S}^k \rightarrow X$ that is simplicial and locally injective for some combinatorial triangulation of \mathbb{S}^k . \square

Definition 3.3. We say that a simplicial complex X has the *disc local injectivity property* up to dimension d if the following holds for all $k \leq d$. Let $f: \mathbb{S}^k \rightarrow X$ be a map that is simplicial and locally injective for some combinatorial triangulation of \mathbb{S}^k . Then there exists a combinatorial triangulation of \mathbb{D}^{k+1} that extends our given triangulation of \mathbb{S}^k and a locally injective simplicial map $F: \mathbb{D}^{k+1} \rightarrow X$ extending f . \square

Example 3.4. All simplicial complexes have the sphere local injectivity property up to dimension 0. A simplicial complex has the disc local injectivity property up to dimension -1 if it is nonempty, and has the disc local injectivity property up to dimension 0 if it is nonempty, connected, and not just a single point. \square

Example 3.5. It will follow from our discussion of weakly Cohen–Macaulay complexes in §4 that if M^n is a combinatorial triangulation of an n -manifold that is $(d - 1)$ -connected for some $d \leq n$, then M^n has the sphere local injectivity property up to dimension d and disc local injectivity property up to dimension $d - 1$. See Example 4.3 and Lemma 4.4. \square

Remark 3.6. By the simplicial approximation theorem, any continuous map $f: \mathbb{S}^k \rightarrow X$ is homotopic to a map that is simplicial with respect to a triangulation of \mathbb{S}^k that can be obtained by subdividing any given triangulation of \mathbb{S}^k . Since the class of combinatorial triangulations of a manifold is closed under subdivisions, by starting with a combinatorial triangulation of \mathbb{S}^k (for instance, the boundary of a $(k + 1)$ -simplex) we can ensure that the resulting map $\mathbb{S}^k \rightarrow X$ is simplicial with respect to a combinatorial triangulation of \mathbb{S}^k . Thus the local injectivity part of the sphere local injectivity property is the key content of that definition. A similar remark applies to the disc local injectivity property, though here you need Zeeman’s version [36] of the relative simplicial approximation theorem.⁵ \square

We have the following.

Lemma 3.7. *Let X be a simplicial complex that has both the sphere and disc local injectivity property up to dimension d . Then X is d -connected.*

Proof. For $k \leq d$, first use the sphere local injectivity property to homotope a given $f: \mathbb{S}^k \rightarrow X$ to a locally injective map, and then use the disc local injectivity property to extend f over \mathbb{D}^{k+1} . \square

⁵There are subtle issues with relative simplicial approximation, and the standard version as found in e.g. [35, Theorem 3.4.8] is not strong enough for what we are doing. See the introduction to [36] for a discussion of this.

3.3. Joins. The sphere and disc local injectivity properties behave well under joins. We will need this for the disc local injectivity property, so we focus on that:

Lemma 3.8. *Let X be a simplicial complex with the disc local injectivity property up to dimension d and Y be a simplicial complex with the disc local injectivity property up to dimension e . Then $X * Y$ has the disc local injectivity property up to dimension $d + e + 2$.*

Proof. For some $k \leq d + e + 2$, fix a combinatorial triangulation of \mathbb{S}^k and let $f: \mathbb{S}^k \rightarrow X * Y$ be a locally injective simplicial map. Our goal is to extend f to a locally injective map $F: \mathbb{D}^{k+1} \rightarrow X * Y$ for some combinatorial triangulation of \mathbb{D}^{k+1} agreeing with our triangulation of \mathbb{S}^k on the boundary. We divide the proof of this into three steps.

Step 1. *We modify f such that $\dim(\sigma) \leq d$ for all simplices σ of \mathbb{S}^k with $f(\sigma) \subset X$.*

Consider a simplex σ of \mathbb{S}^k with $f(\sigma) \subset X$. Pick σ such that its dimension is maximal with this property, and assume that $\dim(\sigma) \geq d + 1$. Then $\text{Link}_{\mathbb{S}^k}(\sigma)$ is a combinatorial triangulation of a $(k - \dim(\sigma) - 1)$ sphere, and by the maximality of the dimension of σ we have $f(\text{Link}_{\mathbb{S}^k}(\sigma)) \subset Y$. Note that

$$k - \dim(\sigma) - 1 \leq (d + e + 2) - (d + 1) - 1 = e.$$

Since Y has the disc local injectivity property up to dimension e and f is locally injective, we can find a combinatorial triangulation of $\mathbb{D}^{k-\dim(\sigma)}$ agreeing with that of $\text{Link}_{\mathbb{S}^k}(\sigma) \cong \mathbb{S}^{k-\dim(\sigma)-1}$ on $\partial\mathbb{D}^{k-\dim(\sigma)}$ and a locally injective map $G: \mathbb{D}^{k-\dim(\sigma)} \rightarrow Y$ extending the restriction of f to $\text{Link}_{\mathbb{S}^k}(\sigma)$. We have

$$\sigma * \text{Link}_{\mathbb{S}^k}(\sigma) \cong \mathbb{D}^{\dim(\sigma)} * \mathbb{S}^{k-\dim(\sigma)-1} \cong \mathbb{D}^k$$

and

$$\partial\sigma * \mathbb{D}^{k-\dim(\sigma)} \cong \mathbb{S}^{\dim(\sigma)-1} * \mathbb{D}^{k-\dim(\sigma)} \cong \mathbb{D}^k,$$

and also

$$\partial(\sigma * \text{Link}_{\mathbb{S}^k}(\sigma)) = \partial\sigma * \text{Link}_{\mathbb{S}^k}(\sigma)$$

and

$$\partial(\partial\sigma * \mathbb{D}^{k-\dim(\sigma)}) = \partial\sigma * \partial\mathbb{D}^{k-\dim(\sigma)} = \partial\sigma * \text{Link}_{\mathbb{S}^k}(\sigma).$$

In all of these $=$ means equality of simplicial complexes and \cong means homeomorphism. It follows that we can homotope f so as to replace

$$f|_{\sigma * \text{Link}_{\mathbb{S}^k}(\sigma)}: \sigma * \text{Link}_{\mathbb{S}^k}(\sigma) \rightarrow X * Y \quad \text{with} \quad (f|_{\partial\sigma}) * G: \partial\sigma * \mathbb{D}^{k-\dim(\sigma)} \rightarrow X * Y.$$

This eliminates σ , and it is easy to see that it is enough to prove the lemma for this new f . Repeating this over and over, we can ensure that $\dim(\sigma) \leq d$ for all simplices σ of \mathbb{S}^k with $f(\sigma) \subset X$.

Step 2. *For a topological space W , let $\text{Cone}(W)$ denote the cone on W . Set $Z = f^{-1}(X) \subset \mathbb{S}^k$. We construct a triangulation of $\text{Cone}(Z)$ and a locally injective simplicial map $F: \text{Cone}(Z) \rightarrow X$ extending $f|_Z$.*

Let p_0 be the cone point of $\text{Cone}(Z)$, and define $F(p_0)$ to be some arbitrary vertex of X . By the previous step, Z has dimension at most d . We can now use the disc local injectivity property of X up to dimension d to extend F over $\sigma * p_0$ for each simplex σ of Z : first over the 0-simplices, then the 1-simplices, etc. At each step, we have already defined F on a subdivision of

$$\sigma \cup_{\partial\sigma} (\partial\sigma * p_0) \cong \mathbb{D}^{\dim(\sigma)} \cup_{\mathbb{S}^{\dim(\sigma)-1}} (\mathbb{S}^{\dim(\sigma)-1} * p_0) \cong \mathbb{S}^{\dim(\sigma)},$$

and we use the disc local injectivity property to extend F to a locally injective map on the interior of this simplex (after a further subdivision).

Step 3. We extend F to the rest of \mathbb{D}^{k+1} .

We have defined F on a subdivision of $\text{Cone}(Z) \subset \text{Cone}(\mathbb{S}^k)$, and F takes $\text{Cone}(Z)$ to X . Subdivide $\text{Cone}(\mathbb{S}^k)$ by subdividing exactly the same simplices subdivided to form our subdivision of $\text{Cone}(Z)$ on which F was defined. The simplices of this subdivision are thus all of the form $\sigma * \tau$, where σ is a simplex of $\text{Cone}(Z)$ with $F(\sigma) \subset X$ and τ is a simplex of \mathbb{S}^k with $f(\tau) \subset Y$. From this, we see that F can be extended over all these simplices to a map with values in $X * Y$, as desired. \square

4. COHEN–MACAULAY COMPLEXES

We now discuss an important class of simplicial complexes.

4.1. Weakly Cohen–Macaulay complexes. We start with the following definition:

Definition 4.1. Let X be a simplicial complex and let $d \geq 0$. We say that X is *weakly Cohen–Macaulay* of dimension d if X is $(d - 1)$ -connected, and for all simplices σ of X the subcomplex $\text{Link}_X(\sigma)$ is $(d - \dim(\sigma) - 2)$ -connected.⁶ If in addition to this X is d -dimensional and for all simplices σ of X the subcomplex $\text{Link}_X(\sigma)$ is $(d - \dim(\sigma) - 1)$ -dimensional, then we say that X is *Cohen–Macaulay* of dimension d . \square

Remark 4.2. If X is weakly Cohen–Macaulay of dimension d and σ is a simplex of X , then $\text{Link}_X(\sigma)$ is weakly Cohen–Macaulay of dimension $(d - \dim(\sigma) - 1)$. A similar remark applies if X is Cohen–Macaulay of dimension d . \square

Example 4.3. Let M^n be a combinatorial triangulation of an n -dimensional manifold that is $(d - 1)$ -connected for some $d \leq n$. Then M^n is weakly Cohen–Macaulay of dimension d . Indeed, it is $(d - 1)$ -connected by assumption, and for a simplex σ of M^n we have that $\text{Link}_{M^n}(\sigma)$ is a combinatorial triangulation of an $(n - \dim(\sigma) - 1)$ -sphere, and in particular is

$$n - \dim(\sigma) - 2 \geq d - \dim(\sigma) - 2$$

connected. \square

4.2. Cohen–Macaulay and local injectivity. Weakly Cohen–Macaulay complexes provide examples of the sphere and disc local injectivity property:

Lemma 4.4.⁷ *Let X be a simplicial complex that is weakly Cohen–Macaulay of dimension d . Then X has the sphere local injectivity property up to dimension d and the disc local injectivity property up to dimension $(d - 1)$.*

Proof. The proof will be by induction on d . The base case $d = 0$ is trivial,⁸ so assume that $d > 0$ and that the lemma is true whenever d is smaller. We prove the two parts of the lemma separately.

Step 1. *The simplicial complex X has the sphere local injectivity property up to dimension d .*

⁶If we considered the empty set to be a (-1) -dimensional simplex whose link is the whole complex X , then we could combine the two hypotheses and just say that for all simplices σ of X the subcomplex $\text{Link}_X(\sigma)$ is $(d - \dim(\sigma) - 2)$ -connected.

⁷See Galatius–Randal-Williams [14, Theorem 2.4] for a similar result.

⁸See Remark 3.4, and note that a simplicial complex X that is weakly Cohen–Macaulay of dimension 0 is (-1) -connected, i.e., nonempty.

For some $k \leq d$, let $f: \mathbb{S}^k \rightarrow X$ be a continuous map. Using simplicial approximation (see Remark 3.6), we can assume that f is simplicial for some combinatorial triangulation of \mathbb{S}^k . We will prove that f can be homotoped to a locally injective simplicial map.

Assume that f is not locally injective. Consider a simplex σ of \mathbb{S}^k with $\dim(f(\sigma)) < \dim(\sigma)$. Pick σ such that its dimension is maximal with this property. Then $\text{Link}_{\mathbb{S}^k}(\sigma)$ is a combinatorial triangulation of a $(k - \dim(\sigma) - 1)$ sphere, and by the maximality of the dimension of σ we have $f(\text{Link}_{\mathbb{S}^k}(\sigma)) \subset \text{Link}_X(f(\sigma))$. This maximality also implies that the restriction of f to $\text{Link}_{\mathbb{S}^k}(\sigma)$ is locally injective. Note that

$$k - \dim(\sigma) - 1 \leq d - \dim(f(\sigma)) - 2.$$

Since X is weakly Cohen–Macaulay of dimension d , it follows that $\text{Link}_X(f(\sigma))$ is weakly Cohen–Macaulay of dimension $d - \dim(f(\sigma)) - 1$. In particular, by our inductive hypothesis it has the disc local injectivity property up to dimension $d - \dim(f(\sigma)) - 2$. We can thus find a combinatorial triangulation of $\mathbb{D}^{k - \dim(\sigma)}$ agreeing with that of $\text{Link}_{\mathbb{S}^k}(\sigma) \cong \mathbb{S}^{k - \dim(\sigma) - 1}$ on $\partial\mathbb{D}^{k - \dim(\sigma)}$ and a locally injective map $G: \mathbb{D}^{k - \dim(\sigma)} \rightarrow X$ extending the restriction of f to $\text{Link}_{\mathbb{S}^k}(\sigma)$. We have

$$\sigma * \text{Link}_{\mathbb{S}^k}(\sigma) \cong \mathbb{D}^{\dim(\sigma)} * \mathbb{S}^{k - \dim(\sigma) - 1} \cong \mathbb{D}^k$$

and

$$\partial\sigma * \mathbb{D}^{k - \dim(\sigma)} \cong \mathbb{S}^{\dim(\sigma) - 1} * \mathbb{D}^{k - \dim(\sigma)} \cong \mathbb{D}^k,$$

and also

$$\partial(\sigma * \text{Link}_{\mathbb{S}^k}(\sigma)) = \partial\sigma * \text{Link}_{\mathbb{S}^k}(\sigma)$$

and

$$\partial(\partial\sigma * \mathbb{D}^{k - \dim(\sigma)}) = \partial\sigma * \partial\mathbb{D}^{k - \dim(\sigma)} = \partial\sigma * \text{Link}_{\mathbb{S}^k}(\sigma).$$

In all of these $=$ means equality of simplicial complexes and \cong means homeomorphism. It follows that we can homotope f so as to replace

$$f|_{\sigma * \text{Link}_{\mathbb{S}^k}(\sigma)}: \sigma * \text{Link}_{\mathbb{S}^k}(\sigma) \rightarrow X$$

with

$$(f|_{\partial\sigma}) * G: \partial\sigma * \mathbb{D}^{k - \dim(\sigma)} \rightarrow f(\sigma) * \text{Link}_X(f(\sigma)) \subset X.$$

This eliminates σ . Repeating this over and over, we can ensure that f is locally injective, as desired.

Step 2. *The simplicial complex X has the disc local injectivity property up to dimension $d - 1$.*

For some $k \leq d - 1$, let $f: \mathbb{S}^k \rightarrow X$ be a map that is locally injective with respect to some combinatorial triangulation of \mathbb{S}^k . Since X is weakly Cohen–Macaulay of dimension d , it is $(d - 1)$ -connected. It follows (see Remark 3.6) that we can extend f to a map $F: \mathbb{D}^{k+1} \rightarrow X$ that is simplicial with respect to a combinatorial triangulation of \mathbb{D}^{k+1} . Using an argument identical to the previous step, we can modify F without changing it on the boundary to ensure that it is locally injective, as desired. \square

5. BAD SIMPLEX ARGUMENTS

Let Y be a simplicial complex and let $X \subset Y$ be a subcomplex. In this section, we discuss two results that let us relate the topological properties of X and Y by studying a set of “bad simplices” that characterize simplices that are in Y but not in X .

5.1. Connectivity and bad simplices. The first of the two results that we will need is due to Hatcher–Vogtmann [21]. To state it, we must first give a definition.

Definition 5.1. Let Y be a simplicial complex and let \mathcal{B} be a set of simplices of Y . For $\sigma \in \mathcal{B}$, let $\mathcal{L}(\sigma, \mathcal{B})$ be the subcomplex of Y consisting of simplices τ satisfying the following two conditions:

- τ is in $\text{Link}_Y(\sigma)$, so $\sigma * \tau$ is a simplex of Y .
- If σ' is a face of $\sigma * \tau$ with $\sigma' \in \mathcal{B}$, then σ' is a face of σ . □

We can now state Hatcher–Vogtmann’s theorem as follows. In it, you should regard \mathcal{B} as the set of “bad simplices” of Y :

Proposition 5.2 ([21, Proposition 2.1]). *Let Y be a simplicial complex and let $X \subset Y$ be a subcomplex. Assume that there exists a set \mathcal{B} of simplices of Y with the following properties for some $d \geq 0$:*

- (i) *A simplex of Y lies in X if and only if none of its faces are in \mathcal{B} . In particular, since simplices are faces of themselves no simplices of \mathcal{B} lie in X .*
- (ii) *If $\sigma, \sigma' \in \mathcal{B}$ are such that $\sigma \cup \sigma'$ is a simplex of Y , then $\sigma \cup \sigma' \in \mathcal{B}$. Note that σ and σ' might share vertices, so $\sigma \cup \sigma'$ might not be $\sigma * \sigma'$.*
- (iii) *For all $\sigma \in \mathcal{B}$, the subcomplex $\mathcal{L}(\sigma, \mathcal{B})$ is $(d - \dim(\sigma) - 1)$ -connected.*

Then the pair (Y, X) is d -connected, i.e., we have $\pi_k(Y, X) = 0$ for $0 \leq k \leq d$.

Both to make this paper more self-contained and to motivate the second and more technical bad simplex argument we discuss below, we include a proof.

Proof of Proposition 5.2. For some $0 \leq k \leq d$, let $f: (\mathbb{D}^k, \partial\mathbb{D}^k) \rightarrow (Y, X)$ be a map of pairs that is simplicial with respect to some combinatorial triangulation of \mathbb{D}^k . Our goal is to homotope f rel boundary such that its image lies in X . So assume that the image of f does not lie in X .

By (i), there exists a simplex σ of \mathbb{D}^k such that $f(\sigma) \in \mathcal{B}$. Pick σ such that its dimension is maximal among all simplices with $f(\sigma) \in \mathcal{B}$. Consider a simplex τ of $\text{Link}_{\mathbb{D}^k}(\sigma)$. We claim that $f(\tau) \in \mathcal{L}(f(\sigma), \mathcal{B})$. Indeed, the maximality of $\dim(\sigma)$ implies that $f(\tau)$ is in $\text{Link}_Y(f(\sigma))$. If $f(\tau)$ does not lie in $\mathcal{L}(f(\sigma), \mathcal{B})$, then $f(\sigma) * f(\tau)$ contains a face η with $\eta \in \mathcal{B}$ and with η containing vertices of $f(\tau)$. But then (ii) implies that $f(\sigma) \cup \eta \in \mathcal{B}$, contradicting the maximality of $\dim(\sigma)$.

Since f takes $\partial\mathbb{D}^k$ to X , the simplex σ does not lie in $\partial\mathbb{D}^k$. It follows that $\text{Link}_{\mathbb{D}^k}(\sigma)$ is a combinatorial triangulation of a $(k - \dim(\sigma) - 1)$ sphere, and by (iii) the complex $\mathcal{L}(f(\sigma), \mathcal{B})$ is $(d - \dim(\sigma) - 1)$ -connected. Since

$$k - \dim(\sigma) - 1 \leq d - \dim(\sigma) - 1,$$

we can find a combinatorial triangulation of $\mathbb{D}^{k - \dim(\sigma)}$ agreeing with that of $\text{Link}_{\mathbb{D}^k}(\sigma) \cong \mathbb{S}^{k - \dim(\sigma) - 1}$ on $\partial\mathbb{D}^{k - \dim(\sigma)}$ and a simplicial map $g: \mathbb{D}^{k - \dim(\sigma)} \rightarrow \mathcal{L}(f(\sigma), \mathcal{B})$ extending the restriction of f to $\text{Link}_{\mathbb{D}^k}(\sigma)$. We have

$$\sigma * \text{Link}_{\mathbb{D}^k}(\sigma) \cong \mathbb{D}^{\dim(\sigma)} * \mathbb{S}^{k - \dim(\sigma) - 1} \cong \mathbb{D}^k$$

and

$$\partial\sigma * \mathbb{D}^{k - \dim(\sigma)} \cong \mathbb{S}^{\dim(\sigma) - 1} * \mathbb{D}^{k - \dim(\sigma)} \cong \mathbb{D}^k,$$

and also

$$\partial(\sigma * \text{Link}_{\mathbb{D}^k}(\sigma)) = \partial\sigma * \text{Link}_{\mathbb{D}^k}(\sigma)$$

and

$$\partial(\partial\sigma * \mathbb{D}^{k - \dim(\sigma)}) = \partial\sigma * \partial\mathbb{D}^{k - \dim(\sigma)} = \partial\sigma * \text{Link}_{\mathbb{D}^k}(\sigma).$$

In all of these $=$ means equality of simplicial complexes and \cong means homeomorphism. It follows that we can homotope f so as to replace

$$f|_{\sigma * \text{Link}_{\mathbb{D}^k}(\sigma)} : \sigma * \text{Link}_{\mathbb{D}^k}(\sigma) \rightarrow Y$$

with

$$(f|_{\partial\sigma}) * g : \partial\sigma * \mathbb{D}^{k-\dim(\sigma)} \rightarrow f(\sigma) * \mathcal{L}(f(\sigma), \mathcal{B}) \subset Y.$$

This eliminates σ , and by (ii) and the definition of $\mathcal{L}(f(\sigma), \mathcal{B})$ it does not introduce any new simplices of dimension at least $\dim(\sigma)$ mapping to simplices in \mathcal{B} . Repeating this over and over, we can ensure that there are no simplices σ of \mathbb{D}^k with $f(\sigma) \in \mathcal{B}$, so by (i) we have $f(\mathbb{D}^k) \subset X$, as desired. \square

5.2. Disc local-injectivity and bad simplices. We now discuss a more technical version of Proposition 5.2 for proving the disc local injectivity property.

Proposition 5.3. *Let Y be a simplicial complex with the disc local injectivity property up to dimension d , and let $X \subset Y$ be a subcomplex. Assume that there exists a set \mathcal{B} of simplices of Y with the following properties:*

- (i) *A simplex of Y lies in X if and only if none of its faces are in \mathcal{B} . In particular, since simplices are faces of themselves no simplices of \mathcal{B} lie in X .*
- (ii) *If $\sigma, \sigma' \in \mathcal{B}$ are such that $\sigma \cup \sigma'$ is a simplex of Y , then $\sigma \cup \sigma' \in \mathcal{B}$. Note that σ and σ' might share vertices, so $\sigma \cup \sigma'$ might not be $\sigma * \sigma'$.*
- (iii) *For all $\sigma \in \mathcal{B}$, there exists a subcomplex $\widehat{\mathcal{L}}(\sigma, \mathcal{B})$ of Y with $\mathcal{L}(\sigma, \mathcal{B}) \subset \widehat{\mathcal{L}}(\sigma, \mathcal{B})$ such that the following holds:*
 - a. *We have $\partial\sigma * \widehat{\mathcal{L}}(\sigma, \mathcal{B}) \subset Y$.*
 - b. *All simplices of $\partial\sigma * \widehat{\mathcal{L}}(\sigma, \mathcal{B})$ that are in \mathcal{B} lie in $\partial\sigma$.*
 - c. *The complex $\widehat{\mathcal{L}}(\sigma, \mathcal{B})$ has the disc local injectivity property up to dimension $(d - \dim(\sigma))$.*

Then X has the disc local injectivity property up to dimension d .

Remark 5.4. We will use the disc local injectivity property to prove that spaces are highly connected. However, it will play an essential role. Indeed, if we change the hypothesis (resp. conclusion) of Proposition 5.3 from Y (resp. X) having the disc local injectivity property up to dimension d to Y (resp. X) being d -connected, then the proof would not work. We will highlight the issue in the proof below with a footnote. \square

Proof of Proposition 5.3. For some $k \leq d$, let $f: \mathbb{S}^k \rightarrow X$ be a locally injective simplicial map with respect to some combinatorial triangulation of \mathbb{S}^k . Since Y has the disc local injectivity property up to dimension d , there exists a combinatorial triangulation of \mathbb{D}^{k+1} agreeing with our given triangulation of \mathbb{S}^k on the boundary and an extension of f to a locally injective simplicial map $F: \mathbb{D}^{k+1} \rightarrow Y$. Assume that the image of F does not lie in X .

By (i), there exists a simplex σ of \mathbb{D}^{k+1} such that $F(\sigma) \in \mathcal{B}$. Pick σ such that its dimension is maximal among all simplices with $F(\sigma) \in \mathcal{B}$. Just like in the proof of Proposition 5.2, we can use (ii) along with the maximality of $\dim(\sigma)$ to deduce that F takes $\text{Link}_{\mathbb{D}^{k+1}}(\sigma)$ to $\mathcal{L}(F(\sigma), \mathcal{B})$.

Since F takes $\partial\mathbb{D}^{k+1}$ to X , the simplex σ does not lie in $\partial\mathbb{D}^{k+1}$. Thus the complex $\text{Link}_{\mathbb{D}^{k+1}}(\sigma)$ is a combinatorial triangulation of a $(k - \dim(\sigma))$ sphere, and by (iii).c the complex $\widehat{\mathcal{L}}(F(\sigma), \mathcal{B})$ has the disc local injectivity property up to dimension $(d - \dim(F(\sigma)))$. Since F is locally injective and

$$k - \dim(\sigma) \leq d - \dim(\sigma),$$

we can find a combinatorial triangulation of $\mathbb{D}^{k-\dim(\sigma)+1}$ agreeing with that of $\text{Link}_{\mathbb{D}^{k+1}}(\sigma) \cong \mathbb{S}^{k-\dim(\sigma)}$ on $\partial\mathbb{D}^{k-\dim(\sigma)+1}$ and a locally injective map $G: \mathbb{D}^{k-\dim(\sigma)+1} \rightarrow \widehat{\mathcal{L}}(F(\sigma), \mathcal{B})$ extending the restriction of F to $\text{Link}_{\mathbb{D}^{k+1}}(\sigma)$.

We have

$$\sigma * \text{Link}_{\mathbb{D}^{k+1}}(\sigma) \cong \mathbb{D}^{\dim(\sigma)} * \mathbb{S}^{k-\dim(\sigma)} \cong \mathbb{D}^{k+1}$$

and

$$\partial\sigma * \mathbb{D}^{k-\dim(\sigma)+1} \cong \mathbb{S}^{\dim(\sigma)-1} * \mathbb{D}^{k-\dim(\sigma)+1} \cong \mathbb{D}^{k+1},$$

and also

$$\partial(\sigma * \text{Link}_{\mathbb{D}^{k+1}}(\sigma)) = \partial\sigma * \text{Link}_{\mathbb{D}^{k+1}}(\sigma)$$

and

$$\partial\left(\partial\sigma * \mathbb{D}^{k-\dim(\sigma)+1}\right) = \partial\sigma * \partial\mathbb{D}^{k-\dim(\sigma)+1} = \partial\sigma * \text{Link}_{\mathbb{D}^{k+1}}(\sigma).$$

In all of these = means equality of simplicial complexes and \cong means homeomorphism. By By (iii).a,⁹ we can modify¹⁰ F so as to replace

$$F|_{\sigma * \text{Link}_{\mathbb{D}^{k+1}}(\sigma)}: \sigma * \text{Link}_{\mathbb{D}^{k+1}}(\sigma) \rightarrow Y$$

with

$$(F|_{\partial\sigma}) * G: \partial\sigma * \mathbb{D}^{k-\dim(\sigma)+1} \rightarrow \partial(F(\sigma)) * \widehat{\mathcal{L}}(F(\sigma), \mathcal{B}) \subset Y.$$

This eliminates σ , and by (ii) and (iii).b it does not introduce any new simplices of dimension at least $\dim(\sigma)$ mapping to simplices in \mathcal{B} . Repeating this over and over, we can ensure that there are no simplices σ of \mathbb{D}^{k+1} with $F(\sigma) \in \mathcal{B}$, so by (i) we have $F(\mathbb{D}^{k+1}) \subset X$, as desired. \square

6. THE COMPLEX OF SPHERES $\mathcal{S}(M_{n,b}, H)$

We now begin our discussion of the topology of the sphere complex and its subcomplexes.

6.1. Basic definitions. We start by recalling the following two definitions from the introduction.

Definition 6.1. Let $M_{n,b}$ be the connect sum of n copies of $S^2 \times S^1$ with b disjoint open balls removed. Our convention is that the connect sum of 0 copies of $S^2 \times S^1$ is the unit S^3 for the connect sum operation, so $M_{0,b}$ is S^3 with b disjoint open balls removed. \square

Definition 6.2. A 2-sphere embedded in $M_{n,b}$ is *essential* if it is not homotopic to $\partial M_{n,b}$ or a point. A *rank- k sphere system* in $M_{n,b}$ is a set $\{S_0, \dots, S_k\}$ of distinct isotopy classes of essential 2-spheres embedded in $M_{n,b}$ that can be realized disjointly. \square

Throughout this paper, we will identify isotopic essential 2-spheres and sphere systems.

⁹This is where local injectivity is key. If F were not locally injective, then F might not take $\partial\sigma$ to $\partial F(\sigma)$, so the map $(F|_{\partial\sigma}) * G$ might not be continuous.

¹⁰Unlike in previous such arguments, we cannot achieve this modification by a homotopy since $\widehat{\mathcal{L}}(F(\sigma), \mathcal{B})$ might not lie in $\text{Link}_Y(F(\sigma))$.

6.2. Compatibility with free factors. We will need to study the relationship between sphere systems and free factors in $\pi_1(M_{n,b})$. We start with the following two definitions.

Definition 6.3. For some $n \geq 0$ and $b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$. Consider a sphere system $\sigma = \{S_0, \dots, S_k\}$ in $M_{n,b}$. Let N be an open regular neighborhood of $S_0 \cup \dots \cup S_k$. The *components of the complement* of σ are the connected components of $M_{n,b} \setminus N$. The component X of the complement with $* \in \partial X$ is the *basepoint-containing component of the complement*. \square

Letting the notation be as in the previous definition, if X is the basepoint-containing component of the complement of σ , then the map $\pi_1(X, *) \rightarrow \pi_1(M_{n,b}, *)$ is injective and its image is a free factor of $\pi_1(M_{n,b}, *) \cong F_n$. We will identify $\pi_1(X, *)$ with its image in $\pi_1(M_{n,b}, *)$. This allows the following definition.

Definition 6.4. For some $n \geq 0$ and $b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. A sphere system σ in $M_{n,b}$ is *H -compatible* if the following holds. Let X be the basepoint-containing component of the complement of σ . We then require that $H \subset \pi_1(X, *)$. \square

6.3. Complex of compatible spheres. This finally brings us to the following key definition.

Definition 6.5. For some $n \geq 0$ and $b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. The *complex of H -compatible spheres* in $M_{n,b}$, denoted $\mathcal{S}(M_{n,b}, H)$, is the simplicial complex whose k -simplices are isotopy classes of H -compatible rank- k sphere systems in $M_{n,b}$. If $H = 1$, then we will sometimes omit it from our notation and just write $\mathcal{S}(M_{n,b})$. \square

This complex was introduced by Hatcher–Vogtmann [19, 20], who proved the following theorem:

Theorem 6.6 (Hatcher–Vogtmann, [20]). *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Assume that $\text{rk}(H) \leq n - 1$. Then $\mathcal{S}(M_{n,b}, H)$ is contractible.*

Remark 6.7. For $H = 1$, this was originally proved by Hatcher [18]. \square

Remark 6.8. In the published version [19] of their paper, Hatcher–Vogtmann only prove the case $b = 1$ of Theorem 6.6. In 2022, they posted the revised version [20] to the arXiv. This version fixes an error (see Remark 7.3 below), and contains the general case of Theorem 6.6. This is done in two steps: [20, Theorem 2.1] proves the case $b = 1$, and [20, Lemma 2.3] proves that $\mathcal{S}(M_{n,b+1}, H)$ deformation retracts to a complex isomorphic to $\mathcal{S}(M_{n,b}, H)$ for $b \geq 1$, proving the case $b \geq 2$. \square

Remark 6.9. Theorem 6.6 fails for $n = 0$. Since $\pi_1(M_{0,b}, *) = 1$, the free factor H is irrelevant here and we will omit it. The complex $\mathcal{S}(M_{0,b})$ is $(b - 4)$ -dimensional, and it turns out that it is homotopy equivalent to a wedge of $(b - 4)$ -dimensional spheres, i.e., is $(b - 5)$ -connected. See¹¹ [22, proof of Theorem 3.1]. \square

6.4. Low complexity cases. When $\dim(H) = n - 1$ and $b = 1$, the complex $\mathcal{S}(M_{n,b}, H)$ has a particularly simple description:

Lemma 6.10. *For some $n \geq 1$, fix a basepoint $* \in \partial M_{n,1}$ and let $H < \pi_1(M_{n,1}, *) \cong F_n$ be a free factor with $\text{rk}(H) = n - 1$. The following then hold:*

¹¹This corresponds to the case $k = 0$ and $C = \emptyset$ of part (1) of the proof of [22, Theorem 3.1]. This theorem concerns a complex of discs and spheres, but since $C = \emptyset$ we are not allowing any discs and it reduces to our complex.

- If $n = 1$ (so $H = 1$), then $\mathcal{S}(M_{1,1})$ consists of a single vertex corresponding to a nonseparating sphere in $M_{1,1}$.
- If $n \geq 2$, then $\mathcal{S}(M_{n,1}, H)$ consists of two vertices joined by an edge, one vertex corresponding to a nonseparating sphere and the other to a separating sphere.

Proof. We start by proving that $\mathcal{S}(M_{1,1})$ consists of a single vertex corresponding to a nonseparating sphere in $M_{1,1}$. First note that by basic 3-manifold topology (in particular, the uniqueness of the connect sum decomposition, see [31] or [24, §3]), we have the following:

- Any separating 2-sphere in $M_{1,1}$ is parallel to $\partial M_{1,1}$, and is thus not essential. It follows that the vertices of $\mathcal{S}(M_{1,1})$ correspond to nonseparating 2-spheres.
- Cutting $M_{1,1}$ open along a nonseparating 2-sphere yields $M_{0,3}$. From this, we see that the mapping class group of $M_{1,1}$ acts transitively on the set of isotopy classes of nonseparating 2-spheres in $M_{1,1}$. For more details, see the discussion of the “change of coordinates” principle from [12, §1.3], which concerns the related case of mapping class groups of surfaces.

Combining these two facts, it is enough to prove that there is a single mapping class group orbit of nonseparating 2-sphere in $M_{1,1}$. As we discussed in the introduction, Laudenbach [26, 27] proved that the action of the mapping class group of $M_{1,1}$ on the set of isotopy classes of 2-spheres factors through $\text{Aut}(F_1)$. We have $\text{Aut}(F_1) = \mathbb{Z}/2$, generated by the automorphism of $F_1 = \mathbb{Z}$ that takes 1 to -1 . This clearly fixes the “core” sphere of

$$M_{1,1} = S^2 \times S^1 \setminus \text{ball},$$

so this is the unique nonseparating 2-sphere up to isotopy, as desired.

We now turn to the case $n \geq 2$. Theorem 6.6 says that $\mathcal{S}(M_{n,1}, H)$ is connected, so to prove the lemma it is enough to prove that each vertex S of $\mathcal{S}(M_{n,1}, H)$ is contained in a unique edge. The vertex S is an essential H -compatible 2-sphere in $M_{n,1}$. If S is a separating 2-sphere, then since $\text{rk}(H) = n - 1$ we must have that S separates $M_{n,1}$ into components X and Y with $X \cong M_{1,1}$ and $Y \cong M_{n-1,2}$. The component Y is the basepoint-containing component and satisfies $\pi_1(Y) = H$. An edge in $\mathcal{S}(M_{n,1}, H)$ must connect S to an essential 2-sphere contained in X , and by the previous paragraph there is a unique such essential 2-sphere. It follows that S is contained in a unique edge, as desired.

If instead S is a nonseparating 2-sphere, then there is a single component Z of the complement of S . We have $Z \cong M_{n-1,3}$, and $\pi_1(Z, *) = H$. An edge in $\mathcal{S}(M_{n,1}, H)$ must connect S to an H -compatible separating 2-sphere T in Z . Letting S' and S'' be the boundary components of Z that are parallel to S in $M_{n,1}$, such a T must be the boundary of a regular neighborhood of $S' \cup S'' \cup \alpha$, where α is an arc connecting S' to S'' . By the lightbulb trick (see, e.g., [33, Exercise 9.F.4]), there is a unique such arc up to isotopy, so there is unique such T , as desired. \square

6.5. Dual graph to vertex. The following notion will be very useful.

Definition 6.11. Let σ be a sphere system in $M_{n,b}$. The *dual graph* of σ , denoted $\Gamma(\sigma)$, is following graph:

- The vertices of $\Gamma(\sigma)$ are the components of the complement of σ .
- The edges of $\Gamma(\sigma)$ are in bijection with the spheres in σ , and the edge corresponding to $S \in \sigma$ connects the vertices corresponding to the components on either side of S . \square

These satisfy the following lemma:

Lemma 6.12. *Let σ be a sphere system in $M_{n,b}$. Let X_1, \dots, X_r be the components of the complement of σ . Write $X_i \cong M_{n_i, b_i}$. Then*

$$n = \text{rk}(\pi_1(\Gamma(\sigma))) + \sum_{i=1}^r n_i.$$

Proof. Immediate. □

6.6. Cohen–Macaulay. Using Theorem 6.6, we can prove the following.

Theorem 6.13. *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Assume that $(n, b) \neq (1, 1)$ and that $\text{rk}(H) \leq n - 1$. Then $\mathcal{S}(M_{n,b}, H)$ is weakly Cohen–Macaulay of dimension $n - \text{rk}(H)$.*

Remark 6.14. The conditions that $(n, b) \neq (1, 1)$ and $\text{rk}(H) \leq n - 1$ are necessary. Indeed, if $\text{rk}(H) = n$ then $\mathcal{S}(M_{n,1}, H)$ is the empty set, and thus is not weakly Cohen–Macaulay of dimension 0. Similarly, by Lemma 6.10 the complex $\mathcal{S}(M_{1,1})$ is a single point. Thus while it is contractible (and hence connected), it is not weakly Cohen–Macaulay of dimension 1. □

Proof of Theorem 6.13. Given our assumptions, Theorem 6.6 implies that $\mathcal{S}(M_{n,b}, H)$ is contractible, and thus is certainly $(n - 1 - \text{rk}(H))$ -connected. Letting $\sigma = \{S_0, \dots, S_k\}$ be a k -simplex of $\mathcal{S}(M_{n,b}, H)$, we must prove that $\text{Link}_{\mathcal{S}(M_{n,b}, H)}(\sigma)$ is $(n - 2 - k - \text{rk}(H))$ -connected. What we will prove is that either

- $\text{Link}_{\mathcal{S}(M_{n,b}, H)}(\sigma)$ is contractible, or
- $n = k + 1 + \text{rk}(H)$ and $\text{Link}_{\mathcal{S}(M_{n,b}, H)}(\sigma)$ is nonempty, or
- $n \leq k + \text{rk}(H)$.

This will imply that $\text{Link}_{\mathcal{S}(M_{n,b}, H)}(\sigma)$ is always $(n - 2 - k - \text{rk}(H))$ -connected.

Let the components of the complement of σ be X_1, \dots, X_r with X_1 the basepoint-containing component. We then have that

$$\text{Link}_{\mathcal{S}(M_{n,b}, H)}(\sigma) = \mathcal{S}(X_1, H) * \mathcal{S}(X_2) * \dots * \mathcal{S}(X_r).$$

Write $X_i = M_{n_i, b_i}$. If $\text{rk}(H) \leq n_1 - 1$, then Theorem 6.6 implies that $\mathcal{S}(X_1, H)$ is contractible, so $\text{Link}_{\mathcal{S}(M_{n,b}, H)}(\sigma)$ is contractible. Similarly, if $n_i \geq 1$ for any $2 \leq i \leq r$, then Theorem 6.6 implies that $\mathcal{S}(X_i)$ is contractible, so $\text{Link}_{\mathcal{S}(M_{n,b}, H)}(\sigma)$ is contractible. We can therefore assume without loss of generality that $n_1 = \text{rk}(H)$ and that $n_i = 0$ for $2 \leq i \leq r$. Our goal then reduces to proving that $n \leq k + 1 + \text{rk}(H)$, and that if $n = k + 1 + \text{rk}(H)$ then $\text{Link}_{\mathcal{S}(M_{n,b}, H)}(\sigma)$ is nonempty.

Letting $\Gamma(\sigma)$ be the dual graph of σ , Lemma 6.12 implies that

$$n = \text{rk}(\pi_1(\Gamma(\sigma))) + \sum_{i=1}^r n_i = \text{rk}(\pi_1(\Gamma(\sigma))) + \text{rk}(H).$$

Since σ contains $(k + 1)$ spheres, the graph $\Gamma(\sigma)$ has $(k + 1)$ edges. It follows that $\text{rk}(\pi_1(\Gamma(\sigma))) \leq k + 1$, so

$$n \leq k + 1 + \text{rk}(H).$$

This is half of what we are trying to prove. What remains is to show that if this inequality is an equality, then $\text{Link}_{\mathcal{S}(M_{n,b}, H)}(\sigma)$ is nonempty.

For our inequality to be an equality, we must have $\text{rk}(\pi_1(\Gamma(\sigma))) = k + 1$. Since $\Gamma(\sigma)$ has $(k + 1)$ edges, this can only happen if $\Gamma(\sigma)$ has a single vertex and all the edges are self-loops. In other words, there is only one component of the complement of σ , namely the basepoint-containing component X_1 . We thus have

$$X_1 \cong M_{n-k-1, b+2k+2} \quad \text{and} \quad H = \pi_1(X_1).$$

If either $b \geq 2$ or $k \geq 1$, then X_1 has at least 4 boundary components, so X_1 has an H -compatible separating¹² sphere that cuts off two boundary spheres and thus

$$\text{Link}_{\mathcal{S}(M_{n,b},H)}(\sigma) \cong \mathcal{S}(X_1, H) \neq \emptyset.$$

If instead $b = 1$ and $k = 0$, then $X_1 \cong M_{n-1,3}$. This is where we finally invoke our assumption that $(n, b) \neq (1, 1)$, which implies that X_1 is not just a 3-holed sphere, so X_1 has an H -compatible separating sphere that cuts off two boundary spheres and

$$\text{Link}_{\mathcal{S}(M_{n,b},H)}(\sigma) \cong \mathcal{S}(X_1, H) \neq \emptyset. \quad \square$$

7. THE NONSEPARATING COMPLEX OF SPHERES $\mathcal{N}(M_{n,b}, H)$

The next step is to consider several subcomplexes of $\mathcal{S}(M_{n,b}, H)$.

7.1. Nonseparating simplices. We start with the following.

Definition 7.1. For some $n \geq 0$ and $b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. The *nonseparating complex of H -compatible spheres* in $M_{n,b}$, denoted $\mathcal{N}(M_{n,b}, H)$, is the subcomplex of $\mathcal{S}(M_{n,b}, H)$ whose k -simplices are H -compatible rank- k sphere systems $\{S_0, \dots, S_k\}$ in $M_{n,b}$ such that $S_0 \cup \dots \cup S_k$ does not separate $M_{n,b}$. \square

7.2. High connectivity. The following theorem of Hatcher–Vogtmann [19, 20] says that these complexes are highly connected:

Theorem 7.2 (Hatcher–Vogtmann [20]). *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Then, $\mathcal{N}(M_{n,b}, H)$ is $(n - \text{rk}(H) - 2)$ -connected.*

Remark 7.3. Hatcher–Vogtmann’s published proof of Theorem 7.2 in [19] has a mistake: [19, Lemma 2.3] claims that $\mathcal{N}(M_{n,b+1}, H)$ deformation retracts to $\mathcal{N}(M_{n,b}, H)$, which is false (e.g., $\mathcal{N}(M_{1,1})$ is a single point, but $\mathcal{N}(M_{1,2})$ is an infinite discrete set). The issue is that their deformation retraction uses simplices that do not lie in $\mathcal{N}(M_{n,b+1}, H)$. In 2022, they posted a revised version [20] of their paper to the arXiv correcting this mistake. Theorem 7.2 is [20, Theorem 2.5]. \square

7.3. Cohen–Macaulay. Using Theorem 7.2, we can prove the following.

Theorem 7.4. *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Then $\mathcal{N}(M_{n,b}, H)$ is weakly Cohen–Macaulay of dimension $n - \text{rk}(H) - 1$.*

Proof. Immediate from Theorem 7.2 along with the fact that if $\sigma = \{S_0, \dots, S_k\}$ is a k -simplex of $\mathcal{N}(M_{n,b}, H)$, then

$$\text{Link}_{\mathcal{N}(M_{n,b},H)}(\sigma) \cong \mathcal{N}(M_{n-k-1,b+2k+2}, H).$$

The point here is that the only connected component of the complement of σ is homeomorphic to $M_{n-k-1,b+2k+2}$. \square

8. THE AUGMENTED NONSEPARATING COMPLEX OF SPHERES $\mathcal{AN}(M_{n,b}, H)$ I: SPHERE LOCAL INJECTIVITY

We now add some simplices to the nonseparating complex of spheres $\mathcal{N}(M_{n,b}, H)$ to increase its connectivity by one.

¹²For later use, note that if $k \geq 1$ then we can choose the H -compatible separating sphere in X_1 such that it becomes nonseparating in the larger 3-manifold $M_{n,b}$.

8.1. Separating core. This requires the following definition.

Definition 8.1. Let σ be a sphere system on $M_{n,b}$ with dual graph $\Gamma(\sigma)$. The *separating core* of σ is the face σ' of σ consisting of all $S \in \sigma$ such that the edge of $\Gamma(\sigma)$ corresponding to S is not a loop. The *nonseparating periphery* of σ is the face of σ consisting of all $S \in \sigma$ such that the edge of $\Gamma(\sigma)$ corresponding to S is a loop. \square

Another way to think about this is as follows. Let σ be a sphere system on $M_{n,b}$ with separating core σ' and nonseparating periphery σ'' . As simplices of $\mathcal{S}(M_{n,b})$, we have $\sigma = \sigma' * \sigma''$. The following hold:

- Any proper face of σ' has strictly fewer components in its complement than σ' .
- Let X_1, \dots, X_r be the components of the complement of σ' . The link of σ' in $\mathcal{S}(M_{n,b})$ is thus

$$\mathcal{S}(X_1) * \dots * \mathcal{S}(X_r).$$

Then σ'' lies in the subcomplex

$$\mathcal{N}(X_1) * \dots * \mathcal{N}(X_r) \subset \mathcal{S}(X_1) * \dots * \mathcal{S}(X_r).$$

8.2. Augmented complex. The definition of our complex is as follows.

Definition 8.2. For some $n \geq 0$ and $b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. The *augmented nonseparating complex of H -compatible spheres* in $M_{n,b}$, denoted $\mathcal{AN}(M_{n,b}, H)$, is the subcomplex of $\mathcal{S}(M_{n,b}, H)$ whose k -simplices are H -compatible rank- k sphere systems $\sigma = \{S_0, \dots, S_k\}$ in $M_{n,b}$ with the following properties:

- Each S_i is a nonseparating sphere.
- Let σ' be the separating core of σ . Then either σ' is empty (so σ does not separate $M_{n,b}$), or σ' has two components X and Y in its complement. Moreover, in the latter case if X is the basepoint-containing component then $\pi_1(Y) = 1$. \square

8.3. Connectivity theorem. Recall that Theorem 7.2 says that $\mathcal{N}(M_{n,b}, H)$ is $(n - 2 - \text{rk}(H))$ -connected. Our main theorem about the augmented complexes is as follows:

Theorem 8.3. *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Assume that $\text{rk}(H) \leq n - 1$. Then $\mathcal{AN}(M_{n,b}, H)$ is $(n - 1 - \text{rk}(H))$ -connected.*

8.4. A stronger result. Some degenerate cases of Theorem 8.3 are immediate:

- If $b = 1$ and $\text{rk}(H) = n - 1$, then by Lemma 6.10 the complex $\mathcal{AN}(M_{n,b}, H)$ consists of a single vertex corresponding to a nonseparating curve. It is thus contractible, and in particular is $(n - 1 - \text{rk}(H)) = 0$ connected.

We thus can exclude these cases. In the remaining ones, we will actually prove the following stronger result.

Theorem 8.4. *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Assume that $\text{rk}(H) \leq n - 1$ and that if $\text{rk}(H) = n - 1$, then $b \geq 2$. Then $\mathcal{AN}(M_{n,b}, H)$ has the sphere and disc local injectivity properties up to dimension $(n - 1 - \text{rk}(H))$.*

By Lemma 3.7, this will imply that $\mathcal{AN}(M_{n,b}, H)$ is $(n - 1 - \text{rk}(H))$ -connected.

Remark 8.5. Theorem 8.4 is weaker than saying that these complexes are weakly Cohen–Macaulay. We do not know if this strong condition holds. \square

We divide the proof of Theorem 8.4 into two parts: in the rest of this section, we prove sphere local injectivity (see Lemma 8.6), and in §10 – §11 we prove disc local injectivity (see Lemma 11.3).

8.5. Sphere local injectivity. The following result takes care of the sphere local injectivity part of Theorem 8.4. We note that the hypothesis that $b \geq 2$ if $\text{rk}(H) = n - 1$ is not needed here, but will be needed for the disc local injectivity property.

Lemma 8.6. *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Assume that $\text{rk}(H) \leq n - 1$. Then $\mathcal{AN}(M_{n,b}, H)$ has the sphere local injectivity property up to dimension $(n - 1 - \text{rk}(H))$.*

Proof. Theorem 7.4 says that $\mathcal{N}(M_{n,b}, H)$ is weakly Cohen–Macaulay of dimension $(n - 1 - \text{rk}(H))$, so by Lemma 4.4 it has the sphere local injectivity property up to dimension $(n - 1 - \text{rk}(H))$. It follows that it is enough to prove that for $k \leq n - 1 - \text{rk}(H)$, every map $\mathbb{S}^k \rightarrow \mathcal{AN}(M_{n,b}, H)$ is homotopic to one whose image lies in $\mathcal{N}(M_{n,b}, H)$. Letting $d = n - 1 - \text{rk}(H)$, this is equivalent to proving that the inclusion

$$\mathcal{N}(M_{n,b}, H) \hookrightarrow \mathcal{AN}(M_{n,b}, H)$$

is d -connected. This will follow from Proposition 5.2 once we verify its hypotheses. The input to Proposition 5.2 is a set \mathcal{B} of “bad simplices”, which for us will be as follows:

- The set \mathcal{B} consist of all simplices $\sigma = \{S_0, \dots, S_k\}$ of $\mathcal{AN}(M_{n,b}, H)$ such that the union of the S_i separates $M_{n,b}$, but no proper subset of σ separates $M_{n,b}$. In other words, σ is its own separating core.

We now verify each hypothesis of Proposition 5.2 in turn.

Condition (i) says that a simplex of $\mathcal{AN}(M_{n,b}, H)$ lies in $\mathcal{N}(M_{n,b}, H)$ if and only if none of its faces are in \mathcal{B} , which is immediate from the definitions.

Condition (ii) says that if $\sigma, \sigma' \in \mathcal{B}$ are such that $\sigma \cup \sigma'$ is a simplex of $\mathcal{AN}(M_{n,b}, H)$, then $\sigma \cup \sigma' \in \mathcal{B}$. In fact, since a simplex of $\mathcal{AN}(M_{n,b}, H)$ can separate $M_{n,b}$ into at most 2 components, the only way that $\sigma \cup \sigma'$ can be a simplex of $\mathcal{AN}(M_{n,b}, H)$ is for $\sigma = \sigma'$, so this condition is trivial.

Condition (iii) says that for all $\sigma \in \mathcal{B}$, the complex $\mathcal{L}(\sigma, \mathcal{B})$ defined in Definition 5.1 is

$$d - \dim(\sigma) - 1 = (n - 1 - \text{rk}(H)) - \dim(\sigma) - 1 = n - 2 - \text{rk}(H) - \dim(\sigma)$$

connected. For our \mathcal{B} , the complex $\mathcal{L}(\sigma, \mathcal{B})$ has the following concrete description. Let X and Y be the components of the complement of σ , with X the basepoint-containing component. Since $\pi_1(Y) = 1$, all 2-spheres in Y separate Y . We thus have

$$(8.1) \quad \mathcal{L}(\sigma, \mathcal{B}) = \mathcal{N}(Y) * \mathcal{N}(X, H) = \mathcal{N}(X, H).$$

Write $X \cong M_{n',b'}$. Theorem 7.2 says that $\mathcal{L}(\sigma, \mathcal{B}) = \mathcal{N}(X, H)$ is $(n' - 2 - \text{rk}(H))$ -connected, so to prove that it is $(n - 2 - \text{rk}(H) - \dim(\sigma))$ -connected we must prove that $n = n' + \dim(\sigma)$.

The dual graph $\Gamma(\sigma)$ has two vertices corresponding to X and Y and $\dim(\sigma) + 1$ edges, so

$$\pi_1(\Gamma(\sigma)) \cong F_{\dim(\sigma)}.$$

Lemma 6.12 now says that

$$n = \text{rk}(\pi_1(\Gamma)) + \text{rk}(\pi_1(X)) + \text{rk}(\pi_1(Y)) = \dim(\sigma) + n' + 0,$$

as desired. □

9. THE COMPLEX OF NONSEPARATING SPHERES $\mathcal{S}_{\text{NS}}(M_{n,b}, H)$

Before we can prove disc local injectivity for the augmented nonseparating sphere complex, we must study the subcomplex of the sphere complex where vertices are nonseparating, but where higher-dimensional simplices can separate.

9.1. Nonseparating spheres, absolute version. The definition is as follows.

Definition 9.1. For some $n \geq 0$ and $b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. The *complex of H -compatible nonseparating spheres* in $M_{n,b}$, denoted $\mathcal{S}_{\text{ns}}(M_{n,b}, H)$, is the full subcomplex of $\mathcal{S}(M_{n,b}, H)$ whose vertices are the isotopy classes of H -compatible essential nonseparating spheres in $M_{n,b}$. If $H = 1$, then we will sometimes omit it from our notation and just write $\mathcal{S}_{\text{ns}}(M_{n,b})$. \square

9.2. Relative version. Our main theorem about $\mathcal{S}_{\text{ns}}(M_{n,b}, H)$ is that it is weakly Cohen–Macaulay of dimension $n - \text{rk}(H)$, just like $\mathcal{S}(M_{n,b}, H)$ (c.f. Theorem 6.13). A technical issue that will arise when studying links in $\mathcal{S}_{\text{ns}}(M_{n,b}, H)$ is that a sphere can separate a submanifold of a 3-manifold without separating the whole manifold. We thus make the following definition:

Definition 9.2. For some $n \geq 0$ and $b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Let M be another connected 3-manifold with boundary such that $M_{n,b}$ is a submanifold of M . The *complex of H -compatible M -nonseparating spheres* in $M_{n,b}$, denoted $\mathcal{S}_{\text{ns}}(M_{n,b}, M, H)$, is the full subcomplex of $\mathcal{S}(M_{n,b}, H)$ whose vertices are the isotopy classes of H -compatible essential spheres in $M_{n,b}$ that do not separate M . If $H = 1$, then we will sometimes omit it from our notation and just write $\mathcal{S}_{\text{ns}}(M_{n,b}, M)$. \square

9.3. Contractibility. The following is the analogue for $\mathcal{S}_{\text{ns}}(M_{n,b}, M, H)$ of Theorem 6.6 for $\mathcal{S}(M_{n,b}, H)$:

Theorem 9.3. *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Assume that $\text{rk}(H) \leq n - 1$, and let M be another connected 3-manifold with boundary such that $M_{n,b}$ is a submanifold of M . Then $\mathcal{S}_{\text{ns}}(M_{n,b}, M, H)$ is contractible.*

Proof. The proof will be by induction on the pair (n, b) , ordered lexicographically. For the base case $(n, b) = (1, 1)$, our assumption on H means that $H = 1$. By Lemma 6.10, the complex

$$\mathcal{S}_{\text{ns}}(M_{1,1}, M) = \mathcal{S}(M_{1,1})$$

is a single vertex represented by a nonseparating sphere, and is thus contractible.

Assume, therefore, that $(n, b) \neq (1, 1)$ and that the theorem is true whenever (n, b) is smaller. Theorem 6.6 says that $\mathcal{S}(M_{n,b}, H)$ is contractible, so it is enough to prove that the inclusion

$$\mathcal{S}_{\text{ns}}(M_{n,b}, M, H) \subset \mathcal{S}(M_{n,b}, H)$$

is d -connected for all $d \geq 0$. This will follow from Proposition 5.2 once we verify its hypotheses. The input to Proposition 5.2 is a set \mathcal{B} of “bad simplices”, which for us will be as follows:

- Let \mathcal{B} be the set of all simplices σ of $\mathcal{S}(M_{n,b}, H)$ such that each vertex of σ separates the ambient manifold M (and thus is not a vertex of $\mathcal{S}_{\text{ns}}(M_{n,b}, M, H)$).

We now verify each hypothesis of Proposition 5.2 in turn.

Condition (i) says that a simplex of $\mathcal{S}(M_{n,b}, H)$ lies in $\mathcal{S}_{\text{ns}}(M_{n,b}, M, H)$ if and only if none of its faces are in \mathcal{B} , which is immediate from the definitions.

Condition (ii) says that if $\sigma, \sigma' \in \mathcal{B}$ are such that $\sigma \cup \sigma'$ is a simplex of $\mathcal{S}(M_{n,b}, H)$, then $\sigma \cup \sigma' \in \mathcal{B}$. Again, this is immediate from the definitions.

Condition (iii) says that for all $\sigma \in \mathcal{B}$, the complex $\mathcal{L}(\sigma, \mathcal{B})$ defined in Definition 5.1 is $(d - \dim(\sigma) - 1)$ -connected for all $d \geq 0$, i.e., is contractible. For our \mathcal{B} , the complex $\mathcal{L}(\sigma, \mathcal{B})$ has the following concrete description. Let X_1, \dots, X_r be the components of the complement of σ in $M_{n,b}$, with X_1 the basepoint-containing component. We then have

$$(9.1) \quad \mathcal{L}(\sigma, \mathcal{B}) = \mathcal{S}_{\text{ns}}(X_1, M, H) * \mathcal{S}_{\text{ns}}(X_2, M) * \cdots * \mathcal{S}_{\text{ns}}(X_r, M).$$

To prove that $\mathcal{L}(\sigma, \mathcal{B})$ is contractible, it is enough to prove that at least one term in this join is contractible.

Write $X_i \cong M_{n_i, b_i}$. Recall that we are inducting on the pair (n, b) , ordered lexicographically. Since X_i is a submanifold of $M_{n, b}$, we must have $n_i \leq n$. Moreover, if $n_i = n$ then since at least one component of ∂X_i is an essential separating¹³ 2-sphere in $M_{n, b}$ and none of the components of ∂X_i bound balls in $M_{n, b}$, it follows that we must have $b_i < b$. In other words, (n_i, b_i) is strictly less than (n, b) in the lexicographic ordering.

This might lead the reader to think that our inductive hypothesis applies to each X_i , and thus that all the terms in (9.1) are contractible. However, there is an issue: this only works if the corresponding complex satisfies the hypotheses of our theorem, and this might not hold. Since we only need one term in (9.1) to be contractible, it is enough to prove that one of the following two things hold:

(†) For some $2 \leq i \leq r$, we have $n_i \geq 1$. Then the term

$$\mathcal{S}_{\text{ns}}(X_i, M) \cong \mathcal{S}_{\text{ns}}(M_{n_i, b_i}, M)$$

of (9.1) satisfies the the hypothesis of our theorem, so by our inductive hypothesis is contractible.

(††) We have $n_1 \geq 1$ and $\text{rk}(H) \leq n_1 - 1$. Then the term

$$\mathcal{S}_{\text{ns}}(X_1, M, H) \cong \mathcal{S}_{\text{ns}}(M_{n_1, b_1}, M, H)$$

of (9.1) satisfies the the hypothesis of our theorem, so by our inductive hypothesis is contractible.

Assume that (†) does not hold, so $n_i = 0$ for $2 \leq i \leq r$. We will prove that (††) then holds. Since each sphere in σ separates the manifold M , it also separates $M_{n, b}$. The dual graph $\Gamma(\sigma)$ is thus a tree. Lemma 6.12 therefore says that

$$n = \text{rk}(\pi_1(\Gamma(\sigma))) + \sum_{i=1}^r n_i = 0 + n_1 + \sum_{i=2}^r 0 = n_1.$$

We thus have $n_1 = n$, so $n_1 \geq 1$ and $\text{rk}(H) \leq n_1 - 1$ by assumption and (††) holds, as desired. \square

9.4. Cohen–Macaulay. Using Theorem 9.3, we can prove the following.

Theorem 9.4. *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n, b}$ and let $H < \pi_1(M_{n, b}, *) \cong F_n$ be a free factor. Assume that $\text{rk}(H) \leq n - 1$, and if $\text{rk}(H) = n - 1$ then assume that $b \geq 2$. Then $\mathcal{S}_{\text{ns}}(M_{n, b}, H)$ is weakly Cohen–Macaulay of dimension $n - \text{rk}(H)$.*

Remark 9.5. The conditions that $\text{rk}(H) \leq n - 1$ and that if $\text{rk}(H) = n - 1$ then $b \geq 2$ are necessary. Indeed, if $\text{rk}(H) = n$ then $\mathcal{S}_{\text{ns}}(M_{n, 1}, H)$ is the empty set, and thus is not weakly Cohen–Macaulay of dimension 0. Similarly, if $\text{rk}(H) = n - 1$ then Lemma 6.10 implies that the complex $\mathcal{S}_{\text{ns}}(M_{n, 1})$ is a single point. Thus while it is contractible (and hence connected), it is not weakly Cohen–Macaulay of dimension 1. \square

Proof of Theorem 9.4. The proof is almost identical to that of Theorem 6.13, so we just indicate the necessary changes:

- Use Theorem 9.3 in place of Theorem 6.6.
- The “relative” version of our complex arises as follows. Consider a simplex σ of $\mathcal{S}_{\text{ns}}(M_{n, b}, H)$. We wish to understand the link of σ . Let the components of the

¹³In fact, it has to separate the ambient manifold M , which is even stronger.

complement of σ be X_1, \dots, X_r with X_1 the basepoint-containing component. We then have

$$\text{Link}_{\mathcal{S}_{\text{ns}}(M_{n,b}, H)}(\sigma) = \mathcal{S}_{\text{ns}}(X_1, M_{n,b}, H) * \mathcal{S}(X_2, M_{n,b}) * \cdots * \mathcal{S}(X_r, M_{n,b})$$

since vertices of the link must not separate $M_{n,b}$ (though they can separate the X_i).

- Finally, the last paragraph of the proof must be adjusted to ensure that the sphere that arises does not separate $M_{n,b}$. The key point here is that the final sentence (where $b = 1$ and $k = 0$) is not needed due to our assumption that $b \geq 2$ if $\text{rk}(H) = n - 1$. \square

10. THE AUGMENTED NONSEPARATING COMPLEX OF SPHERES $\mathcal{AN}(M_{n,b}, H)$ II: EXPANDED DISC LOCAL INJECTIVITY

We now begin studying disc local injectivity for the augmented nonseparating complex of spheres. The first step is to establish this for a slightly larger complex.

10.1. Basic definition. The definition of our complex is as follows.

Definition 10.1. For some $n \geq 0$ and $b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. The *expanded augmented nonseparating complex of H -compatible spheres* in $M_{n,b}$, denoted $\mathcal{AN}_{\text{ex}}(M_{n,b}, H)$, is the subcomplex of $\mathcal{S}(M_{n,b}, H)$ whose k -simplices are H -compatible rank- k sphere systems $\{S_0, \dots, S_k\}$ in $M_{n,b}$ with the following properties:

- Each S_i is a nonseparating sphere.
- The union $S_0 \cup \cdots \cup S_k$ either does not separate $M_{n,b}$, or separates it into exactly two components. \square

We thus have $\mathcal{AN}(M_{n,b}, H) \subset \mathcal{AN}_{\text{ex}}(M_{n,b}, H)$.

10.2. Disc local injectivity property. Lemma 8.6 says that $\mathcal{AN}(M_{n,b}, H)$ has the sphere local injectivity property up to dimension $(n - 1 - \text{rk}(H))$ as long as $\text{rk}(H) \leq n - 1$. The same is true for $\mathcal{AN}_{\text{ex}}(M_{n,b}, H)$, with a very similar proof. We will not need this, but we will need the disc local injectivity property:

Lemma 10.2. *For some $n, b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Assume that $\text{rk}(H) \leq n - 1$ and that if $\text{rk}(H) = n - 1$, then $b \geq 2$. Then $\mathcal{AN}_{\text{ex}}(M_{n,b}, H)$ has the disc local injectivity property up to dimension $(n - 1 - \text{rk}(H))$.*

Proof. Theorem 9.4 says that $\mathcal{S}_{\text{ns}}(M_{n,b}, H)$ is weakly Cohen–Macaulay of dimension $(n - \text{rk}(H))$, so by Lemma 4.4 it has the disc local injectivity property up to dimension $d = (n - 1 - \text{rk}(H))$. To prove the same for $\mathcal{AN}_{\text{ex}}(M_{n,b}, H)$, it is enough to verify the conditions of Proposition 5.3 for the inclusion

$$\mathcal{AN}_{\text{ex}}(M_{n,b}, H) \hookrightarrow \mathcal{S}_{\text{ns}}(M_{n,b}, H).$$

The input to Proposition 5.3 is a set \mathcal{B} of “bad simplices”, which for us will be as follows:

- The set \mathcal{B} consist of all simplices $\sigma = \{S_0, \dots, S_k\}$ of $\mathcal{S}_{\text{ns}}(M_{n,b}, H)$ such that the union of the S_i separates $M_{n,b}$ into $c \geq 3$ components, and all proper subsets of σ separate $M_{n,b}$ into fewer than c components. In particular, σ is its own separating core.

Another way of describing this condition is that the dual graph $\Gamma(\sigma)$ has at least 3 vertices and none of its edges are loops. We now verify each hypothesis of Proposition 5.3 in turn.

Condition (i) says that a simplex of $\mathcal{S}_{\text{ns}}(M_{n,b}, H)$ lies in $\mathcal{AN}_{\text{ex}}(M_{n,b}, H)$ if and only if none of its faces are in \mathcal{B} , which is immediate from the definitions.

Condition (ii) says that if $\sigma, \sigma' \in \mathcal{B}$ are such that $\sigma \cup \sigma'$ is a simplex of $\mathcal{S}_{\text{ns}}(M_{n,b}, H)$, then $\sigma \cup \sigma' \in \mathcal{B}$. This is also immediate from the definitions.

Condition (iii) says that for all $\sigma \in \mathcal{B}$, there exists a subcomplex $\widehat{\mathcal{L}}(\sigma, \mathcal{B})$ of $\mathcal{S}_{\text{ns}}(M_{n,b}, H)$ with $\mathcal{L}(\sigma, \mathcal{B}) \subset \widehat{\mathcal{L}}(\sigma, \mathcal{B})$ that satisfies conditions (iii).a–(iii).c. Here $\mathcal{L}(\sigma, \mathcal{B})$ is the complex defined in Definition 5.1. Fixing some $\sigma \in \mathcal{B}$, we will take $\widehat{\mathcal{L}}(\sigma, \mathcal{B}) = \mathcal{L}(\sigma, \mathcal{B})$. The complex $\mathcal{L}(\sigma, \mathcal{B})$ has the following concrete description. Let X_1, \dots, X_r be the components of the complement of σ , with X_1 the basepoint-containing component. Then

$$\widehat{\mathcal{L}}(\sigma, \mathcal{B}) = \mathcal{L}(\sigma, \mathcal{B}) = \mathcal{N}(X_1, H) * \mathcal{N}(X_2) * \cdots * \mathcal{N}(X_r).$$

We verify (iii).a–(iii).c for this as follows.

Condition (iii).a says that $\partial\sigma * \widehat{\mathcal{L}}(\sigma, \mathcal{B}) \subset \mathcal{S}_{\text{ns}}(M_{n,b}, H)$. In fact, $\widehat{\mathcal{L}}(\sigma, \mathcal{B})$ is contained in the link of σ , so $\sigma * \widehat{\mathcal{L}}(\sigma, \mathcal{B}) \subset \mathcal{S}_{\text{ns}}(M_{n,b}, H)$.

Condition (iii).b says that all simplices of $\partial\sigma * \widehat{\mathcal{L}}(\sigma, \mathcal{B})$ that are in \mathcal{B} lie in $\partial\sigma$. Again, even more is true: all simplices of $\sigma * \widehat{\mathcal{L}}(\sigma, \mathcal{B})$ that are in \mathcal{B} lie in σ .

Finally, Condition (iii).c says that $\widehat{\mathcal{L}}(\sigma, \mathcal{B})$ has the disc local injectivity property up to dimension

$$d - \dim(\sigma) = (n - 1 - \text{rk}(H)) - \dim(\sigma).$$

Let $X_i \cong M_{n_i, b_i}$. Theorem 7.4 says that $\mathcal{N}(X_1, H)$ is weakly Cohen–Macaulay of dimension $n_1 - 1 - \text{rk}(H)$, so by Lemma 4.4 it has the disc local injectivity property up to dimension $n_1 - 2 - \text{rk}(H)$. Similarly, for $2 \leq i \leq r$ the complex $\mathcal{N}(X_i)$ has the disc local injectivity property up to dimension $n_i - 2$. Applying Lemma 3.8, we see that

$$\widehat{\mathcal{L}}(\sigma, \mathcal{B}) = \mathcal{L}(\sigma, \mathcal{B}) = \mathcal{N}(X_1, H) * \mathcal{N}(X_2) * \cdots * \mathcal{N}(X_r)$$

has the disc local injectivity property up to dimension

$$(n_1 - 2 - \text{rk}(H)) + (n_2 - 2) + \cdots + (n_r - 2) + 2(r - 1) = (n_1 + \cdots + n_r) - 2 - \text{rk}(H).$$

We must prove that this is at least $n - 1 - \text{rk}(H) - \dim(\sigma)$, i.e., that

$$n \leq n_1 + \cdots + n_r - 1 + \dim(\sigma).$$

Letting $\Gamma(\sigma)$ be the dual graph of σ , Lemma 6.12 says that

$$n = \text{rk}(\pi_1(\Gamma(\sigma))) + n_1 + \cdots + n_r.$$

Thus reduces us to proving that

$$\text{rk}(\pi_1(\Gamma(\sigma))) \leq \dim(\sigma) - 1.$$

The graph $\Gamma(\sigma)$ has r vertices and $\dim(\sigma) + 1$ edges. It follows that

$$1 - \text{rk}(\pi_1(\Gamma)) = r - \dim(\sigma) - 1 \quad \text{and hence} \quad \text{rk}(\pi_1(\Gamma)) = 2 + \dim(\sigma) - r.$$

This reduces our desired inequality to

$$2 + \dim(\sigma) - r \leq \dim(\sigma) - 1,$$

i.e., $r \geq 3$. This is exactly the defining criterion for simplices in \mathcal{B} : they must have at least 3 components in their complement. The lemma follows. \square

11. THE AUGMENTED NONSEPARATING COMPLEX OF SPHERES $\mathcal{AN}(M_{n,b}, H)$ III: DISC LOCAL INJECTIVITY

We now prove disc local injectivity for the augmented nonseparating complex of spheres $\mathcal{AN}(M_{n,b}, H)$.

11.1. Filtration. Our goal is to deduce this from disc local injectivity of the expanded complex $\mathcal{AN}_{\text{ex}}(M_{n,b}, H)$. To do this, we consider a sequence of complexes that interpolate between $\mathcal{AN}_{\text{ex}}(M_{n,b}, H)$ and $\mathcal{AN}(M_{n,b}, H)$:

Definition 11.1. For some $n \geq 0$ and $b \geq 1$ and $m \geq 0$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. The *rank- m expanded augmented nonseparating complex of H -compatible spheres* in $M_{n,b}$, denoted $\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H)$, is the subcomplex of $\mathcal{S}(M_{n,b}, H)$ whose k -simplices are H -compatible rank- k sphere systems $\sigma = \{S_0, \dots, S_k\}$ in $M_{n,b}$ with the following properties:

- Each S_i is a nonseparating sphere.
- Let σ' be the separating core of σ . Then either σ' is empty (so σ does not separate $M_{n,b}$), or σ' has two components X and Y in its complement. Moreover, in the latter case if X is the basepoint-containing component then $\text{rk}(\pi_1(Y)) \leq m$. \square

We have the following:

Lemma 11.2. *For some $n \geq 0$ and $b \geq 1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Then for all simplices σ of $\mathcal{AN}_{\text{ex}}(M_{n,b}, H)$ and all components Z of the complement of σ , we have $\text{rk}(\pi_1(Z)) \leq n-1$. In particular, $\mathcal{AN}_{\text{ex}}^{n-1}(M_{n,b}, H) = \mathcal{AN}_{\text{ex}}(M_{n,b}, H)$.*

Proof. Since all vertices of σ are nonseparating spheres, removing the edge of the dual graph $\Gamma(\sigma)$ corresponding to any one of them does not separate $\Gamma(\sigma)$. This implies that $\Gamma(\sigma)$ has a nontrivial fundamental group. The lemma now follows from Lemma 6.12. \square

It follows from this that we have

$$\begin{aligned} \mathcal{AN}(M_{n,b}, H) &= \mathcal{AN}_{\text{ex}}^0(M_{n,b}, H) \subset \mathcal{AN}_{\text{ex}}^1(M_{n,b}, H) \\ &\subset \dots \subset \mathcal{AN}_{\text{ex}}^{n-1}(M_{n,b}, H) = \mathcal{AN}_{\text{ex}}(M_{n,b}, H). \end{aligned}$$

11.2. Disc local injectivity. Our result is as follows. The case $m = 0$ of it completes the proof of Theorem 8.4.

Lemma 11.3. *For some $n, b \geq 1$ and $0 \leq m \leq n-1$, fix a basepoint $* \in \partial M_{n,b}$ and let $H < \pi_1(M_{n,b}, *) \cong F_n$ be a free factor. Assume that $\text{rk}(H) \leq n-1$ and that if $\text{rk}(H) = n-1$, then $b \geq 2$. Then $\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H)$ has the disc local injectivity property up to dimension $(n-1 - \text{rk}(H))$.*

Proof. The proof is by reverse induction on m . Lemma 11.2 says that $\mathcal{AN}_{\text{ex}}^{n-1}(M_{n,b}, H) = \mathcal{AN}_{\text{ex}}(M_{n,b}, H)$, so the base case $m = n-1$ is provided by Lemma 10.2. Assume, therefore, that $0 \leq m < n-1$ and that $\mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H)$ has the disc local injectivity property up to dimension $d = n-1 - \text{rk}(H)$. To prove the same for $\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H)$, it is enough to verify the conditions of Proposition 5.3 for the inclusion

$$\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H) \hookrightarrow \mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H).$$

The input to Proposition 5.3 is a set \mathcal{B} of “bad simplices” of $\mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H)$, which for us will be the set of simplices σ satisfying the following:

- σ separates $M_{n,b}$ into two components X and Y , and
- σ equals its own separating core, and
- if X is the basepoint-containing component of the complement, then $\text{rk}(\pi_1(Y)) = m+1$.

We now verify each hypothesis of Proposition 5.3 in turn.

Condition (i) says that a simplex σ of $\mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H)$ lies in $\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H)$ if and only if none of its faces are in \mathcal{B} . Since no simplices of \mathcal{B} lie in $\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H)$, if a face of σ

lies in \mathcal{B} then σ does not lie in $\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H)$. Conversely, assume that no face of σ lies in \mathcal{B} . We will prove that σ lies in $\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H)$. If σ does not separate $M_{n,b}$, then it trivially lies in $\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H)$. If σ does separate, then let σ' be its separating core. Let X and Y be the components of the complement of σ' , with X the basepoint-containing component. Since σ is a simplex of $\mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H)$, we have $\text{rk}(\pi_1(Y)) \leq m + 1$. Moreover, since σ' does not lie in \mathcal{B} this must be a strict inequality, i.e., $\text{rk}(\pi_1(Y)) \leq m$. It follows that σ lies in $\mathcal{AN}_{\text{ex}}^m(M_{n,b}, H)$, as desired.

Condition (ii) says that if $\sigma, \sigma' \in \mathcal{B}$ are such that $\sigma \cup \sigma'$ is a simplex of $\mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H)$, then $\sigma \cup \sigma' \in \mathcal{B}$. In fact, since simplices of $\mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H)$ can separate $M_{n,b}$ into at most 2 components, this can only happen if $\sigma = \sigma'$, so there is nothing to prove.

Condition (iii) says that for all $\sigma \in \mathcal{B}$, there is a subcomplex $\widehat{\mathcal{L}}(\sigma, \mathcal{B})$ of $\mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H)$ with $\mathcal{L}(\sigma, \mathcal{B}) \subset \widehat{\mathcal{L}}(\sigma, \mathcal{B})$ that satisfies conditions (iii).a–(iii).c. Here $\mathcal{L}(\sigma, \mathcal{B})$ is the complex defined in Definition 5.1. Fixing some $\sigma \in \mathcal{B}$, before we can define $\widehat{\mathcal{L}}(\sigma, \mathcal{B})$ we must give a concrete description of $\mathcal{L}(\sigma, \mathcal{B})$. Let X and Y be the components of the complement of σ , with X the basepoint-containing component. The group $\pi_1(Y)$ thus has rank $m + 1$. It follows from the definitions that

$$\mathcal{L}(\sigma, \mathcal{B}) = \text{Link}_{\mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H)}(\sigma) = \mathcal{N}(X, H) * \mathcal{N}(Y).$$

We then define

$$\widehat{\mathcal{L}}(\sigma, \mathcal{B}) = \mathcal{N}(X, H) * \mathcal{AN}_{\text{ex}}^m(Y) = \mathcal{N}(X, H) * \mathcal{AN}_{\text{ex}}^m(Y),$$

where the final equality follows from Lemma 11.2. We verify (iii).a–(iii).c for this as follows.

Condition (iii).a says that $\partial\sigma * \widehat{\mathcal{L}}(\sigma, \mathcal{B}) \subset \mathcal{AN}_{\text{ex}}^{m+1}(M_{n,b}, H)$. This follows immediately from the fact that proper faces of σ do not separate $M_{n,b}$.

Condition (iii).b says that all simplices of $\partial\sigma * \widehat{\mathcal{L}}(\sigma, \mathcal{B})$ that are in \mathcal{B} lie in $\partial\sigma$. Even more is true: since proper faces of σ do not separate $M_{n,b}$, no simplices of $\partial\sigma * \widehat{\mathcal{L}}(\sigma, \mathcal{B})$ lie in \mathcal{B} .

Finally, Condition (iii).c says that $\widehat{\mathcal{L}}(\sigma, \mathcal{B})$ has the disc local injectivity property up to dimension

$$d - \dim(\sigma) = (n - 1 - \text{rk}(H)) - \dim(\sigma).$$

Let $X \cong M_{n',b'}$ and $Y \cong M_{n'',b''}$, so by definition $n'' = m + 1$. It follows from Lemma 10.2 that

$$\mathcal{AN}_{\text{ex}}^m(Y) = \mathcal{AN}_{\text{ex}}(Y)$$

has the disc local injectivity property up to dimension $n'' - 1$. Also, Theorem 7.4 says that $\mathcal{N}(X, H)$ is weakly Cohen–Macaulay of dimension $n' - 1 - \text{rk}(H)$, so by Lemma 4.4 it has the disc local injectivity property up to dimension $n' - 2 - \text{rk}(H)$. Applying Lemma 3.8, we see that

$$\widehat{\mathcal{L}}(\sigma, \mathcal{B}) = \mathcal{N}(X, H) * \mathcal{AN}_{\text{ex}}^m(Y)$$

has the disc local injectivity property up to dimension

$$(n' - 2 - \text{rk}(H)) + (n'' - 1) + 2 = (n' + n'') - 1 - \text{rk}(H).$$

We must prove that this is at least $n - 1 - \text{rk}(H) - \dim(\sigma)$, i.e., that

$$n' + n'' \geq n - \dim(\sigma).$$

In fact this is an equality. To see this note that the dual graph $\Gamma(\sigma)$ has two vertices corresponding to X and Y and $\dim(\sigma) + 1$ edges. It follows that $\text{rk}(\pi_1(\Gamma(\sigma))) = \dim(\sigma)$. Applying Lemma 6.12, we see that

$$n = \text{rk}(\pi_1(\Gamma(\sigma))) + n' + n'' = \dim(\sigma) + n' + n''.$$

The lemma follows. \square

12. A PRESENTATION OF THE STEINBERG MODULE FOR $\text{Aut}(F_n)$

In this section, we use our high connectivity results for the augmented nonseparating sphere complex to construct a presentation for the Steinberg module $\text{St}(F_n)$, whose definition we will recall below.

12.1. Generalities about posets. Let \mathbf{A} be a poset. Let \mathbf{A}^{op} be the opposite poset of \mathbf{A} . Recall that the geometric realization $|\mathbf{A}|$ is the simplicial complex whose k -simplices are chains $a_0 < \cdots < a_k$ in \mathbf{A} . The simplicial complexes $|\mathbf{A}|$ and $|\mathbf{A}^{\text{op}}|$ are isomorphic. A key example is as follows:

Definition 12.1. Let X be a simplicial complex. The *poset of simplices* of X , denoted $\mathcal{P}(X)$, is the poset whose elements are simplices of X ordered by inclusion. \square

For a simplicial complex X , the geometric realization $|\mathcal{P}(X)|$ is the first barycentric subdivision of X . We will also need the following two notions:

Definition 12.2. Let \mathbf{A} be a poset and $a \in \mathbf{A}$. The *height* of a , denoted $\text{ht}(a)$, is the maximal $k \geq 0$ such that there exists a chain $a_0 < a_1 < \cdots < a_k = a$. \square

Example 12.3. If X is a simplicial complex and $\sigma \in \mathcal{P}(X)$, then $\text{ht}(\sigma) = \dim(\sigma)$. \square

Definition 12.4. Let $\phi: \mathbf{A} \rightarrow \mathbf{B}$ be a map of posets and $b \in \mathbf{B}$. The *poset fiber* of ϕ over b , denoted $\phi^{\leq b}$, is the subposet of \mathbf{A} consisting of all $a \in \mathbf{A}$ such that $\phi(a) \leq b$. \square

The following is the case $m = 1$ of Church–Putman [9, Proposition 2.3], which builds on work of Quillen [32]. Recall from Definition 4.1 that a simplicial complex is Cohen–Macaulay of dimension d if it is d -dimensional and weakly Cohen–Macaulay of dimension d .

Proposition 12.5 (Church–Putman [9, Proposition 2.3]). *Fix an abelian group R . Let $\phi: \mathbf{A} \rightarrow \mathbf{B}$ be a map of posets. Assume that $|\mathbf{B}|$ is Cohen–Macaulay of dimension $d \geq 0$ and that for all $b \in \mathbf{B}$, the geometric realization of the poset fiber $|\phi^{\leq b}|$ is $\text{ht}(b)$ -connected. Then $\phi_*: \tilde{H}_i(|\mathbf{A}|; R) \rightarrow \tilde{H}_i(|\mathbf{B}|; R)$ is an isomorphism for $i \leq d$.*

12.2. Relating the free factor complex to the sphere complex. Recall from the introduction that the free factor complex $\text{Tits}(F_n)$ is the geometric realization of the following poset:

Definition 12.6. Let $\mathbf{Tits}(F_n)$ be the poset of proper nontrivial free factors of F_n , ordered by inclusion. \square

Hatcher–Vogtmann [19, 20] proved that $\text{Tits}(F_n) = |\mathbf{Tits}(F_n)|$ is homotopy equivalent to a wedge of $(n - 2)$ -dimensional spheres, and by definition

$$\text{St}(F_n) = \tilde{H}_{n-2}(\text{Tits}(F_n); \mathbb{Q}) = \tilde{H}_{n-2}(|\mathbf{Tits}(F_n)|; \mathbb{Q}).$$

We will relate this to the sphere complex $\mathcal{S}(M_{n,1})$ and its subcomplexes. Fix $*$ $\in \partial M_{n,1}$. Recall from the introduction that the action of $\text{Diff}^+(M_{n,1}, \partial M_{n,1})$ on $\mathcal{S}(M_{n,1})$ factors through

$$\text{Aut}(\pi_1(M_{n,1}, *)) \cong \text{Aut}(F_n).$$

We thus get an action of $\text{Aut}(F_n)$ on $\mathcal{S}(M_{n,1})$. The group $\text{Aut}(F_n)$ also acts on $\text{St}(F_n)$, and while relating $\text{St}(F_n)$ to $\mathcal{S}(M_{n,1})$ and its subcomplexes we will keep track of these group actions.

We introduce the following subcomplex of the augmented nonseparating sphere complex $\mathcal{AN}(M_{n,1})$.

Definition 12.7. For $n \geq 0$, let $\mathcal{AN}'(M_{n,1})$ be the subcomplex of $\mathcal{AN}(M_{n,1})$ consisting of simplices σ such that the basepoint-containing component of the complement of σ has a nontrivial fundamental group. \square

Having done this, we can now make the following key definition:

Definition 12.8. For some $n \geq 0$, fix a basepoint $* \in \partial M_{n,1}$. Identify $\pi_1(M_{n,1}, *)$ with F_n . The *complement map* is the map of posets $\Upsilon: \mathcal{P}(\mathcal{AN}'(M_{n,1})) \rightarrow \mathbf{Tits}(F_n)^{\text{op}}$ defined as follows. Consider $\sigma \in \mathcal{P}(\mathcal{AN}'(M_{n,1}))$, and let X be the basepoint-containing component of the complement of σ . Then

$$\Upsilon(\sigma) = \pi_1(X, *) \subset \pi_1(M_{n,1}, *) = F_n. \quad \square$$

Remark 12.9. This definition makes sense since σ is a simplex of $\mathcal{AN}'(M_{n,1})$, so the basepoint-containing component X of the complement of σ is not simply connected and $\Upsilon(\sigma) = \pi_1(X, *)$ is not trivial. The map Υ cannot be defined on the whole complex $\mathcal{AN}(M_{n,1})$. \square

Remark 12.10. The target of the complement map is the opposite poset $\mathbf{Tits}(F_n)^{\text{op}}$ because if σ and σ' are simplices of $\mathcal{AN}'(M_{n,1})$ with $\sigma \subset \sigma'$, then the basepoint-containing component X of the complement of σ contains the basepoint-containing component X' of the complement of σ' , so

$$\Upsilon(\sigma) = \pi_1(X, *) \supset \pi_1(X', *) = \Upsilon(\sigma'). \quad \square.$$

Remark 12.11. The action of $\text{Aut}(F_n)$ on $\mathcal{S}(M_{n,1})$ preserves the subcomplexes $\mathcal{AN}(M_{n,1})$ and $\mathcal{AN}'(M_{n,1})$, and the map $\Upsilon: \mathcal{P}(\mathcal{AN}'(M_{n,1})) \rightarrow \mathbf{Tits}(F_n)^{\text{op}}$ is $\text{Aut}(F_n)$ -equivariant. \square

Identify

$$\tilde{\mathbf{H}}_{n-2}(\mathcal{AN}'(M_{n,1}); \mathbb{Q}) \quad \text{with} \quad \tilde{\mathbf{H}}_{n-2}(|\mathcal{P}(\mathcal{AN}'(M_{n,1}))|; \mathbb{Q})$$

and

$$\text{St}(F_n) = \tilde{\mathbf{H}}_{n-2}(\mathbf{Tits}(F_n); \mathbb{Q}) \quad \text{with} \quad \tilde{\mathbf{H}}_{n-2}(|\mathbf{Tits}(F_n)|; \mathbb{Q}) = \tilde{\mathbf{H}}_{n-2}(|\mathbf{Tits}(F_n)^{\text{op}}|; \mathbb{Q}).$$

We then have the following:

Lemma 12.12. *For some $n \geq 2$, fix a basepoint $* \in \partial M_{n,1}$. Identify $\pi_1(M_{n,1}, *)$ with F_n . Then the complement map $\Upsilon: \mathcal{P}(\mathcal{AN}'(M_{n,1})) \rightarrow \mathbf{Tits}(F_n)^{\text{op}}$ induces an $\text{Aut}(F_n)$ -equivariant isomorphism*

$$\Upsilon_*: \tilde{\mathbf{H}}_{n-2}(\mathcal{AN}'(M_{n,1}); \mathbb{Q}) \xrightarrow{\cong} \text{St}(F_n).$$

Proof. Since Υ is $\text{Aut}(F_n)$ -equivariant (see Remark 12.11), the map Υ_* is $\text{Aut}(F_n)$ -equivariant. To prove that it is an isomorphism, it is enough to show that the complement map

$$\Upsilon: \mathcal{P}(\mathcal{AN}'(M_{n,1})) \rightarrow \mathbf{Tits}(F_n)^{\text{op}}$$

satisfies the conditions of Proposition 12.5 with $d = n - 2$. The first of these hypotheses is that $|\mathbf{Tits}(F_n)^{\text{op}}|$ is Cohen–Macaulay of dimension $(n - 2)$, which was proved¹⁴ by Hatcher–Vogtmann [19, 20, §4]. The second hypothesis is that for all $H \in \mathbf{Tits}(F_n)^{\text{op}}$, the poset fiber $\Upsilon^{\leq H}$ is $\text{ht}(H)$ -connected. This poset fiber is the poset of simplices σ of $\mathcal{AN}'(M_{n,1})$ such that the fundamental group of the basepoint-containing component of the complement contains H . By definition, this is precisely the poset of simplices of $\mathcal{AN}(M_{n,1}, H)$ – note that there is no $'$ in this since $H \neq 1$, so simplices of $\mathcal{AN}(M_{n,1}, H)$ automatically have a non-simply-connected basepoint-containing component of their complement. Theorem 8.3 says that $\mathcal{AN}(M_{n,1}, H)$ and hence its poset of simplices $\mathcal{P}(\mathcal{AN}(M_{n,1}, H))$ is $(n - 1 - \text{rk}(H))$ connected. Since we are regarding H as an element of the opposite poset $\mathbf{Tits}(F_n)^{\text{op}}$, we have $\text{ht}(H) = n - 1 - \text{rk}(H)$, so the lemma follows. \square

¹⁴They actually proved that $|\mathbf{Tits}(F_n)|$ is Cohen–Macaulay of dimension $(n - 2)$, but the geometric realization of a poset is isomorphic to the geometric realization of its opposite poset.

12.3. Relative homology. We now study the topology of $\mathcal{AN}'(M_{n,1})$. The starting point is the following:

Lemma 12.13. *For $n \geq 2$, we have an $\text{Aut}(F_n)$ -equivariant isomorphism*

$$H_{n-1}(\mathcal{AN}(M_{n,1}), \mathcal{AN}'(M_{n,1}); \mathbb{Q}) \cong \tilde{H}_{n-2}(\mathcal{AN}'(M_{n,1}); \mathbb{Q}).$$

Proof. In this proof, all homology has \mathbb{Q} -coefficients, though we remark that \mathbb{Q} can be replaced by any abelian group. The long exact sequence of the pair $(\mathcal{AN}(M_{n,1}), \mathcal{AN}'(M_{n,1}))$ contains the segment

$$\begin{aligned} \tilde{H}_{n-1}(\mathcal{AN}(M_{n,1})) &\rightarrow H_{n-1}(\mathcal{AN}(M_{n,1}), \mathcal{AN}'(M_{n,1})) \\ &\xrightarrow{\partial} \tilde{H}_{n-2}(\mathcal{AN}'(M_{n,1})) \rightarrow \tilde{H}_{n-2}(\mathcal{AN}(M_{n,1})). \end{aligned}$$

Since $n \geq 2$, Theorem 8.3 implies that

$$\tilde{H}_{n-1}(\mathcal{AN}(M_{n,1})) = \tilde{H}_{n-2}(\mathcal{AN}(M_{n,1})) = 0.$$

The lemma follows. \square

To understand this relative homology group, the key result is as follows. Recall that $\mathcal{AN}(M_{n,1})$ is n -dimensional.

Lemma 12.14. *Fix some $n \geq 2$. The following then hold:*

- $\mathcal{AN}'(M_{n,1})$ is $(n-1)$ -dimensional.
- The $(n-2)$ -skeletons of $\mathcal{AN}'(M_{n,1})$ and $\mathcal{AN}(M_{n,1})$ are equal.
- The only $(n-1)$ -simplices of $\mathcal{AN}(M_{n,1})$ that do not lie in $\mathcal{AN}'(M_{n,1})$ are the $(n-1)$ -simplices of $\mathcal{N}(M_{n,1})$.

Proof. Since $\mathcal{AN}'(M_{n,1})$ is a subcomplex of $\mathcal{AN}(M_{n,1})$, it is enough to characterize which simplices of $\mathcal{AN}(M_{n,1})$ lie in $\mathcal{AN}'(M_{n,1})$. Let σ be a k -simplex of $\mathcal{AN}(M_{n,1})$. Write $\sigma = \{S_0, \dots, S_k\}$. Assume first that the union of the S_i does not separate $M_{n,1}$, so σ is a simplex of $\mathcal{N}(M_{n,1})$. We thus have $k \leq n-1$, and the only component of the complement of σ is homeomorphic to $M_{n-(k+1), 1+2(k+1)}$. This is simply-connected if and only if $k = n-1$, so we see that σ lies in $\mathcal{AN}'(M_{n,1})$ if and only if $k \leq n-2$, as claimed.

Assume next that the union of the S_i does separate $M_{n,1}$. Since σ is a simplex of $\mathcal{AN}(M_{n,1})$, there are two components X and Y of the complement of σ . Moreover, letting X be the basepoint-containing component we have $\pi_1(Y) = 1$. The dual graph $\Gamma(\sigma)$ has two vertices and $(k+1)$ edges, so its fundamental group is F_k . Letting m be such that $\pi_1(X, *) \cong F_m$, Lemma 6.12 says that

$$n = \pi_1(\Gamma(\sigma)) + m + 0 = k + m.$$

The simplex σ lies in $\mathcal{AN}'(M_{n,1})$ if and only if $m \geq 1$, and the above equality shows that this holds if and only if $k \leq n-1$. The lemma follows. \square

This has the following consequence:

Lemma 12.15. *For $n \geq 2$ and $k \geq 0$, we have $\text{Aut}(F_n)$ -equivariant isomorphisms*

$$C_k(\mathcal{AN}(M_{n,1}), \mathcal{AN}'(M_{n,1}); \mathbb{Q}) = \begin{cases} C_{n-1}(\mathcal{N}(M_{n,1}); \mathbb{Q}) & \text{if } k = n-1, \\ C_n(\mathcal{AN}(M_{n,1}); \mathbb{Q}) & \text{if } k = n, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Since $\mathcal{AN}(M_{n,1})$ is n -dimensional, we have

$$C_k(\mathcal{AN}(M_{n,1}), \mathcal{AN}'(M_{n,1}); \mathbb{Q}) = 0 \quad \text{if } k \geq n+1.$$

The remainder of the lemma follows from Lemma 12.14. \square

Combining these results we obtain the following presentation of $\text{St}(F_n)$.

Theorem 12.16. *For $n \geq 2$, we have an exact sequence of $\text{Aut}(F_n)$ -modules*

$$C_n(\mathcal{AN}(M_{n,1}); \mathbb{Q}) \rightarrow C_{n-1}(\mathcal{N}(M_{n,1}); \mathbb{Q}) \rightarrow \text{St}(F_n) \rightarrow 0.$$

Proof. In this proof, all homology has \mathbb{Q} -coefficients. Combining Lemmas 12.12 and 12.13, we get an $\text{Aut}(F_n)$ -equivariant isomorphism

$$H_{n-1}(\mathcal{AN}(M_{n,1}), \mathcal{AN}'(M_{n,1})) \cong \text{St}(F_n).$$

To compute this relative homology group, consider the relevant terms in the simplicial chain complex:

$$\begin{aligned} C_n(\mathcal{AN}(M_{n,1}), \mathcal{AN}'(M_{n,1})) &\longrightarrow C_{n-1}(\mathcal{AN}(M_{n,1}), \mathcal{AN}'(M_{n,1})) \\ &\longrightarrow C_{n-2}(\mathcal{AN}(M_{n,1}), \mathcal{AN}'(M_{n,1})). \end{aligned}$$

Lemma 12.15 shows that these terms are equal to

$$C_n(\mathcal{AN}(M_{n,1})) \rightarrow C_{n-1}(\mathcal{N}(M_{n,1})) \rightarrow 0.$$

The theorem follows. \square

12.4. Flatness. The following lemma shows that this presentation can be used to compute the low-dimensional rational homology of $\text{Aut}(F_n)$ with coefficients in $\text{St}(F_n)$:

Lemma 12.17. *Fix some $n \geq 2$. Then $C_n(\mathcal{AN}(M_{n,1}); \mathbb{Q})$ and $C_{n-1}(\mathcal{N}(M_{n,1}); \mathbb{Q})$ are flat $\mathbb{Z}[\text{Aut}(F_n)]$ -modules.*

Proof. It is enough to prove that if σ is either an n -simplex of $\mathcal{AN}(M_{n,1})$ or an $(n-1)$ -simplex of $\mathcal{N}(M_{n,1})$, then the $\text{Aut}(F_n)$ -stabilizer of σ is finite. See, e.g., [7, Lemma 3.2].¹⁵ Since the action of the mapping class group $\text{Mod}(M_{n,1})$ on these simplicial complexes factors through $\text{Aut}(F_n)$, this will follow if we can prove that the $\text{Mod}(M_{n,1})$ -stabilizer G_σ of σ is finite.

Since every n -simplex of $\mathcal{AN}(M_{n,1})$ contains an $(n-1)$ -simplex of $\mathcal{N}(M_{n,1})$, it is enough to deal with the case where σ is an $(n-1)$ -simplex of $\mathcal{N}(M_{n,1})$. Write $\sigma = \{S_0, \dots, S_{n-1}\}$. Define G'_σ to be the subgroup of $\text{Mod}(M_{n,1})$ consisting of mapping classes f satisfying the following condition:

- For each $0 \leq i \leq n-1$, the mapping class f can be realized by a diffeomorphism of $M_{n,1}$ that fixes a tubular neighborhood of S_i .

The group G'_σ is a finite-index subgroup¹⁶ of G_σ , so it is enough to prove that G'_σ is finite. Cutting $M_{n,1}$ open along the S_i yields $M_{0,2n+1}$. Reversing this, there is a homomorphism $\psi: \text{Mod}(M_{0,2n+1}) \rightarrow \text{Mod}(M_{n,1})$ induced by the map taking diffeomorphisms of $M_{0,2n+1}$ that are the identity on $\partial M_{0,2n+1}$ and gluing boundary components together in pairs. The image of ψ is G'_σ , so we are reduced to proving that $\text{Mod}(M_{0,2n+1})$ is a finite group.

In fact, Laudenbach ([26, 27]; see [25] for an alternate approach using a ‘‘Birman exact sequence’’) proved that $\text{Mod}(M_{0,2n+1}) \cong (\mathbb{Z}/2)^{2n}$, generated by sphere twists about $2n$ of the boundary components. There is no need to take the sphere twist about the remaining boundary component since the product of sphere twists about all the boundary components of $M_{0,2n+1}$ is trivial; see [23, p. 214–215] for a simple proof of this. The lemma follows. \square

¹⁵The key ingredients in the proof are Shapiro’s Lemma and the fact that finite groups have trivial homology in positive degrees with respect to coefficients that are vector spaces over fields of characteristic 0.

¹⁶This uses the fact that homotopy implies ambient isotopy for collections of pairwise non-homotopic disjoint spheres in $M_{n,1}$, which was proved by Laudenbach [26, 27]. The differences between G_σ and G'_σ are that elements of G_σ can permute the S_i and also ‘‘flip them around’’, i.e., reverse the directions of arcs transverse to them.

13. PROOF OF MAIN THEOREM

We close the paper by finally proving Theorem A which asserts that $H_k(\text{Aut}(F_n); \text{St}(F_n)) = 0$ for $k \in \{0, 1\}$ and $n \geq 2$.

Proof of Theorem A. Let $n \geq 2$, and consider the presentation of $\text{St}(F_n)$ from Theorem 12.16:

$$C_n(\mathcal{AN}(M_{n,1}); \mathbb{Q}) \rightarrow C_{n-1}(\mathcal{N}(M_{n,1}); \mathbb{Q}) \rightarrow \text{St}(F_n) \rightarrow 0.$$

By Lemma 12.17, both $C_n(\mathcal{AN}(M_{n,1}); \mathbb{Q})$ and $C_{n-1}(\mathcal{N}(M_{n,1}); \mathbb{Q})$ are flat $\mathbb{Z}[\text{Aut}(F_n)]$ -modules. Complete this to a flat resolution of $\text{St}(F_n)$:

$$\cdots \rightarrow \mathcal{F}_2 \rightarrow C_n(\mathcal{AN}(M_{n,1}); \mathbb{Q}) \rightarrow C_{n-1}(\mathcal{N}(M_{n,1}); \mathbb{Q}) \rightarrow \text{St}(F_n) \rightarrow 0.$$

For a group G and a $\mathbb{Z}[G]$ -module M , write M_G for the G -coinvariants of M . We can use our flat resolution to compute $H_\bullet(\text{Aut}(F_n); \text{St}(F_n))$, which is the homology of the chain complex

$$\cdots \rightarrow (\mathcal{F}_2)_{\text{Aut}(F_n)} \rightarrow (C_n(\mathcal{AN}(M_{n,1}); \mathbb{Q}))_{\text{Aut}(F_n)} \rightarrow (C_{n-1}(\mathcal{N}(M_{n,1}); \mathbb{Q}))_{\text{Aut}(F_n)} \rightarrow 0.$$

To prove the theorem, it suffices to prove that

$$(C_n(\mathcal{AN}(M_{n,1}); \mathbb{Q}))_{\text{Aut}(F_n)} = (C_{n-1}(\mathcal{N}(M_{n,1}); \mathbb{Q}))_{\text{Aut}(F_n)} = 0.$$

The vector space $C_n(\mathcal{AN}(M_{n,1}); \mathbb{Q})$ (resp. $C_{n-1}(\mathcal{N}(M_{n,1}); \mathbb{Q})$) is spanned by oriented n -simplices of $\mathcal{AN}(M_{n,1})$ (resp. $(n-1)$ -simplices of $\mathcal{N}(M_{n,1})$). For such an oriented simplex σ , we will prove below in Lemma 13.1 that there exists some $f \in \text{Aut}(F_n)$ such that $f(\sigma)$ equals σ , but with the opposite orientation. This will imply that the images of σ and $-\sigma$ in the $\text{Aut}(F_n)$ -coinvariants are equal, and thus that the image of 2σ in the $\text{Aut}(F_n)$ -coinvariants is 0. Since we are working over \mathbb{Q} , this implies that the image of σ in the $\text{Aut}(F_n)$ -coinvariants is 0, so these coinvariants are 0, as desired. \square

Lemma 13.1. *Fix some $n \geq 2$. Let σ be either an oriented n -simplex of $\mathcal{AN}(M_{n,1})$ or an oriented $(n-1)$ -simplex of $\mathcal{N}(M_{n,1})$. Then there exists some $f \in \text{Aut}(F_n)$ such that $f(\sigma)$ equals σ , but with the opposite orientation.*

Proof. For $k \geq 2$, we will prove that this holds more generally for all oriented $(k-1)$ -simplices σ of the complex $\mathcal{AN}(M_{n,1})$. The action of the mapping class group $\text{Mod}(M_{n,1})$ on $\mathcal{AN}(M_{n,1})$ factors through $\text{Aut}(F_n)$, so it is enough to find some $\phi \in \text{Mod}(M_{n,1})$ such that $\phi(\sigma)$ equals σ , but with the opposite orientation. Let $\sigma = \{S_1, \dots, S_k\}$, ordered in a way compatible with the orientation of σ .

There are either one or two components in the complement of σ . Assume first that there is one component in the complement, which we call X . We then have that $X \cong M_{n-k, 2k+1}$. Enumerate the components of ∂X as $\{\partial, \beta_1, \beta'_1, \dots, \beta_k, \beta'_k\}$, where ∂ is the component of $\partial M_{n,1}$ and β_i and β'_i are the two components that glue together to form S_i . We can then find an orientation-preserving diffeomorphism $\psi: X \rightarrow X$ such that

$$\psi(\beta_1) = \beta_2 \quad \text{and} \quad \psi(\beta_2) = \beta_1 \quad \text{and} \quad \psi(\beta'_1) = \beta'_2 \quad \text{and} \quad \psi(\beta'_2) = \beta'_1$$

and such that

$$\psi|_{\partial} = \psi|_{\beta_i} = \psi|_{\beta'_i} = \text{id} \quad \text{for } 3 \leq i \leq k.$$

Isotoping ψ if necessary to make sure its behavior on $\beta_1 \cup \beta'_1$ matches up with its behavior on $\beta_2 \cup \beta'_2$, we can glue the boundary components of X back together and get from ψ an induced diffeomorphism $\phi: M_{n,1} \rightarrow M_{n,1}$ that swaps S_1 and S_2 while fixing S_i for $3 \leq i \leq k$. It follows that ϕ takes σ to σ but with the opposite orientation, as desired.

Assume now that there are two components X and Y in the complement of σ , with X the basepoint-containing component. For some $r \geq 3$ and $h \geq 0$ with $r+h=k$, we have

$$Y \cong M_{0,r} \quad \text{and} \quad X \cong M_{n-k+1, 1+r+2h}.$$

The condition $r \geq 3$ holds since all the S_i are nonseparating and pairwise non-isotopic. Enumerate the components of ∂X and ∂Y as

$$\{\partial, \beta_1, \dots, \beta_r, \delta_1, \delta'_1, \dots, \delta_h, \delta'_h\} \quad \text{and} \quad \{\beta'_1, \dots, \beta'_r\},$$

respectively, where the enumeration is chosen such that the following hold when X and Y are glued together to form $M_{n,1}$:

- ∂ is the component of $\partial M_{n,1}$.
- $\beta_i \subset \partial X$ and $\beta'_i \subset \partial Y$ are glued together to form a sphere in $\{S_1, \dots, S_k\}$.
- $\delta_j \subset \partial X$ and $\delta'_j \subset \partial X$ are glued together to form a sphere in $\{S_1, \dots, S_k\}$.

Let S_a (resp. S_b) be the sphere in $\{S_1, \dots, S_k\}$ that is formed when β_1 is glued to β'_1 (resp. β_2 is glued to β'_2). We can then find orientation-preserving diffeomorphisms $\psi_1: X \rightarrow X$ and $\psi_2: Y \rightarrow Y$ such that

$$\psi_1(\beta_1) = \beta_2 \quad \text{and} \quad \psi_1(\beta_2) = \beta_1 \quad \text{and} \quad \psi_2(\beta'_1) = \beta'_2 \quad \text{and} \quad \psi_2(\beta'_2) = \beta'_1$$

and such that

$$\psi_1|_{\partial} = \psi_1|_{\beta_i} = \psi_1|_{\delta_j} = \psi_1|_{\delta'_j} = \psi_2|_{\beta'_i} = \text{id} \quad \text{for } 3 \leq i \leq r \text{ and } 1 \leq j \leq h.$$

Isotoping ψ_1 and ψ_2 if necessary to make sure their behavior on the boundaries match up, we can glue the boundary components of X and Y back up and get from ψ_1 and ψ_2 an induced diffeomorphism $\phi: M_{n,1} \rightarrow M_{n,1}$ that swaps S_a and S_b while fixing S_i for $1 \leq i \leq k$ with $i \neq a, b$. It follows that ϕ takes σ to σ but with the opposite orientation, as desired. \square

REFERENCES

- [1] M. Bestvina and M. Feighn, The topology at infinity of $\text{Out}(F_n)$, *Invent. Math.* 140 (2000), no. 3, 651–692.
- [2] A. Borel and J.-P. Serre, Corners and arithmetic groups, *Comment. Math. Helv.* 48 (1973), 436–491.
- [3] T. Brendle, N. Broaddus, and A. Putman, The high-dimensional cohomology of the moduli space of curves with level structures II: punctures and boundary, to appear in *Israel J. Math.* [arXiv:2003.10913](https://arxiv.org/abs/2003.10913)
- [4] T. Brendle, N. Broaddus, and A. Putman, The mapping class group of connect sums of $S^2 \times S^1$, to appear in *Trans. Amer. Math. Soc.* [arXiv:2012.01529](https://arxiv.org/abs/2012.01529)
- [5] B. Brück and R. Gupta, Homotopy type of the complex of free factors of a free group, *Proc. Lond. Math. Soc.* (3) 121 (2020), no. 6, 1737–1765. [arXiv:1810.09380](https://arxiv.org/abs/1810.09380)
- [6] V. A. Bykovskii, Generating elements of the annihilating ideal for modular symbols, *Funct. Anal. Appl.* 37 (2003), no. 4, 263–272; translated from *Funktional. Anal. i Prilozhen.* 37 (2003), no. 4, 27–38, 95.
- [7] T. Church, B. Farb, and A. Putman, The rational cohomology of the mapping class group vanishes in its virtual cohomological dimension, *Int. Math. Res. Not. IMRN* 2012, no. 21, 5025–5030. [arXiv:1108.0622](https://arxiv.org/abs/1108.0622)
- [8] T. Church, B. Farb, and A. Putman, Integrality in the Steinberg module and the top-dimensional cohomology of $\text{SL}_n \mathcal{O}_K$, *Amer. J. Math.* 141 (2019), no. 5, 1375–1419. [arXiv:1501.01307](https://arxiv.org/abs/1501.01307)
- [9] T. Church and A. Putman, The codimension-one cohomology of $\text{SL}_n \mathbb{Z}$, *Geom. Topol.* 21 (2017), no. 2, 999–1032. [arXiv:1507.06306](https://arxiv.org/abs/1507.06306)
- [10] I. Sadofschi Costa, The complex of partial bases of a free group, *Bull. Lond. Math. Soc.* 52 (2020), no. 1, 109–120. [arXiv:1711.09954](https://arxiv.org/abs/1711.09954)
- [11] P. Deligne, Extensions centrales non résiduellement finies de groupes arithmétiques, *C. R. Acad. Sci. Paris Sér. A-B* 287 (1978), no. 4, A203–A208.
- [12] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, 49, Princeton University Press, Princeton, NJ, 2012.
- [13] N. J. Fullarton and A. Putman, The high-dimensional cohomology of the moduli space of curves with level structures, *J. Eur. Math. Soc. (JEMS)* 22 (2020), no. 4, 1261–1287. [arXiv:1610.03768](https://arxiv.org/abs/1610.03768)
- [14] S. Galatius and O. Randal-Williams, Homological stability for moduli spaces of high dimensional manifolds. I, *J. Amer. Math. Soc.* 31 (2018), no. 1, 215–264. [arXiv:1403.2334](https://arxiv.org/abs/1403.2334)
- [15] F. Gerlits, Invariants in Chain Complexes of Graphs, 2002 PhD thesis.
- [16] F. Gerlits, computer code at <http://pi.math.cornell.edu/m/Research/Dissertations/Gerlits/code>.
- [17] J. L. Harer, The virtual cohomological dimension of the mapping class group of an orientable surface, *Invent. Math.* 84 (1986), no. 1, 157–176.

- [18] A. Hatcher, Homological stability for automorphism groups of free groups, *Comment. Math. Helv.* 70 (1995), no. 1, 39–62.
- [19] A. Hatcher and K. Vogtmann, The complex of free factors of a free group, *Quart. J. Math. Oxford Ser. (2)* 49 (1998), no. 196, 459–468.
- [20] A. Hatcher and K. Vogtmann, The complex of free factors of a free group, 2022 revised version, [arXiv:2203.15602](https://arxiv.org/abs/2203.15602)
- [21] A. Hatcher and K. Vogtmann, Tethers and homology stability for surfaces, *Algebr. Geom. Topol.* 17 (2017), no. 3, 1871–1916. [arXiv:1508.04334](https://arxiv.org/abs/1508.04334)
- [22] A. Hatcher and N. Wahl, Stabilization for the automorphisms of free groups with boundaries, *Geom. Topol.* 9 (2005), 1295–1336. [arXiv:math/0406277](https://arxiv.org/abs/math/0406277)
- [23] A. Hatcher and N. Wahl, Stabilization for mapping class groups of 3-manifolds, *Duke Math. J.* 155 (2010), no. 2, 205–269. [arXiv:0709.2173](https://arxiv.org/abs/0709.2173)
- [24] J. Hempel, *3-Manifolds*, Princeton University Press, Princeton, NJ, 1976.
- [25] J. Landgraf, Cutting and pasting in the Torelli subgroup of $\text{Out}(F_n)$, preprint 2021. [arXiv:2106.14147](https://arxiv.org/abs/2106.14147)
- [26] F. Laudenbach, Sur les 2-sphères d’une variété de dimension 3, *Ann. of Math. (2)* 97 (1973), 57–81.
- [27] F. Laudenbach, *Topologie de la dimension trois: homotopie et isotopie*, Astérisque, No. 12, Société Mathématique de France, Paris, 1974.
- [28] R. Lee and R. H. Szczarba, On the homology and cohomology of congruence subgroups, *Invent. Math.* 33 (1976), no. 1, 15–53.
- [29] J. Miller, R. Nagpal, and P. Patzt, Stability in the high-dimensional cohomology of congruence subgroups, *Compos. Math.* 156 (2020), no. 4, 822–861. [arXiv:1806.11131](https://arxiv.org/abs/1806.11131)
- [30] J. Miller, P. Patzt, and A. Putman, On the top-dimensional cohomology groups of congruence subgroups of $\text{SL}(n, \mathbb{Z})$, *Geom. Topol.* 25 (2021), no. 2, 999–1058. [arXiv:1909.02661](https://arxiv.org/abs/1909.02661)
- [31] J. Milnor, A unique decomposition theorem for 3-manifolds, *Amer. J. Math.* 84 (1962), 1–7.
- [32] D. Quillen, Homotopy properties of the poset of nontrivial p -subgroups of a group, *Adv. Math.* 28 (1978), no. 2, 101–128.
- [33] D. Rolfsen, *Knots and links*, Mathematics Lecture Series, No. 7, Publish or Perish, Inc., Berkeley, CA, 1976.
- [34] L. Solomon, The Steinberg character of a finite group with BN -pair, in *Theory of Finite Groups (Symposium, Harvard Univ., Cambridge, Mass., 1968)*, 213–221, Benjamin, New York.
- [35] E. H. Spanier, *Algebraic topology*, McGraw-Hill Book Co., New York, 1966.
- [36] E. C. Zeeman, Relative simplicial approximation, *Proc. Cambridge Philos. Soc.* 60 (1964), 39–43.

DEPT OF MATHEMATICS; PURDUE UNIVERSITY; 150 N. UNIVERSITY STREET; WEST LAFAYETTE, IN 47907

Email address: himesz@purdue.edu

DEPT OF MATHEMATICS; PURDUE UNIVERSITY; 150 N. UNIVERSITY STREET; WEST LAFAYETTE, IN 47907

Email address: jeremykmiller@purdue.edu

DEPT OF MATHEMATICS; PURDUE UNIVERSITY; 150 N. UNIVERSITY STREET; WEST LAFAYETTE, IN 47907

Email address: snariman@purdue.edu

DEPT OF MATHEMATICS; UNIVERSITY OF NOTRE DAME; 255 HURLEY HALL; NOTRE DAME, IN 46556

Email address: andyp@nd.edu