AGT

Algebraic & Geometric Topology

Volume 25 (2025)

Cutting and pasting in the Torelli subgroup of $\mathrm{Out}(F_n)$

JACOB LANDGRAF

msp

Algebraic & Geometric Topology **25**:1 (2025) 1–38
DOI: 10.2140/agt.2025.25.1
Published: 5 March 2025

msp

Cutting and pasting in the Torelli subgroup of Out(F_n)

JACOB LANDGRAF

Using ideas from 3-manifolds, Hatcher–Wahl defined a notion of automorphism groups of free groups with boundary. We study their Torelli subgroups, adapting ideas introduced by Putman for surface mapping class groups. Our main results show that these groups are finitely generated, and also that they satisfy an appropriate version of the Birman exact sequence.

57K20, 57K30, 57M07

1 Introduction

Let $F_n = \langle x_1, \dots, x_n \rangle$ be the free group on n letters, and let $\mathrm{Out}(F_n)$ be the group of outer automorphisms of F_n. In many ways, $\mathrm{Out}(F_n)$ behaves very similarly to $\mathrm{Mod}(\Sigma_{g,b})$, the mapping class group of the surface $\Sigma_{g,b}$ of genus g with b boundary components. For an overview of some of these similarities, see [7] by Bridson and Vogtmann.

One such connection is that they both contain a Torelli subgroup. In the mapping class group, the Torelli subgroup $\mathcal{I}(\Sigma_{g,b}) \subset \mathrm{Mod}(\Sigma_{g,b})$ is defined to be the kernel of the action on $H_1(\Sigma_{n,b}; \mathbb{Z})$ for $b = 0, 1$. In $\mathrm{Out}(F_n)$, we define a similar subgroup,[1] denoted IO_n, as the kernel of the action of $\mathrm{Out}(F_n)$ on $H_1(F_n; \mathbb{Z}) = \mathbb{Z}^n$.

On surfaces with multiple boundary components, there are many possible definitions one might use to define a Torelli subgroup of $\mathrm{Mod}(\Sigma_{g,b})$. Putman [22] defines a Torelli subgroup $\mathcal{I}(\Sigma_{g,b}, P)$ for $b > 1$ requiring the additional data of a partition P of the boundary components. The goal of the current paper is to mirror Putman's procedure to define an "IO_n with boundary".

Let $M_{n,b} = \#_n(S^1 \times S^2) \setminus (b \text{ open 3-disks})$. For simplicity, we will write M_n if $b = 0$. A key property of $M_{n,b}$ is that it has fundamental group F_n. Fix such an identification. The *mapping class group* $\mathrm{Mod}(M_{n,b})$ is the group of orientation-preserving diffeomorphisms of $M_{n,b}$ fixing the boundary pointwise modulo isotopies fixing the boundary pointwise. Letting $\mathrm{Diff}^+(M_{n,b}, \partial M_{n,b})$ be the topological group of orientation-preserving diffeomorphisms fixing the boundary pointwise, we can also write $\mathrm{Mod}(M_{n,b}) = \pi_0(\mathrm{Diff}^+(M_{n,b}, \partial M_{n,b}))$. By a theorem of Laudenbach [19], there is an exact sequence

$$1 \to (\mathbb{Z}/2)^n \to \mathrm{Mod}(M_n) \to \mathrm{Out}(F_n) \to 1, \tag{1}$$

where the map $\mathrm{Mod}(M_n) \to \mathrm{Out}(F_n)$ is given by the action (up to conjugation) on $\pi_1(M_n)$, and the

[1] It is also common to see this group denoted by IA_n, but we wish to reserve this notation for the analogous subgroup of $\mathrm{Aut}(F_n)$.

$(\mathbb{Z}/2)^n$ is generated by sphere twists about n disjointly embedded 2-spheres (see Section 2 for the definition and relevant properties of sphere twists). Recent work of Brendle, Broaddus, and Putman [6] shows that this sequence actually splits as a semidirect product. This exact sequence implies that, modulo a finite group, $\mathrm{Out}(F_n)$ acts on M_n up to isotopy. Therefore, M_n plays almost the same role for $\mathrm{Out}(F_n)$ that $\Sigma_{g,b}$ plays for $\mathrm{Mod}(\Sigma_{g,b})$.

Adding boundary components From Laudenbach's sequence (1), we see that

$$\mathrm{Out}(F_n) \cong \mathrm{Mod}(M_n)/\mathrm{STwist}(M_n),$$

where $\mathrm{STwist}(M_n) \cong (\mathbb{Z}/2)^n$ is the subgroup of $\mathrm{Mod}(M_n)$ generated by sphere twists. Now that we have related $\mathrm{Out}(F_n)$ to a geometrically defined group, we can start introducing boundary components. Extending the relationship given by Laudenbach's sequence, we define "$\mathrm{Out}(F_n)$ with boundary" as

$$\mathrm{Out}(F_{n,b}) := \mathrm{Mod}(M_{n,b})/\mathrm{STwist}(M_{n,b}).$$

When $b = 1$, Laudenbach [19] also shows that $\mathrm{Out}(F_{n,1}) \cong \mathrm{Aut}(F_n)$. Hatcher and Wahl [14] introduced a more general version of $\mathrm{Out}(F_{n,b})$, which they denoted by $A^s_{n,k}$. The original definition of $A^s_{n,k}$ has to do with classes of self-homotopy equivalences of a certain graph. However, in [14] the authors give an equivalent definition, which says that $A^s_{n,k}$ is the mapping class group of M_n with s spherical and k toroidal boundary components, modulo sphere twists. With this definition, we see that $\mathrm{Out}(F_{n,b}) = A^b_{n,0}$. Similar groups have been examined in the work of Jensen and Wahl [16] and Wahl [26]. Their versions, however, involve only toroidal boundary components, and thus are distinct from $\mathrm{Out}(F_{n,b})$.

Torelli subgroups An important feature of sphere twists (discussed in Section 2) is that they act trivially on homotopy classes of embedded loops, and thus act trivially on $H_1(M_n)$. Therefore, the action of $\mathrm{Mod}(M_{n,b})$ on $H_1(M_{n,b})$ induces an action of $\mathrm{Out}(F_{n,b})$ on $H_1(M_{n,b})$. We can then define the Torelli subgroup $\mathrm{IO}_{n,b} \subset \mathrm{Out}(F_{n,b})$ to be the kernel of this action. However, this definition does not capture all homological information when $b > 1$, especially when $M_{n,b}$ is being embedded in $M_{m,c}$. To see why, consider the scenario depicted in Figure 1, in which $M_{2,2}$ has been embedded into M_4. This embedding induces a homomorphism $\iota_M\colon \mathrm{Mod}(M_{2,2}) \to \mathrm{Mod}(M_4)$ obtained by extending by the identity. This map sends sphere twists to sphere twists, and so we get an induced map $\iota_*\colon \mathrm{Out}(F_{2,2}) \to \mathrm{Out}(F_4)$. However, this does *not* restrict to a map $\mathrm{IO}_{2,2} \to \mathrm{IO}_4$ under this definition of $\mathrm{IO}_{n,b}$ since elements of $\mathrm{IO}_{2,2}$ are not required to fix the homology class of the subarc of α lying inside $M_{2,2}$. To address this issue, we will use a slightly modified homology group.

Definition Fix a partition P of the boundary components of $M_{n,b}$.

(a) Two boundary components ∂_1, ∂_2 of $M_{n,b}$ are *P-adjacent* if there is some $p \in P$ such that $\{\partial_1, \partial_2\} \subset p$.

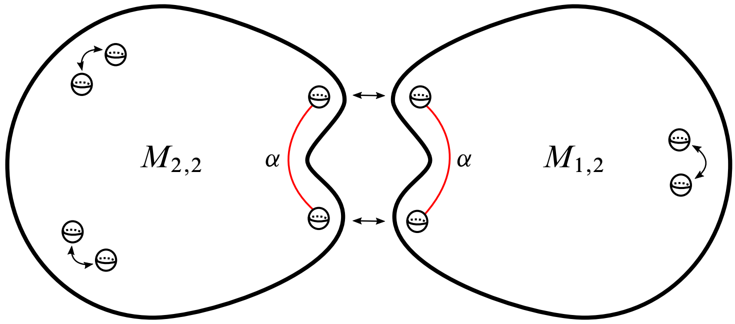

Figure 1: A copy of $M_{2,2}$ and $M_{1,2}$ glued together to obtain M_4. We realize $M_{2,2}$ as a 3-sphere with the six indicated open balls removed, then the boundaries of these removed balls are identified according to the arrows (and similarly for $M_{1,2}$). The class $[\alpha]$ need not be fixed by elements of $\mathrm{IO}_{2,2}$ with the naive definition.

(b) Let $H_1^P(M_{n,b})$ be the subgroup of $H_1(M_{n,b}, \partial M_{n,b})$ spanned by

$$\big\{[h] \in H_1(M_{n,b}, \partial M_{n,b}) \mid \text{ either } h \text{ is a simple closed curve or } h \text{ is a properly embedded arc with endpoints in distinct } P\text{-adjacent boundary components}\big\}.$$

(c) There is a natural action of $\mathrm{Out}(F_{n,b})$ on $H_1^P(M_{n,b})$, and we define the Torelli subgroup $\mathrm{IO}_{n,b}^P \subset \mathrm{Out}(F_{n,b})$ to be the kernel of this action.

Returning to Figure 1, let P be the trivial partition of the boundary components of $M_{2,2}$ with a single P-adjacency class. With this choice of partition, we see that $[\alpha \cap M_{2,2}] \in H_1^P(M_{2,2})$. If $f \in \mathrm{IO}_{2,2}^P$, then it follows that $\iota_*(f) \in \mathrm{Out}(F_4)$ preserves the homology class of α. Therefore, $\iota_*(f) \in \mathrm{IO}_4$, and so ι_* restricts to a map $\mathrm{IO}_{2,2}^P \to \mathrm{IO}_4$.

Restriction As we discussed in the last paragraph, given an embedding $\iota\colon M_{n,b} \hookrightarrow M_m$, we can extend by the identity to get a map $\iota_*\colon \mathrm{Out}(F_{n,b}) \to \mathrm{Out}(F_m)$. In general, ι_* may not be injective. However, it is injective if no connected component of $M_m \setminus M_{n,b}$ is diffeomorphic to D^3 (see Appendix A). Moreover, such an embedding induces a natural partition of the boundary components of $M_{n,b}$ as follows.

Definition Fix an embedding $\iota\colon M_{n,b} \hookrightarrow M_m$. Let N be a connected component of $M_m \setminus \mathrm{int}(M_{n,b})$, and let p_N be the set of boundary components of $M_{n,b}$ shared with N. Then the partition P of the boundary components of $M_{n,b}$ induced by ι is defined to be

$$P = \{p_N \mid N \text{ a connected component of } M_m \setminus M_{n,b}\}.$$

With this definition, one might guess that $\iota_*^{-1}(\mathrm{IO}_n) = \mathrm{IO}_{n,b}^P$. This turns out to be the case, and this is our first main theorem, which we prove in Section 3.

Theorem A (restriction theorem) *Let* $\iota\colon M_{n,b} \hookrightarrow M_m$ *be an embedding,* $\iota_*\colon \mathrm{Out}(F_{n,b}) \to \mathrm{Out}(F_m)$ *the induced map, and* P *the induced partition of the boundary components of* $M_{n,b}$*. Then* $\mathrm{IO}_{n,b}^P = \iota_*^{-1}(\mathrm{IO}_m)$.

Birman exact sequence From here, we move on to exploring the parallels between these Torelli subgroups and those of mapping class groups. There is a well-known relationship between the mapping class groups of surfaces with a different number of boundary components called the Birman exact sequence (see [12]):

$$1 \to \pi_1(\mathrm{UT}(\Sigma_{n,b-1})) \to \mathrm{Mod}(\Sigma_{g,b}) \to \mathrm{Mod}(\Sigma_{g,b-1}) \to 1.$$

Here, $\mathrm{UT}(\Sigma_{n,b-1})$ is the unit tangent bundle of $\Sigma_{n,b-1}$, the map $\pi_1(\mathrm{UT}(\Sigma_{n,b-1})) \to \mathrm{Mod}(\Sigma_{g,b})$ is given by pushing a boundary component around a loop, and the map $\mathrm{Mod}(\Sigma_{g,b}) \to \mathrm{Mod}(\Sigma_{g,b-1})$ is given by attaching a disk onto this boundary component. In Section 4, we will prove versions of the Birman exact sequence for $\mathrm{Mod}(M_{n,b})$ and $\mathrm{Out}(F_{n,b})$, all culminating in the following sequence for $\mathrm{IO}^P_{n,b}$.

Theorem B (Birman exact sequence) *Fix $n, b > 0$ such that $(n, b) \neq (1, 1)$, and let $M_{n,b} \hookrightarrow M_{n,b-1}$ be an embedding obtained by gluing a ball to the boundary component ∂. Fix $x \in M_{n,b-1} \setminus M_{n,b}$. Let P be a partition of the boundary components of $M_{n,b}$, let P' be the induced partition of the boundary components of $M_{n,b-1}$, and let $p \in P$ be the set containing ∂. We then have an exact sequence*

$$1 \to L \to \mathrm{IO}^P_{n,b} \xrightarrow{\iota_*} \mathrm{IO}^{P'}_{n,b-1} \to 1,$$

where L is equal to

(a) $\pi_1(M_{n,b-1}, x) \cong F_n$ *if* $p = \{\partial\}$,

(b) $[\pi_1(M_{n,b-1}, x), \pi_1(M_{n,b-1}, x)] \cong [F_n, F_n]$ *if* $p \neq \{\partial\}$.

Moreover, this sequence splits if $b \geq 2$.

Remark This theorem may seem superficially similar to results proven by Day and Putman [9; 11]. However, we consider a very different notion of "automorphisms with boundary", and so these results are unrelated.

Finite generation Once we have established this version of the Birman exact sequence, in Section 5, we will define a generating set for $\mathrm{IO}^P_{n,b}$. This generating set will be inspired by the generating set for IO_n found by Magnus [21] in 1935.

Theorem 1.1 (Magnus) *Let $F_n = \langle x_1, \dots, x_n \rangle$. The group IO_n is generated by the $\mathrm{Out}(F_n)$-classes of the automorphisms*

$$M_{ij} : x_i \mapsto x_j x_i x_j^{-1}, \quad M_{ijk} : x_i \mapsto x_i [x_j, x_k],$$

for all distinct $i, j, k \in \{1, \dots, n\}$ with $j < k$. Here, the automorphisms are understood to fix x_ℓ for $\ell \neq i$.

Throughout this paper, we will use the convention $[a, b] = aba^{-1}b^{-1}$. Since we defined $\mathrm{IO}^P_{n,b}$ to be a subgroup of $\mathrm{Mod}(M_{n,b})/\mathrm{STwist}(M_{n,b})$, our generators will be defined geometrically rather than algebraically. However, in the case of $b = 0$, they will reduce directly to Magnus's generators. In Section 6, we will show that these elements do indeed generate $\mathrm{IO}^P_{n,b}$.

Theorem C *The group* $\mathrm{IO}_{n,b}^P$ *is finitely generated for* $n \geq 1,\ b \geq 0$.

This is rather striking because the analogous result for Torelli subgroups of mapping class groups with multiple boundary components is still open. We will prove this theorem by using the Birman exact sequence to reduce to $b = 0$ and applying Magnus's theorem. Unfortunately, the tools we have constructed do not seem strong enough to give a novel proof of Magnus's theorem. We will, however, prove a weaker version in Section 7. The original proof of Magnus's Theorem 1.1 comes in two steps: showing that the given automorphisms $\mathrm{Out}(F_n)$-normally generate IO_n, and then showing that the subgroup they generate is normal in $\mathrm{Out}(F_n)$. We will give a proof of the first step in our setting. For alternative proofs of the first step, as well as more information on the second step in this context, see work by Bestvina, Bux and Margalit [5] as well as Day and Putman [10].

Theorem D *The group* IO_n *is* $\mathrm{Out}(F_n)$*-normally generated by the automorphisms* M_{ij} *and* M_{ijk}*, where* $i, j, k \in \{1, \dots, n\}$ *and* $j < k$.

Abelianization Once we have a finite generating set for $\mathrm{IO}_{n,b}^P$, a natural question arises: how does the cardinality of this generating set compare to the rank of $H_1(\mathrm{IO}_{n,b}^P)$? For $b \leq 1$, this question is answered by a result of Andreadakis [1] and Bachmuth [3].

Theorem 1.2 (Andreadakis, Bachmuth) *The abelianization of* $\mathrm{IO}_{n,b}$ *is torsion-free of rank* $n \cdot \binom{n}{2} - n$ *if* $b = 0$*, and rank* $n \cdot \binom{n}{2}$ *if* $b = 1$.

This theorem was proved using a version of the Johnson homomorphism

$$\tau \colon \mathrm{IA}_n \to \mathrm{Hom}\big(H, \textstyle\bigwedge^2 H\big),$$

where $H = H_1(F_n) = \mathbb{Z}^n$. We will recall the definition of this homomorphism in Section 8, along with the proof of Theorem 1.2. We then move on to computing the rank of $H_1(\mathrm{IO}_{n,b}^P)$ for $b > 1$. To do this, we choose an embedding $M_{n,b} \hookrightarrow M_{m,1}$, which induces an injection $\mathrm{IO}_{n,b}^P \to \mathrm{IO}_{m,1} = \mathrm{IA}_m$. Composing this map with τ gives a map $\tau_* \colon \mathrm{IO}_{n,b}^P \to \mathrm{Hom}\big(H, \bigwedge^2 H\big)$. We then compute the image of our generators under τ_*, and use this to count the rank of $\tau_*(\mathrm{IO}_{n,b}^P)$.

Theorem E *The abelianization of* $\mathrm{IO}_{n,b}^P$ *is torsion-free of rank*

$$n \cdot \binom{n}{2} + \left(b \cdot \binom{n}{2} - |P| \cdot \binom{n}{2}\right) + (|P| \cdot n - n).$$

Outline In Section 2, we will give a short overview of sphere twists. We then move on to proving Theorem A in Section 3. We will establish all of our versions of the Birman exact sequence (including Theorem B) in Section 4. In Section 5, we will define our candidate generators for $\mathrm{IO}_{n,b}^P$, and we will prove that they generate (Theorem C) in Section 6 using the Birman exact sequence and Magnus's Theorem 1.1. In Section 7, we will prove Theorem D. We then move on to Section 8, in which we recall the definition of the Johnson homomorphism for IA_n, and use it to compute the rank of the abelianization

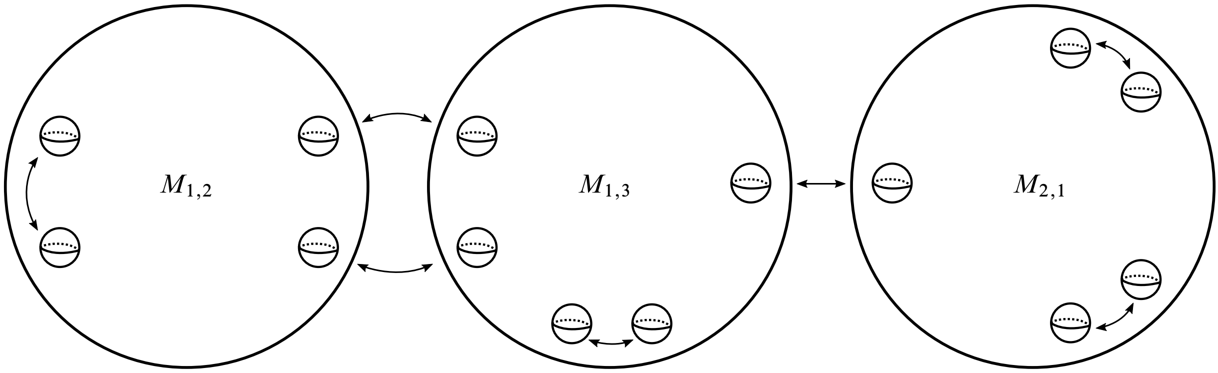

Figure 2: M_5 realized by gluing $M_{1,2}$, $M_{1,3}$, and $M_{2,1}$ together along their boundaries as indicated by the arrows.

of $\mathrm{IO}_{n,b}^P$, proving Theorem E. Finally, we conclude with two appendices. In Appendix A, we provide conditions for a map $\mathrm{Out}(F_{n,b}) \to \mathrm{Out}(F_m)$ induced by an inclusion to be injective, and in Appendix B we prove a lemma which allows us to realize bases of $H_2(M_m)$ as collections of disjoint oriented spheres.

Figure conventions We will frequently direct the reader to figures which are intended to give some geometric intuition for the manifold $M_{n,b}$. In order to assemble $M_{n,b}$, we begin with one or more copies of S^3, remove a collection of open balls, and then glue the resulting boundary components together in pairs. These gluings will be indicated by double-sided arrows connecting the boundary spheres being glued. As an example, see Figure 2.

Acknowledgements I would like to thank my advisor Andy Putman for directing me to $\mathrm{Out}(F_n)$ and its Torelli subgroup, and for his input during the revision process. I would also like to thank Patrick Heslin and Aaron Tyrrell for helpful conversations regarding diffeomorphism groups, as well as Dan Margalit for an enlightening question which resulted in the addition of Section 8.

This material is based upon work supported by the National Science Foundation under grant number DMS-1547292.

2 Preliminaries

Since sphere twists play a fundamental role throughout the remainder of the paper, we will give a brief overview of them here.

Sphere twists Fix a smoothly embedded 2-sphere $S \subset M_{n,b}$, and let $U \cong S \times [0,1]$ be a tubular neighborhood of S. Recall that $\pi_1(\mathrm{SO}(3), \mathrm{id}) \cong \mathbb{Z}/2\mathbb{Z}$, and the nontrivial element $\gamma\colon [0,1] \to \mathrm{SO}(3)$ is given by rotating $\mathbb{R}^3$ one full revolution about any fixed axis through the origin. Fix an identification $S = S^2 \subset \mathbb{R}^3$. Then, we define the sphere twist about S, denoted $T_S \in \mathrm{Mod}(M_{n,b})$, to be the class of the diffeomorphism which is the identity on $M_{n,b} \setminus U$ and is given by $(x,t) \mapsto (\gamma(t)\cdot x, t)$ on $U \cong S \times [0,1]$.

The isotopy class of T_S depends only on the isotopy class of S. In fact, more is true: Laudenbach [19] showed that the class of T_S depends only on the homotopy class of S.

Action on curves and surfaces Since $\pi_1(\mathrm{SO}(3), \mathrm{id}) \cong \mathbb{Z}/2\mathbb{Z}$, we see that sphere twists have order at most two. However, it is tricky to show that sphere twists are actually nontrivial because they act trivially on homotopy classes of embedded arcs and surfaces. To see why this is true, let $S \subset M_{n,b}$ be an embedded 2-sphere, and let $U = S \times [0,1]$ be a tubular neighborhood of S. Suppose that α is either an arc or surface embedded in $M_{n,b}$ (we will handle both cases simultaneously). We can homotope α such that it is either disjoint from U or intersects U transversely. Let $p \in S$ be one of points in S which lies on the axis of rotation used to construct T_S. We can homotope α such that $\alpha \cap U$ collapses into $p \in [0,1]$. Note that this process is not an isotopy, and α is no longer embedded in $M_{n,b}$. This is not an issue because a result of Laudenbach [19] shows that if α is fixed up to homotopy, then it is fixed up to isotopy. Since T_S fixes $p \times [0,1]$ pointwise, it follows that T_S fixes α up to homotopy. The upshot of this is that a more sophisticated invariant must be constructed to detect the nontriviality of T_S. In [18; 19], Laudenbach uses framed cobordisms to show that for $b = 0, 1$, the sphere twist T_S is trivial if and only if S is separating. In the case of no boundary components, Brendle, Broaddus, and Putman [6] give another proof of this fact by showing that sphere twists act nontrivially on a trivialization of the tangent bundle of M_n up to isotopy.

Sphere twist subgroup Let $\mathrm{STwist}(M_{n,b}) \subset \mathrm{Mod}(M_{n,b})$ be the subgroup generated by sphere twists. Given $\mathfrak{f} \in \mathrm{Mod}(M_{n,b})$ and a sphere twists T_S, we have the "change of coordinates" formula

$$\mathfrak{f} T_S \mathfrak{f}^{-1} = T_{\mathfrak{f}(S)}.$$

This shows that $\mathrm{STwist}(M_{n,b})$ is a normal subgroup of $\mathrm{Mod}(M_{n,b})$. In fact, even more is true. Letting $\mathfrak{f} = T_{S'}$ in the above formula and using the fact that sphere twists act trivially on embedded surfaces up to isotopy, we find that

$$T_{S'} T_S T_{S'}^{-1} = T_{T_{S'}(S)} = T_S,$$

which implies $\mathrm{STwist}(M_{n,b})$ is actually abelian. Since nontrivial sphere twists have order two, it follows that $\mathrm{STwist}(M_{n,b})$ is isomorphic to a product of copies of $\mathbb{Z}/2\mathbb{Z}$. For $b = 0, 1$, another result of Laudenbach shows that $\mathrm{STwist}(M_{n,b}) \cong (\mathbb{Z}/2\mathbb{Z})^n$ and is generated by the sphere twists about the n core spheres $* \times S^2$ in each $S^1 \times S^2$ summand. For $b > 1$, one can show that $\mathrm{STwist}(M_{n,b}) \cong (\mathbb{Z}/2\mathbb{Z})^{n+b-1}$. The -1 in the exponent reflects the fact that the product of all the sphere twists about boundary components is trivial. Since we will need this fact later, we include a proof here.

Lemma 2.1 *If $S_1, \dots, S_b \subset M_{n,b}$ be spheres parallel to the b boundary components of $M_{n,b}$, then the element $T_{S_1} \cdots T_{S_b}$ is trivial in* $\mathrm{Mod}(M_{n,b})$.

Proof We will prove this by induction on n. As the base case, consider $M_{0,b}$. The argument in this case follows a proof of Hatcher and Wahl [15, pages 214–215], but we include the proof here as well for

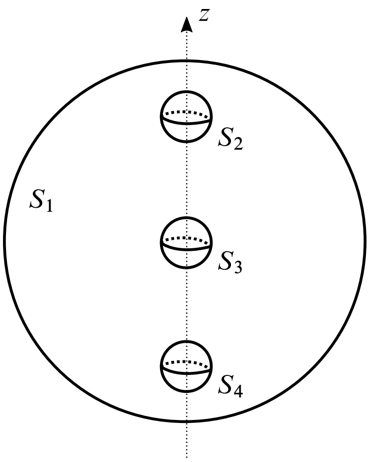

Figure 3: $M_{0,4}$ embedded in $\mathbb{R}^3$.

completeness. If $b = 0$, then the statement is trivial. If $b > 0$, then we can embed $M_{0,b}$ in $\mathbb{R}^3$ as the unit ball with $b-1$ smaller balls removed along the z-axis (see Figure 3). We may then use the z-axis as the axis of rotation for the sphere twists about all the boundary components. Taking S_1 to be the unit sphere, we then see that the product $T_2 \cdots T_b$ is isotopic to T_1. Since sphere twists have order two, this gives the desired relation, and so we have completed the base case.

Next, consider $M_{n,b}$ for $n > 0$. Since $n > 0$, there exists a nonseparating sphere $S \subset M_{n,b}$ which is disjoint from $S_1, \dots, S_b$. Splitting $M_{n,b}$ along S yields a submanifold diffeomorphic to $M_{n-1,b+2}$. Let $\iota_M \colon \mathrm{Mod}(M_{n-1,b+2}) \to \mathrm{Mod}(M_{n,b})$ be the map induced by inclusion. Let $T_1, \dots, T_{b+2}$ be the sphere twists about the boundary components of $M_{n-1,b+2}$, and order them such that $\iota_M(T_j) = T_{S_j}$ for $0 \le j \le b$. With this ordering, notice that $\iota_M(T_{b+1}) = \iota_M(T_{b+2}) = T_S$. Since sphere twists have order two,

$$\iota_M(T_1 \cdots T_{b+2}) = T_{S_1} \cdots T_{S_b} \cdot T_S^2 = T_{S_1} \cdots T_{S_b}.$$

By our induction hypothesis, $T_1 \cdots T_{b+2}$ is trivial in $\mathrm{Mod}(M_{n-1,b+2})$, and so we are done. □

If $b = 1$, this shows that the sphere twist about the boundary component is trivial. However, if $b > 1$, then the sphere twists about boundary components are nontrivial. We will also need this fact, so we prove it here.

Lemma 2.2 *Let $b > 1$, and let ∂ be a boundary component of $M_{n,b}$. Then $T_\partial \in \mathrm{Mod}(M_{n,b})$ is nontrivial.*

Proof Let ∂' be a boundary component of $M_{n,b}$ different from ∂. Then we get an embedding

$$\iota \colon M_{n,b} \hookrightarrow M_{n+1}$$

by attaching ∂ and ∂' with a copy of $S^2 \times I$, and capping off all the remaining boundary components. Let $\iota_M \colon \mathrm{Mod}(M_{n,b}) \to \mathrm{Mod}(M_{n+1})$ be the map induced by ι. Then $\iota_M(T_\partial)$ is a sphere twist about a nonseparating sphere. Earlier in this section, we saw that such sphere twists are nontrivial, and so we conclude that T_∂ is nontrivial as well. □

3 Restriction theorem

Fix $n, b \geq 0$, and let P be a partition of the boundary components of $M_{n,b}$. Recall that we have defined $H_1^P(M_{n,b})$ to be the submodule of $H_1(M_{n,b}, \partial M_{n,b})$ generated by

$$\big\{[h] \in H_1(M_{n,b}, \partial M_{n,b}) \mid \text{ either } h \text{ is a simple closed curve or } h \text{ is a properly embedded arc}$$
$$\text{with endpoints in distinct } P\text{-adjacent boundary components}\big\},$$

and $\mathrm{IO}_{n,b}^P$ is the kernel of the action of $\mathrm{Out}(F_{n,b})$ on $H_1^P(M_{n,b})$ induced by the action of $\mathrm{Mod}(M_{n,b})$.

Remark This version of homology is simpler than the one used in [22]. There are two reasons for this.

- In our case, we can take homology relative to the entire boundary, whereas in [22], homology is taken relative to a set consisting of a single point from each boundary component. This is because in surfaces, the boundary components give nontrivial elements of H_1, and the arcs considered in $H_1^P(\Sigma_{g,b})$ can get "wrapped around" those boundary components. This is not a problem in our setting because loops in boundary components of $M_{n,b}$ are trivial in H_1.
- Next, suppose we have an embedding $\Sigma_{g,b} \hookrightarrow \Sigma_{g'}$ of surfaces. It is possible for a nontrivial element $a \in H_1(\Sigma_{g,b})$ to become trivial in $H_1(\Sigma_{g'})$ (for instance, if a boundary component is capped off). So, there could be elements of $\mathrm{Mod}(\Sigma_{g,b})$ which act trivially on $H_1(\Sigma_{g'})$, but not fix a. In other words, the Torelli group would not be closed under restrictions. To fix this, the author in [22] must mod out by the submodules of $H_1(\Sigma_{g,b})$ spanned by the $p \in P$ (with proper orientation chosen). This is not a problem in the 3-dimensional case however, since an inclusion $M_{n,b} \hookrightarrow M_m$ induces an injection $H_1(M_{n,b}) \to H_1(M_m)$.

Proof of Theorem A Let $\iota\colon M_{n,b} \hookrightarrow M_m$ be an embedding, and let $\iota_*\colon \mathrm{Out}(F_{n,b}) \to \mathrm{Out}(F_m)$ be the induced map. Recall that we must show that $\iota_*^{-1}(\mathrm{IO}_m) = \mathrm{IO}_{n,b}^P$, where P is the partition of the boundary components induced by ι as described in the introduction.

This proof will follow the proof of [22, Theorem 3.3]. Define the following subsets of $H_1(M_m)$ (we use $\cdot$ to denote concatenation of arcs):

$$Q_1 = \{[h] \in H_1(M_m) \mid h \text{ is a simple closed curve in } M_m \setminus M_{n,b}\},$$
$$Q_2 = \{[h] \in H_1(M_m) \mid h \text{ is a simple closed curve in } M_{n,b}\},$$
$$Q_3 = \big\{[h_1 \cdot h_2] \in H_1(M_m) \mid h_1 \text{ is a properly embedded arc in } M_{n,b}$$
$$\text{with endpoints in distinct } P\text{-adjacent boundary components}$$
$$\text{and } h_2 \text{ is a properly embedded arc in } M_m \setminus M_{n,b}$$
$$\text{with the same endpoints as } h_1\big\}.$$

We claim that the homology group $H_1(M_m)$ is spanned by $Q_1 \cup Q_2 \cup Q_3$. To see why, let $[\alpha] \in H_1(M_m)$ be the class of a loop α. If α can be homotoped to lie entirely inside $M_{n,b}$ or $M_m \setminus M_{n,b}$, then we are

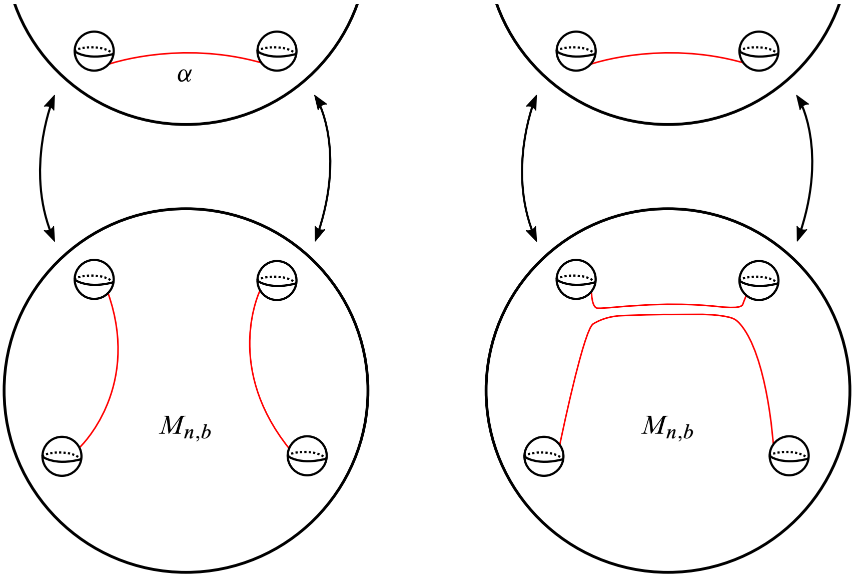

Figure 4: A loop can be surgered into a collection of loops which intersect $\partial M_{n,b}$ exactly twice.

done. On the other hand, suppose that crosses the boundary of $M_{n,b}$. Without loss of generality, we may assume that α crosses the boundary of $M_{n,b}$ exactly twice since any loop can be surgered into a collection of such loops (see Figure 4). It follows that α has the form $\alpha = \gamma \cdot \delta$, where $\gamma \subset M_{n,b}$ is an arc connecting boundary components of $M_{n,b}$, and $\delta \subset M_m \setminus M_{n,b}$ is a arc with the same endpoints as γ. Recall that under the partition P induced by the inclusion ι, two boundary components are P-adjacent if they lie on the same component of $M_m \setminus M_{n,b}$. Therefore, the existence of δ implies that the boundary components intersected by α are P-adjacent, and thus $[\alpha] \in Q_3$. This completes the proof of the claim.

Let $f \in \mathrm{IO}^P_{n,b}$. By the definition of $\mathrm{IO}^P_{n,b}$, the element $\iota_*(f)$ acts trivially on Q_2. Moreover, $\iota_*(f)$ acts trivially on Q_1 by the definition of ι_*. Lastly, suppose that $[h_1 \cdot h_2] \in Q_3$. Then $\iota_*(f)$ fixes the homology class of h_1 since $f \in \mathrm{IO}^P_{n,b}$, and fixes h_2 pointwise by the definition of ι_*. Therefore, $f \in \iota_*^{-1}(\mathrm{IO}_m)$.

Next, suppose that $f \in \iota_*^{-1}(\mathrm{IO}_m)$. By definition, $\iota_*(f)$ acts trivially on $H_1(M_m)$, and thus on Q_2 as well since the map $H_1(M_{n,b}) \to H_1(M_m)$ induced by ι is injective. This implies that f acts trivially on homology classes of simple closed curves in $M_{n,b}$. So, we only need to check that f preserves the homology classes of arcs in M_m connecting distinct P-adjacent boundary components. Suppose there is a class of arcs $[\alpha] \in H^P_1(M_{n,b})$. Since P is the partition of the boundary components induced by ι, $[\alpha]$ can be completed to a homology class $[\alpha \cdot \beta] \in H_1(M_m)$, where β is an arc in $M_m \setminus M_{n,b}$ connecting the endpoints of α. Then since $\iota_*(f) \in \mathrm{IO}_m$ and $\iota_*(f)$ fixes β pointwise, we have

$$0 = ([\alpha \cdot \beta]) - \iota_*(f)([\alpha \cdot \beta]) = [\alpha] - f([\alpha]).$$

This shows that f acts trivially on $[\alpha]$. Therefore, $f \in \mathrm{IO}^P_{n,b}$. □

4 Birman exact sequence

In this section, we give a version of the Birman exact sequence for the groups $\mathrm{IO}^P_{n,b}$. We will start by giving a Birman exact sequence on the level of mapping class groups. We note that Banks has proved a version of the Birman exact sequence for 3-manifolds (see [4]). However, this version involves forgetting a puncture rather than capping a boundary component, so we will prove our own version here. Once we have the sequence for mapping class groups, we will mod out by sphere twists to get a corresponding sequence for $\mathrm{Out}(F_{n,b})$, and finally restrict to get a sequence for $\mathrm{IO}^P_{n,b}$.

Remark In the following theorems, we exclude the case $(n,b)=(1,1)$. This is because boundary drags in $\mathrm{Mod}(M_{1,1})$ are trivial (see the proof of Theorem 4.1 for the definition of boundary drags). In this case, we have isomorphisms

$$\mathrm{Mod}(M_{1,1})\cong\mathrm{Mod}(M_1)\cong\mathbb{Z}/2\times\mathbb{Z}/2,\quad \mathrm{Out}(F_{1,1})\cong\mathrm{Out}(F_1)\cong\mathbb{Z}/2,\quad \mathrm{IO}^{\{\partial\}}_{1,1}\cong\mathrm{IO}_1\cong 1,$$

where one of the generators of $\mathrm{Mod}(M_1)=\mathrm{Mod}(S^1\times S^2)$ is a sphere twist about the sphere $*\times S^2$ and the other is the antipodal map in both coordinates.

Theorem 4.1 *Fix $n,b>0$ such that $(n,b)\neq(1,1)$. Glue a ball to a boundary component of $M_{n,b}$, and let $M_{n,b}\hookrightarrow M_{n,b-1}$ be the resulting embedding. Fix $x\in M_{n,b-1}\setminus M_{n,b}$.*

(a) *If $b>1$, choose a lift $\tilde{x}\in\mathrm{Fr}_x(M_{n,b-1})$ of x, where $\mathrm{Fr}(M_{n,b-1})$ is the oriented frame bundle of $M_{n,b-1}$. We then have an exact sequence*

$$1\to\pi_1(\mathrm{Fr}(M_{n,b-1}),\tilde{x})\to\mathrm{Mod}(M_{n,b})\to\mathrm{Mod}(M_{n,b-1})\to 1.$$

(b) *If $b=1$ and $n>1$, then we have an exact sequence*

$$1\to\pi_1(M_{n,b-1},x)\to\mathrm{Mod}(M_{n,b})\to\mathrm{Mod}(M_{n,b-1})\to 1.$$

Proof Let $\mathrm{Diff}^+(M_{n,b-1})$ denote the space of orientation-preserving diffeomorphisms of $M_{n,b-1}$ which restrict to the identity on $\partial M_{n,b-1}$, and let $\mathrm{Diff}^+(M_{n,b-1},\tilde{x})$ be the subspace of $\mathrm{Diff}^+(M_{n,b-1})$ consisting of diffeomorphisms which fix the framing $\tilde{x}$. It is standard that there is a fiber bundle

$$\mathrm{Diff}^+(M_{n,b-1},\tilde{x})\to\mathrm{Diff}^+(M_{n,b-1})\xrightarrow{p}\mathrm{Fr}(M_{n,b-1}),$$

where the map $p\colon\mathrm{Diff}^+(M_{n,b-1})\to\mathrm{Fr}(M_{n,b-1})$ is given by $\varphi\mapsto d\varphi(\tilde{x})$. Passing to the long exact sequence of homotopy groups associated to this fiber bundle, we find the segment

$$\pi_1(\mathrm{Fr}(M_{n,b-1}))\to\pi_0(\mathrm{Diff}^+(M_{n,b-1},\tilde{x}))\to\pi_0(\mathrm{Diff}^+(M_{n,b-1}))\to\pi_0(\mathrm{Fr}(M_{n,b-1})).\tag{2}$$

Since $\mathrm{Fr}(M_{n,b-1})$ is the *oriented* frame bundle, it is connected, and so $\pi_0(\mathrm{Fr}(M_{n,b-1}))$ is trivial.

Next, we claim that $\pi_0(\mathrm{Diff}^+(M_{n,b-1},\tilde{x}))$ is isomorphic to $\mathrm{Mod}(M_{n,b})$. For a proof of this fact in the surface case, see [12, page 102]. We will give an analogous argument here. Let B be the ball glued to the boundary component of $M_{n,b}$, and let $\mathrm{Diff}^+(M_{n,b-1}, B) \subset \mathrm{Diff}^+(M_{n,b-1})$ be the subspace of diffeomorphisms fixing B pointwise. Then $\pi_0(\mathrm{Diff}^+(M_{n,b-1}, B)) \cong \mathrm{Mod}(M_{n,b})$, where the isomorphism is obtained by simply removing B. Letting $\mathrm{Emb}^+((B, M_{n,b-1}), \tilde{x})$ be the space of smooth, orientation-preserving embeddings $B \hookrightarrow M_{n,b-1}$ which fix the framing $\tilde{x}$, we get a fiber bundle

$$\mathrm{Diff}^+(M_{n,b-1}, B) \hookrightarrow \mathrm{Diff}^+(M_{n,b-1}, \tilde{x}) \to \mathrm{Emb}^+((B, M_{n,b-1}), \tilde{x}),$$

where the second map is given by restriction to B. Again, we pass to the long exact sequence of homotopy groups, and find the segment

$$\pi_1(\mathrm{Emb}^+((B, M_{n,b-1}), \tilde{x})) \to \pi_0(\mathrm{Diff}^+(M_{n,b-1}, B)) \\ \to \pi_0(\mathrm{Diff}^+(M_{n,b-1}, \tilde{x})) \to \pi_0(\mathrm{Emb}^+((B, M_{n,b-1}), \tilde{x})).$$

Since B is contractible, so is $\mathrm{Emb}^+((B, M_{n,b-1}), \tilde{x})$. This gives an isomorphism

$$\pi_0(\mathrm{Diff}^+(M_{n,b-1}, \tilde{x})) \cong \pi_0(\mathrm{Diff}^+(M_{n,b-1}, B)) \cong \mathrm{Mod}(M_{n,b}),$$

as desired.

With these identifications, the sequence (2) then becomes

$$\pi_1(\mathrm{Fr}(M_{n,b-1})) \to \mathrm{Mod}(M_{n,b}) \to \mathrm{Mod}(M_{n,b-1}) \to 1. \tag{3}$$

To get a short exact sequence, we must understand the kernel of the map

$$\widetilde{\mathrm{Push}}\colon \pi_1(\mathrm{Fr}(M_{n,b-1})) \to \mathrm{Mod}(M_{n,b}).$$

We remark here that the map $\widetilde{\mathrm{Push}}$ is given by pushing and rotating a small ball containing x about a loop based at x. This is in analogy with the "disk pushing maps" seen in the case of surfaces. Since $M_{n,b-1}$ is a compact, orientable 3-manifold, it is parallelizable, and hence we have

$$\pi_1(\mathrm{Fr}(M_{n,b-1})) \cong \pi_1(M_{n,b-1}) \times \pi_1(SO(3)) = \pi_1(M_{n,b-1}) \times \mathbb{Z}/2.$$

Consider the map $\mathrm{Mod}(M_{n,b}) \to \mathrm{Aut}(\pi_1(M_{n,b}, y))$, where the basepoint y is on the boundary component ∂ being capped off. As is shown in Figure 5, the composition

$$\pi_1(\mathrm{Fr}(M_{n,b-1})) \cong \pi_1(M_{n,b-1}) \times \mathbb{Z}/2 \xrightarrow{\widetilde{\mathrm{Push}}} \mathrm{Mod}(M_{n,b}) \to \mathrm{Aut}(\pi_1(M_{n,b}, y))$$

is given by conjugation about the loop being pushed around. Since $\mathrm{Aut}(\pi_1(M_{n,b}, y)) \cong \mathrm{Aut}(F_n)$ is centerless for $n > 1$, the entire kernel of $\widetilde{\mathrm{Push}}$ must be contained in $1 \times \mathbb{Z}/2 \subset \pi_1(M_{n,b-1}) \times \mathbb{Z}/2$. However, the image of the generator of this subgroup in $\mathrm{Mod}(M_{n,b})$ is the sphere twist T_∂. By Theorems 2.1 and 2.2, this sphere twist is nontrivial if and only if $b > 1$. If $b > 1$, this shows that Push is injective, and (3) gives us the desired exact sequence. On the other hand, if $b = 1$, then $\ker(\widetilde{\mathrm{Push}}) = 1 \times \mathbb{Z}/2$. Therefore, the image of $\widetilde{\mathrm{Push}}$ in $\mathrm{Mod}(M_{n,b})$ is isomorphic to

$$\pi_1(\mathrm{Fr}(M_{n,b-1}))/\langle T_\partial \rangle \cong \pi_1(M_{n,b-1})$$

as desired. □

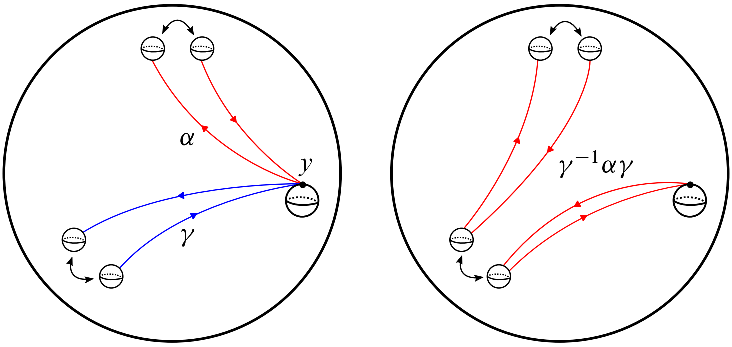

Figure 5: The image of α under $\widetilde{\mathrm{Push}}(\gamma, T)$ is $\gamma^{-1}\alpha\gamma$. Here, T can be either T_∂ or trivial.

Modding out by sphere twists Now that we have a Birman exact sequence for $\mathrm{Mod}(M_{n,b})$, we can mod out by sphere twists to get a Birman exact sequence for $\mathrm{Out}(F_{n,b})$. Consider the map

$$i_M\colon \mathrm{Mod}(M_{n,b}) \to \mathrm{Mod}(M_{n,b-1})$$

given by capping off a boundary component ∂. Since ι_M takes sphere twists to sphere twists, this map descends to a map $\iota_*\colon \mathrm{Out}(F_{n,b}) \to \mathrm{Out}(F_{n,b-1})$. Since ι_M is surjective, ι_* is as well. Let K be the kernel of ι_*, and let $\psi\colon \mathrm{Mod}(M_{n,b}) \to \mathrm{Out}(F_{n,b})$ be the quotient map. If $b > 1$, then the kernel of ι_M is $\pi_1(\mathrm{Fr}(M_{n,b-1}), \tilde{x})$ by Theorem 4.1. Let $\widetilde{\mathrm{Push}}\colon \pi_1(\mathrm{Fr}(M_{n,b-1}), \tilde{x}) \to \mathrm{Mod}(M_{n,b})$ be the map defined in the proof of Theorem 4.1, and fix an identification $\pi_1(\mathrm{Fr}(M_{n,b-1}), \tilde{x}) = \pi_1(M_{n,b-1}, x) \times \mathbb{Z}/2$. Since

$$\iota_*\big(\psi(\widetilde{\mathrm{Push}}(\gamma, T))\big) = \psi\big(\iota_M(\widetilde{\mathrm{Push}}(\gamma, T))\big) = \psi(\mathrm{id}) = \mathrm{id}$$

for all $(\gamma, T) \in \pi_1(M_{n,b-1}, x) \times \mathbb{Z}/2\mathbb{Z}$, the image of $\pi_1(\mathrm{Fr}(M_{n,b-1}), \tilde{x})$ under $\psi \circ \widetilde{\mathrm{Push}}$ is contained in K. In other words, we have the following commutative diagram:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & \pi_1(\mathrm{Fr}(M_{n,b-1}), \tilde{x}) & \xrightarrow{\widetilde{\mathrm{Push}}} & \mathrm{Mod}(M_{n,b}) & \xrightarrow{\iota_M} & \mathrm{Mod}(M_{n,b-1}) & \longrightarrow & 1 \\ & & \downarrow{\scriptstyle \psi_P} & & \downarrow{\scriptstyle \psi} & & \downarrow & & \\ 1 & \longrightarrow & K & \longrightarrow & \mathrm{Out}(F_{n,b}) & \xrightarrow{\iota_*} & \mathrm{Out}(F_{n,b-1}) & \longrightarrow & 1 \end{array}$$

where $\psi_P = \psi \circ \widetilde{\mathrm{Push}}$. Next, we claim that the map $\psi_P\colon \pi_1(\mathrm{Fr}(M_{n,b-1}), \tilde{x}) \to K$ is surjective. To see this, let $f \in K$, and choose a lift $\mathfrak{f} \in \mathrm{Mod}(M_{n,b})$ of f. Since $\iota_*(f) = \mathrm{id}$, the image $\iota_M(\mathfrak{f})$ is a product of sphere twists $T_{S_1} \cdots T_{S_j}$. For each $T_{S_i} \in \mathrm{Mod}(M_{n,b-1})$, choose a preimage $T'_{S_i} \in \mathrm{Mod}(M_{n,b})$ which is also a sphere twist. Then

$$\iota_M(T'_{S_1} \cdots T'_{S_j} \cdot \mathfrak{f}) = \mathrm{id},$$

which implies that $T'_{S_1} \cdots T'_{S_j} \cdot \mathfrak{f} = \widetilde{\mathrm{Push}}(\gamma, T)$ for some $(\gamma, T) \in \pi_1(M_{n,b-1}, x) \times \mathbb{Z}/2\mathbb{Z}$. Moreover, $\psi(T'_{S_1} \cdots T'_{S_j} \cdot \mathfrak{f}) = f$, which verifies our claim that $\psi_P\colon \pi_1(\mathrm{Fr}(M_{n,b-1}), \tilde{x}) \to K$ is surjective.

Now, we wish to identify the kernel of ψ_P. Let $(\gamma, T) \in \pi_1(M_{n,b-1}, x) \times \mathbb{Z}/2\mathbb{Z}$, and fix a basepoint y on the boundary component being capped off. At the end of the proof of Theorem 4.1, we saw that

$\widetilde{\text{Push}}(\gamma, T)$ acts nontrivially on $\pi_1(M_{n,b}, y)$ if and only if γ is trivial. Since sphere twists act trivially on homotopy classes of curves, it follows that $\psi_P(\gamma, T)$ is nontrivial if γ is nontrivial. Therefore, the kernel of ψ_P must lie inside $1 \times \mathbb{Z}/2\mathbb{Z} \subset \pi_1(M_{n,b-1}, x) \times \mathbb{Z}/2\mathbb{Z}$. However, the generator of $1 \times \mathbb{Z}/2\mathbb{Z}$ gets mapped to T_∂ under $\widetilde{\text{Push}}$, which is killed in $\text{Out}(F_{n,b})$. Therefore, $\ker(\psi) = 1 \times \mathbb{Z}/2\mathbb{Z}$, and so it follows that $K \cong \pi_1(M_{n,b-1}, x)$.

On the other hand, if $b = 1$ and $n > 1$, then the kernel of the map $\iota_M\colon \text{Mod}(M_{n,b}) \to \text{Mod}(M_{n,b-1})$ is $\pi_1(M_{n,b-1}, x)$ by Theorem 4.1. Almost exactly the same argument used above shows that the quotient map restricts to a surjective map $\psi_P\colon \pi_1(M_{n,b-1}, x) \to K$. However, in this case, ψ_P is injective since the sphere twist T_∂ has already been killed off. Thus, we find that $K \cong \pi_1(M_{n,b-1}, x)$ in this case as well.

From now on, we will identify the kernel of the map $\iota_*\colon \text{Out}(F_{n,b}) \to \text{Out}(F_{n,b-1})$ with $\pi_1(M_{n,b-1}, x)$. The map $\pi_1(M_{n,b-1}, x) \to \text{Out}(F_{n,b})$ will play a significant role throughout the remainder of the paper, and so we give a formal definition here.

Definition The map $\text{Push}\colon \pi_1(M_{n,b-1}, x) \to \text{Out}(F_{n,b})$ is defined as $\text{Push}(\gamma) = \psi(\widetilde{\text{Push}}(\gamma, T))$, where $T \in \mathbb{Z}/2\mathbb{Z}$ is arbitrary. Since sphere twists become trivial in $\text{Out}(F_{n,b})$, this element depends only on γ.

The upshot of this is that we have proven the Birman exact sequence for $\text{Out}(F_{n,b})$.

Theorem 4.2 *Fix $n, b > 0$ such that $(n, b) \neq (1, 1)$, and let $M_{n,b} \hookrightarrow M_{n,b-1}$ be an embedding obtained by gluing a ball to a boundary component. Fix $x \in M_{n,b-1} \setminus M_{n,b}$. Then the following sequence is exact:*

$$1 \to \pi_1(M_{n,b-1}, x) \xrightarrow{\text{Push}} \text{Out}(F_{n,b}) \xrightarrow{\iota_*} \text{Out}(F_{n,b-1}) \to 1.$$

Restrict to Torelli We now move on to proving Theorem B, which gives a Birman exact sequence for $\text{IO}_{n,b}^P$. We start by recalling its statement. Let P be a partition of the boundary components of $M_{n,b}$, and fix a boundary component ∂. Let $p \in P$ be the set containing ∂, and let $\iota\colon M_{n,b} \hookrightarrow M_{n,b-1}$ be the inclusion obtained by capping off ∂. The partition P induces a partition P' of the boundary components of $M_{n,b-1}$ by removing ∂ from p. With this definition of P', the map $\iota_*\colon \text{Out}(F_{n,b}) \to \text{Out}(F_{n,b-1})$ restricts to a map $\text{IO}_{n,b}^P \to \text{IO}_{n,b-1}^{P'}$, which we will also call ι_*. The sequence from Theorem 4.2 then restricts to

$$1 \to \pi_1(M_{n,b-1}) \cap \text{IO}_{n,b}^P \to \text{IO}_{n,b}^P \xrightarrow{\iota_*} \text{IO}_{n,b-1}^{P'}.$$

Theorem B asserts that ι_* is surjective, and identifies its kernel $\pi_1(M_{n,b-1}) \cap \text{IO}_{n,b}^P$. We start with surjectivity.

Lemma 4.3 *The induced map $\iota_*\colon \text{IO}_{n,b}^P \to \text{IO}_{n,b-1}^{P'}$ is surjective for any embedding $\iota\colon M_{n,b} \hookrightarrow M_{n,b-1}$.*

Proof Consider an element $g \in \text{IO}_{n,b-1}^{P'}$. Our goal is to find some $f \in \text{IO}_{n,b}^P$ such that $\iota_*(f) = g$. There are two cases.

First, suppose that $p=\{\partial\}$. Then the inclusion ι induces an isomorphism $\iota_H : H_1^P(M_{n,b}) \to H_1^{P'}(M_{n,b-1})$ which is equivariant with respect to the actions of $\mathrm{Out}(F_{n,b})$ and $\mathrm{Out}(F_{n,b-1})$. In other words, for any $[h] \in H_1^P(M_{n,b})$ and $f \in \mathrm{Out}(F_{n,b})$, we have

$$\iota_H(f \cdot [h]) = \iota_*(f) \cdot \iota_H([h]). \tag{4}$$

By Theorem 4.2, there exists some $f \in \mathrm{Out}(F_{n,b})$ such that $\iota_*(f) = g$. We claim that $f \in \mathrm{IO}_{n,b}^P$. To see this, let $[h] \in H_1^P(M_{n,b})$. Then, by (4), we see that

$$\iota_H(f \cdot [h]) = \iota_*(f) \cdot \iota_H([h]) = g \cdot \iota_H([h]) = \iota_H([h]).$$

Since ι_H is an isomorphism, this implies that $f \cdot [h] = [h]$, and so $f \in \mathrm{IO}_{n,b}^P$, as desired.

Next, suppose that $p \neq \{\partial\}$. Again, choose $f \in \mathrm{Out}(F_{n,b})$ such that $\iota_*(f) = g$. In this case, there is no longer a well-defined map $H_1^P(M_{n,b}) \to H_1^{P'}(M_{n,b-1})$. However, there is a subgroup of $H_1^P(M_{n,b})$ which projects isomorphically onto $H_1^{P'}(M_{n,b-1})$. Let $A \subset H_1^P(M_{n,b})$ be the subgroup generated by

$$\big\{[a] \in H_1(M_{n,b}, \partial M_{n,b}) \mid \text{ either } a \text{ is a simple closed curve or } a \text{ is a properly embedded arc with neither endpoint on } \partial\big\}.$$

It is clear that $A \cong H_1^{P'}(M_{n,b-1})$.

Let $[k] \in H_1^P(M_{n,b})$ be the class of an arc k which has an endpoint on ∂. We claim that $H_1^P(M_{n,b})$ is generated by A and $[k]$. To establish this claim, it suffices to show that $[\ell] \in \langle A, [k]\rangle$, where $[\ell] \in H_1^P(M_{n,b})$ is the class of any arc with an endpoint on ∂ and the other elsewhere. Such an ℓ exists since $p \neq \{\partial\}$. Fix such a class $[\ell]$, and let $\alpha \subset \partial$ be an arc connecting the endpoints of ℓ and k on ∂. Orient ℓ, α, and k such that the curve $\ell \cdot \alpha \cdot k$ is well-defined.

If the endpoints of ℓ and k which are not on ∂ lie on distinct boundary components, then $\ell \cdot \alpha \cdot k$ is an arc connecting P'-adjacent boundary components. Therefore, $[\ell] + [\alpha] + [k] \in A$. Since $[\alpha] = 0$ in $H_1^P(M_{n,b})$, it follows that $[\ell] \in \langle A, [k]\rangle$. On the other hand, if the endpoints of ℓ and k which are not on ∂ lie on the same boundary component ∂', then we can complete $\ell \cdot \alpha \cdot k$ to a loop $\ell \cdot \alpha \cdot k \cdot \beta$, where $\beta \subset \partial'$ is an arc connecting the endpoints of ℓ and k. Then

$$[\ell] + [k] = [\ell] + [\alpha] + [k] + [\beta] = [\ell \cdot \alpha \cdot k \cdot \beta] \in A,$$

and so $[\ell] \in \langle A, [k]\rangle$. This completes the proof of the claim that $H_1^P(M_{n,b})$ is generated by A and $[k]$.

Since A projects isomorphically onto $H_1^{P'}(M_{n,b-1})$, and this projection is equivariant with respect to the actions of $\mathrm{Out}(F_{n,b})$ and $\mathrm{Out}(F_{n,b-1})$, we have $f \cdot [a] = [a]$. It follows that f acts trivially on A. Therefore, if f fixes $[k]$, then $f \in \mathrm{IO}_{n,b}^P$ by the discussion in the preceding paragraph, and so we are done. On the other hand, if f does not fix $[k]$, then $\gamma = k \cdot f(k)^{-1}$ is a nontrivial loop based at a point on ∂. So, the element $\mathrm{Push}(\gamma)^{-1} \cdot f \in \mathrm{Out}(F_{n,b})$ fixes $[k]$. Moreover, $\mathrm{Push}(\gamma)$ acts trivially on A, and so $\mathrm{Push}(\gamma)^{-1} \cdot f$ does as well. Thus, $\mathrm{Push}(\gamma)^{-1} \cdot f \in \mathrm{IO}_{n,b}^P$. Finally, since $\mathrm{Push}(\gamma) \in \ker(i_*)$, we have that $\iota_*(\mathrm{Push}(\gamma)^{-1} \cdot f) = g$, and so we are done. □

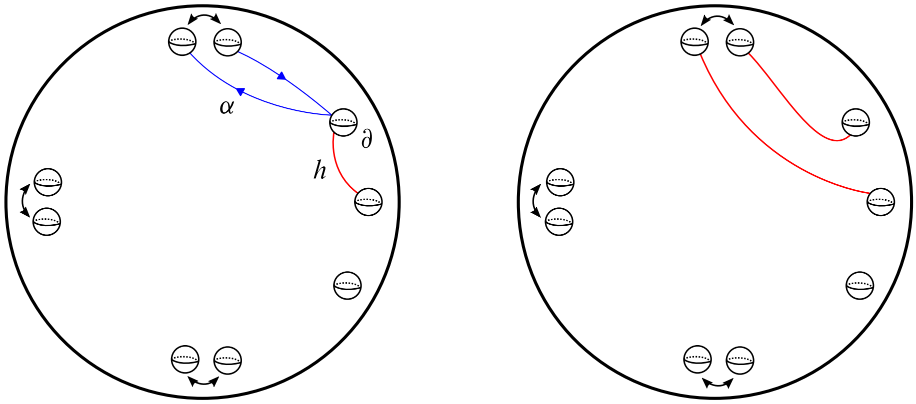

Figure 6: Dragging ∂ around α takes $[h]$ to $[\alpha]+[h]$.

Proof of Theorem B Recall that we want to show that we have an exact sequence

$$1 \to L \xrightarrow{\text{Push}} \mathrm{IO}_{n,b}^{P} \xrightarrow{\iota_*} \mathrm{IO}_{n,b-1}^{P'} \to 1,$$

where L is equal to

(a) $\pi_1(M_{n,b-1},x) \cong F_n$ if $p=\{\partial\}$,

(b) $[\pi_1(M_{n,b-1},x),\pi_1(M_{n,b-1},x)] \cong [F_n,F_n]$ if $p \neq \{\partial\}$.

By Lemma 4.3 and the discussion preceding it, all that is left to show is that $\pi_1(M_{n,b-1}) \cap \mathrm{IO}_{n,b}^{P}$ agrees with the subgroups L given above.

We begin with the case $p=\{\partial\}$. Recall that $\pi_1(M_{n,b-1})$ acts on $M_{n,b}$ by pushing the boundary component ∂ about a given loop. Since ∂ is not P-adjacent to any other boundary components, it follows that $\pi_1(M_{n,b-1})$ acts trivially on $H_1^P(M_{n,b})$. Therefore, $\pi_1(M_{n,b-1}) \subset \mathrm{IO}_{n,b}^{P}$, and so $\pi_1(M_{n,b-1}) \cap \mathrm{IO}_{n,b}^{P} = \pi_1(M_{n,b-1})$. This completes this case.

Next, suppose that $p \neq \{\partial\}$. In this case, not all elements of $\pi_1(M_{n,b})$ are contained in $\mathrm{IO}_{n,b}^{P}$. This is because dragging ∂ about loops may change the homology class of arcs connected to ∂. In particular, if $\gamma \in \pi_1(M_{n,b-1})$ and $[h] \in H_1^P(M_{n,b})$ is the class of arc with an endpoint in ∂ and the other elsewhere, then $\text{Push}(\gamma)$ acts on $[h]$ via

$$\text{Push}(\gamma)\cdot[h] = [\gamma]+[h].$$

See Figure 6 for an illustration. This implies that an element $\text{Push}(\gamma)$ is in $\mathrm{IO}_{n,b}^{P}$ if and only if $[\gamma]=0$ in $H_1(M_{n,b-1})$. Thus,

$$\pi_1(M_{n,b-1}) \cap \mathrm{IO}_{n,b}^{P} = [\pi_1(M_{n,b-1}),\pi_1(M_{n,b-1})],$$

which is what we wanted to show. □

5 Generators

In this section, we will define our generators of $\mathrm{IO}^P_{n,b}$. The definition of these generators will involve splitting and dragging boundary components, so we will discuss these processes in more detail first, then move on to the definitions.

Splitting along spheres Let $S \subset M_{n,b}$ be an embedded 2-sphere. By *splitting along* S, we mean removing an open tubular neighborhood N of S from $M_{n,b}$. If S is nonseparating, the resulting manifold will be diffeomorphic to $M_{n-1,b+2}$ and if S is separating, the result will be diffeomorphic to $M_{m_1,c_1} \sqcup M_{m_2,c_2}$, where $m_1 + m_2 = n$ and $c_1 + c_2 = b + 2$. Notice that the resulting manifold is a submanifold of $M_{n,b}$, and so we get a corresponding map $\mathrm{Mod}(M_{n-1,b+2}) \to \mathrm{Mod}(M_{n,b})$ if S is nonseparating, or $\mathrm{Mod}(M_{m_1,c_1}) \times \mathrm{Mod}(M_{m_2,c_2}) \to \mathrm{Mod}(M_{n,b})$ if S is separating. In either case, this map sends sphere twists to sphere twists, and thus induces a map $\iota_*\colon \mathrm{Out}(F_{n-1,b+2}) \to \mathrm{Out}(F_{n,b})$ or $\iota_*\colon \mathrm{Out}(F_{m_1,c_1}) \times \mathrm{Out}(F_{m_2,c_2}) \to \mathrm{Out}(F_{n,b})$, depending on whether or not S separates $M_{n,b}$.

Dragging boundary components Let ∂ be a boundary component of $M_{n,b}$, and let $\iota\colon M_{n,b} \hookrightarrow M_{n,b-1}$ be the embedding obtained by capping off ∂. By Theorem 4.2, we have an exact sequence

$$1 \to \pi_1(M_{n,b-1}, x) \xrightarrow{\text{Push}} \mathrm{Out}(F_{n,b}) \xrightarrow{\iota_*} \mathrm{Out}(F_{n,b-1}) \to 1,$$

where $x \in M_{n,b-1} \setminus M_{n,b}$. Given $\gamma \in \pi_1(M_{n,b-1}, x)$, recall that the element $\mathrm{Push}(\gamma) \in \mathrm{Out}(M_{n,b})$ is given by pushing ∂ about the loop γ. In the remainder of this section, we will be dragging multiple boundary components at a time. So, from now on we will write $\mathrm{Push}_\partial(\gamma)$ in order to keep track of which boundary component is being pushed.

Magnus generators We now move on to defining our generators for $\mathrm{IO}^P_{n,b}$. In the $b = 0$ case, we have that $\mathrm{IO}^P_{n,0} = \mathrm{IO}_n$, where IO_n is the subgroup of $\mathrm{Out}(F_n)$ acting trivially on homology. In [21], Magnus found the following generating set for IO_n.

Theorem 5.1 (Magnus) *Let $F_n = \langle x_1, \dots, x_n \rangle$. The group IO_n is generated by the $\mathrm{Out}(F_n)$-classes of the automorphisms*

$$M_{ij}\colon x_i \mapsto x_j x_i x_j^{-1}, \quad M_{ijk}\colon x_i \mapsto x_i [x_j, x_k],$$

for all distinct $i, j, k \in \{1, \dots, n\}$ with $j < k$. Here, the automorphisms are understood to fix x_ℓ for $\ell \neq i$.

Our generating set will be inspired by Magnus's, and will indeed reduce to it when $b = 0$. In order to choose a concrete collection of elements, we will need to make some choices. First, fix a basepoint $* \in \mathrm{int}(M_{n,b})$ and a collection of n disjointly embedded oriented 2-spheres $S_1, \dots, S_n \subset M_{n,b} \setminus *$. We will call such a collection a *sphere basis* for $M_{n,b}$. In addition, choose a corresponding *geometric free*

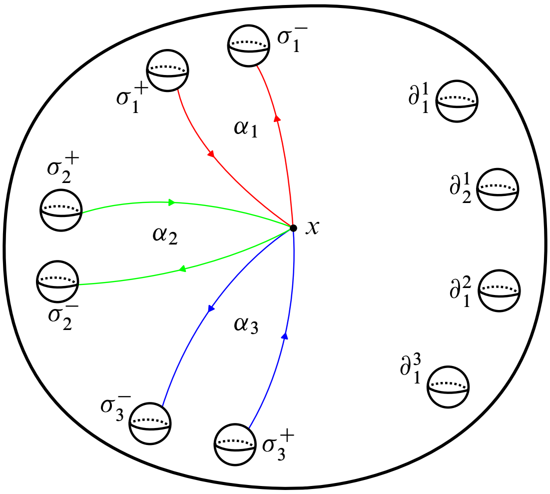

Figure 7: $M_{3,4}$ split along $S_1 \cup S_2 \cup S_3$ with the partition $P = \{\{\partial_1^1, \partial_2^1\}, \{\partial_1^2\}, \{\partial_1^3\}\}$.

basis; that is, a set $\{\alpha_1, \ldots, \alpha_n\}$ of oriented simple closed curves intersecting only at $*$ such that

- α_i intersects S_i exactly once with positive orientation for all i,
- α_i is disjoint from S_j if $i \neq j$.

Notice that the homotopy classes of $\{\alpha_1, \ldots, \alpha_n\}$ necessarily forms a free basis for $\pi_1(M_{n,b}, *)$. Splitting $M_{n,b}$ along the S_i reduces it to a 3-sphere $\mathcal{Z} \subset M_{n,b}$ with $b + 2n$ boundary components. The submanifold $\mathcal{Z}$ will play a significant role throughout the remainder of this section because it will allow all of our choices made in the definitions to be unique. For each S_i, let σ_i^+ and σ_i^- be the boundary components of $\mathcal{Z}$ arising from the split along S_i, where σ_i^+ (resp. σ_i^-) is the component lying on the positive (resp. negative) side of S_i. We will also choose an ordering $P = \{p_1, \ldots, p_{|P|}\}$ and an ordering $p_r = \{\partial_1^r, \ldots, \partial_{b_r}^r\}$ for each $r \in \{1, \ldots, |P|\}$. See Figure 7.

The following lemma will be helpful in showing that our generators lie in $\mathrm{IO}_{n,b}^P$.

Lemma 5.2 *Let $\mathcal{Z}$ be as above, and suppose that $h \subset M_{n,b}$ is a properly embedded oriented arc connecting P-adjacent boundary components of $M_{n,b}$. Then the homology class of $[h] \in H_1^P(M_{n,b})$ has the form*

$$[h] = [\alpha] + [h_0],$$

where α is a loop based at $*$, and h_0 is the unique arc (up to isotopy) in $\mathcal{Z}$ which has the same endpoints as h.

Proof We may homotope h such that it has the form $h = \gamma_s \cdot \alpha \cdot \gamma_e$, where (see Figure 8)

- $\gamma_s \subset \mathcal{Z}$ is the unique arc (up to isotopy) from the initial point of h to the basepoint $*$ of $M_{n,b}$,

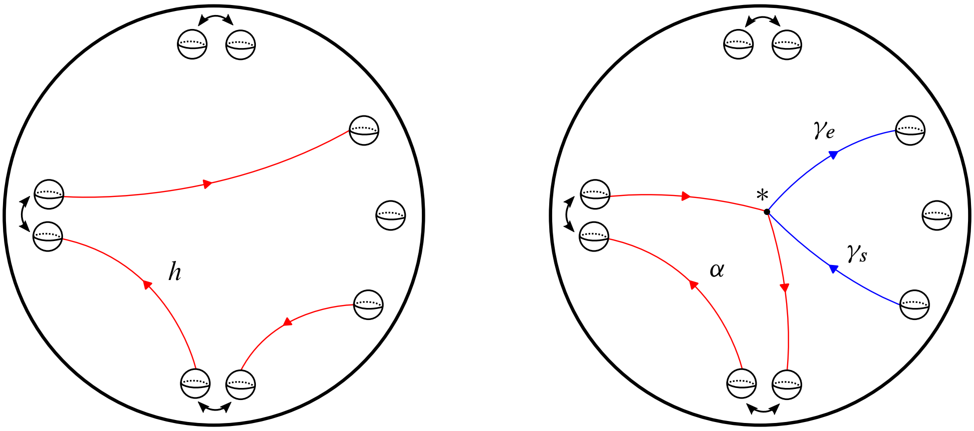

Figure 8: The arc h homotoped to be put in the form $\gamma_s \cdot \alpha \cdot \gamma_e$.

- $\gamma_e \subset \mathcal{Z}$ is the unique arc from $*$ to the endpoint of h,
- $\alpha \in \pi_1(M_{n,b}, *)$.

Then

$$[h] = [\gamma_s \cdot \alpha \cdot \gamma_e] = [\alpha] + [\gamma_s \cdot \gamma_e] = [\alpha] + [h_0],$$

as desired. □

Handle drags Let $i \in \{1, \dots, n\}$, and let h_i be the unique (up to isotopy) properly embedded arc in $\mathcal{Z}$ connecting σ_i^+ and σ_i^- which is disjoint from the α_k. Choose a tubular neighborhood N_i of $\sigma_i^+ \cup h_i \cup \sigma_i^-$ that does not intersect any α_k for $k \neq i$. Let Σ_i be the boundary component of N_i which is not isotopic to σ_i^+ or σ_i^- (notice that Σ_i is diffeomorphic to a 2-sphere). Splitting $M_{n,b}$ along Σ_i yields $M_{n-1,b+1} \sqcup M_{1,1}$. Let $\Sigma_i' \subset \partial M_{n-1,b+1}$ be the boundary component coming from this split, and fix a basepoint $y_i \in \Sigma_i'$. Fix an oriented arc $\delta_i \subset Z$ from y_i to $*$ which only intersects Σ_i' at y_i. Since $\mathcal{Z}$ is a 3-sphere with spherical boundary components, δ_i is unique up to isotopy. The arc δ_i induces an isomorphism $\pi_1(M_{n-1,b+1}, *) \to \pi_1(M_{n-1,b+1}, y_i)$ given by $\gamma \mapsto \delta_i \gamma \delta_i^{-1}$. Define $\beta_j^i = \delta_i \alpha_j \delta_i^{-1}$. Then we define the *handle drag* $\mathrm{HD}_{ij} := \iota_*(\mathrm{Push}_{\Sigma_i'}(\beta_j^i), \mathrm{id}) \in \mathrm{Out}(F_{n,b})$ for $j \neq i$, where ι_* is the map $\mathrm{Out}(F_{n-1,b+1}) \times \mathrm{Out}(F_{1,1}) \to \mathrm{Out}(F_{n,b})$ induced by splitting along Σ_i.

To see that $\mathrm{HD}_{ij} \in \mathrm{IO}_{n,b}^P$, notice that HD_{ij} acts trivially on α_k for $k \neq i$, and acts on α_i via $\alpha_i \mapsto \alpha_j \alpha_i \alpha_j^{-1}$. See Figure 9. This shows that HD_{ij} acts trivially on homology classes of simple closed curves. Additionally, this shows that HD_{ij} reduces to M_{ij} of the Magnus generators if $b = 0$.

Next, suppose that h is an arc connecting P-adjacent boundary components. By Lemma 5.2, we may write $[h] = [\alpha] + [h_0]$, where α is a loop based at $*$, and h_0 is the unique arc (up to isotopy) in $\mathcal{Z}$ which has the same endpoints as h. We have seen that HD_{ij} fixes the homology class of α. Moreover, we may homotope HD_{ij} such that it fixes the arc h_0. Thus, HD_{ij} fixes the homology class of h, and we conclude that $\mathrm{HD}_{ij} \in \mathrm{IO}_{n,b}^P$.

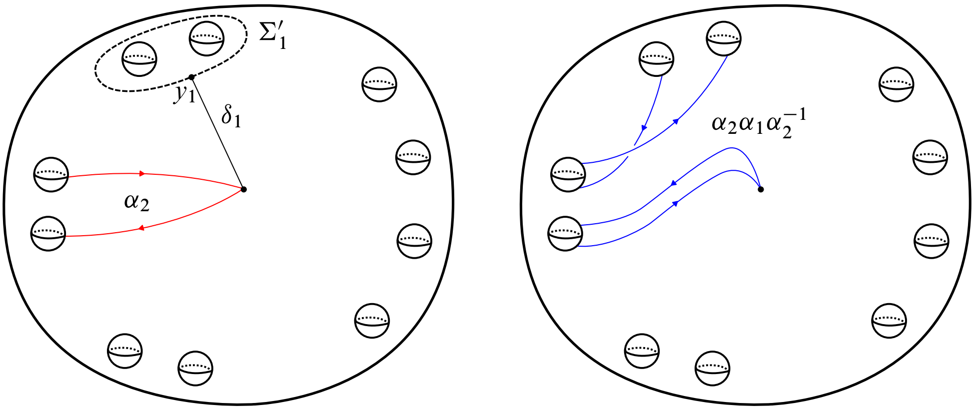

Figure 9: Setup of the handle drag HD_{12} and the image of α_1 under HD_{12}

Commutator drags Let $i, j, k \in \{1, \dots, n\}$ be distinct with $j < k$. Split $M_{n,b}$ along S_i to get $M_{n,b+2}$, where $\mathcal{Z} \subset M_{n,b+2} \subset M_{n,b}$. Fix basepoint $y_i \in \sigma_i^+$ and $z_i \in \sigma_i^-$, and choose oriented arcs $\delta_i, \varepsilon_i \subset \mathcal{Z}$ connecting y_i and z_i to $*$, respectively. Just as in the construction of handle drags, δ_i and ε_i are unique up to isotopy. Let $\beta_\ell^i = \delta_i \alpha_\ell \delta_i^{-1}$ and $\gamma_\ell^i = \varepsilon_i \alpha_\ell \varepsilon_i^{-1}$ for $\ell \in \{j, m\}$. Then, we define the *commutator drags* $\mathrm{CD}_{ijk}^+, \mathrm{CD}_{ijk}^- \in \mathrm{Out}(F_{n,b})$ as $\iota_*(\mathrm{Push}_{\sigma_i^+}([\beta_j^i, \beta_k^i]))$ and $\iota_*(\mathrm{Push}_{\sigma_i^-}([\gamma_j^i, \gamma_k^i]))$, respectively, where $\iota_* \colon \mathrm{Out}(F_{n,b+2}) \to \mathrm{Out}(F_{n,b})$ is the map induced by splitting along S_i. See Figure 10.

Again, we see that $\mathrm{CD}_{ijk}^{\pm}$ acts trivially on α_ℓ for $\ell \neq i$, the commutator drag CD_{ijk}^+ sends α_i to $\alpha_i[\alpha_j, \alpha_k]^{-1}$, and CD_{ijk}^- sends α_i to $[\alpha_j, \alpha_k]\alpha_i$. This shows that CD_{ijk}^+ reduces to M_{ijk}^{-1} of the Magnus generators when $b = 0$.

Now, suppose that h is an arc connecting P-adjacent boundary components of $M_{n,b}$. By Lemma 5.2, we may express $[h]$ in the form $[h] = [\alpha] + [h_0]$. We just saw that $\mathrm{CD}_{ijk}^{\pm}$ fixes $[\alpha]$. We may also homotope $\mathrm{CD}_{ijk}^{\pm}$ such that it fixes h_0. Thus, $\mathrm{CD}_{ijk}^{\pm}$ fixes $[h]$, and so $\mathrm{CD}_{ijk}^{\pm} \in \mathrm{IO}_{n,b}^P$.

Boundary commutator drags Let $p_r \in P$ and $\partial_s^r \in p_r$. Fix $i, j \in \{1, \dots, n\}$ such that $i < j$. Choose a basepoint $y_s^r \in \partial_s^r$. Let $\gamma_s^r \subset \mathcal{Z}$ be the unique arc (up to isotopy) from y_s^r to $*$. Let $\beta_k^{rs} = \gamma_s^r \alpha_k (\gamma_s^r)^{-1}$ for $k \in \{i, j\}$. Then, we define the *boundary commutator drags* $\mathrm{BCD}_{ij}^{rs} = \mathrm{Push}_{\partial_s^r}([\beta_i^{rs}, \beta_j^{rs}]) \in \mathrm{Out}(F_{n,b})$.

It is clear from the definition that BCD_{ij}^{rs} acts trivially on $\alpha_1, \dots, \alpha_n$ and arcs that do not have an endpoint on ∂_s^r. Suppose that h is an oriented arc with an endpoint on ∂_s^r. Without loss of generality, suppose the terminal endpoint of h lies on ∂_s^r. Applying Lemma 5.2, we may write $[h] = [\alpha] + [h_0]$, where α is a loop based at $*$ and $h_0 \subset \mathcal{Z}$ is the unique arc (up to isotopy) which shares endpoints with h. We just saw that BCD_{ij}^{rs} fixes the α_k, and thus fixes the homology class $[\alpha]$. Therefore,

$$\begin{aligned}\mathrm{BCD}_{j\ell m}([h]) = \mathrm{BCD}_{ij}^{rs}([\alpha] + [h_0]) = [\alpha] + \mathrm{BCD}_{ij}^{rs}([h_0]) &= [\alpha] + [h_0] + [\alpha_i \cdot \alpha_j \cdot \alpha_i^{-1} \cdot \alpha_j^{-1}] \\ &= [\alpha] + [h_0] = [h].\end{aligned}$$

So, it follows that $\mathrm{BCD}_{ij}^{rs} \in \mathrm{IO}_{n,b}^P$ as well.

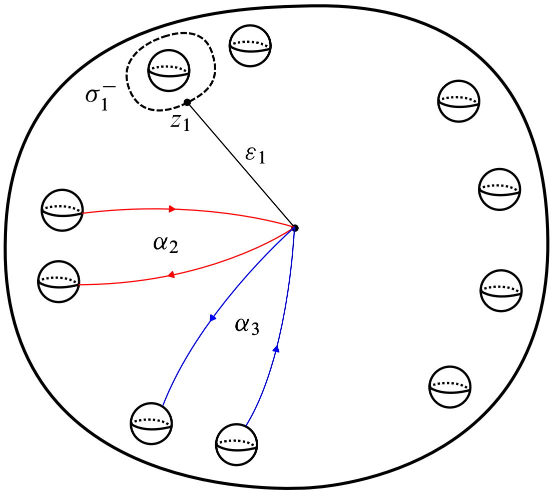

Figure 10: Setup of the commutator drag CD_{123}^-.

P**-drags** The final type of elements we will define are called P-drags, where P is a partition of the boundary components of $M_{n,b}$. Let $p_r \in P$ and $i \in \{1,\dots,n\}$. Let $\Sigma_r \subset \mathcal{Z}$ be the unique 2-sphere (up to isotopy) which separates the boundary components of p_r from the remaining boundary components and the $\sigma_j^{\pm}$. Splitting $M_{n,b}$ along Σ_r gives $M_{n,b-c+1} \sqcup M_{0,c+1}$, where c is the number of boundary components in p. Let $\Sigma_r' \subset \partial M_{n,b-c+1}$ be the boundary component coming from this splitting. Just as in the construction of the other drags, fix a basepoint $y_r \in \Sigma_r'$ and an oriented arc γ_r from y_r to $*$ to get a basis $\{\beta_1^r,\dots,\beta_n^r\}$ of $\pi_1(M_{n,b-c+1},y_r)$. See Figure 11. Then, we define the *P-drag* $\mathrm{PD}_i^r := \iota_*(\mathrm{Push}_{\Sigma_r'}(\beta_i),\mathrm{id}) \in \mathrm{Out}(F_{n,b})$, where $\iota_*\colon \mathrm{Out}(F_{n,b-c+1}) \times \mathrm{Out}(F_{0,c+1}) \to \mathrm{Out}(F_{n,b})$ is the map induced by splitting along Σ_r.

To see why $\mathrm{PD}_i^r \in \mathrm{IO}_{n,b}^P$, first notice that we can isotope PD_i^r to fix all the α_j. Next, if h is an arc connecting P-adjacent boundary components, we write $[h] = [\alpha] + [h_0]$ as in Lemma 5.2. As we just noted, PD_i^r fixes $[\alpha]$, so it suffices to show that PD_i^r fixes the homology class of h_0. If the endpoints of h lie on boundary components in p_r, then we may homotope h_0 such that it never crosses Σ_r. Then, PD_i^r fixes h_0. On the other hand, if the endpoints of h lie on boundary components which are not in p_r, then again we can homotope h_0 such that it does not cross Σ_r, and then homotope PD_i^r such that it fixes h_0. In either case, PD_i^r fixes the homology class of h_0, and so we conclude that $\mathrm{PD}_i^r \in \mathrm{IO}_{n,b}^P$.

Images under capping Suppose we have an embedding $\iota\colon M_{n,b} \hookrightarrow M_{n,b-1}$ given by capping off the boundary component ∂. Let $\iota_*\colon \mathrm{IO}_{n,b}^P \to \mathrm{IO}_{n,b-1}^{P'}$ be the induced map, where P' is the partition of the boundary components of $M_{n,b-1}$ induced by P. Using the sphere basis $\{S_1,\dots,S_n\}$ and geometric free basis $\{\alpha_1,\dots,\alpha_n\}$ for $M_{n,b}$, we get a corresponding sphere basis $\{\iota(S_1),\dots,\iota(S_n)\}$ and geometric free basis $\{\iota(\alpha_1),\dots,\iota(\alpha_n)\}$ for $M_{n,b-1}$. Moreover, the ordering on P (and each $p_r \in P$) induces an ordering on P'. We can repeat the process described throughout this section to define handle drags,

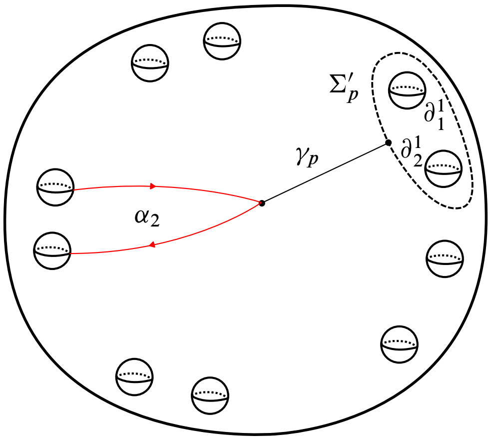

Figure 11: Setup of the P-drag PD_2^p, where $p = \{\partial_1^1, \partial_2^1\} \in P$.

commutator drags, boundary commutator drag, and P'-drags in $\mathrm{IO}_{n,b-1}^{P'}$, which we will denote by $\overline{\mathrm{HD}}_{ij}$, $\overline{\mathrm{CD}}_{ijk}^{\pm}$, $\overline{\mathrm{BCD}}_{ij}^{rs}$, and $\overline{\mathrm{PD}}_i^{r'}$, respectively. With this setup, we find that

$$\begin{aligned}
\iota_*(\mathrm{HD}_{ij}) &= \overline{\mathrm{HD}}_{ij},\\
\iota_*(\mathrm{CD}_{ijk}^{\pm}) &= \overline{\mathrm{CD}}_{ijk}^{\pm},\\
\iota_*(\mathrm{BCD}_{ij}^{rs}) &= \begin{cases} \mathrm{id} & \text{if } \partial = \partial_s^r,\\ \overline{\mathrm{BCD}}_{ij}^{rs} & \text{otherwise,}\end{cases}\\
\iota_*(\mathrm{PD}_i^r) &= \begin{cases} \mathrm{id} & \text{if } p_r = \{\partial_1^r\},\\ \overline{\mathrm{PD}}_i^r & \text{otherwise.}\end{cases}
\end{aligned}$$

6 Finite generation

Now that we have defined our collection of candidate generators for $\mathrm{IO}_{n,b}^{P}$, we now move on to proving that they do in fact generate. The first step in this proof will be an induction on b to reduce to the case of $b = 0$. This induction will rely on the following theorem of Tomaszewski [25] (see [24] for a geometric proof).

Theorem 6.1 (Tomaszewski) *Let F_n be the free group on n letters $\{x_1, \dots, x_n\}$. The commutator subgroup $[F_n, F_n]$ of F_n is freely generated by the set*

$$\big\{[x_i, x_j]^{x_i^{d_i} \cdots x_n^{d_n}},\ 1 \le i < j \le n,\ d_\ell \in \mathbb{Z},\ i \le \ell \le n\big\},$$

where the superscript denotes conjugation.

We will also need the following lemma from group theory.

Lemma 6.2 *Consider an exact sequence of groups*

$$1 \to K \to G \to Q \to 1.$$

Let S_Q be a generating set for Q. Moreover, assume that there are sets $S_K \subset K$ and $S_G \subset G$ such that K is contained in the subgroup of G generated by S_K and S_G. Then G is generated by the set $S_K \cup S_G \cup \widetilde{S}_Q$, where $\widetilde{S}_Q$ is a set consisting of one lift $\widetilde{q} \in G$ for every element $q \in S_q$.

Proof of lemma Let $G' \subset G$ be the subgroup generated by $S_K \cup S_G \cup \widetilde{S}_Q$, and let $K' = G' \cap K$. Then the following diagram commutes and has exact rows:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & K' & \longrightarrow & G' & \longrightarrow & Q & \longrightarrow & 1 \\ & & \downarrow{\scriptstyle\varphi} & & \downarrow & & \| & & \\ 1 & \longrightarrow & K & \longrightarrow & G & \longrightarrow & Q & \longrightarrow & 1 \end{array}$$

The vertical maps are all inclusions, and hence injective. Also, by assumption, the map φ is surjective. Therefore, by the five lemma, all of the vertical maps are isomorphisms, and so we are done. □

We now prove the following result, which implies Theorem C since the collections of handle, commutator, boundary commutator, and P-drags are all finite. In Section 8, we will compute the number of each type of drag, and show how many become redundant in the abelianization of IA_n.

Theorem 6.3 *The group $\mathrm{IO}^P_{n,b}$ is generated by handle, commutator, boundary commutator, and P-drags for $b \geq 0$, $n > 0$.*

Proof As mentioned above, we will prove this by induction on b. The base case $b = 0$ follows directly from Magnus's Theorem 1.1.

If $b > 0$, fix a boundary component ∂ of $M_{n,b}$ and let $p \in P$ be the partition containing ∂. Let $\iota\colon M_{n,b} \hookrightarrow M_{n,b-1}$ be an embedding obtained by capping off ∂, and choose a basepoint $* \in M_{n,b-1} \setminus M_{n,b}$. By Theorem B, there is an exact sequence

$$1 \to L \xrightarrow{\text{Push}} \mathrm{IO}^P_{n,b} \xrightarrow{\iota_*} \mathrm{IO}^{P'}_{n,b-1} \to 1,$$

where $L = \pi_1(M_{n,b}, *)$ if $p = \{\partial\}$ and $L = [\pi_1(M_{n,b}, *), \pi_1(M_{n,b}, *)]$ otherwise. As we saw in the discussion at the end of Section 5, we can define the drags of $\mathrm{IO}^P_{n,b}$ and $\mathrm{IO}^{P'}_{n,b-1}$ in a consistent way; that is, we can define our drags in such a way that ι_* takes handle drags to handle drags, commutator drags to commutator drags, and so on. By induction, $\mathrm{IO}^{P'}_{n,b-1}$ is generated by the desired drags. Therefore, it suffices to show that $\mathrm{Push}(L)$ is generated by our drags as well. If $p = \{\partial\}$, then $\mathrm{Push}(L)$ is precisely the subgroup of $\mathrm{IO}^P_{n,b}$ generated by the P-drags, and so we are done in this case.

The case of $p \neq \{\partial\}$ is less straightforward since the commutator subgroup of a free group is not finitely generated when $n \geq 2$. However, this is not necessary for $\mathrm{IO}^P_{n,b}$ to be finitely generated by our collection of drags. We will appeal to Lemma 6.2. Suppose that $p \neq \{\partial\}$. Then, by Theorem 6.1, the kernel $L = [\pi_1(M_{n,b}, *), \pi_1(M_{n,b}, *)]$ of the Birman exact sequence is generated by elements of the form $[x_i, x_j]^{x_i^{d_i} \cdots x_n^{d_n}}$. First, notice that $\mathrm{Push}([x_i, x_j])$ is the boundary commutator drag BCD^{rs}_{ij}, where $\partial^r_s = \partial$

is the boundary component of $M_{n,b}$ being capped off. Moreover, we have seen that the handle drag $\mathrm{HD}_{k\ell}$ acts on x_k by $x_k \mapsto x_\ell x_k x_\ell^{-1}$. It follows that $\mathrm{HD}_{ik} \cdot \mathrm{HD}_{jk}([x_i, x_j]) = [x_i, x_j]^{x_k}$. Continuing this pattern, we see that

$$[x_i, x_j]^{x_i^{d_i} \cdots x_n^{d_n}} = (\mathrm{HD}_{in} \cdot \mathrm{HD}_{jn})^{d_n} \cdots (\mathrm{HD}_{ii} \cdot \mathrm{HD}_{ji})^{d_i}([x_i, x_j]),$$

where HD_{ii} is taken to be trivial. Let $H = (\mathrm{HD}_{in} \cdot \mathrm{HD}_{jn})^{d_n} \cdots (\mathrm{HD}_{ii} \cdot \mathrm{HD}_{ji})^{d_i}$. Then,

$$\mathrm{Push}([x_i, x_k]^{x_i^{d_i} \cdots x_n^{d_n}}) = \mathrm{Push}(H([x_i, x_j])) = H \cdot \mathrm{Push}([x_i, x_j]) \cdot H^{-1} = H \cdot \mathrm{BCD}_{ij}^{rs} \cdot H^{-1}.$$

This shows that $\mathrm{Push}(L)$ is contained in the subgroup of $\mathrm{IO}_{n,b}^P$ generated by boundary commutator and handle drags. Applying Lemma 6.2 (taking $S_G = \{\text{handle drags}\}$ and $S_K = \{\mathrm{BCD}_{i,j}^{r,s} \mid 1 \le i, j \le n\}$), we conclude that $\mathrm{IO}_{n,b}^P$ is generated by the desired drags. □

7 Partial proof of Magnus's theorem

In this section, we will give a partial proof of Magnus's Theorem 1.1, which constituted the base case in the proof of Theorem 6.3. As stated in the introduction, the original proof of Magnus's theorem involved two steps: showing that the elements M_{ij} and M_{ijk} *normally* generate IO_n, and then showing that the subgroup generated by these elements is normal. We will give a proof of the first step here (Theorem D).

In order to establish this fact, we will examine the action of $\mathrm{IO}_{n,0}^{\{\}} = \mathrm{IO}_n$ on a certain simplicial complex, and apply the following theorem of Armstrong [2]. We say that a group G acts on a simplicial complex X *without rotations* if every simplex s is fixed pointwise by every element of its stabilizer, which we will denote by G_s.

Theorem 7.1 (Armstrong) *Suppose the group G acts on a simply connected simplicial complex X without rotations. If X/G is simply connected, then G is generated by the set*

$$\bigcup_{v \in X^{(0)}} G_v.$$

Here $X^{(0)}$ is the 0-skeleton of X.

Remark In [2], Armstrong proves the converse of this theorem as well. For a modern discussion of the proof of Theorem 7.1, along with some generalizations, we refer the reader to [23, Section 3].

The nonseparating sphere complex The complex to which we will apply Theorem 7.1 will be the *nonseparating sphere complex* $\mathbb{S}_n^{\mathrm{ns}}$. Vertices of $\mathbb{S}_n^{\mathrm{ns}}$ are isotopy classes of smoothly embedded nonseparating 2-spheres in M_n, and $\mathbb{S}_n^{\mathrm{ns}}$ has a k-simplex $\{S_0, \ldots, S_k\}$ if the spheres $S_0, \ldots, S_k$ can be realized pairwise disjointly and their union does not separate M_n. This is a subcomplex of the more ubiquitous *sphere complex*, which was introduced by Hatcher in [13] as a tool to explore the homological stability of $\mathrm{Out}(F_n)$ and $\mathrm{Aut}(F_n)$. In [13, Proposition 3.1], Hatcher proves the following connectivity result about $\mathbb{S}_n^{\mathrm{ns}}$.

Proposition 7.2 (Hatcher) *The complex* $\mathbb{S}_n^{\mathrm{ns}}$ *is* $(n-2)$*-connected.*

In particular, $\mathbb{S}_n^{\mathrm{ns}}$ is simply connected for $n \geq 3$. Recall that sphere twists act trivially on isotopy classes of embedded surfaces, and so we get an action of IO_n on $\mathbb{S}_n^{\mathrm{ns}}$. Notice that spheres in a simplex of $\mathbb{S}_n^{\mathrm{ns}}$ necessarily represent distinct H_2-classes. By Poincaré duality, elements of IO_n act trivially on $H_2(M_n)$, and so this implies that IO_n acts on $\mathbb{S}_n^{\mathrm{ns}}$ without rotations. Thus, in order to apply Theorem 7.1, we must show that $\mathbb{S}_n^{\mathrm{ns}}/\mathrm{IO}_n$ is simply connected.

To do this, we will give a description of $\mathbb{S}_n^{\mathrm{ns}}/\mathrm{IO}_n$ in terms of linear algebra. Fix an identification $H_2(M_n) = \mathbb{Z}^n$. Let $\mathrm{FS}(\mathbb{Z}^n)$ be the simplicial complex whose vertices are rank 1 summands of $\mathbb{Z}^n$, and there is a ℓ-simplex $\{A_0, \dots, A_\ell\}$ if $A_0 \oplus \cdots \oplus A_\ell$ is a rank $\ell + 1$ summand of $\mathbb{Z}^n$. There is a map $\varphi\colon \mathbb{S}_n^{\mathrm{ns}}/\mathrm{IO}_n \to \mathrm{FS}(\mathbb{Z}^n)$ defined as follows. Let $s \in \mathbb{S}_n^{\mathrm{ns}}/\mathrm{IO}_n$ be a vertex, and choose a sphere $S \subset M_n$ which represents s. As noted above, elements of IO_n act trivially on $H_2(M_n)$. Therefore, the homology class $[S] \in H_2(M_n)$ does not depend on the choice of representative S. We then define $\varphi(s)$ to be the span of $[S]$ in $H_2(M_n)$. It is clear that φ extends to simplices.

Lemma 7.3 *The map* $\varphi\colon \mathbb{S}_n^{\mathrm{ns}}/\mathrm{IO}_n \to \mathrm{FS}(\mathbb{Z}^n)$ *is an isomorphism of simplicial complexes.*

Proof Let $\sigma = \{A_0, \dots, A_\ell\}$ be an ℓ-simplex of $\mathrm{FS}(\mathbb{Z}^n)$. We must show that, up to the action of IO_n, there exists a unique ℓ-simplex $\widetilde{\sigma}$ of $\mathbb{S}_n^{\mathrm{ns}}$ which projects to σ.

Let $v_j \in H_2(M_n)$ be a primitive element generating A_j for $0 \leq j \leq \ell$, and extend this to a basis $\{v_0, \dots, v_{n-1}\}$ for $H_2(M_{n,b}) = \mathbb{Z}^n$. In Appendix B, we will prove Lemma B.2, which says that there exists a collection $\{S_0, \dots, S_{n-1}\}$ of disjoint embedded 2-spheres such that $[S_j] = v_j$ for $0 \leq j \leq n-1$. Then the simplex $\widetilde{\sigma} = \{S_0, \dots, S_\ell\}$ of $\mathbb{S}_n^{\mathrm{ns}}$ maps to the σ under the composition

$$\mathbb{S}_n^{\mathrm{ns}} \to \mathbb{S}_n^{\mathrm{ns}}/\mathrm{IO}_n \xrightarrow{\varphi} \mathrm{FS}(\mathbb{Z}^n).$$

We will now show that $\widetilde{\sigma}$ is unique up to the action of IO_n. Suppose that $\widetilde{\sigma}' = \{S_0', \dots, S_\ell'\}$ is another simplex of $\mathbb{S}_n^{\mathrm{ns}}$ which projects to σ. Since $\widetilde{\sigma}$ and $\widetilde{\sigma}'$ both project to σ, we may order and orient the spheres such that $[S_j] = [S_j']$ for $0 \leq j \leq \ell$. Again by Lemma B.2, we can extend $\{S_1, \dots, S_\ell\}$ and $\{S_1', \dots, S_\ell'\}$ to collections of spheres $\{S_1, \dots, S_n\}$ and $\{S_1', \dots, S_n'\}$ such that $[S_j] = [S_j'] = v_j$ for $0 \leq j \leq n-1$. Notice that splitting M_n along either of these collections reduces M_n to a sphere with $2n$ boundary components. Therefore, there exists some $\mathfrak{f} \in \mathrm{Mod}(M_n)$ such that $\mathfrak{f}(S_j) = S_j'$ for all j. Let $f \in \mathrm{Out}(F_n)$ be the image of f. By construction, $f(\widetilde{\sigma}) = \widetilde{\sigma}'$. Furthermore, f fixes a basis for homology, and so $f \in \mathrm{IO}_n$. This completes the proof. □

This description of $\mathbb{S}_n^{\mathrm{ns}}/\mathrm{IO}_n$ is advantageous because $\mathrm{FS}(\mathbb{Z}^n)$ is known to be $(n-2)$-connected, and hence simply connected for $n \geq 3$. The first proof of this fact is due to Maazen [20] in his unpublished thesis (see [8, Theorem E] for a published proof). Thus, we have shown that $\mathbb{S}_n^{\mathrm{ns}}/\mathrm{IO}_n$ is sufficiently connected.

Corollary 7.4 (Maazen) *The complex* $\mathbb{S}_n^{\mathrm{ns}}/\mathrm{IO}_n$ *is simply connected for* $n \geq 3$.

As indicated in Theorem 7.1, the stabilizers of spheres play an important role in the proof of Theorem D, and so we introduce notation for them here. If S is an isotopy class of embedded sphere in M_n, we denote by $\mathrm{Out}(F_n, S)$ the stabilizer of S in $\mathrm{Out}(F_n)$, and define $\mathrm{IO}_n(S) = \mathrm{Out}(F_n, S) \cap \mathrm{IO}_n$. We now move on to the proof of Theorem D.

Proof of Theorem D We will induct on n. The base cases are easy; IO_1 and IO_2 are both trivial. Suppose now that IO_{n-1} is $\mathrm{Out}(F_{n-1})$-normally generated by handle and commutator drags. We must now show that IO_n is $\mathrm{Out}(F_n)$-normally generated by handle and commutator drags as well. By Theorem 7.1, Proposition 7.2, and Corollary 7.4, it suffices to show that $\mathrm{IO}_n(S)$ is generated by $\mathrm{Out}(F_n)$-conjugates of these drags for all S. Let $S_1, \dots, S_n \subset M_{n,b}$ be a sphere basis, and choose a corresponding geometric free basis $\{\alpha_1, \dots, \alpha_n\}$. Identify the homotopy classes of $\alpha_1, \dots, \alpha_n$ with our fixed basis $x_1, \dots, x_n$ for F_n. Use these bases to construct the handle and commutator drags as in Section 5. Recall that handle drags correspond to the automorphisms M_{ij} of Magnus's generators, and commutator drags correspond to M_{ijk}. We will first show that $\mathrm{IO}_n(S_1)$ is $\mathrm{Out}(F_n, S_1)$-normally generated by handle and commutator drags.

Splitting M_n along S_1 yields a copy of $M_{n-1,2}$. Let N be the tubular neighborhood of S_1 removed in this splitting, and let ∂_1 and ∂_2 be the boundary components of $M_{n-1,2}$. Then this splitting induces a surjective map $\mathrm{Out}(F_{n-1,2}) \to \mathrm{Out}(F_n, S_1)$, which restricts to a map $\iota_* \colon \mathrm{IO}^P_{n-1,2} \to \mathrm{IO}_n(S_1)$, where $P = \{p_1\} = \{\{\partial_1, \partial_2\}\}$. This map is also surjective.

Use the bases $\{\alpha_2, \dots, \alpha_n\}$ and $\{S_2, \dots, S_n\}$ to construct the handle, commutator, boundary commutator, and P-drags in $\mathrm{IO}^P_{n-1,2}$. By our induction hypothesis combined with the proof of Theorem 6.3, these drags $\mathrm{Out}(F_{n-1,2})$-normally generate $\mathrm{IO}^P_{n-1,2}$. Notice that with these choices of drags, the map ι_* takes handle and commutator drags to handle and commutator drags. Moreover, ι_* takes boundary commutator drags in $\mathrm{IO}^P_{n-1,2}$ to commutator drags in $\mathrm{IO}_n(S)$, and takes the P-drag PD^P_i to the handle drag HD_{1i}. Thus, $\mathrm{IO}_n(S_1)$ is $\mathrm{Out}(F_n, S_1)$-normally generated by handle and commutator drags.

Finally, let S be an arbitrary vertex of $\mathbb{S}^{\mathrm{ns}}_n$. Since S is nonseparating, there exists some $f \in \mathrm{Out}(F_n)$ such that $f(S_1) = S$. It follows that

$$\mathrm{IO}_n(S) = f \cdot \mathrm{IO}_n(S_1) \cdot f^{-1}.$$

Since $\mathrm{IO}_n(S_1)$ is $\mathrm{Out}(F_n, S_1)$-normally generated by handle and commutator drags, it follows that $\mathrm{IO}_n(S)$ is generated by $\mathrm{Out}(F_n)$-conjugates of handle and commutator drags, which is what we wanted to show. □

8 Computing the abelianization

In this section, we compute the abelianization of the group $\mathrm{IO}^P_{n,b}$, proving Theorem E. For the Torelli subgroup of the mapping class group of a surface, this was done by Johnson [17]. Some key tools used in this computation are the *Johnson homomorphisms*

$$\tau_{\Sigma_{g,1}} \colon \mathcal{I}(\Sigma_{g,1}) \to \textstyle\bigwedge^3 H \quad \text{and} \quad \tau_{\Sigma_g} \colon \mathcal{I}(\Sigma_g) \to \big(\textstyle\bigwedge^3 H\big)/H,$$

where $H = H_1(\Sigma_{g,b})$. Johnson showed that these homomorphisms are the abelianization maps modulo torsion if $g \geq 3$. For $\mathrm{IA}_n = \mathrm{IO}_{n,1}$, Andreadakis [1] and Bachmuth [3] used an analogous homomorphism $\tau\colon \mathrm{IA}_n \to \mathrm{Hom}\big(H, \bigwedge^2 H\big)$ (where now $H = H_1(F_n) = \mathbb{Z}^n$) to show that

$$H_1(\mathrm{IA}_n) \cong \mathrm{Hom}\big(H, \textstyle\bigwedge^2 H\big) \cong \mathbb{Z}^{n\cdot\binom{n}{2}}.$$

We will begin by recalling the definition of τ, along with the computation of the ranks of $H_1(\mathrm{IA}_n)$ and $H_1(\mathrm{IO}_n)$, and then proceed to the case of multiple boundary components.

The Johnson homomorphism Recall that $\mathrm{Out}(F_{n,1}) \cong \mathrm{Aut}(F_n)$, and the subgroup IA_n is precisely those automorphisms which act trivially on $H_1(F_n) = \mathbb{Z}^n$. The goal is to construct a homomorphism $\tau\colon \mathrm{IA}_n \to \mathrm{Hom}\big(H, \bigwedge^2 H\big)$, where $H = H_1(F_n) = \mathbb{Z}^n$.

First, we claim that the group $[F_n, F_n]/[F_n, [F_n, F_n]]$ is isomorphic to $\bigwedge^2 H$, where $[F_n, F_n]$ denotes the commutator subgroup of F_n. To see this, consider the short exact sequence

$$1 \to [F_n, F_n] \to F_n \to \mathbb{Z}^n \to 1.$$

Passing to the five-term exact sequence in homology, we get the sequence

$$0 \to H_2(\mathbb{Z}^n) \to H_1([F_n, F_n])_{\mathbb{Z}^n} \to H_1(F_n) \to H_1(\mathbb{Z}^n) \to 0,$$

where $H_1([F_n, F_n])_{\mathbb{Z}^n} = [F_n, F_n]/[F_n, [F_n, F_n]]$ denotes the group of coinvariants of $H_1([F_n, F_n])$ with respect to the action of $\mathbb{Z}^n$ (induced by the conjugation action of F_n on $[F_n, F_n]$). The map $H_1(F_n) \to H_1(\mathbb{Z}^n)$ is clearly an isomorphism, and so it follows that the map $H_2(\mathbb{Z}^n) \to [F_n, F_n]/[F_n, [F_n, F_n]]$ is an isomorphism as well. This proves our claim because $H_2(\mathbb{Z}^n) \cong \bigwedge^2 \mathbb{Z}^n$. Let $\rho\colon [F_n, F_n] \to \bigwedge^2 \mathbb{Z}^n$ be the projection. Following the definitions above, we see that ρ is defined by

$$\rho([x, y]) = [x] \wedge [y],$$

where $[x]$ and $[y]$ denote the classes of x and y in H, respectively.

Next, let $f \in \mathrm{IA}_n$. Then $f(x)x^{-1}$ is nullhomologous for all $x \in F_n$, and therefore lies in $[F_n, F_n]$. We define the map $\hat{\tau}_f\colon F_n \to \bigwedge^2 H$ via

$$\hat{\tau}_f(x) = \rho(f(x)x^{-1}).$$

We now check that $\hat{\tau}_f$ is a homomorphism. Let $x, y \in F_n$. Applying the relation $ab = [a, b]ba$, we get

$$\begin{aligned}\hat{\tau}_f(xy) &= \rho(f(x)f(y)y^{-1}x^{-1})\\ &= \rho\big([f(x), f(y)y^{-1}]\cdot(f(y)y^{-1})\cdot(f(x)x^{-1})\big)\\ &= [f(x)] \wedge [f(y)y^{-}1] + \hat{\tau}_f(y) + \hat{\tau}_f(y)\\ &= \hat{\tau}_f(y) + \hat{\tau}_f(y),\end{aligned}$$

since $[f(y)y^{-1}] = 0$. This shows that $\hat{\tau}$ is indeed a homomorphism. Furthermore, since $\bigwedge^2 H$ is abelian, the map $\hat{\tau}\colon F_n \to \bigwedge^2 H$ factors through the abelianization, inducing a map $\tau_f\colon H \to \bigwedge^2 H$. Therefore, we have a map

$$\tau\colon \mathrm{IA}_n \to \mathrm{Hom}\big(H, \textstyle\bigwedge^2 H\big)$$

sending f to τ_f. We now check that τ is a homomorphism. Let $f, g \in \mathrm{IA}_n$. Then

$$\tau_{fg}([x]) = \rho(f(g(x))x^{-1}) = \rho\big(f(g(x))(g(x))^{-1}g(x)x^{-1}\big) = \tau_f([g(x)]) + \tau_g([x]) = \tau_f([x]) + \tau_g([x])$$

since g fixes $[x]$. Thus, τ is the desired homomorphism.

We now move on to computing the image of our generators under τ. Since we are dealing with the case of one boundary component, boundary commutator drags are unnecessary since they are products of P-drags. Fix a basepoint $* \in \partial M_{n,1}$, and choose a basis $\{x_1, \dots, x_n\}$ of $\pi_1(M_{n,1}, *)$. Construct the handle, commutator, and P-drags with respect to this basis.

Handle drags Recall that the handle drag HD_{ij} acts on $\pi_1(M_{n,1})$ by sending x_i to $x_j x_i x_j^{-1}$, and fixing the remaining basis elements. Therefore,

$$\tau(\mathrm{HD}_{ij})([x_\ell]) = \rho(\mathrm{HD}_{ij}(x_\ell)x_\ell^{-1}) = \begin{cases} 0 & \text{if } \ell \neq i, \\ \rho(x_j x_i x_j^{-1} x_i^{-1}) & \text{if } \ell = i. \end{cases}$$

Thus, $\tau(\mathrm{HD}_{ij})$ is the homomorphism $[x_i] \mapsto [x_j] \wedge [x_i]$ (and all other generators are sent to 0).

Commutator drags Notice that the product of commutator drags $\mathrm{CD}^+_{ijk} \cdot \mathrm{CD}^-_{ijk}$ is equal to a commutator of handle drags. Therefore, only the CD^-_{ijk} are needed in our generating set, and we can disregard the CD^+_{ijk} from now on. Recall that CD^-_{ijk} acts on $\pi_1(M_{n,1})$ by sending x_i to $[x_j, x_k]x_i$. Therefore,

$$\tau(\mathrm{CD}^-_{ijk})([x_\ell]) = \rho(\mathrm{CD}^-_{ijk}(x_\ell)x_\ell^{-1}) = \begin{cases} 0 & \text{if } \ell \neq i, \\ \rho([x_j, x_k]) & \text{if } \ell = i. \end{cases}$$

It follows that $\tau(\mathrm{CD}^-_{ijk})$ is the map $[x_i] \mapsto [x_j] \wedge [x_k]$.

***P*-drags** Next, we note that the product

(5) $$\mathrm{PD}_j \cdot \mathrm{HD}_{1j} \cdots \mathrm{HD}_{nj}$$

is trivial in IA_n. For a justification of this fact, see the proof of the claim at the end of Theorem A.2. It follows that the P-drags are also redundant in our generating set for IA_n, and can be removed.

Abelianization of IA_n To compute the rank of the abelianization of IA_n, we use the following lemma.

Lemma 8.1 *Let G be a group and S a finite generating set for G. Suppose that $\varphi\colon G \to \mathbb{Z}^{|S|}$ is a surjective homomorphism. Then $H_1(G) \cong \mathbb{Z}^{|S|}$.*

Proof Let $F(S)$ denote the free group on the set S. Since $\mathbb{Z}^{|S|}$ is abelian, the homomorphism φ factors through the abelianization to give a map $\bar{\varphi}\colon H_1(G) \to \mathbb{Z}^n$, which is also surjective. Additionally, by the universal property of free groups, we have a map $\psi\colon F(S) \to G$. Passing to the abelianizations induces a map $\bar{\psi}\colon H_1(F(S)) \to H_1(G)$. Since S is a generating set for G, this map is also surjective. It follows that $\bar{\varphi} \circ \bar{\psi}$ is a surjective map between free abelian groups of equal rank, and is hence an isomorphism. Thus, $\bar{\varphi}$ is an isomorphism as well. □

From the discussion in the preceding paragraphs, we have a generating set for IA_n of size

$$\#(\text{handle drags}) + \#(\text{commutator drags}) = n(n-1) + n \cdot \binom{n-1}{2} = n \cdot \binom{n}{2}$$

(since P-drags can be written as a product of handle drags), and the image of this generating set spans $\mathrm{Hom}\big(H, \bigwedge^2 H\big)$, which also has dimension $n \cdot \binom{n}{2}$. Therefore, by Lemma 8.1, the group $H_1(\mathrm{IA}_n)$ has rank $n \cdot \binom{n}{2}$.

To compute the rank of $H_1(\mathrm{IO}_n)$, consider the quotient map $\mathrm{IA}_n \to \mathrm{IO}_n$, whose kernel is the subgroup of inner automorphisms (or P-drags under our geometric interpretation of IA_n). We compute the image of a P-drag under τ:

$$\tau(\mathrm{PD}_i)([x_\ell]) = \rho(\mathrm{PD}_i(x_\ell)x_\ell^{-1}) = \rho(x_i^{-1}x_\ell x_i x_\ell^{-1}) = [x_\ell] \wedge [x_i].$$

Since $\tau(\mathrm{PD}_i)$ is nontrivial, τ does not descend to a map $\mathrm{IO}_n \to \mathrm{Hom}\big(H, \bigwedge^2 H\big)$. However, the images $\{\tau(\mathrm{PD}_i)\}$ span a subgroup of $\mathrm{Hom}\big(H, \bigwedge^2 H\big)$ isomorphic to H (where the isomorphism is given by $[h] \mapsto ([x_\ell] \mapsto [x_\ell] \wedge [h])$). So, τ induces a map $\mathrm{IO}_n \to \mathrm{Hom}\big(H, \bigwedge^2 H\big)/H$. Just as the element given in (5) is trivial in IA_n, the product $\mathrm{HD}_{1j} \cdots \mathrm{HD}_{nj}$ is trivial in IO_n for all $j \in \{1, \ldots, n\}$. Thus, we may throw out n handle drags from our generating set to obtain a generating set for IO_n of size $n \cdot \binom{n}{2} - n$. Since $\mathrm{Hom}\big(H, \bigwedge^2 H\big)$ has rank $n \cdot \binom{n}{2} - n$, Lemma 8.1 implies that $H_1(\mathrm{IO}_n)$ has rank $n \cdot \binom{n}{2} - n$ as well. This verifies Theorem 1.2 from the introduction.

Multiple boundary components We now move on to the case of multiple boundary components. Just as we did when constructing our drags in Section 5, fix an ordering $P = \{p_1, \ldots, p_{|P|}\}$ and an ordering $p_r = \{\partial_1^r, \ldots, \partial_{b_r}^r\}$ for each $p_r \in P$. We cap off the boundary components of each $p \in P$ as follows:

- If $|p| = 1$, we attach a copy of $M_{1,1}$ to the single boundary component of p.
- If $|p| = k > 1$, we attach a copy of $M_{0,k}$ to these k boundary components.
- If $p = p_1$, we follow the same rules as above, except we introduce an additional boundary component in the piece glued to p.

Capping off the boundary components in this way gives an embedding

$$\iota\colon M_{n,b} \hookrightarrow M_{m,1},$$

where the boundary component of $M_{m,1}$ lies in the piece attached to p_1. This embedding induces a map $\iota_*\colon \mathrm{IO}_{n,b}^P \to IA_m$. We obtain a similar map $\mathrm{IA}_m \to \mathrm{IO}_m$ by attaching a disk to the boundary component of $M_{m,1}$. In Appendix A, we will prove Theorem A.2, which says that the composition

$$\mathrm{IO}_{n,b}^P \xrightarrow{\iota_*} \mathrm{IA}_m \to \mathrm{IO}_m$$

is injective. It follows that ι_* is injective as well. Therefore, to compute the rank of the abelianization of $\mathrm{IO}_{n,b}^P$, it suffices to compute the rank of the abelianization of its image in IA_m. Let $H = H_1(M_{m,1})$, and let $\tau_*\colon \mathrm{IO}_{n,b}^P \to \mathrm{Hom}\big(H, \bigwedge^2 H\big)$ denote the composition

$$\mathrm{IO}_{n,b}^P \xrightarrow{\iota_*} \mathrm{IA}_m \xrightarrow{\tau} \mathrm{Hom}\big(H, \textstyle\bigwedge^2 H\big).$$

Our goal now becomes computing the images of handle, commutator, boundary commutator, and P-drags under τ_*.

Choosing a basis To carry out this computation, it will be helpful to choose bases for $\pi_1(M_{n,b})$ and $\pi_1(M_{m,1})$ carefully. For simplicity, we will assume that $|p_1| > 1$. The case of $|p_1| = 1$ is more straightforward. Fix a basepoint $z_1^1 \in \partial_1^1$, a sphere basis $\{S_1, \dots, S_n\}$, and a geometric free basis $\{x_1, \dots, x_n\}$ for $\pi_1(M_{n,b}, z_1^1)$. We define our drags in $\mathrm{IO}_{n,b}^P$ with respect to these bases.

Next, choose a basepoint $z \in \partial M_{m,1}$, and an oriented arc $\alpha_1 \subset M_{m,1} \setminus \mathrm{int}(M_{n,b})$ from z to z_1^1 (this is possible since the boundary component of $M_{m,1}$ lies on the piece attached to p_1). For $i \in \{1, \dots, n\}$, let $y_i = \alpha_1 x_i \alpha_1^{-1}$. Then $\{y_1, \dots, y_n\}$ is a partial basis for $\pi_1(M_{m,1}, z)$. We wish to extend this to a full basis. Throughout the definition of this extended basis, we encourage the reader to follow along in Figure 12.

For each boundary component ∂_s^r of $M_{n,b}$, fix a point $z_s^r \in \partial_s^r$ (leaving z_1^1 as before). For each $z_s^r \neq z_1^1$, let β_s^r be the unique oriented arc (up to isotopy) in $M_{n,b} \setminus \bigcup S_i$ from z_1^1 to z_s^r. For $s \in \{2, \dots, b_1\}$, define $y_s^1 = \alpha_1 \beta_s^1 \alpha_s^{-1}$. In Figure 12, the loops y_1^1 and y_2^1 are of this form.

Next, let $r > 1$. If $|p_r| = 1$, let γ_1^r be an oriented loop based at z_1^r which generates the fundamental group of the copy of $M_{1,1}$ attached to p_r. Then we define $y_1^r = \alpha_1 \beta_1^r \gamma_1^1 (\beta_1^r)^{-1} (\alpha_1)^{-1}$. In Figure 12, the curve y_1^3 is an example of such a loop.

On the other hand, suppose $|p_r| > 1$. For $s \in \{2, \dots, b_r\}$, let γ_s^r be the unique (up to isotopy) oriented curve in $M_{m,1} \setminus \mathrm{int}(M_{n,b})$ from z_1^r to z_s^r. Then, define

$$y_s^r = \alpha_1 \beta_1^r \gamma_s^r (\beta_s^r)^{-1} (\alpha_1)^{-1}.$$

The curve y_2^2 is an example of this type of loop in Figure 12.

Let $Y = \{y_1 \dots, y_n\}$. For $r \in \{1, \dots, |P|\}$, let $Y_r = \{y_1^r\}$ if $|p_r| = 1$, and $Y_r = \{y_2^r, \dots, y_{b_r}^r\}$ otherwise. Then the collection

$$Y \cup Y_1 \cup \dots \cup Y_{|P|}$$

forms a free basis for $\pi_1(M_{m,1}, z)$.

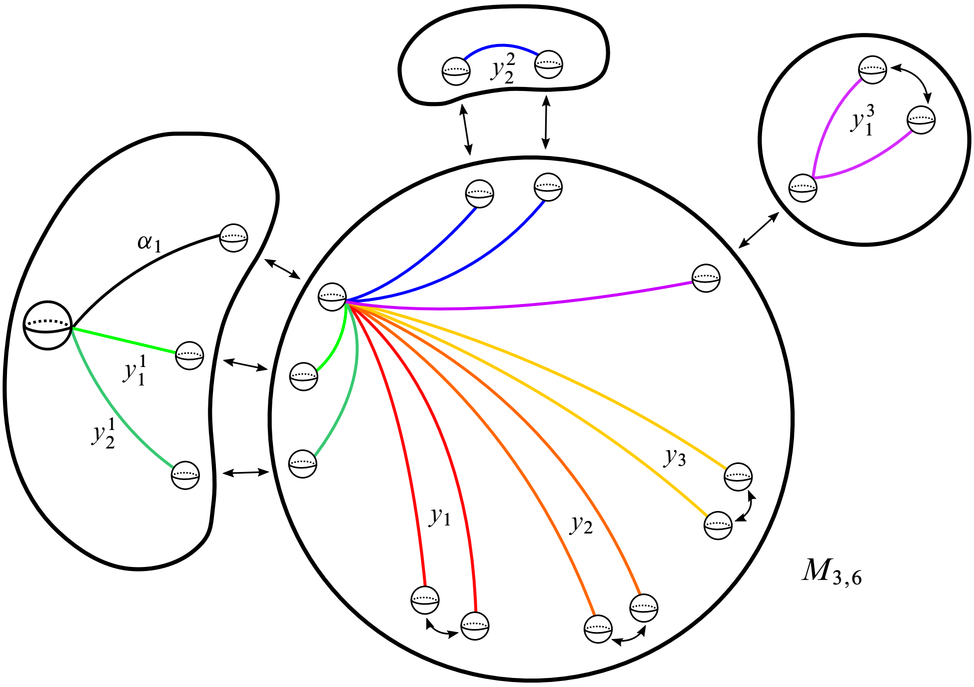

Figure 12: $M_{3,6}$ with the partition $P = \{\{\partial_1^1, \partial_2^1, \partial_3^1\}, \{\partial_1^2, \partial_2^2\}, \{\partial_1^3\}\}$ embedded into $M_{7,1}$. The loops $\{y_1, y_2, y_3\}$ are freely homotopic to a basis for $\pi_1(M_{3,6})$, and this basis has been extended to a basis $\{y_1, y_2, y_3, y_1^1, y_2^1, y_2^2, y_1^3\}$ of $\pi_1(M_{7,1}, z)$.

Computations and relations We now move on to the computation of the images of our collection of drags under τ_*. These computations are straightforward, and are summarized in Table 1. We see from these computations that there is a relation between the images of P-drags and Handle Drags. Namely,

(6) $$\sum_{r=1}^{|P|} \tau_*(\mathrm{PD}_j^r) = -\sum_{i=1}^{n} \tau_*(\mathrm{HD}_{ij})$$

for all $j \in \{1, \dots, n\}$. As we saw in the case of one boundary component, this is because

(7) $$\mathrm{PD}_j^1 \cdots \mathrm{PD}_j^{|P|} \cdot \mathrm{HD}_{1j} \cdot \mathrm{HD}_{nj} = 1$$

in $\mathrm{IO}_{n,b}^P$. Additionally, we see a relation between the image of boundary commutator drags:

(8) $$\sum_{s=1}^{b_r} \tau_*(\mathrm{BCD}_{ij}^{rs}) = 0$$

for all $r \in \{1, \dots, |P|\}$ and $i, j \in \{1, \dots, n\}$ with $i < j$. This relation holds because

(9) $$\mathrm{BCD}_{ij}^{r1} \cdots \mathrm{BCD}_{ij}^{rb_r} = [\mathrm{PD}_i^r, \mathrm{PD}_j^r]$$

in $\mathrm{IO}_{n,b}^P$.

drag	action on π_1	image under τ_*
HD_{ij}	$y_i \mapsto y_j y_i y_j^{-1}$	$[y_i] \mapsto [y_j] \wedge [y_i]$
CD_{ijk}^-	$y_i \mapsto [y_j, y_k] y_i$	$[y_i] \mapsto [y_j] \wedge [y_k]$
BCD_{jk}^{rs} $(r, s > 1)$	$y_s^r \mapsto y_s^r [y_j, y_k]^{-1}$	$[y_s^r] \mapsto [y_k] \wedge [y_j]$
BCD_{jk}^{r1} $(r > 1)$	$y_s^r \mapsto [y_j, y_k] y_s^r$ $(s > 1)$	$[y_s^r] \mapsto [y_j] \wedge [y_k]$ $(s > 1)$
BCD_{jk}^{1s} $(s > 1)$	$y_s^1 \mapsto y_s^1 [y_j, y_k]$	$[y_s^1] \mapsto [y_j] \wedge [y_k]$
BCD_{jk}^{11}	$y \mapsto [y_j, y_k]^{-1} y [y_j, y_k]$ $(y \notin Y_1)$ $y_s^1 \mapsto [y_j, y_k]^{-1} y_s^r$ $(s > 1)$	$[y] \mapsto 0$ $(y \notin Y_1)$ $[y_s^1] \mapsto [y_k] \wedge [y_j]$ $(s > 1)$
PD_j^r $(r > 1)$	$y_s^r \mapsto y_j y_s^r y_j^{-1}$ $(s > 1)$	$[y_s^r] \mapsto [y_j] \wedge [y_s^r]$ $(s > 1)$
PD_j^1	$y \mapsto y_j^{-1} y y_j$ $(y \notin Y_1)$	$[y] \mapsto [y] \wedge [y_j]$ $(y \notin Y_1)$

Table 1: Computing the image of drags under τ_*.

Contributions to abelianization From the computations and relations above, we see that the handle drags and commutator drags together still contribute $n \cdot \binom{n}{2}$ dimensions to the abelianization of $\mathrm{IO}_{n,b}^P$. There are $b \cdot \binom{n}{2}$ boundary commutator drags, but the relations given in (8) kill off $|P| \cdot \binom{n}{2}$ of these in the abelianization (though we can also remove this many elements from our generating set by using (9)). Finally, the number of P-drags is $|P| \cdot n$, but n of them are killed in the abelianization by (6) (and again, we may remove n elements from our generating set by (7)). Adding these all together, we find that the image of $\tau_* \colon \mathrm{IO}_{n,b}^P \to \mathrm{Hom}\big(H, \bigwedge^2 H\big)$ has rank

$$R = n \cdot \binom{n}{2} + \left(b \cdot \binom{n}{2} - |P| \cdot \binom{n}{2}\right) + (|P| \cdot n - n).$$

Moreover, we can reduce our generating set (using (7) and (9)) to a set of size R as well. Thus, by Lemma 8.1, the group $H_1(\mathrm{IO}_{n,b}^P)$ has rank R, which proves Theorem E.

Appendix A Injectivity of the inclusion map

We end this paper with a proof of the following facts, which are surely known to experts, but for which we do not know a reference. They are significant because they allow us to realize the groups $\mathrm{Out}(F_{n,b})$ (and hence $\mathrm{IO}_{n,b}^P$) as subgroups of $\mathrm{Out}(F_m)$. We will begin with a low-genus case.

Lemma A.1 *The induced map $\iota_* \colon \mathrm{Out}(F_{1,1}) \to \mathrm{Out}(F_m)$ is injective for any embedding $\iota \colon M_{1,1} \hookrightarrow M_m$.*

Proof By Laudenbach [18], the group $\mathrm{Out}(F_{1,1}) \cong \mathrm{Aut}(F_1) \cong \mathbb{Z}/2$, where the nontrivial element $f \in \mathrm{Out}(F_{1,1})$ acts on $\pi_1(M_{1,1}, x) \cong \mathbb{Z}$ by inverting the generator. Therefore, $\iota_*(f) \in \mathrm{Out}(F_m)$ is the class of the automorphism

$$\begin{cases} x_1 \mapsto x_1^{-1}, & \\ x_j \mapsto x_j & \text{if } j > 1. \end{cases}$$

This automorphism is not an inner automorphism for any $m \geq 1$, so i_* is injective. □

Theorem A.2 *Fix $n, b \geq 1$ such that $(n, b) \neq (1, 1)$, and let $\iota\colon M_{n,b} \hookrightarrow M_m$ be an embedding. The induced map $\iota_*\colon \mathrm{Out}(F_{n,b}) \to \mathrm{Out}(F_m)$ is injective if and only if no component of $M_m \setminus \mathrm{int}(M_{n,b})$ is diffeomorphic to a 3-disk.*

Proof Suppose first that some component of $M_m \setminus \mathrm{int}(M_{n,b})$ is diffeomorphic to a disk, and let ∂ be the boundary component of $M_{n,b}$ capped off by this disk. By the Birman exact sequence (Theorem 4.2), dragging this boundary component along any nontrivial loop will give a nontrivial element in the kernel of ι_*.

Suppose now that no component of $M_m \setminus \mathrm{int}(M_{n,b})$ is a disk. We will first prove the theorem in the case $b = 1$, and then move on to the general result.

Case 1 Suppose we have an embedding $\iota\colon M_{n,1} \hookrightarrow M_m$. Since no component of $M_m \setminus \mathrm{int}(M_{n,b})$ is a disk, $m > n$. If $n = 1$, then we are done by Lemma A.1, so we may assume that $n > 1$. Fix a basepoint x on the boundary of $M_{n,1}$, and choose a free basis $\{x_1, \dots, x_n\}$ of $\pi_1(M_{n,1}, x)$. The embedding ι induces an injection $\pi_1(M_{n,b}, x) \hookrightarrow \pi_1(M_m, x)$ which identifies $\pi_1(M_{n,1}, x)$ with a free summand of $\pi_1(M_m, x)$. This allows us to extend $\{x_1, \dots, x_n\}$ to a free basis $\{x_1, \dots, x_m\}$ of $\pi_1(M_m, x)$. Given $f \in \mathrm{Out}(F_{n,1}) \cong \mathrm{Aut}(F_n)$, the image $\iota_*(f) \in \mathrm{Out}(F_m)$ is the class of the automorphism $\varphi \in \mathrm{Aut}(F_m)$ generated by

$$\varphi\colon x_i \mapsto \begin{cases} f(x_i) & \text{if } 1 \leq i \leq n, \\ x_i & \text{if } n < i \leq m. \end{cases}$$

Suppose that φ is an inner automorphism. If $m > n + 1$, then φ fixes at least two generators of F_m, and thus must be trivial. It follows that f is trivial as well. On the other hand, if $m = n + 1$, then φ fixes x_m. Since φ is inner, φ must conjugate by a power of x_m. However, if φ conjugates by a nontrivial power of x_m, then f would not act as an automorphism on $\langle x_1, \dots, x_n \rangle \subset F_m$, which is a contradiction. Thus, φ is trivial, and so f is trivial as well.

In summary, we have shown that φ is an inner automorphism if and only if f is trivial, which implies that ι_* is injective.

Case 2 Next, suppose that $\iota\colon M_{n,b} \hookrightarrow M_m$ is an embedding, where $b > 1$. Let $\partial_1, \dots, \partial_b$ be the boundary components of $M_{n,b}$. Let $\Sigma \subset M_{n,b}$ be a 2-sphere which separates $M_{n,b}$ into $M_{n,1}$ and $M_{0,b+1}$ (see Figure 13). Then we have a composition of inclusions

$$M_{n,1} \hookrightarrow M_{n,b} \hookrightarrow M_m.$$

Let $\kappa_*\colon \mathrm{Out}(F_{n,1}) \to \mathrm{Out}(F_{n,b})$ be the map induced by inclusion. By the preceding case, $\iota_* \circ \kappa_*$ is injective. Let $f \in \mathrm{Out}(F_{n,b})$, and suppose that $\iota_*(f) = \mathrm{id}$. By repeated applications of the Birman exact sequence (Theorem 4.2), f has the form $f = p_1 p_2 \cdots p_b \cdot \kappa_*(g)$, where $g \in \mathrm{Out}(F_{n,1}) \cong \mathrm{Aut}(F_n)$ and $p_j \in \mathrm{Out}(F_{n,1})$ is a boundary drag of ∂_j along a loop β_j. Fix a basepoint $x \in \Sigma$, and let $\gamma_j \in \pi_1(M_{n,b}, x)$ be representative of the free homotopy class of β_j. Choose a free basis $\{x_1, \dots, x_n\}$ for $\pi_1(M_{n,1}, x)$. Extend this to a free basis $\{x_1, \dots, x_m\}$ for $\pi_1(M_m, x)$ such that for each $i > n$, the loop x_i intersects

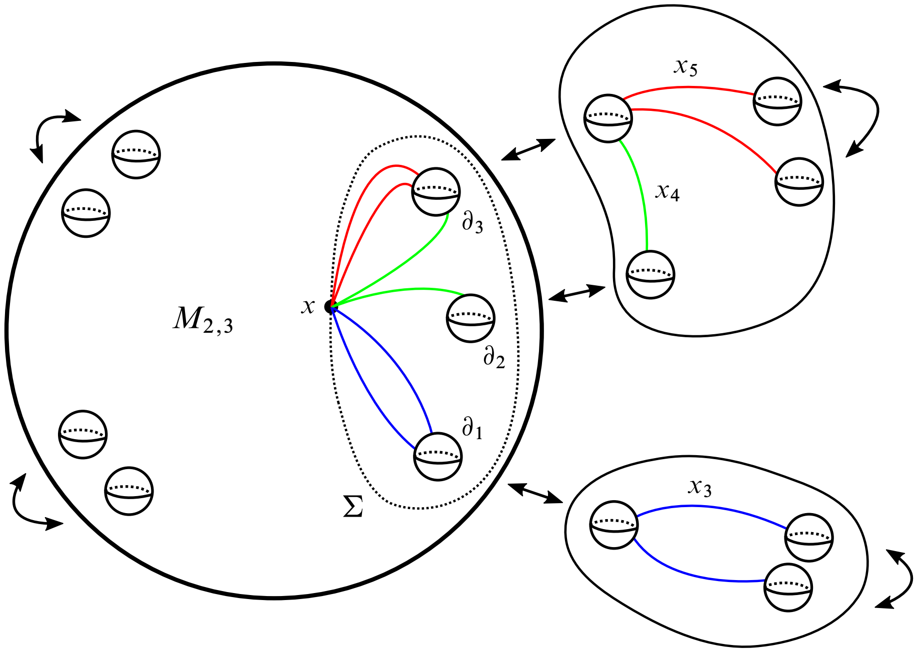

Figure 13: $M_{2,3}$ embedded inside M_5. For clarity, x_1 and x_2 are not shown, but they lie entirely on the opposite side of Σ from x_3, x_4, and x_5.

the set $\bigcup_{j=1}^{b} \partial_j$ exactly twice: once when exiting $M_{n,b}$, and once when reentering (see Figure 13). For $i > n$, let $\partial_{\ell(i)}$ be the boundary component through which α_i leaves $M_{n,b}$, and let $\partial_{r(i)}$ be the boundary component through which it returns. Then $\iota_*(f)$ is the class of the automorphism $\varphi \in \mathrm{Aut}(F_m)$ given by

$$\varphi\colon x_i \mapsto \begin{cases} g(x_i) & \text{for } 1 \le i \le n, \\ \gamma_{\ell(i)} x_i \gamma_{r(i)}^{-1} & \text{for } n < i \le m. \end{cases}$$

By assumption, this automorphism is an inner automorphism. Suppose that φ conjugates by a reduced word w in the x_i. Since g is an automorphism of $\langle x_1, \dots, x_n \rangle \subset F_m$, it follows that $w \in \langle x_1, \dots x_n \rangle$. We will show that this implies that f is trivial by induction on the reduced word length of w.

For the base case, suppose that the word length of w is 0. Then w and φ are both trivial. Since $\iota_* \circ \kappa_*$ is injective, g is trivial as well. Suppose now that some γ_j is not nullhomotopic. Since no component of $M_m \setminus \mathrm{int}(M_{n,b})$ is a disk, there exists some x_i which passes through ∂_j, where $i > n$. In other words, either $\ell(i) = j$ or $r(i) = j$. This is a contradiction because then $\varphi(x_i) = \gamma_{\ell(i)} x_i \gamma_{r(i)}^{-1} \ne x_i$. Thus, all γ_j are nullhomotopic, and so f is trivial. This completes the base case.

Next, suppose that w has positive word length, and let $x_i^{\pm 1}$ be the last letter in the reduced form of w. Then, $w = w' x_i^{\pm 1}$, where the length of w' is less than that of w. To avoid notational complexity, we will assume that $x_i^{\pm 1} = x_1$, but the same argument works for any other x_i. Consider the element

$$h := \mathrm{HD}_{21}\, \mathrm{HD}_{31} \cdots \mathrm{HD}_{n1} \cdot q_1 \cdots q_b \in \mathrm{Out}(F_{n,b}),$$

where HD_{i1} is the handle drag of the i^{th} handle about the first handle (see Section 5) and q_j is obtained by dragging ∂_j about a loop in the free homotopy class of x_1. By construction, $\iota_*(h) \in \mathrm{Out}(F_m)$ is the class of the automorphism which conjugates by x_1. Therefore, $\iota_*(h^{-1}f)$ is the class of the automorphism which conjugates by w'. By our induction hypothesis, this implies that $h^{-1}f$ is trivial.

Claim *The element h is trivial.*

Proof Let $\Sigma' \subset M_{n,b}$ be a 2-sphere which separates $M_{n,b}$ into $M_{1,1}$ and $M_{n-1,b+1}$, where the $M_{1,1}$ is the handle containing x_1. Let $\lambda_*\colon \mathrm{Out}(F_{1,1}) \to \mathrm{Out}(F_{n,b})$ be the map induced by this inclusion. Notice that $h = \lambda_*(q)$, where $q \in \mathrm{Out}(F_{1,1})$ drags the boundary component of $M_{1,1}$ about the nontrivial loop in the positive direction. We saw in the proof of Lemma A.1 that $\mathrm{Out}(F_{1,1}) \cong \mathbb{Z}/2$, and the nontrivial element acts on $\pi_1(M_{1,1})$ by inversion. However, the element q acts trivially on $\pi_1(M_{1,1})$, and is thus trivial itself. It follows that h is trivial as well. □

Combining the claim with the fact that $h^{-1}f$ is trivial, we find that f is trivial. This completes the induction, and so we conclude that ι_* is injective. □

Appendix B Realizing homology classes as spheres

In this section, we prove a result used in the proof of Lemma 7.3 which involves realizing bases of $H_2(M_n)$ as collections of 2-spheres. Recall that $H_2(M_n) = \mathbb{Z}^n$. This identification induces a homomorphism $\eta\colon \mathrm{Mod}(M_n) \to \mathrm{GL}_n(\mathbb{Z})$ which takes a mapping class to its action on homology.

Lemma B.1 *The map $\eta\colon \mathrm{Mod}(M_n) \to \mathrm{GL}_n(\mathbb{Z})$ is surjective.*

Proof First, notice that $H^1(M_n) = \mathbb{Z}^n$. This identification also induces a homomorphism

$$\eta'\colon \mathrm{Mod}(M_n) \to \mathrm{GL}_n(\mathbb{Z})$$

which is well-known to be surjective. Indeed, this map factors as

$$\mathrm{Mod}(M_n) \xrightarrow{q} \mathrm{Out}(F_n) \xrightarrow{\varphi} \mathrm{GL}_n(\mathbb{Z}),$$

where q is the quotient map, and φ sends an automorphism class to its action on H^1. Therefore, if we choose our identifications $H^1(M_n) = \mathbb{Z}^n$ and $H_2(M_n) = \mathbb{Z}^n$ to agree with Poincaré duality, then η and η' are the same map. Thus, η is surjective. □

Lemma B.2 *Let $\{v_1, \dots, v_n\}$ be a basis for $H_2(M_n) = \mathbb{Z}^n$, and let $A = \{S_1, \dots, S_\ell\}$ be a collection of disjoint embedded oriented 2-spheres in M_n which satisfy $[S_j] = v_j$ for $1 \le j \le \ell$. Then A can be extended to a collection $\bar{A} = \{S_1, \dots, S_n\}$ of disjoint embedded oriented 2-spheres such that $[S_j] = v_j$ for $1 \le j \le n$.*

Proof We will induct on n. The base case $n = 0$ is trivial. So assume $n > 0$, and let $\{v_1, \dots, v_n\}$ and $A = \{S_1, \dots, S_\ell\}$ be as stated. There are two cases.

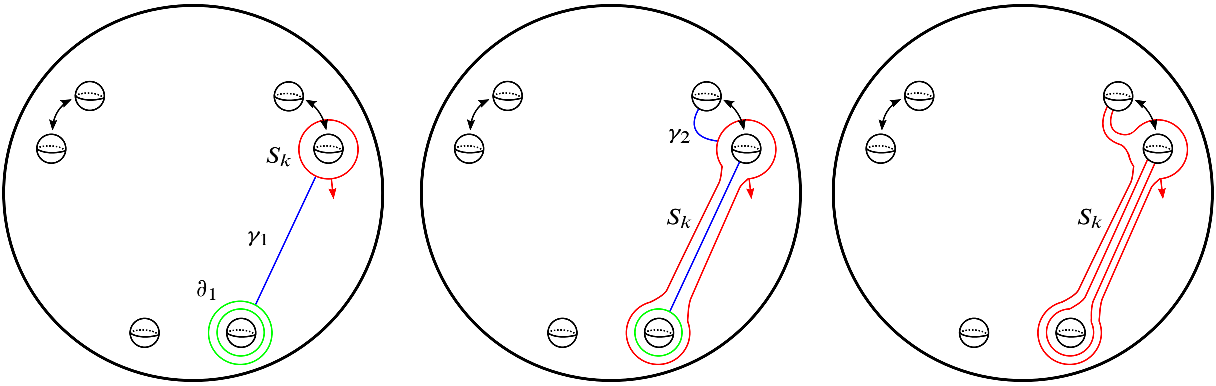

Figure 14: Surgering boundary spheres onto S_k.

First, suppose that $\ell = 0$. If we identity $H_2(M_n)$ with $\mathbb{Z}^n$, then by Lemma B.1 the resulting map $\eta\colon \mathrm{Mod}(M_n) \to \mathrm{GL}_n(\mathbb{Z})$ is surjective. Choose any collection $\Sigma_1, \dots, \Sigma_n \subset M_n$ of disjoint embedded 2-spheres such that $M_n \setminus (\Sigma_1 \cup \dots \cup \Sigma_n)$ is connected. Then $\{[\Sigma_1], \dots, [\Sigma_n]\}$ is a basis for $H_2(M_n)$. Since $\mathrm{GL}_n(\mathbb{Z})$ acts transitively on ordered bases of $\mathbb{Z}^n$ and the map η is surjective (Lemma B.1), there exists some $\mathfrak{f} \in \mathrm{Mod}(M_n)$ such that $\eta(\mathfrak{f}) \cdot [\Sigma_j] = v_j$ for all $1 \le j \le n$. In other words, $[\mathfrak{f}(\Sigma_j)] = v_j$, and so $\{\mathfrak{f}(\Sigma_1), \dots, \mathfrak{f}(\Sigma_n)\}$ is the desired collection of spheres.

Next, suppose that $\ell > 0$. Splitting M_n along S_1 gives an embedding $\iota\colon M_{n-1,2} \hookrightarrow M_n$. Notice that the induced map $\iota_H\colon H_2(M_{n-1,2}) \to H_2(M_n)$ is an isomorphism. Let $w_j = \iota_H^{-1}(v_j)$ for $1 \le j \le n$, and let ∂ and ∂' be the boundary components of $M_{n-1,2}$. Capping the two boundary components of $M_{n-1,2}$ with disks D and D', we get another embedding $\iota'\colon M_{n-1,2} \hookrightarrow M_{n-1}$. This embedding induces a surjective map $\iota'_H\colon H_2(M_{n-1,2}) \to H_2(M_{n-1})$ whose kernel is generated by $[\partial]$. Let $w'_j = \iota'_H(w_j)$ for $2 \le j \le n$, and let $S'_k = \iota'(S_k)$ for $2 \le k \le \ell$. By our induction hypothesis, we can extend the collection $\{S'_2, \dots, S'_\ell\}$ to a collection $\{S'_2, \dots, S'_{n-1}\}$ of disjoint embedded oriented 2-spheres in M_{n-1} such that $[S'_j] = w'_j$ for $2 \le j \le n$. Moreover, since the disks D and D' used to cap the boundary components of $M_{n-1,2}$ are contractible, we may isotope $S'_{\ell+1}, \dots, S'_{n-1}$ such that they are disjoint from D and D'. Let $S_j = (\iota')^{-1}(S'_j)$ for $\ell + 1 \le j \le n$. If $[S_k] = w_k$ for all k, then $\{S_1, \dots, S_n\}$ is the desired collection, and we are done. However, since the kernel of ι'_H is generated by $[\partial]$, we have

$$[S_k] = w_k + c_k[\partial],$$

where $c_k \in \mathbb{Z}$. Note that $c_k = 0$ for $2 \le k \le \ell$. To fix this, we may surger parallel copies of ∂ or ∂' onto S_k such that it has the correct homology class. The process is as follows (see Figure 14):

(i) If $c_k > 0$, take c_k parallel copies of ∂', which we denote by $\partial_1, \dots, \partial_{c_k}$. If instead $c_k < 0$, take $\partial_1, \dots, \partial_{c_k}$ to be parallel copies of ∂. Order the ∂_j such that ∂_1 is furthest from its respective boundary component, then ∂_2, and so on.

(ii) Let γ_1 be a properly embedded arc connecting the positive side of S_k to ∂_1 which does not intersect any of the other S_j or ∂_j.

(iii) Surger S_k and ∂_1 together via a tube running along γ_1.

(iv) Repeat steps (ii) and (iii) for the remaining ∂_j.

Once we have carried out this process for all the S_k, we will have obtained a collection $\{S_2, \dots, S_n\}$ of spheres whose homology classes are exactly $w_2, \dots, w_n$. Thus, $\{S_1, \dots, S_n\}$ is the desired collection of 2-spheres. □

References

[1] **S Andreadakis**, *On the automorphisms of free groups and free nilpotent groups*, Proc. Lond. Math. Soc. 15 (1965) 239–268 MR Zbl

[2] **M A Armstrong**, *On the fundamental group of an orbit space*, Proc. Cambridge Philos. Soc. 61 (1965) 639–646 MR Zbl

[3] **S Bachmuth**, *Induced automorphisms of free groups and free metabelian groups*, Trans. Amer. Math. Soc. 122 (1966) 1–17 MR Zbl

[4] **J E Banks**, *The Birman exact sequence for 3-manifolds*, Bull. Lond. Math. Soc. 49 (2017) 604–629 MR Zbl

[5] **M Bestvina**, **K-U Bux**, **D Margalit**, *Dimension of the Torelli group for* $\mathrm{Out}(F_n)$, Invent. Math. 170 (2007) 1–32 MR Zbl

[6] **T Brendle**, **N Broaddus**, **A Putman**, *The mapping class group of connect sums of* $S^2 \times S^1$, Trans. Amer. Math. Soc. 376 (2023) 2557–2572 MR Zbl

[7] **M R Bridson**, **K Vogtmann**, *Automorphism groups of free groups, surface groups and free abelian groups*, from "Problems on mapping class groups and related topics", Proc. Sympos. Pure Math. 74, Amer. Math. Soc., Providence, RI (2006) 301–316 MR Zbl

[8] **T Church**, **B Farb**, **A Putman**, *Integrality in the Steinberg module and the top-dimensional cohomology of* $\mathrm{SL}_n\mathcal{O}_K$, Amer. J. Math. 141 (2019) 1375–1419 MR Zbl

[9] **M Day**, **A Putman**, *A Birman exact sequence for* $\mathrm{Aut}(F_n)$, Adv. Math. 231 (2012) 243–275 MR Zbl

[10] **M Day**, **A Putman**, *The complex of partial bases for* F_n *and finite generation of the Torelli subgroup of* $\mathrm{Aut}(F_n)$, Geom. Dedicata 164 (2013) 139–153 MR Zbl

[11] **M Day**, **A Putman**, *A Birman exact sequence for the Torelli subgroup of* $\mathrm{Aut}(F_n)$, Int. J. Algebra Comput. 26 (2016) 585–617 MR Zbl

[12] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Math. Ser. 49, Princeton Univ. Press (2012) MR Zbl

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

[14] **A Hatcher**, **N Wahl**, *Stabilization for the automorphisms of free groups with boundaries*, Geom. Topol. 9 (2005) 1295–1336 MR Zbl

[15] **A Hatcher**, **N Wahl**, *Stabilization for mapping class groups of 3-manifolds*, Duke Math. J. 155 (2010) 205–269 MR Zbl

[16] **C A Jensen**, **N Wahl**, *Automorphisms of free groups with boundaries*, Algebr. Geom. Topol. 4 (2004) 543–569 MR Zbl

[17] **D Johnson**, *The structure of the Torelli group, III: The abelianization of $\mathcal{T}$*, Topology 24 (1985) 127–144 MR Zbl

[18] **F Laudenbach**, *Sur les 2-sphères d'une variété de dimension 3*, Ann. of Math. 97 (1973) 57–81 MR

[19] **F Laudenbach**, *Topologie de la dimension trois: homotopie et isotopie*, Astérisque 12, Soc. Math. France, Paris (1974) MR Zbl

[20] **H Maazen**, *Homology stability for the general linear group*, PhD thesis, Utrecht University (1951) Available at https://dspace.library.uu.nl/handle/1874/237657

[21] **W Magnus**, *Über n-dimensionale Gittertransformationen*, Acta Math. 64 (1935) 353–367 MR Zbl

[22] **A Putman**, *Cutting and pasting in the Torelli group*, Geom. Topol. 11 (2007) 829–865 MR Zbl

[23] **A Putman**, *Obtaining presentations from group actions without making choices*, Algebr. Geom. Topol. 11 (2011) 1737–1766 MR Zbl

[24] **A Putman**, *The commutator subgroups of free groups and surface groups*, Enseign. Math. 68 (2022) 389–408 MR Zbl

[25] **W Tomaszewski**, *A basis of Bachmuth type in the commutator subgroup of a free group*, Canad. Math. Bull. 46 (2003) 299–303 MR Zbl

[26] **N Wahl**, *From mapping class groups to automorphism groups of free groups*, J. Lond. Math. Soc. 72 (2005) 510–524 MR Zbl

Department of Mathematics, University of Notre Dame
Notre Dame, IN, United States

landgrafjacob@gmail.com

Received: 18 November 2021 Revised: 15 September 2023

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2025 is US $760/year for the electronic version, and $1110/year (+$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

https://msp.org/

Algebraic & Geometric Topology

Volume 25 Issue 1 (pages 1–644) 2025