

ISOTOPY VERSUS EQUIVARIANT ISOTOPY IN DIMENSIONS THREE AND HIGHER

TRENT LUCAS

ABSTRACT. Given a finite group action on a smooth manifold, we study the following question: if two equivariant diffeomorphisms are isotopic, must they be equivariantly isotopic? Birman–Hilden and Maclachlan–Harvey proved the answer is “yes” for most surfaces. By contrast, we give a general criterion in higher dimensions under which there are many equivariant diffeomorphisms which are isotopic but not equivariantly isotopic. Examples satisfying this criterion include branched covers of split links and “stabilized” branched covers. We prove the result by constructing an invariant valued in the homology of a certain infinite cover of the manifold. We give applications to outer automorphism groups of free products and to group actions on manifolds which fiber over the circle.

1. INTRODUCTION

1.1. The main question and result. Let M be a closed oriented smooth manifold, and G a finite group of orientation-preserving diffeomorphisms of M . We call a diffeomorphism of M *equivariant* if it commutes with G . Given two equivariant diffeomorphisms of M , we can ask whether they are isotopic, or we can ask whether they are *equivariantly* isotopic, i.e. isotopic through equivariant diffeomorphisms. Our goal in this paper is to measure the difference between these two equivalence relations.

More formally, let $\text{Diff}(M)$ denote the group of orientation-preserving diffeomorphisms of M , and let $\text{Diff}(M)^G$ denote the subgroup of equivariant diffeomorphisms. Let $\text{Mod}(M)$ denote $\pi_0(\text{Diff}(M))$, i.e. the group of isotopy classes of orientation-preserving diffeomorphisms, also called the *mapping class group* of M . Let $\Gamma_G(M)$ denote $\pi_0(\text{Diff}(M)^G)$, i.e. the group of equivariant isotopy classes of equivariant diffeomorphisms. Then the inclusion $\text{Diff}(M)^G \hookrightarrow \text{Diff}(M)$ induces a map $\mathcal{P}_G : \Gamma_G(M) \rightarrow \text{Mod}(M)$. We then ask:

Question 1.1. *Is the map $\mathcal{P}_G : \Gamma_G(M) \rightarrow \text{Mod}(M)$ injective? If not, how large is the kernel?*

A remarkable theorem of Birman–Hilden [BH73] and Maclachlan–Harvey [MH75] says that if M is a closed oriented surface of genus $g \geq 2$, then \mathcal{P}_G is injective. By contrast, we showed in previous work [Luc25] that for many finite group actions on closed oriented 3-manifolds, the kernel $\text{Ker}(\mathcal{P}_G)$ is infinite, answering a question of Margalit–Winarski [MW21, Ques 11.4]. Our main theorem strengthens this result and extends it to all dimensions $n \geq 3$.

Theorem A. *Let M be a closed oriented connected smooth manifold of dimension $n \geq 3$, and let G be a finite group of orientation-preserving diffeomorphisms. Let $M^\circ \subseteq M$ denote the set of points whose G -stabilizer is trivial. Assume that:*

- *For some $k > 0$, the quotient manifold M°/G is diffeomorphic to a k -fold connected sum $Q_0 \# \cdots \# Q_{k-1}$ where $\pi_1(Q_i)$ is nontrivial for each i .*
- *For some $g_0 \in G$, the set of fixed points $\text{Fix}(g_0) \subseteq M$ has a nonempty codimension 2 orientable component B whose homology class in $H_{n-2}(M; \mathbb{Z}/\ell\mathbb{Z})$ is trivial for some $\ell > 1$.*

Then if $k \geq 3$:

- (i) *the kernel of the map $\mathcal{P}_G : \Gamma_G(M) \rightarrow \text{Mod}(M)$ contains a virtual free product, and*
- (ii) *the kernel $\text{Ker}(\mathcal{P}_G)$ is not finitely generated.*

In part (i), a *virtual free product* is a group which is virtually isomorphic to a free product of two nontrivial groups; here we say that two groups G_1 and G_2 are *virtually isomorphic* if there are finite index subgroups $H_i \leq G_i$ and finite normal subgroups $F_i \leq H_i$ such that $H_1/F_1 \cong H_2/F_2$. Note that Theorem A reduces to the case $k = 3$, since for higher k the quotient can be written as $Q_0 \# Q_1 \# Q$ where $Q = Q_2 \# \cdots \# Q_{k-1}$.

The surprising aspect of Theorem A is not just that $\text{Ker}(\mathcal{P}_G)$ is nontrivial, but also that $\text{Ker}(\mathcal{P}_G)$ is so large at this level of generality. Indeed, the bulk of this paper is devoted to proving part (ii). The proof requires us to find an invariant that can distinguish many equivariant diffeomorphisms which are all isotopically trivial; we accomplish this with an “obstruction map” valued in the homology of a certain infinite cover of M . This invariant is simple to describe, but finding even one example where this invariant is nontrivial requires a careful and explicit computation. The proof of part (i) is simpler, and serves as a warmup to the proof of part (ii).

We give more details of the proof in Section 1.6 below. Before that, we will describe examples of group actions satisfying the hypotheses of Theorem A, and give two applications of Theorem A.

1.2. Example: Branched covers of links. A simple way to build examples for Theorem A in all dimensions $n \geq 3$ is to take branched covers of links in S^n . We say that a surjective map of smooth manifolds $p : M \rightarrow N$ is a *finite regular branched cover* if there is a compact codimension 2 submanifold $C \subseteq N$ such that

- the restriction of p to $M \setminus p^{-1}(C) \rightarrow N \setminus C$ is a finite-sheeted regular cover, and
- for each $x \in p^{-1}(C)$, we can choose coordinates around x and $p(x)$ valued in $\mathbb{C} \times \mathbb{R}^{n-2}$ in which p is of the form $(z, x) \mapsto (z^m, x)$ for some integer $m \geq 2$.

A *link* is the image of a smooth embedding $S^{n-2} \sqcup \cdots \sqcup S^{n-2} \hookrightarrow S^n$. Given a link L in S^n and an integer $d \geq 2$, there is a closed oriented smooth manifold $M_{L,d}$ and a finite regular branched cover $p_{L,d} : M_{L,d} \rightarrow S^n$ with branch set L and deck group $G = \mathbb{Z}/d\mathbb{Z}$ (see e.g. [Ran98, §27], which constructs $M_{L,d}$ in the topological category but can be easily adapted to the smooth category).

We say a link L is *k-split* if there are k disjoint closed n -disks $D_0, \dots, D_{k-1} \subseteq S^n$ such that $D_i \cap L$ is a nonempty union of components of L . Given a k -split link L , we claim that the action of G on $M = M_{L,d}$ satisfies the hypotheses of Theorem A. Indeed, let Q_i be the n -manifold obtained by capping the boundary of $D_i \setminus L$ with an n -disk. Then $S^n \setminus L$ is diffeomorphic to the connected sum $Q_0 \# \cdots \# Q_{k-1}$. Since $H_1(Q_i)$ is nontrivial for each i by Alexander duality, the first hypothesis of Theorem A is satisfied. For the second hypothesis, we can take B to be the preimage of any component L_i of L ; then B will be homologically trivial in M since its meridian will be nontrivial in $H_1(M \setminus B)$ (see Lemma 3.3). Thus Theorem A implies:

Corollary B. *Let L be a k -split link in S^n for $n \geq 3$, and let $M = M_{L,d}$ be the associated branched cover with deck group $G = \mathbb{Z}/d\mathbb{Z}$. If $k \geq 3$, then $\text{Ker}(\mathcal{P}_G)$ is not finitely generated.*

The simplest case of Corollary B is when $d = 2$ and L is the k -component unlink in S^n (i.e. each component is isotopic to the standard embedding $S^{n-2} \hookrightarrow S^n$). In this case, $M \cong (S^1 \times S^{n-1})^{\#k-1}$ and the action of $G = \mathbb{Z}/2\mathbb{Z}$ on M is a generalization of the hyperelliptic involution on a genus $k-1$ surface (which was the example first studied by Birman and Hilden). In [Luc25], we showed that $\text{Ker}(\mathcal{P}_G)$ is not finitely generated for the double cover of the 3-component unlink in S^3 using an algebraic proof much different from that of Theorem A.

1.3. Example: Stabilized actions and branched covers. Another family of examples for Theorem A come from “stabilizing” group actions and branched covers, as follows.

Suppose P is a closed smooth n -manifold and G is a finite group of diffeomorphisms. For $1 \leq i \leq k$, let Q_i be any closed oriented connected smooth n -manifold with $\pi_1(Q_i)$ nontrivial, and let P_i be the disjoint union of d copies of Q_i , where $d = |G|$. We then define a closed smooth n -manifold M by taking an *equivariant connected sum* of P with each P_i , meaning that we remove an open disk from each component of P_i , remove a free G -orbit of open disks from P , and glue together the

resulting boundary components (some care must be taken to preserve smoothness and orientability, as in Section 2.3). The action of G on P naturally extends to an action on M by permuting the components of each P_i . We call M a k -fold *stabilization* of the action of G on P .

If $q : P \rightarrow Q$ is a finite regular branched cover with deck group G , and M is a k -fold stabilization of the action of G on P , then we get an induced branched cover $p : M \rightarrow N$ where $N = Q \# Q_1 \cdots \# Q_k$. We call the branched cover p a k -fold *stabilization* of the branched cover q . Theorem A implies:

Corollary C. *Let P and Q be closed oriented connected smooth manifolds of dimension $n \geq 3$, and let $q : P \rightarrow Q$ be a finite regular branched cover with branch set $C \subseteq Q$ and deck group G . Assume that $p^{-1}(C)$ represents a trivial element of $H_{n-2}(P; \mathbb{Z}/\ell\mathbb{Z})$ for some $\ell > 1$. Let $p : M \rightarrow N$ be a k -fold stabilization of q . Then if $k \geq 2$, the kernel of $\mathcal{P}_G : \Gamma_G(M) \rightarrow \text{Mod}(M)$ is not finitely generated.*

To deduce Corollary C from Theorem A, we use the fact that $Q \setminus C$ is not simply connected. This is because the unbranched cover $P \setminus q^{-1}(C) \rightarrow Q \setminus C$ is classified by a surjection $\pi_1(Q \setminus C) \rightarrow G$.

For a concrete family of examples in even dimensions, we can take P to be a degree d cyclic cover of $Q = \mathbb{CP}^{n/2}$ branched over a smooth degree m complex hypersurface C , where d is a proper divisor of m . To check that the assumption of Corollary C holds, recall first that $\pi_1(\mathbb{CP}^{n/2} \setminus C) \cong \mathbb{Z}/m\mathbb{Z}$ and is generated by a meridian of C (see [Cog11]). Then $\pi_1(P \setminus q^{-1}(C)) \cong \mathbb{Z}/\ell\mathbb{Z}$, where $\ell = m/d$. Thus $H_1(P \setminus q^{-1}(C); \mathbb{Z}/\ell\mathbb{Z})$ is isomorphic to $\mathbb{Z}/\ell\mathbb{Z}$ and is generated by a meridian of $q^{-1}(C)$. This implies that $q^{-1}(C)$ must represent a trivial homology class in $H_{n-2}(P; \mathbb{Z}/\ell\mathbb{Z})$ (this can be seen by adapting the proof of Lemma 3.3).

1.4. Application: Symmetric automorphism groups of free products. We can apply Theorem A to study certain natural maps arising in group theory.

Let $L = L_1 * \cdots * L_k$ be any free product of groups. We define $\text{SymAut}(L)$ to be the group of *symmetric* automorphisms of L , i.e. the group of automorphisms of L sending each L_i to a conjugate of some L_j , and we define $\text{SymOut}(L)$ to be the image of $\text{SymAut}(L)$ in $\text{Out}(L)$. In the case that each L_i is freely indecomposable and not infinite cyclic, $\text{SymAut}(L)$ is in fact the full automorphism group $\text{Aut}(L)$ by the Kurosh subgroup theorem. In the case that L is the free group F_k , the group $\text{SymAut}(F_k)$ arises naturally in the study of the braid group B_k , as B_k acts on the fundamental group of the punctured disk by symmetric automorphisms.

Given integers $k, d > 1$, we define the free product

$$H_{k,d} := \underbrace{(\mathbb{Z}/d\mathbb{Z}) * \cdots * (\mathbb{Z}/d\mathbb{Z})}_{k \text{ times}}.$$

Then the natural surjection $F_k \rightarrow H_{k,d}$ induces maps

$$\begin{aligned} \widehat{\mathcal{Q}}_{k,d} : \text{SymAut}(F_k) &\rightarrow \text{SymAut}(H_{k,d}) \\ \mathcal{Q}_{k,d} : \text{SymOut}(F_k) &\rightarrow \text{SymOut}(H_{k,d}). \end{aligned}$$

As a surprising application of their work on Question 1.1, Birman–Hilden [BH73] showed that the map $\widehat{\mathcal{Q}}_{k,d}$ restricts to an embedding $B_k \hookrightarrow \text{SymAut}(H_{k,d})$, answering a question of Magnus¹ (see also [MW21, §7]). Apart from this, it appears that the kernels of the maps $\widehat{\mathcal{Q}}_{k,d}$ and $\mathcal{Q}_{k,d}$ are not well-studied. These kernels contain some obvious elements (e.g. conjugating one generator by the d th power of another), but it is not clear which elements generate these kernels or what finiteness properties these kernels possess. Using Theorem A, we show:

Theorem D. *Let $k \geq 3$ and $d \geq 2$. Then the groups $\text{Ker}(\widehat{\mathcal{Q}}_{k,d})$ and $\text{Ker}(\mathcal{Q}_{k,d})$ are not finitely generated.*

¹Magnus later applied this result in [Mag80].

Theorem D is sharp. For $d = 1$, the maps $\widehat{\mathcal{Q}}_{k,d}$ and $\mathcal{Q}_{k,d}$ are trivial, and the groups $\text{SymAut}(F_k)$ and $\text{SymOut}(F_k)$ are finitely generated (see e.g. [McC86]). For $k = 1$, we have that $\text{SymAut}(\mathbb{Z})$ and $\text{SymOut}(\mathbb{Z})$ are isomorphic to $\mathbb{Z}/2\mathbb{Z}$. For $k = 2$, the group $\text{Inn}(F_2)$ has finite index in $\text{SymAut}(F_2)$ (again see e.g. [McC86]), which implies that $\text{Ker}(\widehat{\mathcal{Q}}_{k,d})$ is not finitely generated but $\text{Ker}(\mathcal{Q}_{k,d})$ is finite.

Following our previous work in [Luc25], we give a topological proof of Theorem D. The key idea is to apply Theorem A to finite covers $p_{k,d} : M_{k,d} \rightarrow S^3$ branched over the k -component unlink C_k with deck group $G = \mathbb{Z}/d\mathbb{Z}$. Let $\text{Diff}(S^3, C_k)$ denote the group orientation-preserving diffeomorphisms of S^3 preserving C_k setwise, and let $\text{Mod}(S^3, C_k) = \pi_0(\text{Diff}(S^3, C_k))$. Then the relationship between the cover $p = p_{k,d}$ and the group $\text{SymOut}(F_k)$ comes from a theorem of Wattenberg [Wat72], which says that the action of $\text{Mod}(S^3, C_k)$ on $\pi_1(S^3 \setminus C_k)$ induces an isomorphism

$$\text{Mod}(S^3, C_k) \cong \text{SymOut}(F_k)$$

(the analogous result holds in the topological category by work of Goldsmith [Gol81], and both results build off the work of Dahm [Dah62]). The cover p induces an orbifold structure \mathcal{O}_p on S^3 , and $\text{Mod}(S^3, C_k)$ also acts on the orbifold fundamental group $\pi_1^{\text{orb}}(\mathcal{O}_p) \cong H_{k,d}$. This yields a relationship between $\text{Ker}(\mathcal{Q}_{k,d})$ and $\text{Ker}(\mathcal{P}_G)$ which we can use to prove Theorem D.

In [Luc25], we computed a finite *normal* generating set of $\text{Ker}(\mathcal{Q}_{k,2})$ using results of McCullough–Miller [MM96] on the groups $\text{SymOut}(F_k)$ and $\text{SymOut}(H_{k,d})$. It would be interesting to determine whether the groups $\text{Ker}(\mathcal{Q}_{k,d})$ have a finite normal generating set for all k and d using our topological framework.

1.5. Application: Trivial bundles with a nontrivial fiberwise action. As explained in [Luc25], Theorem A provides many examples of fiber bundles $E \rightarrow S^1$ equipped with a fiberwise group action such that E is smoothly trivial, but not equivariantly trivial.

Let M be a closed oriented n -manifold and G a finite group of diffeomorphisms satisfying the hypotheses of Theorem A. Then the non-injectivity of the map \mathcal{P}_G implies that the map

$$\widehat{\mathcal{P}}_G : \pi_1(B \text{Diff}(M)^G) \rightarrow \pi_1(B \text{Diff}(M))$$

is not injective, where B denotes the classifying space. Conjugacy classes in $\pi_1(B \text{Diff}(M))$ correspond to isomorphism classes of M -bundles over S^1 . Conjugacy classes in $\pi_1(B \text{Diff}(M)^G)$ correspond to isomorphism classes of M -bundles over S^1 with structure group $\text{Diff}(M)^G$; these are bundles for which the action of G on M extends to a fiberwise action on the total space. Thus, nontrivial elements in $\text{Ker}(\widehat{\mathcal{P}}_G)$ yield bundles which are smoothly trivial, but not equivariantly trivial, i.e. not equivariantly isomorphic to the product $S^1 \times M$ equipped with the standard action of G in each fiber.

The same reasoning tells us that elements in the kernel of the map $\pi_k(\text{Diff}(M)^G) \rightarrow \pi_k(\text{Diff}(M))$ correspond to M -bundles over S^{k+1} which are smoothly but not equivariantly trivial. Our construction of elements in the kernel for $k = 0$ does not obviously extend to higher k , so we naturally pose:

Problem 1.2. *Let G be a finite group of diffeomorphisms of a closed oriented smooth n -manifold M . For $k \geq 1$, construct elements in the kernel of the map $\pi_k(\text{Diff}(M)^G) \rightarrow \pi_k(\text{Diff}(M))$ (or show this map is injective).*

1.6. Proof idea. As mentioned above, the surprising aspect of Theorem A is that $\text{Ker}(\mathcal{P}_G)$ is so large. In fact, if one only wishes to show that $\text{Ker}(\mathcal{P}_G)$ is nontrivial, there is a simple strategy to do so. Namely, for any $g_0 \in G$, any equivariant diffeomorphism must preserve the fixed set $\text{Fix}(g_0) \subseteq M$ setwise, and an element of $\text{Ker}(\mathcal{P}_G)$ may restrict to a nontrivial mapping class of $\text{Fix}(g_0)$ (for example, this occurs for the hyperelliptic involution on the 2-torus [MW21, §2]). However, if the mapping class group of $\text{Fix}(g_0)$ is finite (such as for branched covers of links in S^n for $n \leq 5$), then this obstruction only gives a finite lower bound on the size of $\text{Ker}(\mathcal{P}_G)$. The elements of $\text{Ker}(\mathcal{P}_G)$

that we construct in the proof of Theorem A will in fact act trivially on the fixed points of any $g_0 \in G$.

The elements of $\text{Ker}(\mathcal{P}_G)$ that we use to prove Theorem A are lifts of “summand slides,” defined as follows. Choose a splitting of M°/G as a connected sum $N_1 \# N_2$, and let $\Sigma \subseteq M^\circ/G$ be a separating $(n-1)$ -sphere realizing this decomposition. Let δ be a loop in N_1 based at Σ . Informally, a summand slide is a diffeomorphism of M°/G obtained by dragging Σ along δ , similar to a point-pushing diffeomorphism of a surface.

Part (i) of Theorem A has a fairly direct proof. We find loops δ_1 and δ_2 in M°/G which are nontrivial in $\pi_1(M^\circ/G)$, but lift to trivial elements of $\pi_1(M)$. We then let f_i be a summand slide of M°/G along δ_i , and let \tilde{f}_i be a lift of f_i to an equivariant diffeomorphism of M . A key property of summand slides is that up to finite order, they are determined by their action on π_1 . We show that f_1 and f_2 generate a free product in $\text{Out}(\pi_1(M^\circ/G))$, which implies that \tilde{f}_1 and \tilde{f}_2 generate a virtual free product in $\Gamma_G(M)$, but \tilde{f}_1 and \tilde{f}_2 act trivially on $\pi_1(M)$, which implies that they generate a finite subgroup of $\text{Mod}(M)$.

Part (ii) is more difficult, and its proof comprises the majority of this paper. Let $\text{Mod}(M, B) = \pi_0(\text{Diff}(M, B))$, where $\text{Diff}(M, B)$ is the group of orientation-preserving diffeomorphisms that preserve B setwise. Let $\text{Mod}_G(M, B) \leq \text{Mod}(M, B)$ and $\text{Mod}_G(M) \leq \text{Mod}(M)$ each denote the subgroup of mapping classes with an equivariant representative. Then the natural forgetful map $\text{Mod}(M, B) \rightarrow \text{Mod}(M)$ restricts to a map

$$\mathcal{F}_G : \text{Mod}_G(M, B) \rightarrow \text{Mod}_G(M).$$

Then a finite index subgroup of $\text{Ker}(\mathcal{P}_G)$ will surject onto $\text{Ker}(\mathcal{F}_G)$, so it’s enough to show that $\text{Ker}(\mathcal{F}_G)$ is not finitely generated.

We construct an obstruction map

$$\varphi : \text{Ker}(\mathcal{F}_G) \rightarrow H_{n-2}(\tilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$$

where \tilde{M} is a certain infinite-sheeted cover of M and \tilde{B} is a component of the preimage of B in \tilde{M} . The homology group $H_{n-2}(\tilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$ is non-finitely generated, so the strategy of the proof is to find an infinite linearly independent subset of the image of φ .

The map φ can be computed as follows: suppose $\alpha \in \text{Ker}(\mathcal{F}_G)$ is represented by an equivariant diffeomorphism $f : M \rightarrow M$ that preserves B setwise. Then there is an isotopy $f \simeq \text{id}_M$. By tracking B under this isotopy, we get a map $B \times [0, 1] \rightarrow M$ mapping $B \times \{0\}$ and $B \times \{1\}$ diffeomorphically onto B . We show that this lifts to a map $B \times [0, 1] \rightarrow \tilde{M}$ mapping $B \times \{0\}$ and $B \times \{1\}$ diffeomorphically onto \tilde{B} . This lifted map yields a homology class in $H_{n-1}(\tilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$.

The most difficult step is to find at least one element $\alpha_0 \in \text{Ker}(\mathcal{F}_G)$ such that $\varphi(\alpha_0)$ is nonzero. We can then obtain an infinite linearly independent subset of $H_{n-2}(\tilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$ by conjugating α_0 . To find α_0 , we construct an explicit equivariant diffeomorphism f and an explicit isotopy from f to id_M . As in the proof of part (i), the diffeomorphism f is a lift of a summand slide.

The obstruction map φ can be viewed as an invariant defined on the homology of the space of unparametrized embeddings $B \hookrightarrow M$ (see Remark 5.1). It would be interesting to adapt the map φ to study the homology of embedding spaces in a broader context.

1.7. Generalizations of Theorem A. Theorem A is stated and proved in the smooth category. By slightly modifying the proof, one can show that if M is a smooth manifold and G is a finite group of diffeomorphisms satisfying the hypotheses of Theorem A, then the kernel of the map $\pi_0(\text{Homeo}(M)^G) \rightarrow \pi_0(\text{Homeo}(M))$ is not finitely generated. One can also adapt the proof of Theorem A to the case that M is noncompact or has nonempty boundary, as long as one still assumes the submanifold B is compact and without boundary.

1.8. Sharpness of Theorem A in dimension 3. In dimension $n = 3$, the assumptions of Theorem A are fairly mild restrictions. The Prime Decomposition Theorem implies that M°/G splits uniquely as a connected sum of prime 3-manifolds, and the Poincaré conjecture implies that each summand is not simply connected. Therefore, the first assumption reduces to the statement that M°/G has k prime factors. For the second assumption, we note that if $\text{Fix}(g_0)$ is nonempty, then any component will be an embedded loop in M , and the assumption will be satisfied if any such loop is a torsion element (e.g. zero) in $H_1(M; \mathbb{Z})$. Under these assumptions, Theorem A is sharp in dimension 3. Namely, by [Luc25, Cor D], there is an action of $G = \mathbb{Z}/2\mathbb{Z}$ on $M = S^1 \times S^2$ satisfying the assumptions of Theorem A with $k = 2$, and for this action, $\text{Ker}(\mathcal{P}_G)$ is finite.

1.9. Question: Normal generation. In [Luc25, Thm C], we show that for a certain action of $G = \mathbb{Z}/2\mathbb{Z}$ on $M = (S^1 \times S^2)^{\#k-1}$, the group $\text{Ker}(\mathcal{P}_G)$ is the $\Gamma_G(M)$ -normal closure of a finite set (the theorem is stated in the topological category, but can easily be adapted to the smooth category). The proof is special to this case; it exploits the isomorphism $\text{Mod}(S^3, C_k) \cong \text{SymOut}(F_k)$ described in Section 1.4 and appeals to the action of $\text{SymOut}(F_k)$ on a certain contractible simplicial complex constructed by McCullough–Miller [MM96]. We therefore ask:

Question 1.3. *Let M be a closed oriented smooth manifold of dimension $n \geq 3$, and G a finite group of orientation-preserving diffeomorphisms. Is $\text{Ker}(\mathcal{P}_G)$ the $\Gamma_G(M)$ -normal closure of a finite set?*

1.10. Free actions. For non-free group actions, the fixed points of G provide a natural obstruction to equivariant isotopy. In the case of free actions, it is less obvious how to obstruct equivariant isotopy, and indeed, there are classes of free actions for which $\text{Ker}(\mathcal{P}_G)$ is known to be trivial.

For instance, suppose M is a closed surface of genus $g \geq 1$ and G acts freely. Then by studying the action of diffeomorphisms on π_1 and appealing to $K(\pi, 1)$ -theory, one can show that $\text{Ker}(\mathcal{P}_G)$ is trivial (see e.g. [MW21, §9], which gives an argument following Birman–Hilden [BH73] and Aramayona–Leininger–Souto [ALS09]). This argument easily adapts to the case that M is a closed hyperbolic 3-manifold; the main technical ingredient is that homotopic diffeomorphisms are isotopic, which was proved by Gabai–Meyerhoff–Thurston [GMT03].

Beyond these examples, it appears little is known about $\text{Ker}(\mathcal{P}_G)$ in dimension $n \geq 3$ when G acts freely. A recent theorem of Raman [Ram25] uses Hodge theory to prove that $\text{Ker}(\mathcal{P}_G)$ is trivial if M is an irreducible hyperkähler manifold and G is a finite group of automorphisms acting freely. This provides examples of free actions for which $\text{Ker}(\mathcal{P}_G)$ is trivial in dimension $4n$ for $n \geq 1$; the simplest example comes from a free involution on a K3 surface (which has real dimension 4). Vogt [Vog77] proved that $\text{Ker}(\mathcal{P}_G)$ is trivial for certain free actions on Seifert fibered 3-manifolds, and applied this result to the study of foliations of certain 5-manifolds.

1.11. Outline. In Section 2, we describe certain diffeomorphisms of manifolds that we call *disk slides* and *summand slides*; these are the diffeomorphisms that we use to prove Theorem A. In Section 3, we prove part (i) of Theorem A. Sections 4 through 7 are dedicated to the proof of part (ii) of Theorem A. In Section 4, we give local and global models for the action of G on M , and construct a certain infinite cover \widetilde{M} whose homology is not finitely generated. In Section 5, we construct the obstruction map valued in the homology of \widetilde{M} . In Section 6, we construct an equivariant diffeomorphism which maps to a nontrivial homology class under the obstruction map. In Section 7, we complete the proof of part (ii) of Theorem A by using the element in Section 6 to build an infinite linearly independent subset of the image of the obstruction map. In Section 8, we use Theorem A to prove Theorem D.

1.12. Acknowledgements. We thank Bena Tshishiku for introducing us to this problem, for helpful comments on an earlier draft, and for many helpful conversations throughout this project.

We also thank Tom Goodwillie for a helpful discussion about the obstruction map constructed in Section 5.

2. DISK SLIDES AND SUMMAND SLIDES

In this section, we define certain “slide diffeomorphisms” which we will use to construct elements of $\text{Ker}(\mathcal{P}_G)$. We first define the notion of a “disk slide” in Section 2.1, and study their action on the fundamental group in Section 2.2. We then use disk slides to define “summand slides” in Section 2.3.

2.1. Disk slides. Let M be a compact oriented smooth manifold of dimension $n \geq 3$ (possibly with boundary). An *(ordered) multidisk* in M is a tuple $\Delta = (D_1, \dots, D_d)$, where D_1, \dots, D_d are disjoint closed smooth n -disks in the interior $\text{Int}(M)$. Given a multidisk Δ , we define $\text{Diff}_\partial(M \text{ rel } \Delta)$ to be the group of diffeomorphisms of M which fix ∂M and Δ pointwise, and we define $\text{Mod}(M \text{ rel } \Delta) := \pi_0(\text{Diff}_\partial(M \text{ rel } \Delta))$. Observe that we have a natural forgetful map $\text{Mod}(M \text{ rel } \Delta) \rightarrow \text{Mod}(M)$.

Definition 2.1. Given a smooth manifold M and a multidisk Δ in M , we define the group

$$\text{DS}_\Delta(M) := \text{Ker}(\text{Mod}(M \text{ rel } \Delta) \rightarrow \text{Mod}(M)),$$

and we define a *disk slide* to be any diffeomorphism of M that represents an element of DS_Δ .

Our goal in this subsection is to characterize disk slides in terms of the fundamental group of a certain bundle. First, let $\text{Conf}_d(M)$ denote the configuration space of d ordered points in the interior of M , i.e.

$$\text{Conf}_d(M) := \{(x_1, \dots, x_d) \in \text{Int}(M)^d \mid x_i \neq x_j \text{ for all } 1 \leq i \neq j \leq d\}.$$

Then we define the *multiframe bundle of degree d* to be the principal $\text{GL}_n(\mathbb{R})^d$ -bundle $\text{Fr}_d(M) \rightarrow \text{Conf}_d(M)$ where the fiber over (x_1, \dots, x_n) is the set of tuples $(\omega_1, \dots, \omega_n)$, where ω_i is a basis of the tangent space $T_{x_i}M$. In the case $d = 1$, $\text{Fr}_1(M)$ is simply the frame bundle $\text{Fr}(M) \rightarrow \text{Int}(M)$. To construct $\text{Fr}_d(M)$ formally, one can start with the direct product bundle $(\text{Fr}(M))^d$ on M^d and restrict it to the subspace $\text{Conf}_d(M) \subseteq M^d$. Alternatively, one can construct $\text{Fr}_d(M)$ as the subspace of $\text{Conf}_d(\text{Fr}(M))$ consisting of tuples of frames over distinct points.

Suppose now Δ is a *framed* multidisk, meaning that $\Delta = (D_1, \dots, D_d)$ is a multidisk equipped with a choice of frame ω_i at the center of each disk D_i . Then we let $\text{Fr}_\Delta(M)$ denote the frame bundle $\text{Fr}_d(M)$, equipped with the base point $(\omega_1, \dots, \omega_d)$. The main result of this subsection is then the following.

Proposition 2.2. *Let Δ be a framed multidisk in M . Then there is natural surjection*

$$\mathcal{DS}_\Delta : \pi_1(\text{Fr}_\Delta(M)) \twoheadrightarrow \text{DS}_\Delta(M).$$

We prove Proposition 2.2 with two standard lemmas about embedding spaces of disks. Given a multidisk $\Delta = (D_1, \dots, D_d)$, we let $\text{Emb}(\Delta, M)$ denote the space of embeddings $D_1 \sqcup \dots \sqcup D_d \hookrightarrow \text{Int}(M)$. Then the following two lemmas follow from the work of Cerf [Cer61].

Lemma 2.3. *The natural map $\text{Diff}_\partial(M) \rightarrow \text{Emb}(\Delta, M)$ is a locally trivial fibration, where $\text{Diff}_\partial(M)$ is the group of diffeomorphisms fixing ∂M pointwise.*

Lemma 2.3 is given in [Cer61, §II.2.2.2].

Lemma 2.4. *The natural map $\text{Emb}(\Delta, M) \rightarrow \text{Fr}_\Delta(M)$ is a homotopy equivalence.*

Lemma 2.4 is deduced in [Cer61, §II.5.1.5] in the case of single disk, but the proof can be adapted to the case of a multidisk with minor modifications.

Now, we can prove Proposition 2.2.

Proof of Proposition 2.2. By Lemma 2.3, we get an exact sequence

$$\pi_1(\text{Emb}(\Delta, M)) \rightarrow \text{Mod}(M \text{ rel } \Delta) \rightarrow \text{Mod}(M).$$

By exactness, the image of the first map is precisely $\text{DS}_\Delta(M)$. By Lemma 2.4, we have an isomorphism $\pi_1(\text{Fr}_\Delta(M)) \rightarrow \pi_1(\text{Emb}(\Delta, M))$. Thus we define \mathcal{DS}_Δ as the composition

$$\pi_1(\text{Fr}_\Delta(M)) \xrightarrow{\cong} \pi_1(\text{Emb}(\Delta, M)) \rightarrow \text{DS}_\Delta(M).$$

□

2.2. Disk slides, sphere twists, and inner automorphisms. Next, given a framed multidisk $\Delta = (D_1, \dots, D_d)$ in M , we can study the behavior of disks slides in terms of $\pi_1(\text{Fr}_\Delta(M))$. Observe first that since $\text{Fr}_\Delta(M)$ is a principal $\text{GL}_n(\mathbb{R})^d$ -bundle, we have an exact sequence

$$\pi_1(\text{GL}_n(\mathbb{R})^d) \rightarrow \pi_1(\text{Fr}_\Delta(M)) \rightarrow \pi_1(\text{Conf}_d(M)) \rightarrow \pi_0(\text{GL}_n(\mathbb{R})^d).$$

Since M is orientable, the last map is in fact trivial. So, we can understand disk slides in terms of $\pi_1(\text{GL}_n(\mathbb{R})^d)$ and $\pi_1(\text{Conf}_d(M))$. We will summarize our findings as Proposition 2.6 below. For a similar discussion in the case of 3-manifolds, see [HW10, Rem 2.4].

2.2.1. Sphere twists. First, recall that $\pi_1(\text{GL}_n(\mathbb{R})) \cong \mathbb{Z}/2\mathbb{Z}$, and it is generated by a 2π -rotation in the xy -plane. Then $\pi_1(\text{GL}_n(\mathbb{R})^d) \cong (\mathbb{Z}/2\mathbb{Z})^d$; let e_1, \dots, e_d be the images of standard generators in $\pi_1(\text{Fr}_\Delta(M))$. If we let p_i denote the center of D_i , then e_i is a 2π -rotation of our chosen frame at p_i . It follows that $\mathcal{DS}_\Delta(e_i)$ is the isotopy class of *sphere twist* about ∂D_i . This is a diffeomorphism of M supported on a collar neighborhood $\partial D_i \times [0, 1] \subseteq M - \text{Int}(D_i)$ which acts on the slice $\partial D_i \times \{t\}$ by a $2\pi t$ -rotation. In particular, $\mathcal{DS}_\Delta(e_i)$ acts trivially on $\pi_1(M, p_i)$. We emphasize that a sphere twist may be isotopically trivial, i.e. the composition $(\mathbb{Z}/2\mathbb{Z})^d \rightarrow \pi_1(\text{Fr}_\Delta(M)) \rightarrow \text{DS}_\Delta(M)$ need not be injective (in particular, if e_i is in the image of the composition

$$\pi_1(\text{Diff}_\partial(M)) \rightarrow \pi_1(\text{Emb}(\Delta, M)) \xrightarrow{\cong} \pi_1(\text{Fr}_\Delta(M)),$$

then $\mathcal{DS}_\Delta(e_i)$ will be trivial).

2.2.2. Inner automorphisms. Next, we can study disk slides in terms of $\pi_1(\text{Conf}_d(M))$. We let p_i denote the center of D_i , and fix (p_1, \dots, p_d) as a base point of $\text{Conf}_d(M)$. First, we can compute $\pi_1(\text{Conf}_d(M))$.

Lemma 2.5. *The natural map*

$$\pi_1(\text{Conf}_d(M)) \rightarrow \prod_{i=1}^d \pi_1(M, p_i)$$

is an isomorphism.

Proof. Fix local coordinates x^1, \dots, x^n on $\text{Int}(M)$. Then $\text{Int}(M)^d$ has local coordinates of the form

$$x_1^1, \dots, x_1^n, \dots, x_d^1, \dots, x_d^n,$$

and $\text{Conf}_d(M)$ is given locally as the complement of the linear subspaces

$$V_{i,j} = \{(x_1^1, \dots, x_1^n, \dots, x_d^1, \dots, x_d^n) \in \mathbb{R}^{nd} \mid x_i^k = x_j^k \text{ for all } 1 \leq k \leq n\}$$

for $1 \leq i \neq j \leq d$. Since the equation $x_i^k = x_j^k$ defines a codimension 1 subspace, it follows that $V_{i,j}$ has codimension n .

Thus, since $n \geq 3$, the space $\text{Conf}_d(M)$ is obtained from $\text{Int}(M)^d$ by removing a transverse intersection of closed submanifolds of codimension at least 3, which means that the map $\pi_1(\text{Conf}_d(M)) \rightarrow \pi_1(M^d)$ induces an isomorphism on π_1 . □

By Lemma 2.5, we have a natural map

$$\rho : \pi_1(\text{Fr}_\Delta(M)) \rightarrow \prod_{i=1}^d \pi_1(M, p_i).$$

We saw above that \mathcal{DS}_Δ sends elements in $\text{Ker}(\rho)$ to sphere twists, which act trivially on $\pi_1(M)$. Thus, we will use the map ρ to study the action of disk slides on $\pi_1(M)$.

Observe that for each $j \in \{1, \dots, d\}$, any element of $\text{Mod}(M \text{ rel } \Delta)$ fixes the point $p_j \in D_j$, and hence $\text{Mod}(M \text{ rel } \Delta)$ acts on $\pi_1(M, p_j)$. If $\alpha \in \text{DS}_\Delta(M)$ is represented by a disk slide f , then since f is isotopic to id_M , α must induce an inner automorphism on $\pi_1(M, p_j)$. Indeed, if we choose an isotopy h_t with $h_0 = \text{id}$ and $h_1 = f$, then f will conjugate $\pi_1(M, p_j)$ by the path of p_j under h_t .

Now, consider the composition

$$\pi_1(\text{Fr}_\Delta(M)) \xrightarrow{\mathcal{DS}_\Delta} \text{Mod}(M \text{ rel } \Delta) \rightarrow \text{Aut}(\pi_1(M, p_j)).$$

It follows from the construction of \mathcal{DS}_Δ and the discussion above that for $\gamma \in \pi_1(\text{Fr}_\Delta(M))$, the image of γ in $\text{Aut}(\pi_1(M, p_j))$ is the inner automorphism $\text{Inn}(\bar{\gamma})$, where $\bar{\gamma}$ is the image of γ in $\pi_1(M, p_j)$. Indeed, γ determines a loop in $\text{Emb}(\Delta, M)$, which we can extend to an ambient isotopy h_t with $h_0 = \text{id}_M$. Then $\mathcal{DS}_\Delta(\gamma)$ is represented by h_1 . By the discussion above, $\mathcal{DS}_\Delta(\gamma)$ conjugates $\pi_1(M, p_j)$ by the path traveled by p_j under h_t , which is precisely the loop $\bar{\gamma}$.

2.2.3. Conclusion. All together, we can summarize the discussion in this subsection as the following proposition.

Proposition 2.6. *Let $\Delta = (D_1, \dots, D_d)$ be a framed multidisk in M , and let p_i be the center of the disk D_i . Then there is an exact sequence*

$$(\mathbb{Z}/2\mathbb{Z})^d \rightarrow \pi_1(\text{Fr}_\Delta(M)) \rightarrow \prod_{i=1}^d \pi_1(M, p_i) \rightarrow 1.$$

Moreover:

- (i) *If e_i is the image of the i th generator under the map $(\mathbb{Z}/2\mathbb{Z})^d \rightarrow \pi_1(\text{Fr}_\Delta(M))$, then $\mathcal{DS}_\Delta(e_i)$ is the isotopy class of a sphere twist about ∂D_i .*
- (ii) *For each $j \in \{1, \dots, d\}$, the composite map*

$$\pi_1(\text{Fr}_\Delta(M)) \xrightarrow{\mathcal{DS}_\Delta} \text{Mod}(M \text{ rel } \Delta) \rightarrow \text{Aut}(\pi_1(M, p_j))$$

is equal to the composition

$$\pi_1(\text{Fr}_\Delta(M)) \rightarrow \prod_{i=1}^d \pi_1(M, p_i) \rightarrow \pi_1(M, p_j) \rightarrow \text{Inn}(\pi_1(M, p_j)).$$

2.3. Summand slides. Next, we will use disk slides to define *summand slides*, which are the diffeomorphisms we will use to prove Theorem A. The main result of this subsection is Proposition 2.8 below, which computes the action of a summand slide on π_1 .

First, we introduce the notion of gluing manifolds along multidisks (following [KM63]). Suppose M_1 and M_2 are two oriented smooth n -manifolds, and $\Delta_i = (D_{i,1}, \dots, D_{i,d})$ is a multidisk in the interior of M_i . Let $p_{i,j}$ denote the center of the disk $D_{i,j}$. We choose parametrizations $k_{i,j} : D^n \rightarrow D_{i,j}$, where D^n is the standard closed n -disk in \mathbb{R}^n . We require each $k_{1,j}$ to be orientation-preserving and $k_{2,j}$ to be orientation-reversing. Define

$$M_i^* := M_i \setminus \{p_{i,1}, \dots, p_{i,d}\}.$$

To glue M_1 and M_2 along Δ_1 and Δ_2 , we start with the disjoint union $M_1^* \sqcup M_2^*$. Then, using our parametrizations $k_{i,j}$, we identify $\text{Int}(D_{1,j}) \setminus \{p_{1,j}\}$ with $\text{Int}(D_{2,j}) \setminus \{p_{2,j}\}$ via the diffeomorphism of $\text{Int}(D^n) \setminus \{0\}$ defined by $tu \mapsto (1-t)u$ for each unit vector $u \in S^{n-1}$ and $t \in (0, 1)$. The resulting

manifold M is well-defined independent of the choice of $k_{i,j}$, but if M_1 and M_2 are both oriented, then M may depend on the choice of orientation of M_1 and M_2 (see [Pal60]). We will frequently identify M_i^* with its image in M .

With this convention set, we can define summand slides.

Definition 2.7. Suppose M is obtained by gluing two manifolds M_1 and M_2 along multidisks Δ_1 and Δ_2 . A *summand slide* is a diffeomorphism of M which acts on M_1^* by the restriction of a disk slide, and which acts on M_2^* by $\text{id}_{M_2^*}$.

Note that any summand slide is well-defined since a disk slide on M_1 fixes the multidisk Δ_1 pointwise.

An important special case is that $d = 1$, so M a connected sum $M_1 \# M_2$. In this case, we can easily compute the action of a summand slide on $\pi_1(M)$. Namely, suppose M_1 and M_2 are glued along disks $D_1 \subseteq M_1$ and $D_2 \subseteq M_2$. Let p_i be the center of D_i . Fix points $q_i \in D_i$ away from the center so that each q_i corresponds to the same point $q \in M$. By choosing arcs in each D_i from p_i to q_i , we get a natural free product decomposition

$$\pi_1(M, q) \cong \pi_1(M_1, p_1) * \pi_1(M_2, p_2).$$

Since any summand slide fixes D_1 and D_2 pointwise, it will induce well-defined automorphisms of $\pi_1(M, q)$ and $\pi_1(M_i, p_i)$ (in particular, it will act trivially on $\pi_1(M_2, p_2)$). Then we get the following.

Proposition 2.8. *Suppose $M = M_1 \# M_2$ as above. Choose $\gamma \in \pi_1(\text{Fr}(M_1))$ and let $\bar{\gamma} \in \pi_1(M_1)$ be its image. Suppose f is a summand slide of M where the action on M_1^* represents $\mathcal{DS}_{D_1}(\gamma)$. Then, with respect to the identification*

$$\pi_1(M, q) \cong \pi_1(M_1, p_1) * \pi_1(M_2, p_2),$$

the diffeomorphism f acts on $\pi_1(M)$ by conjugating $\pi_1(M_1)$ by $\bar{\gamma}$ and fixing $\pi_1(M_2)$ pointwise.

Proof. This follows directly from the definition of f and Proposition 2.6 above. \square

3. PROOF OF PART (I) OF THEOREM A

In this section, we prove part (i) of Theorem A. The proof is a relatively straightforward application of Propositions 2.6 and 2.8, and is essentially a warmup to part (ii) of Theorem A.

In Section 3.1, we establish some notation and a model of the group action that we will use throughout this section. In Sections 3.2 and 3.3, we describe how disk slides lift along the quotient map $M \rightarrow M/G$. In Section 3.4, we reinterpret the assumption on the homology class of B in terms of $\pi_1(M \setminus B)$. In Section 3.5 we construct the candidate subgroup of $\text{Ker}(\mathcal{P}_G)$, and we complete the proof in Section 3.6.

3.1. Preliminaries. We begin by establishing some notation. Let M and G be a manifold and group respectively which satisfy the hypotheses of Theorem A with $k = 3$. Let $R \subseteq M$ denote the set of points whose G -stabilizer is nontrivial. Then R is a deformation retract of a G -invariant open submanifold $U \subseteq M$. Indeed, by [Ill00], we can choose an equivariant triangulation of M with R as a G -invariant subcomplex, and then take U to be the interior of a regular neighborhood of R . Let $X = M \setminus U$, so X is a manifold with boundary on which G acts freely. Let $Y = X/G$. Note that if we let $M^\circ \subseteq M$ denote the set of points whose G -stabilizer is trivial, then X is a deformation retract of M° and Y is a deformation retract of M°/G .

Next, we fix connected sum decompositions of X and Y as follows. By assumption, we have a connected sum decomposition $M^\circ/G \cong Q_0 \# Q_1 \# Q_2$ where $\pi_1(Q_i)$ is nontrivial for each i . It follows that there is a connected sum decomposition $Y \cong Y_1 \# Y_2$ where $\pi_1(Y_1) \cong \pi_1(Q_0) * \pi_1(Q_1)$ and $\pi_1(Y_2) \cong \pi_1(Q_2)$. Following the convention of Section 2.3, we choose n -disks $D_1 \subseteq Y_1$ and $D_2 \subseteq Y_2$ and view Y as obtained by gluing Y_1 and Y_2 along these disks. We let Y_i^* (resp. D_i^*) denote Y_i (resp. D_i) minus the center of D_i , and we identify Y_i^* with its image in Y . Let $X_i^* \subseteq X$ denote the

preimage of Y_i^* , and let Δ_i^* denote the preimage of D_i^* . Then let X_i and Δ_i be obtained by filling in the punctures. Thus we can view X as obtained by gluing together X_1 and X_2 along Δ_1 and Δ_2 (note that each X_i may be disconnected). Note that we have regular (unbranched) G -coverings $X_i \rightarrow Y_i$ and $X_i^* \rightarrow Y_i^*$.

Recall that by assumption, M has a distinguished codimension 2 compact orientable submanifold B which is fixed by an element $g_0 \in G$. The submanifold B is contained in R . Let U_B be the component of U containing B . Then ∂U_B is a boundary component of X ; we assume without loss of generality that this boundary component lies on X_1^* . In particular, X_1 and Y_1 both have nonempty boundary.

Observe that the inclusion $X \hookrightarrow M$ induces a way to fill in each boundary component of X_1 and X_2 . We call the resulting filled manifolds M_1 and M_2 respectively, so M is obtained by gluing M_1 and M_2 along the multidisks Δ_1 and Δ_2 . See Figure 1 for an illustration.




FIGURE 1. A schematic illustration of the manifolds X and Y . The trefoils and unknots (in magenta) represent the set R and its image in M/G ; the manifolds X and Y are the complement of a tubular neighborhood of these curves.

Finally, we establish some notation for the various mapping class groups that arise throughout this section. Given an oriented manifold W , we let $\text{Diff}_\partial(W)$ denote the group of orientation-preserving diffeomorphisms fixing ∂W pointwise, and let $\text{Mod}(W) = \pi_0(\text{Diff}_\partial(W))$. Similarly, given a closed submanifold $Z \subseteq W$, we let $\text{Diff}_\partial(W \text{ rel } Z)$ denote the subgroup of diffeomorphisms fixing Z pointwise, and let $\text{Mod}(W \text{ rel } Z) = \pi_0(\text{Diff}_\partial(W \text{ rel } Z))$. Suppose $W' \rightarrow W$ is a finite cover. Then we let $\text{LDiff}_\partial(W) \leq \text{Diff}_\partial(W)$ and $\text{LDiff}_\partial(W \text{ rel } Z) \leq \text{Diff}_\partial(W \text{ rel } Z)$ denote the finite index

subgroups of diffeomorphisms that lift to a diffeomorphism of W' that fixes $\partial W'$ pointwise, and define $\text{LMod}(W) = \pi_0(\text{LDiff}_\partial(W))$ and $\text{LMod}(W \text{ rel } Z) = \pi_0(\text{Diff}_\partial(W \text{ rel } Z))$.

3.2. Lifting disk slides. Our next goal is to describe lifts of summand slides of Y on the level of frame bundles. It will be enough to describe lifts of disk slides on Y_1 .

Fix an arbitrary framing on D_1 . Since the covering map $X_1 \rightarrow Y_1$ is a local diffeomorphism, this induces framings on Δ_1 . Recall from Proposition 2.6 that we have an exact sequence

$$\mathbb{Z}/2\mathbb{Z} \rightarrow \pi_1(\text{Fr}_{D_1}(Y_1)) \rightarrow \pi_1(Y_1) \rightarrow 1.$$

Here $\text{Fr}_{D_1}(Y_1)$ is simply the frame bundle $\text{Fr}(Y_1)$ with a base point at our chosen frame at the center of D_1 . Let $\pi_1(Y_1)^{\text{lift}} \leq \pi_1(Y_1)$ be the subgroup of loops that lift to loops in X_1 (equivalently, $\pi_1(Y_1)^{\text{lift}}$ is the image of the fundamental group of any component of X_1). Let $\pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}} \leq \pi_1(\text{Fr}_{D_1}(Y_1))$ be the preimage of $\pi_1(Y_1)^{\text{lift}}$.

Now, observe that there's a natural map $\text{Fr}_{\Delta_1}(X_1) \rightarrow \text{Fr}_{D_1}(Y_1)$ which maps a tuple $(\omega_1, \dots, \omega_d)$ to the image of ω_1 in Y_1 . Then we get the following.

Lemma 3.1. *The map $\pi_1(\text{Fr}_{\Delta_1}(X_1)) \rightarrow \pi_1(\text{Fr}_{D_1}(Y_1))$ above admits a section*

$$\pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}} \rightarrow \pi_1(\text{Fr}_{\Delta_1}(X_1))$$

which fits into the following commutative square:

$$\begin{array}{ccc} \pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}} & \longrightarrow & \pi_1(\text{Fr}_{\Delta_1}(X_1)) \\ \mathcal{DS}_{D_1} \downarrow & & \downarrow \mathcal{DS}_{\Delta_1} \\ \text{LMod}(Y_1 \text{ rel } D_1) & \dashrightarrow & \text{Mod}(X_1 \text{ rel } \Delta_1) \end{array}$$

Here, the dashed arrow is defined on a finite index subgroup containing the image of $\pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}}$, and sends a diffeomorphism to its unique lift that fixes ∂X_1 .

Before proving the lemma, we will have to introduce *unordered* embedding spaces and multiframe bundles.

Let $d = |G|$, so $X_1 \rightarrow Y_1$ is a degree d cover. Note that since Δ_1 is simply the preimage of D_1 in X_1 , we can identify $\text{Emb}(\Delta_1, X_1)$ with the space of embeddings $\bigsqcup_{i=1}^d D_1 \hookrightarrow \text{Int}(X_1)$. In particular, $\text{Emb}(\Delta_1, X_1)$ admits a natural action of the symmetric group S_d by precomposition; we define the quotient

$$\text{UEmb}(\Delta_1, X_1) := \text{Emb}(\Delta_1, X_1)/S_d,$$

so $\text{UEmb}(\Delta_1, X_1)$ is the space of unordered embeddings.

Similarly, we have a natural action of S_d on $\text{Conf}_d(X_1)$; we let $\text{UConf}_d(X_1)$ denote the quotient by this action, so $\text{UConf}_d(X_1)$ is the configuration space of d unordered points in $\text{Int}(X_1)$. Then the multiframe bundle $\text{Fr}_{\Delta_1}(X_1)$ descends to a bundle

$$\text{UFr}_{\Delta_1}(X_1) \rightarrow \text{UConf}_d(X_1),$$

where the fiber over an unordered set of points $\{x_1, \dots, x_d\}$ is an unordered set $\{\omega_1, \dots, \omega_d\}$ where ω_i is a basis for $T_{x_i}X_1$.

Note that we have the following commutative diagram with exact rows:

$$(3.1) \quad \begin{array}{ccccccc} 1 & \longrightarrow & \pi_1(\text{Emb}(\Delta_1, X_1)) & \longrightarrow & \pi_1(\text{UEmb}(\Delta_1, X_1)) & \longrightarrow & S_d \longrightarrow 1 \\ & & \downarrow & & \downarrow & & \downarrow \\ 1 & \longrightarrow & \pi_1(\text{Fr}_{\Delta_1}(X_1)) & \longrightarrow & \pi_1(\text{UFr}_{\Delta_1}(X_1)) & \longrightarrow & S_d \longrightarrow 1 \end{array}$$

By the 5 Lemma, we conclude that $\pi_1(\text{UEmb}(\Delta_1, X_1)) \cong \pi_1(\text{UFr}_{\Delta_1}(X_1))$.

Now, we can prove the lemma.

Proof of Lemma 3.1. Observe that by lifting an embedding $D_1 \hookrightarrow Y_1$, we get a natural map

$$\text{Emb}(D_1, Y_1) \rightarrow \text{UEmb}(\Delta_1, X_1).$$

This map fits into the following commutative diagram:

$$\begin{array}{ccccc} \text{LDiff}_{\partial}(Y_1 \text{ rel } D_1) & \longrightarrow & \text{LDiff}_{\partial}(Y_1) & \longrightarrow & \text{Emb}(D_1, Y_1) \\ \downarrow & & \downarrow & & \downarrow \\ \text{Diff}_{\partial}(X_1 \text{ rel}_{S_d} \Delta_1) & \longrightarrow & \text{Diff}_{\partial}(X_1) & \longrightarrow & \text{UEmb}(\Delta_1, X_1) \end{array}$$

Here, the middle vertical map sends a liftable diffeomorphism to the unique lift that fixes ∂X_1 (recall that X_1 and Y_1 both have a boundary component coming from B). The group $\text{Diff}_{\partial}(X_1 \text{ rel}_{S_d} \Delta_1)$ is the group of diffeomorphisms of X_1 fixing ∂X_1 pointwise and fixing the natural embedding $\Delta_1 \hookrightarrow X_1$ up to precomposition by an element of S_d , and the left vertical map is the restriction of the middle vertical map. By Lemma 2.3 (replacing M with X_1 and Y_1 respectively), each row in this diagram is a fibration.

Now, observe that by mapping a frame in Y_1 to the set of frames above it in X_1 , we have a natural map

$$\text{Fr}_{D_1}(Y_1) \rightarrow \text{UFr}_{\Delta_1}(X_1).$$

Combining this map with the diagram (3.1) and the long exact sequences of the fibrations above, we get the following commutative diagram:

$$\begin{array}{ccc} \pi_1(\text{Fr}_{D_1}(Y_1)) & \longrightarrow & \pi_1(\text{UFr}_{\Delta_1}(X_1)) \\ \cong \downarrow & & \downarrow \cong \\ \pi_1(\text{Emb}(D_1, Y_1)) & \longrightarrow & \pi_1(\text{UEmb}(\Delta_1, Y_1)) \\ \downarrow & & \downarrow \\ \text{LMod}(Y_1 \text{ rel } \Delta_1) & \longrightarrow & \text{Mod}(X_1 \text{ rel}_{S_d} \Delta_1) \end{array}$$

Here $\text{Mod}(X_1 \text{ rel}_{S_d} \Delta_1) = \pi_0(\text{Diff}_{\partial}(X_1 \text{ rel}_{S_d} \Delta_1))$.

Now, we can use the above diagram to deduce the desired commutative square. Following the proof of Proposition 2.2, we see that the composition of the left vertical maps is precisely \mathcal{DS}_{D_1} . Next, recall from (3.1) that $\pi_1(\text{Fr}_{\Delta_1}(X_1))$ is naturally a subgroup of $\pi_1(\text{UFr}_{\Delta_1}(X_1))$ (and the analogous statement holds for Emb and UEmb). Composing the right vertical maps in the above diagram yields a map $\pi_1(\text{UFr}_{\Delta_1}(X_1)) \rightarrow \text{Mod}(X_1 \text{ rel}_{S_d} \Delta)$, and the restriction to the subgroup $\pi_1(\text{Fr}_{\Delta_1}(X_1))$ is precisely the map \mathcal{DS}_{Δ_1} . Finally, it follows by the definition of $\pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}}$ that the image of $\pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}}$ under the map $\pi_1(\text{Fr}_{D_1}(Y_1)) \rightarrow \pi_1(\text{UFr}_{\Delta_1}(X_1))$ lands in the subgroup $\pi_1(\text{Fr}_{\Delta_1}(X_1))$. This yields the desired commutative square, where the dashed arrow is defined on the preimage of the finite index subgroup $\text{Mod}(X_1 \text{ rel } \Delta_1) \leq \text{Mod}(X_1 \text{ rel}_{S_d} \Delta_1)$. The fact that $\pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}} \rightarrow \pi_1(\text{Fr}_{\Delta_1}(X_1))$ is a section follows by construction. \square

3.3. Lifting and including. Now, we have the following commutative diagram:

$$(3.2) \quad \begin{array}{ccccc} \pi_1(\text{Fr}_{\Delta_1}(Y_1))^{\text{lift}} & \longrightarrow & \pi_1(\text{Fr}_{\Delta_1}(X_1)) & \longrightarrow & \pi_1(\text{Fr}_{\Delta_1}(M_1)) \\ \downarrow \mathcal{DS}_{D_1} & & \downarrow \mathcal{DS}_{\Delta_1} & & \downarrow \mathcal{DS}_{\Delta_1} \\ \text{LMod}(Y_1 \text{ rel } D_1) & \dashrightarrow & \text{Mod}(X_1 \text{ rel } \Delta_1) & \longrightarrow & \text{Mod}(M_1 \text{ rel } \Delta_1) \\ \downarrow & & \downarrow & & \downarrow \\ \text{LMod}(Y) & \longrightarrow & \text{Mod}(X) & \longrightarrow & \text{Mod}(M) \end{array}$$

The top-left square comes directly from Lemma 3.1. The top right square is induced by the inclusion $X_1 \hookrightarrow M_1$. In the bottom right square, the horizontal maps are induced by the inclusions $X_1 \hookrightarrow M_1$ and $X \hookrightarrow M$, and the vertical maps are induced by the inclusions $X_1^* \hookrightarrow X$ and $M_1^* \hookrightarrow M$. In the bottom left square, the bottom horizontal map is induced by sending a diffeomorphism to its unique lift that fixes ∂X , and the vertical maps are again induced by inclusions.

3.4. A meridian of the fixed set. Next, we will identify a loop in M that we will use to construct the summand slides needed to prove part (i) of Theorem A. This loop will be a *meridian* of B , i.e. the boundary of a fiber of a tubular neighborhood of B .

Lemma 3.2. *The submanifold B has a meridian μ which lies in X and represents a nontrivial element of $\pi_1(X)$.*

To prove Lemma 3.2, we appeal to the following general fact about meridians of codimension 2 submanifolds.

Lemma 3.3. *Let W be a closed oriented smooth manifold of dimension $n \geq 3$, and let $Z \subseteq W$ be a compact oriented codimension 2 submanifold without boundary. Let μ be a meridian of Z . Then for any prime p , the meridian μ represents a nontrivial element of $H_1(W \setminus Z; \mathbb{Z}/p\mathbb{Z})$ if and only if the homology class $[Z] \in H_{n-2}(W; \mathbb{Z}/p\mathbb{Z})$ is trivial.*

Proof. Let $\nu(Z)$ denote the normal bundle of Z , so $\nu(Z)$ is a D^2 -bundle over Z . By excision, we have an isomorphism

$$H_2(W, W \setminus Z; \mathbb{Z}/p\mathbb{Z}) \cong H_2(\nu(Z), \partial\nu(Z); \mathbb{Z}/p\mathbb{Z}).$$

By the Thom isomorphism and Poincaré duality, we have an isomorphism $H_2(\nu(Z), \partial\nu(Z); \mathbb{Z}/p\mathbb{Z}) \cong \mathbb{Z}/p\mathbb{Z}$, and a homology class in $H_2(\nu(Z), \partial\nu(Z); \mathbb{Z}/p\mathbb{Z})$ is determined by its $\mathbb{Z}/p\mathbb{Z}$ -intersection number with Z . This implies in particular that $H_2(\nu(Z), \partial\nu(Z); \mathbb{Z}/p\mathbb{Z})$ is generated by a fiber of $\nu(Z)$.

Now, from the long exact sequence of the pair $(W, W \setminus Z)$, we have an exact sequence

$$H_2(W; \mathbb{Z}/p\mathbb{Z}) \rightarrow H_2(\nu(Z), \partial\nu(Z); \mathbb{Z}/p\mathbb{Z}) \rightarrow H_1(W \setminus Z; \mathbb{Z}/p\mathbb{Z}).$$

The second map sends a fiber of $\nu(Z)$ to its boundary, i.e. to the homology class of μ . The first map sends a homology class of W to its $\mathbb{Z}/p\mathbb{Z}$ -intersection number with Z . Thus, the homology class $[Z] \in H_2(W; \mathbb{Z}/p\mathbb{Z})$ is trivial if and only if the first map is trivial, which occurs if and only if the second map is injective, which occurs if and only if the meridian μ is nontrivial in $H_1(W \setminus Z; \mathbb{Z}/p\mathbb{Z})$. \square

Now, we can prove Lemma 3.2.

Proof of Lemma 3.2. Fix $b \in B$, and let G_b denote the G -stabilizer of b . We may assume that G acts by isometries with respect to some Riemannian metric on M . Then the exponential map $T_b M \rightarrow M$ is a G_b -equivariant diffeomorphism near the origin. Since G acts by orientation-preserving diffeomorphisms, we know that G_b does not fix any codimension 1 subspace of $T_b M$. Thus near b , the preimage of M° in $T_b M$ is the complement of a finite union of subspaces of codimension at least 2. It follows that we can find a meridian μ of B contained in M° , and up to homotopy, we may assume μ lies in X . Then by Lemma 3.3 (applied to some prime p dividing ℓ), the class of μ in $\pi_1(X)$ maps to a nontrivial element under the composition

$$\pi_1(X) \rightarrow \pi_1(M \setminus B) \rightarrow H_1(M \setminus B; \mathbb{Z}) \rightarrow H_1(M \setminus B; \mathbb{Z}/\ell\mathbb{Z}),$$

and is therefore nontrivial. \square

3.5. Constructing the subgroup. Our next goal is to construct a subgroup $\Lambda \leq \pi_1(\text{Fr}_{\Delta_1}(Y_1))^{\text{lift}}$ which maps onto a virtual free product in $\text{LMod}(Y)$, and which maps to a finite subgroup in $\text{Mod}(M)$. From here, we will be able to deduce part (i) of Theorem A.

We construct the subgroup Λ as follows. Let μ be a meridian of B from Lemma 3.2. Since B lies in M_1 , we may assume without loss of generality that μ is contained in X_1 and represents a nontrivial conjugacy class in $\pi_1(X_1)$. Let $\bar{\gamma} \in \pi_1(X_1)$ be an element in this conjugacy class. Let $\bar{\delta} \in \pi_1(Y_1)$ be the image of $\bar{\gamma}$; by construction, $\bar{\delta}$ lies in the subgroup $\pi_1(Y_1)^{\text{lift}}$. Let $\delta \in \pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}}$ be an element in the preimage of $\bar{\delta}$. Recall that we have a free splitting

$$\pi_1(Y_1) \cong \pi_1(Q_0) * \pi_1(Q_1).$$

Without loss of generality, we may assume $\bar{\delta}$ lies in the subgroup $\pi_1(Q_0)$. Let $\bar{\delta}'$ be a conjugate of $\bar{\delta}$ by a nontrivial element of $\pi_1(Q_1)$, and let $\delta' \leq \pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}}$ be an element in the preimage of $\bar{\delta}'$. We define $\Lambda \leq \pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}}$ to be the subgroup generated by δ and δ' . See Figure 2 for an illustration of the curves $\bar{\gamma}$ and $\bar{\delta}$.

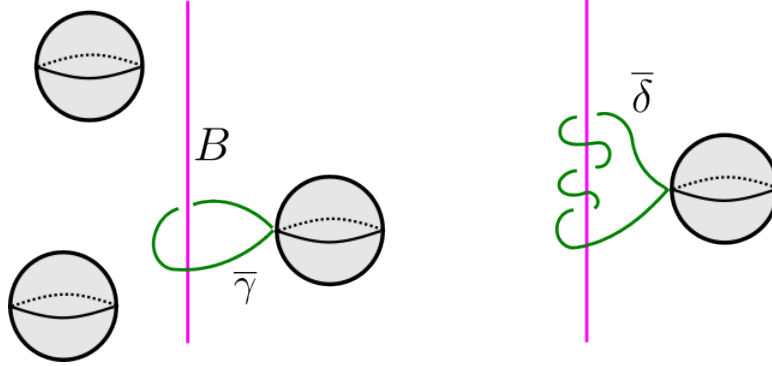


FIGURE 2. The curves $\bar{\gamma}$ and $\bar{\delta}$. The 3-disk on the right represents D_1 , and the 3-disks on the left represent a subset of Δ_1 .

Now, we can prove that Λ has the desired properties using the commutative diagram (3.2).

Lemma 3.4. *Let $\Lambda' \leq \text{LMod}(Y)$ be the image of Λ , and let $\Lambda'' \leq \text{Out}(\pi_1(Y))$ be the image of Λ' . Then Λ'' is a nontrivial free product of cyclic groups, and the surjection $\Lambda' \rightarrow \Lambda''$ has a finite kernel.*

Proof. Let α_δ and $\alpha_{\delta'}$ denote the images of δ and δ' in $\text{LMod}(Y)$, so Λ' is generated by α_δ and $\alpha_{\delta'}$, and Λ'' is generated by the images of α_δ and $\alpha_{\delta'}$ in $\text{Out}(\pi_1(Y))$.

First, we can identify the group Λ'' more explicitly. Let $\text{Inn}(\bar{\delta})$ and $\text{Inn}(\bar{\delta}')$ denote the inner automorphisms of $\pi_1(Y_1)$ determined by $\bar{\delta}$ and $\bar{\delta}'$. We have a nontrivial free splitting

$$\pi_1(Y) \cong \pi_1(Y_1) * \pi_1(Y_2).$$

Each element of $\text{Inn}(\pi_1(Y_1))$ extends to an automorphism of $\pi_1(Y)$ which acts trivially on $\pi_1(Y_2)$, and these automorphisms are never inner (except for the identity), so we get an embedding $\text{Inn}(\pi_1(Y_1)) \hookrightarrow \text{Out}(\pi_1(Y))$. Now, α_δ and $\alpha_{\delta'}$ are summand slides, and they act on Y_1 by $\mathcal{DS}_{D_1}(\bar{\delta})$ and $\mathcal{DS}_{D_1}(\bar{\delta}')$ respectively. Thus by Proposition 2.8, we see that Λ'' is the image of $\langle \text{Inn}(\bar{\delta}), \text{Inn}(\bar{\delta}') \rangle$ under the embedding $\text{Inn}(\pi_1(Y_1)) \hookrightarrow \text{Out}(\pi_1(Y))$.

Now, we can complete the proof. Since $\pi_1(Y_1)$ is a free product, it has trivial center, and hence the map $\pi_1(Y_1) \rightarrow \text{Inn}(\pi_1(Y_1))$ is an isomorphism. We therefore know that Λ'' is a free product since $\bar{\delta}$ and $\bar{\delta}'$ generate a free product in $\pi_1(Y_1)$. It remains just to check that the map $\Lambda' \rightarrow \Lambda''$ has a finite kernel. But this follows because the map $\Lambda \rightarrow \Lambda''$ has a finite kernel by Proposition 2.6 (applied to $M = Y_1$ and $\Delta = D_1$). \square

Lemma 3.5. *The group Λ maps onto a finite subgroup in $\text{Mod}(M)$.*

Proof. From the diagram (3.2), it's enough to show that Λ maps to a finite subgroup of $\pi_1(\text{Fr}_{\Delta_1}(M_1))$. Let $\tilde{\delta}$ and $\tilde{\delta}'$ be the images of δ and δ' in $\pi_1(\text{Fr}_{\Delta_1}(X_1))$. Let $d = |G|$. By Proposition 2.6, we have an exact sequence

$$(\mathbb{Z}/2\mathbb{Z})^d \rightarrow \pi_1(\text{Fr}_{\Delta_1}(X_1)) \rightarrow \pi_1(X_1)^d \rightarrow 1$$

(note that in the product $\pi_1(X_1)^d$, each factor has a different basepoint, which we omit for simplicity). By construction of the map $\pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}} \rightarrow \pi_1(\text{Fr}_{\Delta_1}(X_1))$, the image of $\tilde{\delta}$ in $\pi_1(X_1)^d$ is a d -tuple comprised of the G -orbit of $\bar{\gamma}$, and in particular it is the G -orbit of a loop in the free homotopy class of μ . Similarly, since $\tilde{\delta}'$ is a conjugate of $\tilde{\delta}$, the image of $\tilde{\delta}'$ in $\pi_1(X_1)^d$ is a d -tuple comprised of the G -orbit of a loop in the free homotopy class of μ .

Proposition 2.6 also gives us an exact sequence

$$(\mathbb{Z}/2\mathbb{Z})^d \rightarrow \pi_1(\text{Fr}_{\Delta_1}(M_1)) \rightarrow \pi_1(M_1)^d \rightarrow 1.$$

Since μ is homotopically trivial in M_1 , it follows from the discussion above that under the map $\pi_1(\text{Fr}_{\Delta_1}(X_1)) \rightarrow \pi_1(\text{Fr}_{\Delta_1}(M_1))$, the elements $\tilde{\delta}$ and $\tilde{\delta}'$ are sent to the kernel of the map $\pi_1(\text{Fr}_{\Delta_1}(M_1)) \rightarrow \pi_1(M_1)^d$, and hence lie in the image of $(\mathbb{Z}/2\mathbb{Z})^d$ in $\pi_1(\text{Fr}_{\Delta_1}(M_1))$. \square

3.6. Completing the proof. Finally, we can complete the proof of part (i) of Theorem A as follows.

By Lemmas 3.4 and 3.5, the subgroup $\Lambda \leq \pi_1(\text{Fr}_{D_1}(Y_1))^{\text{lift}}$ maps to a subgroup $\Lambda' \leq \text{LMod}(Y)$ which is a finite extension of a free product, and Λ maps to a finite group in $\text{Mod}(M)$. Then Λ' has a finite index subgroup Ω which maps to the identity in $\text{Mod}(M)$.

We have a lifting map $\text{LDiff}_{\partial}(Y) \rightarrow \text{Diff}_{\partial}(X)$ that maps a diffeomorphism to the unique lift fixing ∂X (which is nonempty). The image of this map is contained in the normalizer of G in $\text{Diff}_{\partial}(X)$, which contains the centralizer $\text{Diff}_{\partial}(X)^G$ as a finite index subgroup. Let $\text{LDiff}_{\partial}^G(Y)$ denote the preimage of $\text{Diff}_{\partial}(X)^G$. Then $\pi_0(\text{LDiff}_{\partial}^G(Y))$ has finite index in $\text{LMod}(Y)$. Thus, up to passing to a further finite index subgroup of Ω , we may assume that each diffeomorphism of Y representing an element of Ω lifts to an equivariant diffeomorphism of X . This implies that Ω embeds into $\pi_0(\text{Diff}_{\partial}(X)^G)$; indeed, if f represents an element of Ω and \tilde{f} is its equivariant lift, then any equivariant isotopy $\tilde{f} \simeq \text{id}_X$ will descend to an isotopy $f \simeq \text{id}_Y$.

We have a natural inclusion $\text{Diff}_{\partial}(X) \hookrightarrow \text{Diff}(M)$ by extending a diffeomorphism by the identity, and this restricts to an inclusion $\text{Diff}_{\partial}(X)^G \hookrightarrow \text{Diff}(M)^G$. We claim that the composition

$$\Omega \hookrightarrow \pi_0(\text{Diff}_{\partial}(X)^G) \rightarrow \pi_0(\text{Diff}(M)^G)$$

has a finite kernel. To see this, suppose $f \in \text{Diff}_{\partial}(Y)$ represents an element of Ω and $\tilde{f} \in \text{Diff}_{\partial}(X)^G$ is its lift. Let $\tilde{f}' \in \text{Diff}(M)^G$ be the natural extension of \tilde{f} (so \tilde{f}' represents the image of $[f]$ in $\pi_0(\text{Diff}(M)^G)$). Then \tilde{f}' restricts to a diffeomorphism of M° , and thus descends to a diffeomorphism f' of M°/G which extends f (recall that Y is a deformation retract of M°/G). Suppose now that f represents an element in the kernel of the map $\Omega \rightarrow \pi_0(\text{Diff}(M)^G)$, meaning there is an equivariant isotopy $\tilde{f}' \simeq \text{id}_M$. This isotopy must preserve the set $R \subseteq M$ of points with nontrivial G -stabilizers, and thus this isotopy restricts to M° and descends to an isotopy $f' \simeq \text{id}_{M^\circ/G}$. Since f' is an extension of f and Y is a deformation retract of M°/G , this implies that f is homotopic to id_Y . Then f must represent an element of the kernel of $\Omega \rightarrow \text{Out}(\pi_1(Y))$, which we know is finite by Lemma 3.4.

Thus Ω virtually embeds into the kernel of the natural map

$$\mathcal{P}_G : \pi_0(\text{Diff}(M)^G) \rightarrow \pi_0(\text{Diff}(M)),$$

which completes the proof.

4. A MODEL OF THE ACTION, AND AN INFINITE COVER

We now begin towards the proof of part (ii) of Theorem A. Fix a manifold M and a group G satisfying the assumptions of Theorem A with $k = 3$. By assumption, we have a codimension 2 submanifold $B \subseteq M$ which is fixed by some element of G and which represents a trivial homology class in $H_{n-2}(M; \mathbb{Z}/\ell\mathbb{Z})$ for some $\ell > 1$. In this section, we construct local and global models of the action of G on M , and use them to construct an infinite sheeted cover $\pi : \widetilde{M} \rightarrow M$ that we will use in the proof of Theorem A.

In Section 4.1, we describe the local action of G near a generic point of B . In Section 4.2, we construct a global model of the action by decomposing M according to the connected summands of M°/G . In Sections 4.3 and 4.4, we describe $\pi_1(M)$ in terms of this global model. In Section 4.5, we construct the cover \widetilde{M} using this description of $\pi_1(M)$, and in Section 4.6, we show that the homology of \widetilde{M} is infinite dimensional.

4.1. A local model of the action. We begin by describing the local behavior of the action of G near a generic point of B .

Lemma 4.1. *There exists a point $b_0 \in B$ whose stabilizer $G_0 := \text{Stab}_G(b_0)$ is cyclic and fixes B pointwise. Moreover, each point b in the G -orbit of b_0 has a closed neighborhood Z_b which is invariant under the stabilizer $G_b := \text{Stab}_G(b)$ and satisfies the following properties:*

- (i) *There is a diffeomorphism $Z_b \cong D^{n-2} \times D^2$ under which G_b fixes the D^{n-2} -factor pointwise and acts on the D^2 -factor by a $\frac{2\pi}{m}$ -rotation, where $m = |G_b|$.*
- (ii) *For any two distinct points b and b' in the G -orbit of b_0 , the neighborhoods Z_b and $Z_{b'}$ are disjoint.*
- (iii) *For any $g \in G$, $g(Z_b) = Z_{g(b)}$.*

Note that Lemma 4.1 requires the hypothesis that G acts by orientation-preserving diffeomorphisms. Indeed, the lemma fails if we let $M = S^3$, let $G = \mathbb{Z}_2 \times \mathbb{Z}_2$ where the two generators act by reflections in the xy - and xz -planes (viewing S^3 as $\mathbb{R}^3 \cup \{\infty\}$), and let B be the x -axis, which is the fixed set of the element $g_0 = (1, 1)$.

Proof of Lemma 4.1. Given any $x \in M$, we let G_x denote its stabilizer $\text{Stab}_G(x)$. Assume G fixes some Riemannian metric on M . Then for any $x \in M$, the stabilizer G_x acts on the tangent space $T_x M$ as a finite subgroup of $\text{SO}(n)$. Moreover, the exponential map $T_x M \rightarrow M$ is a G_x -equivariant local diffeomorphism at $0 \in T_x M$. The submanifold B is a closed totally geodesic submanifold of M , since it is a component of $\text{Fix}(g_0)$ (see [Kob72, Thm 5.1]).

Let $H \leq G$ be the subgroup of elements that fix B pointwise. Let T be an H -invariant closed tubular neighborhood of B (see [Bre72, Thm VI.2.2] for the existence of such a neighborhood). Since B has codimension 2, each point $x \in B$ has a closed neighborhood Z with a diffeomorphism $Z \cong D^{n-2} \times D^2$ taking $Z \cap B$ to $D^{n-2} \times \{0\}$, and taking each fiber of T to a fiber $\{y\} \times D^2$.

For $x \in B$, the action of H on $T_x M$ fixes a codimension 2 subspace $W \subseteq T_x M$, and the exponential map yields an isomorphism between the action of H on the D^2 -fibers of T and the linear action of H on a small 2-disk $D \subseteq W^\perp$. Note that H must act faithfully on $T_x M$; indeed, if $h \in H$ fixes $T_x M$ pointwise, then the fixed set $\text{Fix}(h) \subseteq M$ must be a closed submanifold of codimension 0, and thus $\text{Fix}(h) = M$. Since H acts trivially on W , it must be that H acts faithfully on W^\perp . Since G is orientation-preserving, we know that no element of H can act on W^\perp by a reflection, and thus H acts on W^\perp by a finite subgroup of $\text{SO}(2)$. Thus H must be cyclic and act on the D^2 -factor of Z by rotations.

Now, we claim that there exists a point $b_0 \in B$ whose stabilizer $G_0 := G_{b_0}$ is equal to H . Assuming this claim, we can prove the lemma. Indeed, let Z_{b_0} be the neighborhood Z described above, and for $b = gb_0$, we define $Z_b := g(Z_{b_0})$. This is well-defined since Z_{b_0} is G_0 -invariant. Then the fact that

G_0 is cyclic and statement (i) follow from the discussion above, statement (ii) follows by choosing Z_{b_0} sufficiently small, and statement (iii) is automatic.

Thus, it remains to show that there exists a point $b_0 \in B$ such that $G_{b_0} = H$. Choose any $x \in B$. Observe that $H \leq G_x$ automatically, so suppose that there exists $g \in G_x \setminus H$. Let $W \subseteq T_x M$ be the codimension 2 subspace spanned by $\exp^{-1}(V)$, where $V \subseteq B$ is a small neighborhood of x . Note that g cannot fix W pointwise; otherwise g would fix an open subset of B pointwise, which would imply that g fixes B pointwise (as B is totally geodesic and g is an isometry), contradicting that $g \notin H$. Thus if we let $Q \subseteq T_x M$ denote the fixed set of g , the intersection $W \cap Q$ has positive codimension in W . Thus, there is an open subset $U \subseteq B$ such that $g \notin G_y$ for any $y \in U$. We can repeat the argument inductively; choosing $y \in U$, if there exists $g' \in G_y \setminus H$, we can pass to a smaller neighborhood where g' does not stabilize any point. Since G is finite, we must eventually find a point $b_0 \in B$ such that $G_{b_0} \setminus H$ is empty. \square

4.2. A global model of the action. Next, we can describe how G acts on the connected summands of M . This model will be slightly different from the one in Section 3.1. Recall that we let M° denote the set of points whose G -stabilizer is trivial. Let $N = M^\circ/G$. By assumption, $N \cong Q_0 \# Q_1 \# Q_2$ where $\pi_1(Q_i)$ is nontrivial for each i .

We begin by constructing a model of N . Let D_1^+ and D_2^+ be embedded n -disks in Q_0 . Similarly, for $i \in \{1, 2\}$, let D_i^- be an embedded n -disk in Q_i . Then, we view N as the manifold obtained by gluing each D_i^+ to D_i^- (see Section 2.3 for our gluing conventions). We let Q_0^* denote Q_0 minus the centers of each D_i^+ , and for $i \in \{1, 2\}$ we let Q_i^* denote Q_i minus the center of D_i^- . For $i \in \{0, 1, 2\}$ we identify Q_i^* with its image in N . We let $(D_i^+)^* = D_i^+ \cap Q_0^*$ and $(D_i^-)^* = D_i^- \cap Q_i^*$.

Next, we can construct a model of M . For each $i \in \{0, 1, 2\}$, we let $(P_i^\circ)^* \subseteq M^\circ$ denote the preimage of Q_i^* in M° , and we let P_i° be the (possibly disconnected) manifold obtained by filling in the punctures. Similarly, we let $(\Delta_i^+)^* \subseteq (P_0^\circ)^*$ denote the preimage of $(D_i^+)^*$ and let $(\Delta_i^-)^* \subseteq (P_i^\circ)^*$ denote the preimage of $(D_i^-)^*$, and we let $\Delta_i^+ \subseteq P_0^\circ$ and $\Delta_i^- \subseteq P_i^\circ$ denote the multidisks obtained by filling in the punctures. Then we can view M° as the manifold obtained by gluing P_0° to P_i° along the multidisks Δ_i^\pm for each $i \in \{1, 2\}$. Finally, the inclusion $M^\circ \hookrightarrow M$ induces a way to fill in the ends of P_i° (resp. $(P_i^\circ)^*$) to obtain a (possibly disconnected) manifold P_i (resp. P_i^*), so M is obtained by gluing P_0 to P_i for $i \in \{1, 2\}$ along the multidisks Δ_i^\pm . Note that the action of G on M induces an action of G on each P_i .

Note that for each $i \in \{1, 2\}$, the group G acts freely and transitively on the components of Δ_i^+ and Δ_i^- . We can therefore label the components of Δ_i^\pm by the elements of G , i.e. $\Delta_i^\pm = (D_{i,g}^\pm)_{g \in G}$. We choose these labels so that $gD_{i,h}^\pm = D_{i,gh}^\pm$, and so that the disk $D_{i,g}^+$ is glued to $D_{i,g}^-$. We call the disks $D_{i,g}^\pm$ the *gluing disks* of M . We define $\Sigma_{i,g} := \partial D_{i,g}^+$, so $\Sigma_{i,g}$ is an $(n-2)$ -sphere in P_0 .

Fix a component M_0 of P_0 , and let $M_0^* \subseteq M$ denote the corresponding component of P_0^* . We assume without loss of generality that B lies on M_0 , and that M_0 contains the gluing disks $D_{i,\text{id}}^+$ for $i \in \{1, 2\}$. We fix a basepoint $b_0 \in B$ as in Lemma 4.1, and let $G_0 \leq G$ denote the stabilizer of b_0 .

See Figure 3 for a schematic illustration of M .

4.3. The fundamental group. Our next goal is to describe the fundamental group of M in terms of the pieces P_i . It will be convenient to describe $\pi_1(M, b_0)$ in terms of a graph of groups; we will refer to [Ser80, §I.5] and [Hig76] for background material on graphs of groups. We adapt the convention that a *graph* \mathcal{G} consists of

- a *vertex set* $V(\mathcal{G})$, and
- an *edge set* $E(\mathcal{G}) \subseteq V(\mathcal{G}) \times V(\mathcal{G})$ such that for each edge $e = (v_1, v_2) \in E(\mathcal{G})$, the set $E(\mathcal{G})$ contains the unique *reverse edge* $\bar{e} = (v_2, v_1)$.




FIGURE 3. A schematic illustration of M and N . Here the trefoils and unknots (in magenta) represent the points with nontrivial stabilizers and their images in M/G , so N and M° are the complement of these curves.

The *valence* of a vertex v is the number of edges of the form (v, v') for some $v' \in V(\mathcal{G})$. A *graph of groups* is a graph equipped with a group \mathcal{G}_v for each vertex $v \in V(\mathcal{G})$ (typically one allows edge groups as well, but this will not be relevant for us).

For $i \in \{0, 1, 2\}$, let \mathcal{C}_i denote the set of components of P_i , and let $\mathcal{C} = \mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2$. For each $C \in \mathcal{C}$, choose a point $p_C \in C$ away from any gluing disks; we assume that $p_{M_0} = b_0$. For each $i \in \{1, 2\}$ and $g \in G$, we also fix a point $p_{i,g} \in M(D_{i,g}^+)^* \subseteq M$ (we can equivalently view $p_{i,g}$ as lying in $(D_{i,g}^-)^*$). We define \mathcal{G} to be the graph whose vertex set $V(\mathcal{G})$ is the set of all the points p_C and $p_{i,g}$, and whose edge set $E(\mathcal{G})$ consists of pairs $(C, p_{i,g})$ and $(p_{i,g}, C)$ where $p_{i,g} \in C$. Note that each vertex $p_{i,g}$ has valence 2. We equip \mathcal{G} with the structure of a graph of groups where the vertex group at p_C is $\mathcal{G}_C = \pi_1(C, p_C)$ and whose vertex group at $p_{i,g}$ is trivial.

Following Higgins [Hig76], we can define the fundamental groupoid of $\Pi(\mathcal{G})$ of \mathcal{G} . Namely, $\Pi(\mathcal{G})$ is the groupoid whose objects are precisely the elements of $V(\mathcal{G})$, and whose morphisms are generated by the edges of \mathcal{G} and the vertex groups of \mathcal{G} , subject to the relation that for each $e \in \mathcal{G}$, the morphism associated to \bar{e} is the inverse of the morphism associated to e . We define the fundamental group $\pi_1(\mathcal{G}, M_0)$ to be the automorphism group of the object $b_0 = p_{M_0}$ of $\Pi(\mathcal{G})$.

There is a more concrete description of $\pi_1(\mathcal{G}, M_0)$ following Serre [Ser80]. Given an edge $e = (v, v')$, we write $i(e) = v$ and $t(e) = v'$. We define a *path* in \mathcal{G} to be a pair (\hat{e}, \hat{r}) , where either $\hat{e} = (e_1, \dots, e_m)$ is a tuple of edges with $t(e_j) = i(e_{j+1})$ and $\hat{r} = (r_0, \dots, r_m)$ is a tuple with $r_0 \in \mathcal{G}_{i(e_0)}$ and $r_j \in \mathcal{G}_{t(e_j)}$ for $j \geq 1$, or $\hat{e} = \emptyset$ and $\hat{r} = (r_0)$ where $r_0 \in \mathcal{G}_v$ for some $v \in V(\mathcal{G})$. A *loop* in \mathcal{G} based at a vertex v is a path (\hat{e}, \hat{r}) such that either $\hat{e} = \emptyset$ and $r_0 \in \mathcal{G}_v$, or $i(e_1) = t(e_m) = v$. Given a path (\hat{e}, \hat{r}) , we get

an associated morphism in $\Pi(\mathcal{G})$ given by

$$|(\widehat{e}, \widehat{r})| := r_0 e_1 r_1 \cdots e_m r_m.$$

Then $\pi_1(\mathcal{G}, M_0)$ is the group generated by elements of the form $|(\widehat{e}, \widehat{r})|$, where $(\widehat{e}, \widehat{r})$ is a loop based at the vertex M_0 . Note that we can naturally view $\pi_1(M_0, b_0)$ as the subgroup of $\pi_1(\mathcal{G}, M_0)$ generated by paths with $\widehat{e} = \emptyset$.

There is yet another characterization which will be useful for us. Let T be a spanning tree of the graph underlying \mathcal{G} . Let $\pi(\mathcal{G})$ be the group obtained by identifying all the objects of $\Pi(\mathcal{G})$ (equivalently, $\pi(\mathcal{G})$ is the free product of the groups $\pi_1(C, p_C)$ with the free group on $E(\mathcal{G})$, subject to the relations $\bar{e} = e^{-1}$). Then we can define $\pi_1(\mathcal{G}, T)$ to be the group obtained by starting with the group $\pi(\mathcal{G})$ and adding the relation $e = 1$ for each edge in T . By [Ser80, Prop I.20], the natural map $\pi_1(\mathcal{G}, M_0) \rightarrow \pi_1(\mathcal{G}, T)$ is an isomorphism for any spanning tree T .

Now, we can describe $\pi_1(M, b_0)$ in terms of \mathcal{G} . Given a component C of P_i , we let $C^* \subseteq M$ denote the corresponding component of P_i^* . For each edge $e = (C, p_{i,g})$ of \mathcal{G} , choose an oriented arc a_e in C^* joining p_C to $p_{i,g}$; we define $a_{\bar{e}}$ by reversing the arc a_e . Then, given a path $(\widehat{e}, \widehat{r})$ in \mathcal{G} , we can construct a path $\gamma_{(\widehat{e}, \widehat{r})}$ in M by concatenating representatives of the elements r_j with the arcs a_{e_j} . The path $\gamma_{(\widehat{e}, \widehat{r})}$ is well-defined up to homotopy rel the points p_C and $p_{i,g}$. Then we have the following.

Lemma 4.2. *There is an isomorphism*

$$\pi_1(\mathcal{G}, M_0) \cong \pi_1(M, b_0)$$

which sends the element $|(\widehat{e}, \widehat{r})|$ to the loop $\gamma_{(\widehat{e}, \widehat{r})}$ for any loop $(\widehat{e}, \widehat{r})$ in \mathcal{G} based at M_0 .

Proof. Let $\mathbf{p} \subseteq M$ be the set of points p_C for $C \in \mathcal{C}$ and $p_{i,g}$ for $i \in \{1, 2\}$ and $g \in G$. Let $\Pi(M, \mathbf{p})$ be the groupoid whose objects are the points in \mathbf{p} and whose morphisms are homotopy classes of paths with endpoints in \mathbf{p} . We can apply the groupoid version of van Kampen's theorem (see [BS84]) to show that there is an isomorphism $\Pi(M, \mathbf{p}) \cong \Pi(\mathcal{G})$ which induces the desired isomorphism on fundamental groups. Namely, for each $C \in \mathcal{C}$, we have a natural map $\Pi(C^*, \mathbf{p} \cap C^*) \rightarrow \Pi(\mathcal{G})$ which maps the fundamental group $\pi_1(C^*, p_C)$ to the vertex group at p_C , and which maps each arc a_e for $e = (C, p_{i,g})$ to the morphism corresponding to e . Each 2-fold intersection of the sets C^* for $C \in \mathcal{C}$ is a disjoint union of simply connected subsets, and each 3-fold intersection is empty. Van Kampen's theorem then implies that the maps $\Pi(C^*, \mathbf{p} \cap C^*) \rightarrow \Pi(\mathcal{G})$ assemble into a map $\Pi(M, \mathbf{p}) \rightarrow \Pi(\mathcal{G})$. This map has a natural inverse, and is therefore an isomorphism. \square

4.4. The invariant subgroup of the fundamental group. Recall that we let $G_0 \leq G$ denote the stabilizer of the basepoint b_0 . We have a well-defined action of G_0 on $\pi_1(M, b_0)$. Let $\pi_1(M, b_0)^{G_0} \leq \pi_1(M, b_0)$ denote the fixed points of this action. The following lemma constrains the subgroup $\pi_1(M, b_0)^{G_0}$. We will use this result in Section 5 in the construction of our obstruction map (see Lemma 5.2).

Lemma 4.3. *The subgroup $\pi_1(M, b_0)^{G_0}$ is contained in the subgroup $\pi_1(M_0, b_0) \leq \pi_1(M, b_0)$.*

Proof. Let $\widehat{M} \rightarrow M$ denote the universal cover of M , and fix a base point $\widehat{b}_0 \in \widehat{M}$ above b_0 . Let $\widehat{P}_i^* \subseteq \widehat{M}$ denote the preimage of P_i^* , and let \widehat{P}_i be obtained by filling in the punctures. We call each component of each \widehat{P}_i a *sheet* of \widehat{M} . Let \widehat{M}_0 denote the sheet containing \widehat{b}_0 (so \widehat{M}_0 is a component of \widehat{P}_0 , and is diffeomorphic to the universal cover of M_0).

We will prove the lemma by contrapositive. Suppose $\gamma \in \pi_1(M, b_0)$ is an element that is not contained in the subgroup $\pi_1(M_0, b_0) \leq \pi_1(M, b_0)$, and let $g \in G_0$ be a generator (recall that G_0 is cyclic). Our goal is then to show that $g(\gamma) \neq \gamma$. We will show that a lift of γ to \widehat{M} travels through a unique sequence of sheets up to homotopy, and then show that a lift of g cannot preserve this sequence of sheets.

Let $\widehat{\gamma}$ denote the lift of γ to \widehat{M} based at \widehat{b}_0 . Since \widehat{M} is simply connected, we may assume up to homotopy rel endpoints that for any sheet W of \widehat{M} , the intersection $\widehat{\gamma} \cap W$ is connected, i.e. $\widehat{\gamma}$ travels through each sheet at most once. We claim that $\widehat{\gamma}$ must end on a different sheet than \widehat{M}_0 . Indeed, suppose for the sake of contradiction that $\widehat{\gamma}$ ends at a point on \widehat{M}_0 . Then we can assume that $\widehat{\gamma}$ is entirely contained in the sheet \widehat{M}_0 . Then $\widehat{\gamma}$ projects to a loop in M entirely contained within M_0 , contradicting that γ lies outside of the subgroup $\pi_1(M_0, b_0)$. Thus, $\widehat{\gamma}$ travels through a sequence of distinct sheets $\widehat{M}_0 = W_0, W_1, \dots, W_m$ with $m > 0$, where W_i is adjacent to W_{i+1} and W_m is a component of \widehat{P}_0 (since the endpoint of $\widehat{\gamma}$ must live above b_0 , which lies in M_0).

Next, we claim that this sequence of sheets is unique; that is, if $\widehat{\gamma}$ is homotopic rel endpoints to an arc $\widehat{\gamma}'$ traveling through a sequence of distinct sheets $\widehat{M}_0 = W_0, W'_1, \dots, W'_{m'}$, then $m' = m$ and $W_j = W'_j$ for all j . Suppose for contradiction that for some $k < m$, we have that $W_j = W'_j$ for $j \leq k$ but $W_{k+1} \neq W'_{k+1}$. Then the curves $\widehat{\gamma}$ and $\widehat{\gamma}'$ leave through two different disks on the sheet $W_k = W'_k$, say D and D' . Let $\Sigma = \partial D$ and $\Sigma' = \partial D'$. Equip Σ and Σ' with orientations such that $\widehat{\gamma}$ and $\widehat{\gamma}'$ intersect them positively, and let $\overline{\Sigma}'$ denote Σ' equipped with the opposite orientation. Let δ be the loop obtained by concatenating $\widehat{\gamma}$ with the reverse of $\widehat{\gamma}'$. Then the loop δ and the submanifold $\Sigma \sqcup \overline{\Sigma}'$ have a positive signed intersection number. Since this is a homotopy invariant of δ , it follows that the loop δ is not homotopic to a point, contradicting that \widehat{M} is simply connected.

Now, lift g to a diffeomorphism $\widehat{g} : \widehat{M} \rightarrow \widehat{M}$ that fixes the point \widehat{b}_0 . Since g preserves $M_0^* \subseteq M$ (as it fixes b_0), it follows that \widehat{g} preserves the sheet \widehat{M}_0 . However, since g does not fix any of the gluing disks in M_0 , it follows that \widehat{g} does not preserve any of the sheets adjacent to \widehat{M}_0 . Thus, $\widehat{g}(\widehat{\gamma})$ travels through a different sequence of sheets than $\widehat{\gamma}$. By the uniqueness statement above, this means that $\widehat{g}(\widehat{\gamma})$ is not homotopic rel endpoints to $\widehat{\gamma}$, and thus we can conclude that $g(\gamma) \neq \gamma$. \square

4.5. An infinite cover. Next, we will define the cover \widetilde{M} that we will use to construct our obstruction map. Let \mathcal{G} denote the graph of groups defined above, and let $\overline{\mathcal{G}}$ denote the graph of groups obtained from \mathcal{G} by replacing the vertex group $\mathcal{G}_{M_0} = \pi_1(M_0, b_0)$ with the trivial group. By Lemma 4.2, the natural map $\Pi(\mathcal{G}) \rightarrow \Pi(\overline{\mathcal{G}})$ induces a surjection

$$\pi_1(M, b_0) \rightarrow \pi_1(\overline{\mathcal{G}}, M_0).$$

The kernel of this map is the normal closure $\langle\langle \pi_1(M_0, b_0) \rangle\rangle \leq \pi_1(M, b_0)$ (this is easily seen using the isomorphisms $\pi_1(\mathcal{G}, M_0) \cong \pi_1(\mathcal{G}, T)$ and $\pi_1(\overline{\mathcal{G}}, M_0) \cong \pi_1(\overline{\mathcal{G}}, T)$, where T is a spanning tree of the graphs underlying \mathcal{G} and $\overline{\mathcal{G}}$). We define $\pi : \widetilde{M} \rightarrow M$ to be cover associated to the kernel of this map. We fix a basepoint $\widetilde{b}_0 \in \widetilde{M}$ above $b_0 \in M$, so $\pi_1(\widetilde{M}, \widetilde{b}_0) \cong \langle\langle \pi_1(M_0, b_0) \rangle\rangle$, and the deck group of \widetilde{M} is the group $\Lambda := \pi_1(\overline{\mathcal{G}}, M_0)$.

Observe that the cover \widetilde{M} is trivial over M_0^* , i.e. each component of $\pi^{-1}(M_0^*)$ projects diffeomorphically onto M_0^* . On the other hand, if C^* is a component of P_1^* or P_2^* , or if C^* is a component of $P_0^* \setminus M_0^*$, then each component of $\pi^{-1}(C^*)$ is the universal cover of C^* . Note that it is possible that $\pi_1(M_0, b_0)$ is trivial, in which case $\overline{\mathcal{G}} = \mathcal{G}$ and \widetilde{M} is simply the universal cover of M .

The cover \widetilde{M} will be infinite-sheeted; we can construct an infinite subgroup of the deck group $\Lambda = \pi_1(\overline{\mathcal{G}}, M_0)$ as follows. Fix $i \in \{1, 2\}$. Let $\overline{\mathcal{G}}'_i$ be the induced subgraph of $\overline{\mathcal{G}}$ whose vertices are the points p_C for $C \in \mathcal{C}_0 \cup \mathcal{C}_i$ together with the points $p_{i,g}$ lying on each $C \in \mathcal{C}_0 \cup \mathcal{C}_i$, and we define $\overline{\mathcal{G}}_i$ to be the component of $\overline{\mathcal{G}}'_i$ containing the vertex $p_{M_0} = b_0$ (we equip $\overline{\mathcal{G}}_i$ with the same vertex groups as $\overline{\mathcal{G}}$). We define $\Lambda_i := \pi_1(\overline{\mathcal{G}}_i, M_0)$, and observe that Λ_i is naturally a subgroup of Λ .

Lemma 4.4. *At least one of the following is true:*

- $\overline{\mathcal{G}}_i$ has at least two nontrivial vertex groups \mathcal{G}_C for $C \in \mathcal{C}_i$, or
- $\overline{\mathcal{G}}_i$ contains a cycle.

In particular, the subgroup $\Lambda_i \leq \Lambda$ is infinite.

Proof. Observe first that each component of P_i contains the same number of gluing disks. Thus, we analyze two cases: the case that each component has a single gluing disk, and the case that each component has at least two gluing disks.

Suppose first that each component of P_i has a single gluing disk. In this case, we claim that $\bar{\mathcal{G}}_i$ contains at least two nontrivial vertex groups \mathcal{G}_C for $C \in \mathcal{C}_i$. First, observe that since each component has a single gluing disk, the cover $P_i^\circ \rightarrow Q_i$ must be trivial. This means that G acts freely on the components of P_i° and each component is diffeomorphic to Q_i . Then G must act freely on P_i , so $P_i = P_i^\circ$. Thus each component of P_i is not simply connected, since Q_i is not simply connected. To complete the proof, we need to check that $\bar{\mathcal{G}}_i$ contains at least two vertices of the form p_C for $C \in \mathcal{C}_i$. But this follows since each component of P_0 contains at least two of the gluing disks $\Delta_{i,g}^+$ for $g \in G$, since the basepoint $b_0 \in M_0$ has a nontrivial G -stabilizer. In particular, we see that Λ_i is infinite since it contains a free product $\pi_1(C) * \pi_1(C')$ for two distinct components $C, C' \in P_i$.

Otherwise, suppose each component of P_i has at least two gluing disks. In this case, we claim that $\bar{\mathcal{G}}_i$ contains a cycle; the fact that Λ_i is infinite then follows from the isomorphism $\pi_1(\bar{\mathcal{G}}_i, M_0) \cong \pi_1(\bar{\mathcal{G}}_i, T)$ for any spanning tree T . As in the previous case, we observe that each component of P_0 also contains at least two of the gluing disks $\Delta_{i,g}^+$ for $g \in G$. Therefore, since each component of P_0 and P_i contains at least two gluing disks, each vertex of $\bar{\mathcal{G}}_i$ must have valence at least 2. Thus $\bar{\mathcal{G}}_i$ must contain a cycle, as any tree contains a vertex of valence 0 or 1. \square

4.6. The homology of the cover. The key property of the cover \widetilde{M} is that the homology group $H_{n-1}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z})$ is not finitely generated. In fact, we will describe an infinite dimensional subspace of $H_{n-1}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z})$. In order to do so, we establish some notation.

Let $G_{M_0} \leq G$ denote the stabilizer of M_0 (with respect to the action of G on the components of P_0). Then for each $i \in \{1, 2\}$, since M_0 contains the gluing disk $D_{i,\text{id}}^+$, it follows that the disks $D_{i,g}^+$ for $g \in G_{M_0}$ are precisely the gluing disks contained in M_0 . Note that $G_0 \leq G_{M_0}$.

Let W_{id} be the component of $\pi^{-1}(M_0^*) \subseteq \widetilde{M}$ containing the base point \widetilde{b}_0 . Then, for $u \in \Lambda$, let W_u be the image of W_{id} under u . We will call each W_u a *sheet* of \widetilde{M} . For each $u \in \Lambda$, the cover $\pi : \widetilde{M} \rightarrow M$ restricts to a diffeomorphism $W_u \rightarrow M_0^*$. Recall that for $g \in G_{M_0}$, we let $\Sigma_{i,g} \subseteq M_0$ denote the boundary sphere of the disk $D_{i,g}^+$. On W_u , let $\Sigma_{i,g,u}$ be the copy of $\Sigma_{i,g}$ for $i \in \{1, 2\}$ and $g \in G_{M_0}$.

Now, we can construct an infinite linearly independent subset of $H_{n-1}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z})$ indexed by the subgroup Λ_2 defined above.

Proposition 4.5. *For any two distinct pairs $(g, u), (g', u') \in G_{M_0} \times \Lambda_2$, the spheres $\Sigma_{1,g,u}$ and $\Sigma_{1,g',u'}$ represent nonequal classes in $H_{n-1}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z})$. Moreover, the set*

$$\mathcal{S} = \{[\Sigma_{1,g,u}] \in H_{n-1}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z}) \mid g \in G_{M_0}, u \in \Lambda_2\}$$

is an infinite linearly independent subset of $H_{n-1}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z})$.

To prove Proposition 4.5, we need the following preliminary lemma.

Lemma 4.6. *For any $i \in \{1, 2\}$, $g \in G_{M_0}$, and $u \in \Lambda$, the complement $\widetilde{M} \setminus \Sigma_{i,g,u}$ has two components, and each component is noncompact.*

Proof. We begin by showing that $\widetilde{M} \setminus \Sigma_{i,g,u}$ is disconnected (and hence has two components). If $\widetilde{M} \setminus \Sigma_{i,g,u}$ were connected, then by cutting \widetilde{M} along $\Sigma_{i,g,u}$, we could construct a simple closed curve that intersects $\Sigma_{i,g,u}$ once. Since mod 2 intersection numbers are a homotopy invariant, it's therefore enough to show that every loop in \widetilde{M} has an even intersection number with $\Sigma_{i,g,u}$.

Let γ be a loop in \widetilde{M} . Up to homotopy, we may assume that γ is based at \widetilde{b}_0 . Let $\bar{\gamma}$ be the image of γ in M , so $\bar{\gamma}$ is a loop based at b_0 . Since $\pi_1(\widetilde{M}, \widetilde{b}_0)$ is the normal closure $\langle\langle \pi_1(M_0, b_0) \rangle\rangle$, we know

that $\bar{\gamma}$ is homotopic to a concatenation of loops $\bar{\gamma}_1, \dots, \bar{\gamma}_m$ where each $\bar{\gamma}_j$ is freely homotopic into the subspace $M_0^* \subseteq M$. This means that γ is homotopic to a concatenation of loops $\gamma_1, \dots, \gamma_m$ where each γ_j is freely homotopic into W_{u_j} for some $u_j \in \Lambda$. This implies that the intersection number of γ with any sphere of the form $\Sigma_{i,g,u}$ must be even as desired.

Now, we can show that each component is noncompact. First, we claim that each component $\widetilde{M} \setminus \Sigma_{i,g,u}$ contains at least one sheet $W_{u'}$ for some $u' \in \Lambda$. To prove this, it's enough to find a loop γ in M based at b_0 which has a lift to an arc in \widetilde{M} that passes through $\Sigma_{i,g,u}$ once. Let C be the component of P_i containing $D_{i,g}^-$. If C is not simply connected, then we can construct γ by conjugating a nontrivial element of $\pi_1(C, p_C)$ with a path from b_0 to p_C . Otherwise, assume each component of P_i is simply connected. Then it's enough to show that the sphere $\Sigma_{i,g}$ is non-separating in M , since then we can take γ to be a loop that intersects $\Sigma_{i,g}$ once. To prove this, suppose for contradiction that $\Sigma_{i,g}$ separates M . Then for any $g' \in G$, the sphere $\Sigma_{i,g'g}$ is also separating. This implies that every edge of the graph $\bar{\mathcal{G}}_i$ is separating, which implies that the graph underlying $\bar{\mathcal{G}}_i$ is a tree, contradicting Lemma 4.4.

Now, we can show that each component of $\widetilde{M} \setminus \Sigma_{i,g,u}$ contains infinitely-many sheets, and is therefore noncompact. From above, we know that each component contains a sheet $W_{u'}$ for some $u' \in \Lambda$. Then we can choose infinitely-many distinct elements of Λ_j (where $j \neq i$), represent them by loops in M which are disjoint from the spheres $\Sigma_{i,g}$ for $g \in G$, and lift these loops to arcs based at the sheet $W_{u'}$. Then the endpoint of each lift will be a distinct sheet in the same component of $\widetilde{M} \setminus \Sigma_{i,g,u}$. □

Proof of Proposition 4.5. Throughout this proof, we use homology with $\mathbb{Z}/\ell\mathbb{Z}$ -coefficients. Lemma 4.4 says that the group Λ_2 is infinite, so it's enough to show that $\Sigma_{1,g,u}$ and $\Sigma_{1,g',u'}$ represent distinct classes and that \mathcal{S} is linearly independent. In particular, we must show that for any finite subset

$$\{(g_1, u_1), \dots, (g_m, u_m)\} \subseteq G_{M_0} \times \Lambda_2,$$

the homology classes $[\Sigma_{1,g_1,u_1}], \dots, [\Sigma_{1,g_m,u_m}]$ in $H_{n-1}(\widetilde{M})$ are linearly independent.

Let $V \subseteq M_0$ denote the complement of $\text{Int}(D_{i,g}^+)$ for each $i \in \{1, 2\}$ and $g \in G_{M_0}$, so in particular $V \subseteq M_0^*$. For $u \in \Lambda$, let $V_u = W_u \cap \pi^{-1}(V)$, so V_u is a compact oriented n -manifold with

$$\partial V_u = \bigcup_{g \in G_{M_0}} \Sigma_{1,g,u} \cup \Sigma_{2,g,u}.$$

Now, let $X' \subseteq \widetilde{M}$ be the union of the sets V_{u_j} for $1 \leq j \leq m$. Then X' is a (possibly disconnected) compact oriented n -manifold with boundary

$$\partial X' = \bigcup_{j=1}^m \bigcup_{g \in G_{M_0}} \Sigma_{1,g,u_j} \cup \Sigma_{2,g,u_j}.$$

Now, we claim that for each $1 \leq j \leq m$, there exists an arc $a_j : [0, 1] \rightarrow \widetilde{M}$ with $a_j(0) = \tilde{b}_0$ and $a_j(1) \in V_{u_j}$, and such that a_j is disjoint from the spheres $\Sigma_{1,g,u}$ for all $(g, u) \in G_{M_0} \times \Lambda$. To construct a_j , let $\gamma \in \pi_1(M, b_0)$ be an element mapping to the deck transformation $u_j \in \Lambda_2$. By the definition of Λ_2 , we can represent γ by a loop in M which is disjoint from the gluing disks $D_{1,g}^\pm$ for all $g \in G$. Then we can take a_j to be the lift of this loop based at \tilde{b}_0 .

Let X be the union of X' with a tubular neighborhood of each arc a_j . Then X is a connected compact oriented n -manifold with boundary. Each sphere Σ_{1,g_j,u_j} for $1 \leq j \leq m$ is a boundary component of X . Let $Y \subseteq \widetilde{M}$ be the closure of $\widetilde{M} \setminus X$, so Y is a (possibly disconnected) non-compact oriented n -manifold with $\partial Y = \partial X = X \cap Y$. Let $L \subseteq H_{n-1}(X \cap Y)$ denote submodule spanned by the classes $[\Sigma_{1,g_1,u_1}], \dots, [\Sigma_{1,g_m,u_m}]$, so L is free submodule of rank m .

Let $L_X \subseteq H_{n-1}(X)$ and $L_Y \subseteq H_{n-1}(Y)$ denote the images of L induced by the inclusions $X \cap Y \hookrightarrow X$ and $X \cap Y \hookrightarrow Y$. We claim that the maps $L \rightarrow L_X$ and $L \rightarrow L_Y$ are isomorphisms. Assuming this claim, we can prove the proposition as follows. From Mayer-Vietoris, we have an exact sequence

$$H_{n-1}(X \cap Y) \rightarrow H_{n-1}(X) \oplus H_{n-1}(Y) \rightarrow H_{n-1}(\widetilde{M}).$$

Since $L \cong L_X$ and $L \cong L_Y$, the submodule $L_X \oplus 0 \subseteq H_{n-1}(X) \oplus H_{n-1}(Y)$ intersects the image of $H_{n-1}(X \cap Y)$ trivially. Thus, $L_X \oplus 0$ embeds in $H_{n-1}(\widetilde{M})$, proving the proposition.

Now, we can prove the claim. To show that $L \cong L_X$, it's enough to show that the classes $[\Sigma_{1,g_j,u_j}]$ for $1 \leq j \leq m$ are linearly independent in $H_{n-1}(X)$. Consider the exact sequence

$$H_n(X, \partial X) \rightarrow H_{n-1}(\partial X) \rightarrow H_{n-1}(X).$$

The middle term has a basis given by the components of ∂X . The image of the first map is the submodule spanned by the sum of all the components of ∂X . The spheres Σ_{1,g_j,u_j} for $1 \leq j \leq m$ are not the only boundary components of X , as X has at least one other boundary component coming from either one of the arcs a_j or one of the spheres Σ_{2,g,u_j} for some $g \in G_{M_0}$. Thus, the submodule of $L \subseteq H_{n-1}(X \cap Y) = H_{n-1}(\partial X)$ intersects the kernel of the second map trivially, and hence L embeds in $H_{n-1}(X)$.

We can show that $L \cong L_Y$ similarly. We again have an exact sequence

$$H_n(Y, \partial Y) \rightarrow H_{n-1}(\partial Y) \rightarrow H_{n-1}(Y).$$

It's enough to show that $L \subseteq H_{n-1}(\partial Y) = H_{n-1}(X \cap Y)$ intersects the image of the first map trivially. Let Y_1, \dots, Y_m denote the components of Y . Then $H_n(Y, \partial Y) \cong \bigoplus_{k=1}^m H_n(Y_k, \partial Y_k)$. If Y_k is compact, then the first map sends the fundamental class in $H_n(Y_k, \partial Y_k)$ to the sum of the boundary components of Y_k . Otherwise, if Y_k is non-compact, then the group $H_n(Y_k, \partial Y_k)$ is trivial. Then we note that by Lemma 4.6 that each Σ_{1,g_j,u_j} is the boundary of some noncompact component Y_{k_j} of Y . Thus we can conclude L intersects the image of the first map trivially as desired. \square

5. THE OBSTRUCTION MAP

In this section, we will construct the obstruction map that we use to prove part (ii) of Theorem A. Fix a manifold M and a group G satisfying the assumptions of Theorem A. Then by assumption we have an oriented submanifold $B \subseteq M$ which is fixed by some element of G , and whose homology class in $H_{n-2}(M; \mathbb{Z}/\ell\mathbb{Z})$ is trivial for some $\ell > 1$. We fix a basepoint $b_0 \in B$ as in Lemma 4.1. Let $\pi : \widetilde{M} \rightarrow M$ be the infinite-sheeted cover described in Section 4.5, and let \widetilde{b}_0 be a basepoint above b_0 . Recall that the cover \widetilde{M} is trivial over the subset $M_0^* \subseteq M$, and so in particular it is trivial over B . We let \widetilde{B} be the component of $\pi^{-1}(B)$ containing the basepoint \widetilde{b}_0 (so π restricts to a diffeomorphism $\widetilde{B} \cong B$).

Our obstruction map is not defined on $\text{Ker}(\mathcal{P}_G)$, but a closely related group. Let $\text{Diff}(M, B)$ denote the group of orientation-preserving diffeomorphisms of M that preserve B setwise (the orientation of B need not be preserved), and let $\text{Mod}(M, B) = \pi_0(\text{Diff}(M, B))$. Let $\text{Mod}_G(M, B) \leq \text{Mod}(M, B)$ and $\text{Mod}_G(M) \leq \text{Mod}(M)$ denote the subgroups of mapping classes with an equivariant representative. Then there is a forgetful map $\mathcal{F} : \text{Mod}(M, B) \rightarrow \text{Mod}(M)$ which restricts to a map

$$\mathcal{F}_G : \text{Mod}_G(M, B) \rightarrow \text{Mod}_G(M).$$

Our goal in this section is construct a map

$$\varphi : \text{Ker}(\mathcal{F}_G) \rightarrow \mathcal{H}_{n-1}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z}).$$

We give a simple description of the map φ in Section 5.1; the rest of this section is devoted to proving this map is well-defined. In Section 5.2, we describe how to obtain a homology class of \widetilde{M} from an isotopy. In Sections 5.3 and 5.4, we describe how a subgroup of isotopies on M lift to \widetilde{M} .

We formally construct φ in Section 5.5. Throughout this section, we implicitly use homology with $\mathbb{Z}/\ell\mathbb{Z}$ -coefficients, but our arguments apply verbatim for other coefficients.

5.1. A simple description of the map. The obstruction map φ works as follows. Let $\alpha \in \text{Ker}(\mathcal{F}_G)$, and let f be an equivariant diffeomorphism of M that represents α . Then f preserves B setwise, and there is an isotopy $h : M \times [0, 1] \rightarrow M$ with $h_0 = \text{id}_M$ and $h_1 = f$. This induces a map

$$\begin{aligned} r : B \times [0, 1] &\rightarrow M \\ (b, t) &\mapsto h_t(b). \end{aligned}$$

In particular, $r|_{B \times \{0\}} = \text{id}_B$ and $r|_{B \times \{1\}}$ is a diffeomorphism of B . Now, lift r to a map

$$\tilde{r} : B \times [0, 1] \rightarrow \tilde{M}$$

which sends $B \times \{0, 1\}$ to \tilde{B} . Then by pushing forward the fundamental class of $B \times [0, 1]$, the map \tilde{r} yields an element $x \in H_{n-1}(\tilde{M}, \tilde{B})$. We define $\varphi(\alpha) = x$.

It is not obvious that the map φ is well-defined with this description. One has to show that φ is independent of the choice of representative f and the choice of isotopy h . One also has to show that the map r lifts as desired; we can choose a lift \tilde{r} sending $B \times \{0\}$ to \tilde{B} , but it's not clear that this lift will map $B \times \{1\}$ to \tilde{B} (as opposed to some other component of $\pi^{-1}(B)$).

Thus, in this section we will construct φ more formally so as to ensure that it is well-defined. The above description will still be useful to compute φ in practice (which we will do in Section 6).

5.2. Homology classes from isotopies. We begin by describing how an isotopy of \tilde{B} in \tilde{M} gives rise to a homology class.

There is a natural evaluation map

$$\begin{aligned} \text{ev}_{\tilde{B}} : \text{Diff}(\tilde{M}) \times \tilde{B} &\rightarrow \tilde{M} \\ (f, b) &\mapsto f(b). \end{aligned}$$

This map takes the subspace $\text{Diff}(\tilde{M}, \tilde{B}) \times \tilde{B}$ to \tilde{B} . Applying the relative Künneth formula (see e.g. [Spa81, Thm 5.3.10]), we get a map

$$H_1\left(\text{Diff}(\tilde{M}), \text{Diff}(\tilde{M}, \tilde{B})\right) \otimes H_{n-2}(\tilde{B}) \rightarrow H_{n-1}(\tilde{M}, \tilde{B}).$$

By pairing an element of $H_1\left(\text{Diff}(\tilde{M}), \text{Diff}(\tilde{M}, \tilde{B})\right)$ with the fundamental class of \tilde{B} , we get a map

$$H_1\left(\text{Diff}(\tilde{M}), \text{Diff}(\tilde{M}, \tilde{B})\right) \rightarrow H_{n-1}(\tilde{M}, \tilde{B}).$$

Now, we define the group

$$\text{Iso}(\tilde{M}, \tilde{B}) := \pi_1\left(\text{Diff}(\tilde{M}), \text{Diff}(\tilde{M}, \tilde{B}), \text{id}_{\tilde{M}}\right).$$

So, an element of $\text{Iso}(\tilde{M}, \tilde{B})$ is the homotopy class of an isotopy $h : \tilde{M} \times [0, 1] \rightarrow \tilde{M}$ where $h_0 = \text{id}_{\tilde{M}}$ and h_1 restricts to a diffeomorphism of \tilde{B} . A priori, $\text{Iso}(\tilde{M}, \tilde{B})$ is just a (pointed) set, but it inherits a group structure from $\text{Diff}(\tilde{M})$ (cf. [Gol81]): given two elements $[f], [g] \in \text{Iso}(\tilde{M}, \tilde{B})$ represented by isotopies $f, g : \tilde{M} \times [0, 1] \rightarrow \tilde{M}$ with $f_0 = g_0 = \text{id}_{\tilde{M}}$ and $f_1(\tilde{B}) = g_1(\tilde{B}) = \tilde{B}$, the product $[f] \cdot [g]$ is represented by the isotopy

$$(f \cdot g)_t := \begin{cases} f_{2t} & 0 \leq t \leq 1/2, \\ g_{2t-1} \circ f_1 & 1/2 \leq t \leq 1. \end{cases}$$

We define $\psi : \text{Iso}(\tilde{M}, \tilde{B}) \rightarrow H_{n-1}(\tilde{M}, \tilde{B})$ to be the composition

$$\text{Iso}(\tilde{M}, \tilde{B}) \rightarrow H_1\left(\text{Diff}(\tilde{M}), \text{Diff}(\tilde{M}, \tilde{B})\right) \rightarrow H_{n-1}(\tilde{M}, \tilde{B}).$$

We can interpret the map ψ as follows. Let $[h] \in \text{Iso}(\widetilde{M}, \widetilde{B})$, so $h : \widetilde{M} \times [0, 1] \rightarrow \widetilde{M}$ is an isotopy with $h_0 = \text{id}_{\widetilde{M}}$ and $h_1(\widetilde{B}) = \widetilde{B}$. Since $\pi : \widetilde{M} \rightarrow M$ restricts to a diffeomorphism $\widetilde{B} \rightarrow B$, there is a well-defined inverse $\pi^{-1} : B \rightarrow \widetilde{B}$. We then get a map

$$\begin{aligned} \tilde{r} : B \times [0, 1] &\rightarrow \widetilde{M} \\ (b, t) &\mapsto h_t(\pi^{-1}(b)). \end{aligned}$$

We have an induced map on relative homology

$$\tilde{r}_* : H_{n-1}(B \times [0, 1], B \times \{0, 1\}) \rightarrow H_{n-1}(\widetilde{M}, \widetilde{B}).$$

Then $\psi([h])$ is the image of the fundamental class under \tilde{r}_* .

Remark 5.1. Let $\text{Emb}(\widetilde{B}, \widetilde{M})$ denote the space of embeddings $\widetilde{B} \hookrightarrow \widetilde{M}$, and let $\text{Sub}(\widetilde{B}, \widetilde{M})$ denote quotient of $\text{Emb}(\widetilde{B}, \widetilde{M})$ by $\text{Diff}(\widetilde{B})$ (which acts by precomposition). Then a theorem of Cerf [Cer61] says that the map $\text{Diff}(\widetilde{M}) \rightarrow \text{Emb}(\widetilde{B}, \widetilde{M})$ is a fibration, and a theorem of Binz-Fischer [BF81] says that the map $\text{Emb}(\widetilde{B}, \widetilde{M}) \rightarrow \text{Sub}(\widetilde{B}, \widetilde{M})$ is a fibration. Thus the map $\text{Diff}(\widetilde{M}) \rightarrow \text{Sub}(\widetilde{B}, \widetilde{M})$ is a fibration with fiber $\text{Diff}(\widetilde{M}, \widetilde{B})$. It follows that $\text{Iso}(\widetilde{M}, \widetilde{B}) \cong \pi_1(\text{Sub}(\widetilde{M}, \widetilde{B}))$, so the map ψ can be viewed as an invariant of loops in $\text{Sub}(\widetilde{M}, \widetilde{B})$. Indeed, this was our original motivation in defining the map ψ , but for the arguments in this section it is generally simpler to work with the pair $(\text{Diff}(\widetilde{M}), \text{Diff}(\widetilde{M}, \widetilde{B}))$ instead.

5.3. Isotopies and equivariant diffeomorphisms. Next, we will identify the relationship between $\text{Ker}(\mathcal{F}_G)$ and isotopies of B in M .

Similar to above, we define the group

$$\text{Iso}(M, B) := \pi_1(\text{Diff}(M), \text{Diff}(M, B), \text{id}_M).$$

From the long exact sequence of the pair $(\text{Diff}(M), \text{Diff}(M, B))$, we have an exact sequence

$$\pi_1(\text{Diff}(M)) \rightarrow \text{Iso}(M, B) \rightarrow \text{Mod}(M, B) \xrightarrow{\mathcal{F}} \text{Mod}(M).$$

Equivalently, we have an exact sequence

$$\pi_1(\text{Diff}(M)) \rightarrow \text{Iso}(M, B) \rightarrow \text{Ker}(\mathcal{F}) \rightarrow 1.$$

The group $\text{Ker}(\mathcal{F}_G)$ is a subgroup of $\text{Ker}(\mathcal{F})$; let $\text{Iso}(M, B)_G \leq \text{Iso}(M, B)$ denote the preimage of $\text{Ker}(\mathcal{F}_G)$. Thus we have an exact sequence

$$\pi_1(\text{Diff}(M)) \rightarrow \text{Iso}(M, B)_G \rightarrow \text{Ker}(\mathcal{F}_G) \rightarrow 1.$$

The map $\text{Iso}(M, B)_G \rightarrow \text{Ker}(\mathcal{F}_G)$ is given by $[h] \mapsto [h_1]$, where $h : M \times [0, 1] \rightarrow M$ is an isotopy with $h_0 = \text{id}_M$ and $h_1(B) = B$. Any element of $\text{Iso}(M, B)_G$ can be represented by an isotopy h such that h_1 is equivariant.

5.4. Lifting isotopies. Our next goal is to show that every element of the group $\text{Iso}(M, B)_G$ lifts to an element of $\text{Iso}(\widetilde{M}, \widetilde{B})$.

Lemma 5.2. *There is a well-defined map $\lambda_G : \text{Iso}(M, B)_G \rightarrow \text{Iso}(\widetilde{M}, \widetilde{B})$ such that $\lambda_G([h])$ is represented by an isotopy $\tilde{h} : \widetilde{M} \times [0, 1] \rightarrow \widetilde{M}$ such that $\pi \circ \tilde{h}_t = h_t$.*

Proof. Let $h : M \times [0, 1] \rightarrow M$ be an isotopy with $h_0 = \text{id}_M$ and $h_1(B) = B$. Then h lifts to an isotopy $\tilde{h} : \widetilde{M} \times [0, 1] \rightarrow \widetilde{M}$ with $\tilde{h}_0 = \text{id}_{\widetilde{M}}$. Since \tilde{h} is a lift of h , we know that \tilde{h}_1 preserves the set $\pi^{-1}(B)$. Moreover, a homotopy of h lifts to a homotopy of \tilde{h} . We therefore have a map

$$\text{Iso}(M, B) \rightarrow \pi_1\left(\text{Diff}(\widetilde{M}), \text{Diff}(\widetilde{M}, \pi^{-1}(B))\right).$$

Now, let $\text{Diff}(\widetilde{M}, \pi^{-1}(B), \widetilde{B})$ be the group of diffeomorphisms preserving $\pi^{-1}(B)$ setwise, and preserving the component \widetilde{B} setwise. Then we have a natural inclusion

$$\pi_1 \left(\text{Diff}(\widetilde{M}), \text{Diff}(\widetilde{M}, \pi^{-1}(B), \widetilde{B}) \right) \hookrightarrow \pi_1 \left(\text{Diff}(\widetilde{M}), \text{Diff}(\widetilde{M}, \pi^{-1}(B)) \right).$$

We claim that the image of $\text{Iso}(M, B)_G$ is contained in the image of this inclusion. Assuming this, we have a map

$$\text{Iso}(M, B)_G \rightarrow \pi_1 \left(\text{Diff}(\widetilde{M}), \text{Diff}(\widetilde{M}, \pi^{-1}(B), \widetilde{B}) \right),$$

and we can obtain the map λ_G by post-composing with the natural map

$$\pi_1 \left(\text{Diff}(\widetilde{M}), \text{Diff}(\widetilde{M}, \pi^{-1}(B), \widetilde{B}) \right) \rightarrow \text{Iso}(\widetilde{M}, \widetilde{B}).$$

So, it remains to prove our claim. Let $h : M \times [0, 1] \rightarrow M$ represent an element of $\text{Iso}(M, B)_G$. This means that h_1 is an equivariant diffeomorphism $f : M \rightarrow M$ that preserves B . Let $a : [0, 1] \rightarrow M$ be the path $a(t) = h_t(b_0)$, i.e. a is the path traveled by the basepoint b_0 under h . Then \widetilde{h} restricts to a lift $\widetilde{a} : [0, 1] \rightarrow \widetilde{M}$ with $\widetilde{a}(0) = \widetilde{b}_0$. To check that \widetilde{h} preserves the component \widetilde{B} , it enough to check that $\widetilde{a}(1) \in \widetilde{B}$. We will first prove this in the case that a is a loop, i.e. $f(b_0) = b_0$. Then we will address the general case.

So, assume that $f(b_0) = b_0$. Our goal is to show that the loop a lifts to a loop in \widetilde{M} . It's enough to show that a represents an element of the subgroup $\pi_1(M_0, b_0) \leq \pi_1(M, b_0)$. Let $G_0 \leq G$ denote the stabilizer of b_0 (see Lemma 4.1 for a description of G_0). Then by Lemma 4.3, it's enough to show that G_0 fixes the element $[a] \in \pi_1(M, b_0)$. Since $\pi_1(M, b_0)$ splits as a free product by Lemma 4.2, it has a trivial center, and therefore it's enough to show that $[a]^{-1}g_0([a])$ is central for any $g_0 \in G_0$.

We can show that $[a]^{-1}g_0([a])$ is central as follows (we adapt an argument of Birman–Hilden [BH73, Lem 1.3]). Let $c : S^1 \rightarrow M$ be any loop based at b_0 . Applying the isotopy h to c , we get that $c \simeq a \cdot f(c) \cdot \bar{a}$, where \cdot denotes concatenation and $\overline{(-)}$ denotes the reverse path. Applying g_0 , we conclude

$$g_0(c) \simeq g_0(a) \cdot g_0(f(c)) \cdot \overline{g_0(a)} \simeq g_0(a) \cdot f(g_0(c)) \cdot \overline{g_0(a)}.$$

On the other hand, applying the isotopy h directly to the element $g_0(c)$, we get

$$g_0(c) \simeq a \cdot f(g_0(c)) \cdot \bar{a}.$$

Thus, we conclude that

$$(\bar{a} \cdot g_0(a)) \cdot f(g_0(c)) \cdot \overline{(\bar{a} \cdot g_0(a))} \simeq f(g_0(c)).$$

Since any element of $\pi_1(M, b_0)$ is represented by $f(g_0(c))$ for some c (as $f \circ g_0$ induces an automorphism of $\pi_1(M, b_0)$), we conclude that $[a]^{-1}g_0([a])$ is central as desired.

Finally, we must address the case that $f(b_0) \neq b_0$. Note that in the argument above, we only used the fact that f was G_0 -equivariant. So, it's enough to find a G_0 -equivariant ambient isotopy of M taking $f(b_0)$ to b_0 ; we can then concatenate h with this isotopy to reduce to the case $f(b_0) = b_0$.

To build this isotopy, fix a G_0 -invariant tubular neighborhood of B (see [Bre72, Thm VI.2.2]). Fix an $(n-2)$ -disk $D \subseteq B$ containing b_0 and $f(b_0)$, and choose an isotopy $k : D \times [0, 1] \rightarrow D$ such that $k_0 = \text{id}$, $k_1(f(b_0)) = b_0$, and $k_t(x) = x$ for all x in a neighborhood of ∂D and $t \in [0, 1]$. As in Lemma 4.1, over D the tubular neighborhood is diffeomorphic to $D \times D^2$, and G_0 acts by rotations on each D^2 -fiber. Choose a smooth function $\mu : [0, 1] \rightarrow [0, 1]$ which equals 1 in a neighborhood of 0 and equals 0 in a neighborhood of 1. Then we define an isotopy $\widehat{k} : D \times D^2 \times [0, 1] \rightarrow D \times D^2$ by $\widehat{k}_t(x, re^{i\theta}) = (k_{\mu(r)t}(x), re^{i\theta})$. Then \widehat{k}_t is G_0 -equivariant since it is independent of θ . Moreover, since \widehat{k}_t is the identity on a neighborhood of $\partial(D \times D^2)$, it extends to a G_0 -equivariant ambient isotopy of M . \square

5.5. Defining the obstruction map. Finally, we can define our obstruction map by combining the maps λ_G and ψ . For emphasis, we make the $\mathbb{Z}/\ell\mathbb{Z}$ -coefficients explicit moving forward.

We define the map $\widehat{\varphi} : \text{Iso}(M, B)_G \rightarrow H_{n-1}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z})$ to be the composition

$$\text{Iso}(M, B)_G \xrightarrow{\lambda_G} \text{Iso}(\widetilde{M}, \widetilde{B}) \xrightarrow{\psi} H_{n-1}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z})$$

The main result of this section is then the following.

Proposition 5.3. *The map $\widehat{\varphi} : \text{Iso}(M, B)_G \rightarrow H_{n-1}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z})$ factors through a map*

$$\varphi : \text{Ker}(\mathcal{F}_G) \rightarrow H_{n-1}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z}).$$

Proof. Recall that we have an exact sequence

$$\pi_1(\text{Diff}(M)) \rightarrow \text{Iso}(M, B)_G \rightarrow \text{Ker}(\mathcal{F}_G) \rightarrow 1.$$

Let $[h] \in \text{Iso}(M, B)_G$ be an element in the image of $\pi_1(\text{Diff}(M))$. Our goal is to show that $\widehat{\varphi}([h])$ is trivial. Applying the map λ_G , we lift $[h]$ to an element $[\tilde{h}] \in \text{Iso}(\widetilde{M}, \widetilde{B})$. Then $\widehat{\varphi}([h]) = \psi([\tilde{h}])$, so it's enough to show that $\psi([\tilde{h}]) = 0$.

First, since $[h] \in \text{Iso}(M, B)_G$ lies in the image of $\pi_1(\text{Diff}(M))$, we may assume that $h_1 = \text{id}_M$. Then, the lift \tilde{h}_1 will be a deck transformation of \widetilde{M} . Since h_1 preserves the component $\widetilde{B} \subseteq \pi^{-1}(B)$, and the deck group of \widetilde{M} acts freely on the components of $\pi^{-1}(B)$ (as $B \subseteq M_0^*$ and the cover is trivial over M_0^*), it must be that $\tilde{h}_1 = \text{id}_{\widetilde{M}}$. Thus, \tilde{h} in fact defines an element of $H_1(\text{Diff}(\widetilde{M}); \mathbb{Z}/\ell\mathbb{Z})$.

Recall from Section 5.2 that we have an evaluation map

$$\text{ev}_{\widetilde{B}} : \text{Diff}(\widetilde{M}) \times \widetilde{B} \rightarrow \widetilde{M},$$

and since subspace $\text{Diff}(\widetilde{M}, \widetilde{B}) \times \widetilde{B}$ maps into \widetilde{B} , the relative Künneth formula gives us a pairing

$$H_1(\text{Diff}(\widetilde{M}), \text{Diff}(\widetilde{M}, \widetilde{B}); \mathbb{Z}/\ell\mathbb{Z}) \otimes H_{n-2}(\widetilde{B}; \mathbb{Z}/\ell\mathbb{Z}) \rightarrow H_1(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z}).$$

On the other hand, we have an evaluation map

$$\text{ev}_{\widetilde{M}} : \text{Diff}(\widetilde{M}) \times \widetilde{M} \rightarrow \widetilde{M},$$

and the Künneth formula gives us a pairing

$$H_1(\text{Diff}(\widetilde{M}); \mathbb{Z}/\ell\mathbb{Z}) \otimes H_{n-2}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z}) \rightarrow H_1(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z}).$$

So, we have the following commutative diagram:

$$\begin{array}{ccccc} H_1(S^1; \mathbb{Z}/\ell\mathbb{Z}) & \longrightarrow & H_1(\text{Diff}(\widetilde{M}); \mathbb{Z}/\ell\mathbb{Z}) & \longrightarrow & H_{n-1}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z}) \\ \uparrow \cong & & & & \downarrow \\ H_1([0, 1], \{0, 1\}; \mathbb{Z}/\ell\mathbb{Z}) & \longrightarrow & H_1(\text{Diff}(\widetilde{M}), \text{Diff}(\widetilde{M}, \widetilde{B}); \mathbb{Z}/\ell\mathbb{Z}) & \longrightarrow & H_{n-1}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z}) \end{array}$$

The upper-left and lower-left horizontal maps are both induced by the isotopy \tilde{h} . The upper right map comes from pairing a class with $[\widetilde{B}] \in H_{n-2}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z})$ via the map $\text{ev}_{\widetilde{M}}$, while the lower right map comes from pairing a class with the fundamental class $[\widetilde{B}] \in H_{n-2}(\widetilde{B}; \mathbb{Z}/\ell\mathbb{Z})$ via $\text{ev}_{\widetilde{B}}$.

The key observation is that the upper-right horizontal map is in fact the zero map. This is because the class $[B] \in H_{n-2}(M; \mathbb{Z}/\ell\mathbb{Z})$ is trivial by assumption, which implies that the class $[\widetilde{B}] \in H_{n-2}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z})$ is trivial (to see this, use the fact that $\pi : \widetilde{M} \rightarrow M$ restricts to a diffeomorphism $W_{\text{id}} \cong M_0^*$, and apply Mayer-Vietoris to the subspaces W_{id} and $\widetilde{M} \setminus W_{\text{id}}$). On the other hand, the composition of the two lower horizontal maps is $\psi([\tilde{h}])$ by definition. Thus, we conclude that $\psi([\tilde{h}])$ is zero as desired. \square

6. AN ELEMENT WITH NONTRIVIAL OBSTRUCTION

In Section 5, we defined an obstruction map

$$\varphi : \text{Ker}(\mathcal{F}_G) \rightarrow H_{n-2}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z}).$$

Our next goal is to construct an element $\alpha_0 \in \text{Ker}(\mathcal{F}_G)$ for which $\varphi(\alpha_0)$ is nontrivial.

First, we recall some notation from Section 4.6. Let $G_{M_0} \leq G$ denote the stabilizer of M_0 (with respect to the action of G on the components of P_0). Then the gluing disks $D_{i,g}^+$ for $g \in G_{M_0}$ are precisely the gluing disks that lie on M_0 . We let $\Sigma_{i,g} \subseteq M$ denote the boundary of the gluing disk $D_{i,g}^+$, and for $g \in G_{M_0}$ and $u \in \Lambda$, and we let $\Sigma_{i,g,u}$ denote the copy of $\Sigma_{i,g}$ on the sheet W_u above M_0 .

Let $G_B \leq G$ denote the setwise stabilizer of B , so $G_B \leq G_{M_0}$. For $(g, u) \in G_{M_0} \times \Lambda$, we let $[\Sigma_{1,g,u}]$ denote the homology class of $\Sigma_{1,g,u}$ in $H_{n-1}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z})$. Since the map $H_{n-1}(\widetilde{M}; \mathbb{Z}/\ell\mathbb{Z}) \rightarrow H_{n-1}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z})$ is injective (as \widetilde{B} has dimension $n-2$), we can equivalently view $[\Sigma_{1,g,u}]$ as an element of $H_{n-1}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z})$. Then we have the following.

Proposition 6.1. *There exists an element $\alpha_0 \in \text{Ker}(\mathcal{F}_G)$ such that*

$$\varphi(\alpha_0) = \sum_{g \in G_B} [\Sigma_{1,g,\text{id}}].$$

We will compute $\varphi(\alpha_0)$ using the description of φ in Section 5.1.

6.1. An informal description of the construction. In order to motivate the construction of this section, we start with an informal picture of the mapping class α_0 . Assume for simplicity that $G_B = G$. Assume also that the gluing disks on M_0 are positioned in a tubular neighborhood of B . We will construct an equivariant diffeomorphism f as a summand slide which rotates all the gluing disks $D_{1,g}^+$ around B , and which fixes the remaining gluing disks. The path traveled by the center of each $D_{1,g}^+$ is a meridian of B .

We can then (non-equivariantly) isotope to f to id_M by modifying the path traveled by each $D_{1,g}^+$. In particular, we pull the path across B and then shrink it down to the trivial loop at the center of each $D_{1,g}^+$, as in Figure 4 (in our actual construction, this will isotope f to a product of sphere twists, so we actually define α_0 by multiplying f with a product of sphere twists).




FIGURE 4. We isotope f to id_M by sliding the disks $D_{1,g}^+$ along shorter and shorter loops.

This gives us an isotopy f_t with $f_0 = f$ and $f_1 = \text{id}_M$, where each f_t is a summand slide. For some range of times $t_1 \leq t \leq t_2$, the summand slide f_t will push the gluing disk through B (this occurs as

we pull the sliding path across B). The end result is that in the image of the map $B \times [0, 1] \rightarrow M$ given by $(b, t) \mapsto f_t(b)$, we get a “bubble” around each gluing disk $D_{1,g}^+$ (see Figure 5). Thus in homology, the map $B \times [0, 1] \rightarrow M$ yields the sum of all the classes $[\Sigma_{1,g}]$ for $g \in G$. Since the cover $\pi : \widetilde{M} \rightarrow M$ restricts to a diffeomorphism $W_{\text{id}} \rightarrow M_0$, we see the same phenomenon when we lift this isotopy to \widetilde{M} .




FIGURE 5. During the isotopy f_t , the sliding path of the disks $D_{1,g}^+$ runs into B , creating a “bubble” around each disk (cf. Figures 8 and 9).

6.2. A family of diffeomorphisms of the 2-disk. Let D^2 denote the closed unit 2-disk in \mathbb{R}^2 , and let C denote the group generated by a $2\pi/m$ -rotation. We begin by defining a family of C -equivariant diffeomorphisms $F_s : D^2 \rightarrow D^2$ for $s \in [0, 1]$. In our construction of α_0 , we will apply these diffeomorphisms to the fibers of a tubular neighborhood of B .

First, for $i \in \{1, 2\}$ and $j \in \{1, \dots, m\}$, let $E_{i,j}$ be a smaller 2-disk in D^2 such that for each i , the disks $E_{i,1}, \dots, E_{i,m}$ are permuted by C . We choose the disks small enough so that $E_{i,j}$ is disjoint from $E_{i',j'}$ for all distinct pairs (i, j) and (i', j') , and so that no $E_{i,j}$ contains the origin. We arrange the disks so that for each j , the centers of the disks $E_{1,j}$ and $E_{2,j}$ lie on the same radial line in D^2 , with $E_{1,j}$ closer to the origin and $E_{2,j}$ farther. See Figure 6 for an illustration.

Let (r, θ) denote polar coordinates on $D^2 - \{0\}$. Choose $r_1 < r_2$ such that the annulus $X := [r_1, r_2] \times [0, 2\pi)$ contains each disk $E_{1,j}$ for $1 \leq j \leq m$ within its interior, and such that the disks $E_{2,j}$ for $1 \leq j \leq m$ are disjoint from X (see Figure 6). Choose $\varepsilon > 0$ such that the regions $[r_1, r_1 + \varepsilon] \times [0, 2\pi)$ and $[r_2 - \varepsilon, r_2] \times [0, 2\pi)$ are disjoint from all the disks $E_{i,j}$. Let $\xi : [0, \varepsilon] \rightarrow [0, 1]$ be a smooth function which equals 0 in a neighborhood of 0 and equals 1 in a neighborhood of ε . Finally, let $S_r \subseteq D^2 - \{0\}$ be the circle of radius $r \in (0, 1]$. Then for $t \in [0, 1]$, we define a C -equivariant diffeomorphism $F_s : D^2 \rightarrow D^2$ as follows:

- For $r \in [r_1, r_1 + \varepsilon]$, the map F_s acts by a $\xi(r - r_1)2\pi s$ -rotation on S_r .
- For $r \in [r_1 + \varepsilon, r_2 - \varepsilon]$, the map F_s acts by a $2\pi s$ rotation on S_r .
- For $r \in [r_2 - \varepsilon, r_2]$, the map F_s acts by a $\xi(r_2 - r)2\pi s$ -rotation.
- Outside of the annulus X , the map F_s acts by id_{D^2} .

In particular, $F_0 = \text{id}_{D^2}$ and F_1 is a left Dehn twist about the inner boundary component of X composed with a right Dehn twist about the outer boundary component of X .

6.3. An equivariant diffeomorphism: Step one. Our next goal is to use the family of diffeomorphisms F_s to construct an equivariant diffeomorphism $f : M \rightarrow M$. The diffeomorphism f will be a summand slide supported on P_0 , but we will need to carefully describe it in coordinates to compute φ . As a first step, we will construct a disk slide $f' : M_0 \rightarrow M_0$.

First, apply Lemma 4.1 to the action of G_{M_0} on M_0 to obtain a point $b_0 \in B$ and associated closed neighborhood $Z_{b_0} \subseteq M_0$. We let G_0 denote the G_{M_0} -stabilizer of b_0 , so G_0 is cyclic. In the following discussion, we will implicitly identify Z_{b_0} with $D^{n-2} \times D^2$.

Without loss of generality, we assume that the gluing disks $D_{i,\text{id}}^+$ for $i \in \{1, 2\}$ lie in the interior of Z_{b_0} , and thus so do the gluing disks $D_{i,g}^+$ for $g \in G_0$. Without loss of generality, we may also assume that for all $g \in G_0$ and $i \in \{1, 2\}$, there is a unique $j \in \{1, \dots, m\}$ such that the gluing disk $D_{i,g}^+$ is contained in the interior of $D^{n-2} \times E_{i,j}$.




FIGURE 6. The disks $E_{i,j}$ and the annulus X . The submanifold B intersects D^2 at the origin.

Fix a closed $(n-2)$ -disk $K \subseteq \text{Int}(D^{n-2})$ such that under the projection map $Z_{b_0} \rightarrow D^{n-2}$, the image of each gluing disk $D_{i,g}^+$ lies in $\text{Int}(K)$. Choose a larger closed disk $K' \subseteq \text{Int}(D^{n-2})$ such that $K \subseteq \text{Int}(K')$. Let $\nu : D^{n-2} \rightarrow [0, 1]$ be a smooth function which is identically 1 on K and which is identically 0 outside of K' . Then, we define $f' : M_0 \rightarrow M_0$ to be the following G_0 -equivariant diffeomorphism:

- on Z_{b_0} , f' acts on the fiber $\{x\} \times D^2$ by the diffeomorphism $F_{\nu(x)}$,
- outside of Z_{b_0} , f' acts by id_{M_0} .

Note that since F_1 fixes the disks $E_{i,j} \subseteq D^2$ pointwise, it follows that f' fixes all of the gluing disks in M_0 . In fact, f' represents a disk slide in $\text{Mod}(M_0, \Delta^0)$, where Δ^0 the union of all the gluing disks on M_0 . Indeed, we can build an isotopy from f' to id_{M_0} by smoothly homotoping the function ν to the constant 0 map. Under this isotopy, the 2-disks $E_{1,1}, \dots, E_{1,m}$ move a full rotation about the origin in each D^2 -fiber over K .

6.4. An equivariant diffeomorphism: Step two. Next, we will use f' to construct the desired equivariant diffeomorphism $f : M \rightarrow M$. To do so, we will first define a disk slide f'' on P_0 , which will then induce a summand slide on M .

Let $\{b_0, \dots, b_\ell\} \subseteq P_0$ denote the G -orbit of b_0 , and choose $g_j \in G$ such that $g_j b_0 = b_j$. Let $Z_{b_j} = g_j Z_{b_0}$ (this set is independent of the choice of g_j). Then we define $f'' : P_0 \rightarrow P_0$ to be the following diffeomorphism:

- on Z_{b_j} , f'' acts by $g_j \circ f' \circ g_j^{-1}$,
- outside of $\bigcup_{j=0}^{\ell} Z_{b_j}$, f acts by id_{P_0} .

This definition is independent of the choice of g_j . Moreover, one can directly check that f is G -equivariant.

The diffeomorphism f'' represents a disk slide in $\text{Mod}(P_0, \Delta_1^+ \cup \Delta_2^+)$ (recall that Δ_i^+ is comprised of the disks $D_{i,g}^+$ for $g \in G$). Indeed, the isotopy $f' \simeq \text{id}_{M_0}$ induces an isotopy $f'' \simeq \text{id}_{P_0}$.

Since f'' fixes each gluing disk pointwise, it extends a diffeomorphism $f : M \rightarrow M$ which acts by the identity outside of P_0^* . By definition, the diffeomorphism f is a summand slide. Moreover, since f'' is equivariant, it follows that f is equivariant. Note that the isotopy $f'' \simeq \text{id}_{P_0}$ does not fix the gluing disks at all times, and therefore does *not* induce an isotopy from f to id_M .

6.5. Identifying a product of sphere twists. As discussed above, the diffeomorphism f is a summand slide, and is induced by a disk slide f'' on P_0 . Let $d_{i,g}$ denote the center of the gluing disk $D_{i,g}^+$. We can see from our explicit description of the map f' that the path traveled by $d_{i,g}$ under the isotopy $f'' \simeq P_0$ is null-homotopic in P_0 . Therefore, by Proposition 2.6, the mapping class $[f] \in \text{Mod}(M)$ lies in the subgroup generated by sphere twists about the spheres $\Sigma_{i,g} = \partial D_{i,g}^+$ for $(i, g) \in \{1, 2\} \times G$.

In fact, we can explicitly determine the class $[f]$ as a product of sphere twists. For each point b_j , we have a diffeomorphism $Z_{b_j} \cong D^{n-2} \times D^2$. Under this identification, each tangent space in Z_{b_j} is naturally isomorphic to $\mathbb{R}^{n-2} \times \mathbb{R}^2$. Fix a frame at $d_{i,g}$ corresponding to the standard basis of $\mathbb{R}^{n-2} \times \mathbb{R}^2$. For $i = 2$, the frame is fixed throughout the isotopy $f'' \simeq \text{id}_{P_0}$. For $i = 1$, the basis vectors in \mathbb{R}^{n-2} remain fixed, while the basis vectors in \mathbb{R}^2 trace out a 2π -rotation. Thus, we conclude that in $\text{Mod}(M)$, we have an equality

$$[f] = \prod_{g \in G} \tau_{1,g},$$

where $\tau_{1,g} \in \text{Mod}(M)$ is the isotopy class of a sphere twist about $\Sigma_{i,g} = \partial D_{1,g}^+$.

6.6. An explicit isotopy. Next, we construct an isotopy $f_t : M \rightarrow M$ for $t \in [0, 1]$ with $f_0 = f$.

First, we construct an isotopy of the diffeomorphisms $F_s : D^2 \rightarrow D^2$. Choose an ambient isotopy $k : D^2 \times [0, 1] \rightarrow D^2$ with $k_0 = \text{id}_{D^2}$ which fixes the disks $E_{i,j}$ pointwise at all times and which pushes the annulus X across the origin, so that both boundary components of $k_1(X)$ have winding number 0 about the origin (see Figure 7). Then, for $t \in [0, 1]$, define $F_{s,t} := k_t \circ F_s \circ k_t^{-1}$. Thus $F_{s,0} = F_s$ for all s , $F_{0,t}$ is the identity for all t , and $F_{1,t}$ is a product of opposite Dehn twists about the two boundary components of $k_t(X)$.




FIGURE 7. The annus X at time $t = 0$ and $t = 1$.

Next, we can build the isotopy of f . For $t \in [0, 1]$, let $f'_t : M_0 \rightarrow M_0$ be the diffeomorphism defined the same as f' , except on each fiber $\{x\} \times D^2$ of Z_{b_0} , f'_t acts by the diffeomorphism $F_{\nu(x),t}$.

instead of $F_{\nu(x)}$. Then we build a disk slide f_t'' on P_0 exact as how we built f'' from f' , and f_t'' induces a summand slide $f_t : M \rightarrow M$.

Note that for each fixed t we have an isotopy $f_t' \simeq \text{id}_{M_0}$ constructed the same way as the isotopy $f' \simeq \text{id}_{M_0}$, and this induces an isotopy $f_t'' \simeq \text{id}_{P_0}$. The key observation is that the path traveled by the center $d_{i,g}$ of $D_{i,g}^+$ under the isotopy $f_1' \simeq \text{id}_{P_0}$ is null-homotopic in $P_0 \setminus B$ (note that when we remove B from P_0 , this removes the origin from each D^2 -fiber of each Z_b). Thus in the mapping class group $\text{Mod}(M \setminus B)$, we have the equality

$$(6.1) \quad [f_1] = \prod_{g \in G} \tau_{1,g}.$$

6.7. Constructing the element. Now, we will construct the desired element $\alpha_0 \in \text{Ker}(\mathcal{F}_G)$. In fact, we will construct an isotopy $h : [0, 1] \rightarrow \text{Diff}(M)$ such that $h_0 = \text{id}_M$ and h_1 is equivariant and preserves B setwise; we can then let $\alpha_0 = [h_1]$. The isotopy f_{1-t} almost works for our purposes, but the problem is that the diffeomorphism $f = f_0$ is isotopic to a product of sphere twists, and thus may not be isotopic to id_M .

Let $T_{i,g} : M \rightarrow M$ be a sphere twist about $\partial D_{i,g}^-$ supported in a collar neighborhood $\partial D_{i,g}^- \times [0, 1] \subseteq P_i \setminus \text{Int}(D_{i,g}^-)$. Since $\partial D_{i,g}^+$ and $\partial D_{i,g}^-$ are isotopic spheres in $M \setminus B$, we know that $T_{i,g}$ represents the mapping class $\tau_{i,g}$ (in both $\text{Mod}(M)$ and $\text{Mod}(M \setminus B)$). Define $h' : [0, 1] \rightarrow \text{Diff}(M)$ to be the isotopy

$$h_t' := f_{1-t} \circ \prod_{g \in G} T_{1,g}.$$

From the equality (6.1) in $\text{Mod}(M \setminus B)$, it follows that there is an isotopy $\text{id}_M \simeq h_0'$ that fixes B at all times. Thus, we construct h by concatenating the isotopy $\text{id}_M \simeq h_0'$ with the isotopy h' .

6.8. Analyzing the image. Let $r : B \times [0, 1] \rightarrow M$ be the map $r(b, t) = h_t(b)$. Our next goal is to describe the image of r . Recall that b_0, \dots, b_ℓ are the points in the G -orbit of b_0 . Assume without loss of generality that for some $k \leq \ell$, the points b_0, \dots, b_k lie on B , and the points b_{k+1}, \dots, b_ℓ do not lie on B . Let $Z_{b_j}^* = Z_{b_j} \cap P_0^* \subseteq M$. Note that $Z_{b_j}^*$ is diffeomorphic to $D^{n-2} \times D^2$ with a finite number of points removed (namely, the center of each gluing disk in Z_{b_j}).

Now, let $b \in B$; then we can describe the loop $r : \{b\} \times [0, 1] \rightarrow M$ as follows. Suppose first that b lies in $Z_{b_0}^*$. Let $K, K' \subseteq D^{n-2}$ be the subsets described in the construction of f' ; we will identify K and K' as subsets of B . If b lies in the $(n-2)$ -disk K , then following the description of the isotopy $F_{\nu(b),t} = F_{1,t}$, the restriction $r : \{b\} \times [0, 1] \rightarrow M$ is a closed curve in the fiber $(\{b\} \times D^2) \cap P_0^* \subseteq Z_{b_0}^*$. One side of this loop contains the disks $E_{1,j}$ for $j \in \{1, \dots, m\}$, and the other side contains the disks $E_{2,j}$ for $j \in \{1, \dots, m\}$ (see Figure 8). As we vary from b from ∂K towards $\partial K'$, this loop shrinks down towards the origin, and for b outside of K' , the loop $r : \{b\} \times [0, 1] \rightarrow M$ is the constant loop at b (see Figure 9).

Next, observe that if b lies in $Z_{b_j}^*$ for $1 \leq j \leq k$, then by applying the element g_j to Z_{b_0} , we see the same picture of the loop $r : \{b\} \times [0, 1] \rightarrow M$ as in the case that $b \in Z_{b_0}^*$. Finally, if b lies outside of the sets Z_{b_j} for each $0 \leq j \leq k$, then $r(b, t) = b$ for all $t \in [0, 1]$.

6.9. Computing the obstruction. Finally, we can compute $\varphi(\alpha_0)$ following the description of φ given in Section 5.1. Let $h : [0, 1] \rightarrow \text{Diff}(M)$ be the isotopy constructed above, so $h_0 = \text{id}_M$ and h_1 represents α_0 . Let $r : B \times [0, 1] \rightarrow M$ be the map $r(b, t) = h_t(b)$ as above. Note that the image of r is contained in $M_0^* \subseteq M$.

Throughout the remainder of this section, we use homology with $\mathbb{Z}/\ell\mathbb{Z}$ -coefficients. To prove Proposition 6.1, it's enough to show that the homology class $r_*([B \times [0, 1]]) \in H_{n-1}(M_0^*, B)$ is equal to the sum of the classes $[\Sigma_{1,g}]$ for $g \in G_B$. Indeed, the image of the lift $\tilde{r} : B \times [0, 1] \rightarrow \tilde{M}$ is




FIGURE 8. On the left is the loop $r : \{b\} \times [0, 1] \rightarrow M$ for $b \in K$. On the right is a slight perturbation of this loop for clarity.




FIGURE 9. The loop $r : \{b\} \times [0, 1] \rightarrow M$ for $b \in K' \setminus K$. On the left is the case that b is closer to ∂K , on the right is the case that b is closer to $\partial K'$. For some values of $b \in K' \setminus K$, the loop will intersect the disks $E_{1,j}$.

contained in the sheet W_{id} . The diffeomorphism $W_{\text{id}} \cong M_0^*$ induces an isomorphism

$$H_{n-1}(M_0^*, B) \cong H_{n-1}(W_{\text{id}}, \tilde{B})$$

which takes $[\Sigma_{1,g}]$ to $[\Sigma_{1,g,\text{id}}]$. Then by Mayer–Vietoris, the natural map

$$H_{n-1}(W_{\text{id}}, \tilde{B}) \rightarrow H_{n-1}(\tilde{M}, \tilde{B})$$

is an embedding (cf. the proof of Proposition 4.5).

Moreover, since h_1 fixes B pointwise, the map r descends to a map $\bar{r} : B \times S^1 \rightarrow M_0^*$. We have the following commutative diagram:

$$\begin{array}{ccc} H_{n-1}(B \times [0, 1], B \times \{0, 1\}) & \xrightarrow{\cong} & H_{n-1}(B \times S^1, B \times \{*\}) \xleftarrow{\cong} H_{n-1}(B \times S^1) \\ \downarrow r_* & & \downarrow \bar{r}_* \\ H_{n-1}(M_0^*, B) & \xleftarrow{\quad\quad\quad} & H_{n-1}(M_0^*) \end{array}$$

Therefore, it's enough to show that the map \bar{r}_* takes the fundamental class of $B \times S^1$ to the sum of the classes $[\Sigma_{1,g}] \in H_{n-1}(M_0^*)$ for $g \in G_B$.

For $j \in \{0, \dots, k\}$, let K_j denote the copy of the $(n-2)$ -disk K in Z_{b_j} , so each K_j is an $(n-2)$ -disk in B (recall that b_0, \dots, b_k are the points in the G -orbit of b_0 that lie on B). Let B° be the closure of $B \setminus \bigcup_{j=0}^k K_j$. For each $b \in \partial K_j$, we saw above that the loop $r : \{b\} \times [0, 1] \rightarrow \{b\} \times D^2$ is a simple closed curve γ_b that bounds a 2-disk C_b containing $E_{1,1}, \dots, E_{1,m}$. Let $\sigma_j \in H_{n-1}(M_0^*)$ be the homology class obtained by taking $r(K_j \times [0, 1])$ and capping the loop γ_b with the 2-disk C_b for each $b \in \partial K_j$. Let $\sigma^\circ \in H_{n-1}(M_0^*)$ be the homology class obtained capping the boundary of $r(B^\circ \times [0, 1])$ in the same way. Then we see that the homology class $\bar{r}_*([B \times S^1]) \in H_{n-1}(M_0^*; \mathbb{Z}/\ell\mathbb{Z})$ is equal to the sum $\sigma^\circ + \sigma_0 + \dots + \sigma_m$.

The class σ° is trivial because the map $\bar{r} : B^\circ \times S^1 \rightarrow M_0^*$ extends to a map $B^\circ \times D^2 \rightarrow M_0^*$. Indeed, if $b \in B^\circ$ lies outside of each Z_{b_j} , then the map $\bar{r} : \{b\} \times S^1 \rightarrow M_0^*$ is the constant loop at b . Otherwise, if $b \in Z_{b_j} \cap B^\circ$, then the loop $\bar{r} : \{b\} \times S^1 \rightarrow M_0^*$ is either the constant loop at b , or a closed curve in a D^2 -fiber of Z_{b_j} which bounds 2-disk in the same fiber. ²

We also see that

$$\sigma_j = \sum_{g \in G_0} [\Sigma_{1,g_j g}].$$

for each $0 \leq j \leq k$. Indeed, after we cap off $r(K_j \times S^1)$, we get a manifold with corners that cobounds an n -dimensional submanifold of M with the spheres $\Sigma_{1,g_j g}$ for $g \in G_0$. Thus the proposition follows.

7. PROOF OF PART (II) OF THEOREM A

In this section, we will complete the proof of part (ii) of Theorem A. Fix a manifold M and a finite group G satisfying the hypotheses of Theorem A. Recall that we let $\text{Mod}_G(M, B) \leq \text{Mod}(M, B)$ and $\text{Mod}_G(M) \leq \text{Mod}(M)$ denote the subgroups of mapping classes with an equivariant representative, and we have a forgetful map $\mathcal{F}_G : \text{Mod}_G(M, B) \rightarrow \text{Mod}_G(M)$. In Section 5, we constructed a map

$$\varphi : \text{Ker}(\mathcal{F}_G) \rightarrow H_{n-1}(\widetilde{M}, \widetilde{B}; \mathbb{Z}/\ell\mathbb{Z}),$$

where $\pi : \widetilde{M} \rightarrow M$ is the cover constructed in Section 4 and \widetilde{B} is a component of $\pi^{-1}(B)$. In Section 6, we constructed an element $\alpha_0 \in \text{Ker}(\mathcal{F}_G)$ for which $\varphi(\alpha_0)$ is nontrivial.

We prove Theorem A as follows. In Section 7.1, we show that it's enough to prove the group $\text{Ker}(\mathcal{F}_G)$ is not finitely generated. In Section 7.2, we show that the obstruction map φ is equivariant with respect to the subgroup $L \leq \text{Mod}_G(M, B)$ of mapping classes which lift to the cover \widetilde{M} . In Section 7.3, we show that the L -orbit of $\varphi(\alpha_0)$ contains an infinite linearly independent set of homology classes, which completes the proof.

7.1. Reduction to the forgetful kernel. First, we will show that it's enough to prove the kernel $\text{Ker}(\mathcal{F}_G)$ is not finitely generated.

Lemma 7.1. *If $\text{Ker}(\mathcal{F}_G)$ is not finitely generated, then $\text{Ker}(\mathcal{P}_G)$ is not finitely generated.*

Proof. Let $\Gamma_G(M, B) = \pi_0(\text{Diff}(M, B)^G)$, where $\text{Diff}(M, B)^G$ is the group of orientation-preserving diffeomorphisms of M that preserve B setwise and commute with G . Then we have the following commutative diagram:

$$\begin{array}{ccc} \Gamma_G(M, B) & \xrightarrow{\Psi} & \text{Mod}_G(M, B) \\ \Phi \downarrow & & \downarrow \mathcal{F}_G \\ \Gamma_G(M) & \xrightarrow{\mathcal{P}_G} & \text{Mod}_G(M) \end{array}$$

²Note that this argument does not apply if $b \in K_j$, since then the loop $\bar{r} : \{b\} \times S^1 \rightarrow M_0^*$ might lie on a D^2 -fiber of Z_{b_j} that contains the center of a gluing disk, and thus does not bound a 2-disk in the intersection of that fiber with M_0^* .

Here Φ is induced by the inclusion $\text{Diff}(M, B)^G \hookrightarrow \text{Diff}(M)^G$, and Ψ is induced by the inclusion $\text{Diff}(M, B)^G \hookrightarrow \text{Diff}(M, B)$. The maps Ψ and \mathcal{P}_G are surjective by definition of $\text{Mod}_G(M, B)$ and $\text{Mod}_G(M)$. To see that the map Φ is injective, note first that if f is an equivariant diffeomorphism that preserves B , then any equivariant isotopy of f must also preserve B setwise, since B is a component of the fixed set of $g_0 \in G$. Thus, if f represents an element of $\text{Ker}(\Phi)$, then there is an equivariant isotopy $f \simeq \text{id}_M$, and since this isotopy must preserve B , this implies that f represents the identity in $\Gamma_G(M, B)$.

Suppose now $\text{Ker}(\mathcal{F}_G)$ is not finitely generated. Since Ψ is surjective, we have a surjection $\text{Ker}(\mathcal{F}_G \circ \Psi) \rightarrow \text{Ker}(\mathcal{F}_G)$, and so $\text{Ker}(\mathcal{F}_G \circ \Psi) = \text{Ker}(\mathcal{P}_G \circ \Phi)$ is not finitely generated. Since Φ is injective, we have that $\text{Ker}(\mathcal{P}_G \circ \Phi) \cong \text{Ker}(\mathcal{P}_G) \cap \text{Im}(\Phi)$. Thus it's enough to show that $\text{Im}(\Phi)$ has finite index in $\Gamma_G(M)$. But this follows since B is a component of $\text{Fix}(g_0)$, and any equivariant diffeomorphism will permute the components of $\text{Fix}(g_0)$, and $\text{Fix}(g_0)$ has only finitely many components since $\text{Fix}(g_0)$ is a closed submanifold and M is compact. \square

7.2. Equivariance of the obstruction map. Next, we will show that the obstruction map φ satisfies a certain equivariance property. Let $\text{Diff}(M, B, b_0)$ denote the group of orientation-preserving diffeomorphisms of M that preserve B setwise and fix the basepoint $b_0 \in B$, and let $\text{Mod}(M, B, b_0) = \pi_0(\text{Diff}(M, B, b_0))$. Let $\text{Mod}_G(M, B, b_0) \leq \text{Mod}_G(M, B, b_0)$ denote the subgroup of mapping classes with an equivariant representative. Then, we let $L_* \leq \text{Mod}_G(M, B, b_0)$ denote the subgroup of mapping classes that lift along the cover $\pi : \widetilde{M} \rightarrow M$, and we let $L \leq \text{Mod}_G(M, B)$ denote the image of L_* .

Since every $\beta \in L_*$ fixes the base point $b_0 \in M$, any representative f has a distinguished lift \tilde{f} which fixes the base point $\tilde{b}_0 \in \widetilde{M}$. Furthermore, since f preserves B setwise, it follows that \tilde{f} preserves \tilde{B} setwise. Thus we get a well-defined map $L_* \rightarrow \text{Mod}(\widetilde{M}, \tilde{B})$. Post-composing with the action on homology, we have an action of L_* on $H_{n-1}(\widetilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$.

First, we can show that the basepoint is irrelevant to this action.

Lemma 7.2. *The action $L_* \rightarrow \text{Aut}(H_{n-1}(\widetilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z}))$ factors through L .*

Proof. Suppose $\beta \in \text{Ker}(L_* \rightarrow L)$, and choose a representative f of β . Then f is isotopic to id_M through an isotopy preserving B setwise. This lifts to an isotopy from \tilde{f} to a deck transformation. Since this lifted isotopy preserves \tilde{B} , and the deck group of \widetilde{M} acts faithfully on the components of $\pi^{-1}(B)$, we in fact have that \tilde{f} is isotopic to $\text{id}_{\widetilde{M}}$ through an isotopy preserving \tilde{B} setwise. Thus \tilde{f} acts trivially on $H_{n-1}(\widetilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$. \square

Thus, we have a well-defined action of L on $H_{n-1}(\widetilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$. On the other hand, L acts on the normal subgroup $\text{Ker}(\mathcal{F}_G) \triangleleft \text{Mod}_G(M, B)$ by conjugation. We can show that φ is equivariant with respect to these two actions, at least up to finite index. Let $L^+ \leq L$ denote the index 2 subgroup of mapping classes which restrict to an orientation-preserving diffeomorphism of B . Then we have the following;

Lemma 7.3. *The map $\varphi : \text{Ker}(\mathcal{F}_G) \rightarrow H_{n-1}(\widetilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$ is L^+ -equivariant.*

Proof. Take any $\alpha \in \text{Ker}(\mathcal{F}_G)$ and $\beta \in L^+$. Our goal is to show that $\varphi(\beta\alpha\beta^{-1}) = \beta \cdot \varphi(\alpha)$.

First, choose an isotopy $h : M \times [0, 1] \rightarrow M$ such that $h_0 = \text{id}$ and h_1 represents α . Let $r : B \times [0, 1] \rightarrow M$ be the map $r(b, t) = h_t(b)$, and let $\tilde{r} : B \times [0, 1] \rightarrow \widetilde{M}$ be the lift mapping $B \times \{0, 1\}$ to \tilde{B} . So, $\varphi(\alpha)$ is the image of the fundamental class under the natural map

$$\tilde{r}_* : H_{n-1}(B \times [0, 1], B \times \{0, 1\}; \mathbb{Z}/\ell\mathbb{Z}) \rightarrow H_{n-1}(\widetilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z}).$$

Next, let f be a representative of β ; without loss of generality we may assume that f fixes the basepoint $b_0 \in B$. Define the isotopy $h^f : M \times [0, 1] \rightarrow M$ to be $f \circ h \circ (f^{-1} \times \text{id})$, so $h_t^f = fh_t f^{-1}$. Then the diffeomorphism h_1^f represents $\beta\alpha\beta^{-1}$. Let $r^f : B \times [0, 1] \rightarrow M$ be the map $r^f(b, t) = h_t^f(b)$,

and let $\tilde{r}^f : B \times [0, 1] \rightarrow \tilde{M}$ be the lift mapping $B \times \{0, 1\}$ to \tilde{B} . Then $\varphi(\beta\alpha\beta^{-1})$ is the image of the fundamental class under the map

$$\tilde{r}_*^f : H_{n-1}(B \times [0, 1], B \times \{0, 1\}; \mathbb{Z}/\ell\mathbb{Z}) \rightarrow H_{n-1}(\tilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z}).$$

Now, let \tilde{f} be the lift of f along the cover $\pi : \tilde{M} \rightarrow M$ that fixes \tilde{b}_0 . We claim that $\tilde{r}^f = \tilde{f} \circ \tilde{r} \circ (f^{-1} \times \text{id})$. Indeed, we see that

$$\begin{aligned} \pi \circ \left(\tilde{f} \circ \tilde{r} \circ (f^{-1} \times \text{id}) \right) &= f \circ \pi \circ \tilde{r} \circ (f^{-1} \times \text{id}) \\ &= f \circ r \circ (f^{-1} \times \text{id}) \\ &= r^f, \end{aligned}$$

so $\tilde{r}^f = \tilde{f} \circ \tilde{r} \circ (f^{-1} \times \text{id})$ by the uniqueness of lifts.

So, we have the following commutative diagram:

$$\begin{array}{ccc} H_{n-1}(B \times [0, 1], B \times \{0, 1\}; \mathbb{Z}/\ell\mathbb{Z}) & \xrightarrow{(f^{-1} \times \text{id})_*} & H_{n-1}(B \times [0, 1], B \times \{0, 1\}; \mathbb{Z}/\ell\mathbb{Z}) \\ & \searrow \tilde{r}_*^f & \downarrow \tilde{r}_* \\ & & H_{n-1}(\tilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z}) \\ & & \downarrow \tilde{f}_* \\ & & H_{n-1}(\tilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z}) \end{array}$$

Since f^{-1} is an orientation-preserving diffeomorphism of B , the map $f^{-1} \times \text{id}$ acts by the identity on $H_{n-1}(B \times [0, 1], B \times \{0, 1\}; \mathbb{Z}/\ell\mathbb{Z})$. Thus, the lemma follows. \square

7.3. Completing the proof. By Proposition 6.1, we know that

$$\varphi(\alpha_0) = \sum_{g \in G_B} [\Sigma_{1,g,\text{id}}],$$

where $G_B \leq G$ is the subgroup of G preserving B setwise. By Lemma 7.3, the map φ is L^+ -equivariant. Thus, to prove Theorem A, it's enough to show that the L^+ orbit of $\sum_{g \in G_B} [\Sigma_{1,g,\text{id}}]$ contains an infinite linearly independent subset of $H_{n-1}(\tilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$.

Recall from Proposition 4.5 that the set

$$\mathcal{S} = \{[\Sigma_{1,g,u}] \in H_{n-1}(\tilde{M}; \mathbb{Z}/\ell\mathbb{Z}) \mid g \in G_{M_0}, u \in \Lambda_2\}$$

is infinite and linearly independent, where G_{M_0} is the G -stabilizer of M_0 (so $G_B \leq G_{M_0}$) and $\Lambda_2 \leq \pi_1(M)$ is the subgroup constructed in Section 4.5. The set \mathcal{S} embeds into $H_{n-1}(\tilde{M}, \tilde{B}; \mathbb{Z}/\ell\mathbb{Z})$. Thus, we can complete the proof of Theorem A with the following lemma.

Lemma 7.4. *For each $u \in \Lambda_2$, there exists $\beta_u \in L^+$ such that $\beta_u([\Sigma_{1,\text{id},\text{id}}]) = [\Sigma_{1,\text{id},u}]$, and for all $g \in G_{M_0}$, $\beta_u([\Sigma_{1,g,\text{id}}]) = [\Sigma_{1,g,u}]$ for some $u' \in \Lambda_2$.*

Proof. Throughout this proof, we will use the notation of Sections 2 and 4. Recall that we define Λ_2 as $\pi_1(\bar{\mathcal{G}}_2, M_0)$, and by Lemma 4.2, we can represent elements of Λ_2 by loops in M based at b_0 which are disjoint from the gluing disks $D_{1,g}^\pm$ for all $g \in G$. Let M' denote the manifold obtained from gluing P_0 and P_2 along the multidisks $\Delta_{2,g}^\pm$, and let M'_0 denote the component of M' containing M_0 . Then we can represent elements of Λ_2 with loops in M'_0 .

Now, fix $u \in \Lambda_2$. For $g \in G$, let $d_{1,g}$ denote the center of the gluing disk $D_{1,g}^+$ in P_0 . By choosing an arc in M_0 from b_0 to $d_{1,\text{id}}$, we get an identification $\pi_1(M'_0, b_0) \cong \pi_1(M'_0, d_{1,\text{id}})$; let $\gamma_{\text{id}} : S^1 \rightarrow M'_0$ be a loop based at $d_{1,\text{id}}$ representing the element u . Then for $g \in G$, let $\gamma_g = g\gamma_{\text{id}}$, so γ_g is a loop in M' based at $d_{1,g}$.

Without loss of generality, we may assume that the loops γ_g are all disjoint. Choose an element $\delta \in \pi_1(\text{Fr}_{\Delta_1^+}(M'))$ which maps onto the tuple $(\gamma_g)_{g \in G}$ under the map $\pi_1(\text{Fr}_{\Delta_1^+}(M')) \rightarrow \prod_{g \in G} \pi_1(M', d_{1,g})$. Then δ determines a disk slide $\mathcal{DS}_{\Delta_1^+}(\delta)$ on M' ; since the loops γ_g are disjoint and permuted by G , we may choose an equivariant representative f' of $\mathcal{DS}_{\Delta_1^+}(\delta)$. Since f' fixes the multidisks Δ_1^\pm pointwise, it extends to an equivariant diffeomorphism $f : M \rightarrow M$. Since B has codimension 2, we may assume that the loops γ_g are disjoint from B , and hence we may assume that f fixes B pointwise.

Let $\beta_u \in \text{Mod}_G(M, B)$ be the element represented by f . Note that β lies in the subgroup $L^+ \leq \text{Mod}_G(M, B)$; indeed, β is liftable since f acts trivially on the subgroup $\pi_1(M_0, b_0) \leq \pi_1(M, b_0)$ and hence preserves the normal closure $\langle\langle \pi_1(M_0, b_0) \rangle\rangle$. Let $\tilde{f} : \tilde{M} \rightarrow \tilde{M}$ be the lift fixing \tilde{b}_0 . Then it follows by construction that \tilde{f} preserves the Λ_2 -orbit of each $[\Sigma_{1,g,\text{id}}]$, and that β_u maps $[\Sigma_{1,\text{id},\text{id}}]$ to $[\Sigma_{1,\text{id},u}]$. \square

8. PROOF OF THEOREM D

In this section, we prove Theorem D using Theorem A. For this proof, it will be simpler to work with *homeomorphisms*, rather than diffeomorphisms. This is not a significant difference; if we let M be a smooth manifold and G a finite group of diffeomorphisms satisfying the hypotheses of Theorem A, then a slight modification of the proof of Theorem A shows that the kernel of the map $\pi_0(\text{Homeo}(M)^G) \rightarrow \pi_0(\text{Homeo}(M))$ is not finitely generated (essentially, one can just replace each instance of the word “diffeomorphism” with “homeomorphism” in Section 5 to extend our obstruction map to the topological setting).

We prove Theorem D as follows. In Section 8.1, we recall the relationship between symmetric automorphisms and branched covers of the unlink in S^3 . In Section 8.2, we define a “lifting map” whose kernel is virtually isomorphic to (the topological version of) $\text{Ker}(\mathcal{P}_G)$. In Sections 8.3 and 8.4, we study the orbifold fundamental group associated to the branched cover. We complete the proof in Section 8.5 by using the orbifold fundamental group to show that $\text{Ker}(\mathcal{Q}_{k,d})$ is virtually isomorphic to the kernel of the lifting map.

8.1. Branched covers of the unlink and symmetric automorphisms. Fix integers $k, d > 1$, and let $p_{k,d} : M_{k,d} \rightarrow S^3$ be the finite regular cover branched over the k -component unlink $C_k \subseteq S^3$ and with deck group $G_d = \mathbb{Z}/d\mathbb{Z}$. For notational simplicity, we let $p = p_{k,d}$, $M = M_{k,d}$, and $G = G_d$. We note that the manifold M is homeomorphic to a connected sum of $S^1 \times S^2$'s, though we will not use this fact.³

Let $\text{Homeo}(S^3, C_k)$ denote the group of orientation-preserving homeomorphisms of S^3 preserving C_k setwise (not necessarily preserving the orientation of C_k), and let $\text{Mod}(M, C_k) = \pi_0(\text{Homeo}(M, C_k))$. We have an isomorphism $\pi_1(S^3 \setminus C_k) \cong F_k$, where F_k is the free group of rank k . We fix a basis $\{x_1, \dots, x_k\}$ of F_k where each x_i corresponds to a loop that winds once around the i th component of C_k . As mentioned in Section 1.4, a theorem of Goldsmith [Gol81], building off the work of Dahm [Dah62], says that the action of $\text{Mod}(S^3, C_k)$ on $\pi_1(S^3 \setminus C_k)$ induces an isomorphism

$$\text{Mod}(S^3, C_k) \cong \text{SymOut}(F_k),$$

where $\text{SymOut}(F_k)$ is the group of outer automorphisms sending each x_i to a conjugate of some $x_j^{\pm 1}$ (Goldsmith actually proves an analogous result about the \mathbb{R}^3 , but the above isomorphism can easily be deduced, see e.g. [Luc25, Prop 5.1]). We let $\text{Mod}^+(S^3, C_k)$ denote the finite index subgroup of homeomorphisms preserving the orientation of C_k , and let $\text{SymOut}^+(F_k)$ denote its image in $\text{SymOut}(F_k)$. Then $\text{SymOut}^+(F_k)$ is the group of outer automorphisms sending each x_i to a conjugate of some x_j .

³This can be seen by viewing $S^3 \setminus C_k$ as a k -fold connected sum $(S^3 \setminus C_1)^{\#k}$ and constructing a model of M as in Section 4.

8.2. The lifting map. Let $M^\circ = M \setminus p^{-1}(C)$. Then the unbranched cover $M^\circ \rightarrow S^3 \setminus C_k$ is classified by the natural map $F_k \rightarrow \mathbb{Z}/d\mathbb{Z}$ sending each x_i to 1. Since precomposition by an element of $\text{SymOut}^+(F_k)$ preserves this homomorphism, it follows that every element of $\text{Mod}^+(S^3, C_k)$ lifts along the cover p to an equivariant homeomorphism of M (see [Luc25, §2.4]). Let $\text{Mod}(M) = \pi_0(\text{Homeo}(M))$, and let $\text{Mod}_G(M)$ be the subgroup of mapping classes represented by an equivariant homeomorphism. Since any lift is well-defined up to a deck transformation, we get a *lifting map*

$$\mathcal{L}_p : \text{Mod}^+(S^3, C_k) \rightarrow \text{Mod}_G(M)/G.$$

which sends a mapping class to a lift.

Let $\Gamma_G(M) = \pi_0(\text{Homeo}(M)^G)$. Observe that the kernel $\text{Ker}(\mathcal{L}_p)$ is virtually isomorphic to the kernel of the map $\mathcal{P}_G : \Gamma_G(M) \rightarrow \text{Mod}_G(M)$. Indeed, since any equivariant (topological) isotopy of M descends to an isotopy of S^3 preserving C_k setwise, we get a map $\Gamma_G(M) \rightarrow \text{Mod}(S^3, C_k)$ with a finite kernel and a finite-index image. This induces a virtual isomorphism between $\text{Ker}(\mathcal{L}_p)$ and $\text{Ker}(\mathcal{P}_G)$. As mentioned at the beginning of this section, the proof of Theorem A easily adapts to the setting of homeomorphisms. Thus, we conclude that $\text{Ker}(\mathcal{L}_p)$ is not finitely generated for $k \geq 3$.

8.3. The orbifold fundamental group. The branched cover p induces an orbifold structure \mathcal{O}_p on S^3 , where each point of C_k is an orbifold point of order d . Then we have an isomorphism

$$\pi_1^{\text{orb}}(\mathcal{O}_p) \cong \underbrace{\mathbb{Z}/d\mathbb{Z} * \cdots * \mathbb{Z}/d\mathbb{Z}}_{k \text{ times}} =: H_{k,d}.$$

Indeed, the orbifold fundamental group is the quotient of $\pi_1(S^3 \setminus C_k)$ obtained by adding the relations $x_i^d = 1$ for each i (see e.g. [HD84]). We let z_i denote the image of x_i in $\pi_1^{\text{orb}}(\mathcal{O}_p)$. The branched cover p gives us the following commutative diagram with exact rows:

$$(8.1) \quad \begin{array}{ccccccc} 1 & \longrightarrow & \pi_1(M^\circ) & \longrightarrow & \pi_1(S^3 \setminus C_k) & \longrightarrow & G \longrightarrow 1 \\ & & \downarrow & & \downarrow & & \parallel \\ 1 & \longrightarrow & \pi_1(M) & \longrightarrow & \pi_1^{\text{orb}}(\mathcal{O}_p) & \longrightarrow & G \longrightarrow 1 \end{array}$$

where the map $\pi_1^{\text{orb}}(\mathcal{O}_p) \rightarrow G = \mathbb{Z}/d\mathbb{Z}$ maps each z_i to 1.

Since $\text{Mod}^+(S^3, C_k)$ preserves the normal closure of the elements x_1^d, \dots, x_k^d , the action of $\text{Mod}^+(S^3, C_k)$ on $\pi_1(S^3 \setminus C_k)$ descends to an action of $\text{Mod}^+(S^3, C_k)$ on $\pi_1^{\text{orb}}(\mathcal{O}_p)$ by outer automorphisms. Let $\text{SymOut}(H_{k,d})$ denote the group of outer automorphisms sending each z_i to a conjugate of some z_j^\pm , and let $\text{SymOut}^+(H_{k,d})$ denote the finite index subgroup sending each z_i to a conjugate of some z_j . The projection $F_k \rightarrow H_{k,d}$ induces a natural map

$$\mathcal{Q}_{k,d} : \text{SymOut}(F_k) \rightarrow \text{SymOut}(H_{k,d}).$$

It follows that the action of $\text{Mod}^+(S^3, C_k)$ on $\pi_1^{\text{orb}}(\mathcal{O}_p)$ factors through the restriction

$$\mathcal{Q}_{k,d}^+ : \text{SymOut}^+(F_k) \rightarrow \text{SymOut}^+(H_{k,d}).$$

8.4. The restriction map. Let $\text{Aut}(\pi_1^{\text{orb}}(\mathcal{O}_p), \pi_1(M))$ denote the group of automorphisms of $\pi_1^{\text{orb}}(\mathcal{O}_p)$ that preserve the subgroup $\pi_1(M)$. Since $\pi_1(M)$ is a normal subgroup of $\pi_1^{\text{orb}}(\mathcal{O}_p)$, we know that $\text{Inn}(\pi_1^{\text{orb}}(\mathcal{O}_p)) \leq \text{Aut}(\pi_1^{\text{orb}}(\mathcal{O}_p), \pi_1(M))$, and so we define

$$\text{Out}(\pi_1^{\text{orb}}(\mathcal{O}_p), \pi_1(M)) := \text{Aut}(\pi_1^{\text{orb}}(\mathcal{O}_p), \pi_1(M)) / \text{Inn}(\pi_1^{\text{orb}}(\mathcal{O}_p)).$$

Let G_* denote the image of G in $\text{Out}(\pi_1(M))$, and let $\text{Out}_{G_*}(\pi_1(M))$ denote its normalizer. Then we claim there is a map

$$\bar{r} : \text{Out}(\pi_1^{\text{orb}}(\mathcal{O}_p), \pi_1(M)) \rightarrow \text{Out}_{G_*}(\pi_1(M)) / G_*.$$

To construct the map \bar{r} , we start with the restriction map

$$r : \text{Aut}(\pi_1^{\text{orb}}(\mathcal{O}_p), \pi_1(M)) \rightarrow \text{Aut}(\pi_1(M)),$$

and post-compose to obtain a map

$$r' : \text{Aut}(\pi_1^{\text{orb}}(\mathcal{O}_p), \pi_1(M)) \rightarrow \text{Out}(\pi_1(M)).$$

From the short exact sequence $\pi_1(M) \hookrightarrow \pi_1^{\text{orb}}(\mathcal{O}_p) \rightarrow G$, we see that the map r' takes $\text{Inn}(\pi_1^{\text{orb}}(\mathcal{O}_p))$ to the group G_* . Since $\text{Inn}(\pi_1^{\text{orb}}(\mathcal{O}_p))$ is normal in $\text{Aut}(\pi_1^{\text{orb}}(\mathcal{O}_p), \pi_1(M))$, it follows that G_* is normal in the image of r' . Thus the map r' descends to the map \bar{r} .

Now, applying the isomorphism $\pi_1^{\text{orb}}(\mathcal{O}_p) \cong H_{k,d}$, we can view $\text{SymOut}^+(H_{k,d})$ as a subgroup of $\text{Out}(\pi_1^{\text{orb}}(\mathcal{O}_p), \pi_1(M))$. Then \bar{r} restricts a map

$$\bar{R} : \text{SymOut}^+(H_{k,d}) \rightarrow \text{Out}_{G_*}(\pi_1(M))/G_*.$$

Lemma 8.1. *The map $\bar{R} : \text{SymOut}^+(H_{k,d}) \rightarrow \text{Out}_{G_*}(\pi_1(M))/G_*$ has a finite kernel.*

Proof. Let $\text{PSymOut}(H_{k,d})$ be the finite index subgroup of $\text{SymOut}(H_{k,d})$ consisting of outer automorphisms that send each z_i to a conjugate of itself. Then it's enough to show that \bar{R} is injective on $\text{PSymOut}(H_{k,d})$.

Suppose f represents an element of $\text{PSymOut}(H_{k,d})$ and $\bar{R}([f])$ is trivial. Our goal is to show that f is an inner automorphism of $H_{k,d}$. Since $\bar{R}([f])$ is trivial, this means there exists $\iota \in \text{Inn}(H_{k,d})$ such that $\iota \circ f$ restricts to the identity map on $\pi_1(M)$. It's enough to show that $\iota \circ f$ is inner. Thus, we assume without loss of generality that f itself acts trivially on $\pi_1(M)$.

We will view elements of $H_{k,d}$ as freely reduced words on the elements z_1, \dots, z_k . For each $1 \leq i \leq k$, we have that $f(z_i) = w_i z_i w_i^{-1}$ for some reduced word w_i . Without loss of generality, we may assume that w_i does not end with a power of z_i .

First, we claim that if w_i is trivial for some i , then w_i must be trivial for all $1 \leq i \leq k$, in which case we are done. To see this, suppose w_i is trivial and let $1 \leq j \leq k$ with $j \neq i$. Then the element $z_i z_j^{-1}$ lies in the subgroup $\pi_1(M)$, and so since f is the identity on $\pi_1(M)$, we have that

$$z_i z_j^{-1} = f(z_i z_j^{-1}) = z_i w_j z_j^{-1} w_j^{-1}.$$

This is possible only if w_j is trivial.

Otherwise, assume that w_i is nontrivial for all i . Let i and j be arbitrary. Then we have that

$$z_i z_j^{-1} = f(z_i z_j^{-1}) = w_i z_i w_i^{-1} w_j z_j w_j^{-1}.$$

Given a reduced word w on the elements z_1, \dots, z_k , we say that a reduced word w' is a *left subword* (resp. *right subword*) of w if there exists a reduced word v such that $w = w'v$ and the last letter of w' is not a power of the first letter of v (resp. $w = vw'$ and the last letter of v is not a power of the first letter of w'). We say the left (resp. right) subword is *proper* if v is nontrivial. Then the above equality implies that

$$w_i^{-1} w_j = w$$

where w is either a proper left subword of w_i^{-1} or a proper right subword of w_j . Now, the equality

$$z_i z_j^{-1} = w_i z_i w_i^{-1} w_j z_j w_j^{-1}$$

implies that w is nontrivial, and either

$$w_i z_i w_i^{-1} w_j^{-1} = w'_i z_j^{-1} w_j^{-1}$$

where w'_i is a left subword of w_i , or

$$w_i z_i w_i^{-1} w_j^{-1} = w_i z_i (w_j^{-1})'$$

where $(w_j^{-1})'$ is a right subword of w_j^{-1} . Then we have that

$$z_i z_j^{-1} = w'_i z_j^{-1} w_j^{-1} \text{ or } z_i z_j^{-1} = w_i z_i (w_j^{-1})',$$

and either case implies that w_i starts with z_i or w_j starts with z_j . But then from the equality

$$w_i^{-1} w_j = w,$$

we see that if w_i starts with z_i , then w_j must start with z_i , and if w_j starts with z_j , then w_i must start with z_j . Thus, we conclude that w_i and w_j must start with the same letter. Since i and j were arbitrary, we conclude that the words w_1, \dots, w_k all must start with the same letter z . Then if we let $f' = \text{Inn}(z) \circ f$, we have that $f'(z_i) = v_i z_i v_i^{-1}$ where v_i is a reduced word of shorter length than w_i . By repeating the argument inductively, we obtain a sequence of inner automorphisms ι_1, \dots, ι_m such that $\iota_1 \circ \dots \circ \iota_m \circ f$ is trivial, and thus f is inner as desired. \square

8.5. Proof of Theorem D. Now, we can prove Theorem D. First, note that if $\text{Ker}(\mathcal{Q}_{k,d})$ is not finitely generated, then $\text{Ker}(\widehat{\mathcal{Q}}_{k,d})$ must also be not finitely generated. Indeed, we have the following commutative diagram:

$$\begin{array}{ccc} \text{SymAut}(F_k) & \xrightarrow{\widehat{\mathcal{Q}}_{k,d}} & \text{SymAut}(H_{k,d}) \\ \Phi \downarrow & & \downarrow \Psi \\ \text{SymOut}(F_k) & \xrightarrow{\mathcal{Q}_{k,d}} & \text{SymOut}(H_{k,d}) \end{array}$$

Here Φ and Ψ are both the quotient map by the group of inner automorphisms. If $\text{Ker}(\mathcal{Q}_{k,d})$ is not finitely generated, then since $\text{Ker}(\mathcal{Q}_{k,d} \circ \Phi)$ surjects onto $\text{Ker}(\mathcal{Q}_{k,d})$, we conclude that $\text{Ker}(\mathcal{Q}_{k,d} \circ \Phi) = \text{Ker}(\Psi \circ \widehat{\mathcal{Q}}_{k,d})$ is not finitely generated. Now, since the map $F_k \rightarrow H_{k,d}$ is surjective, it follows that $\widehat{\mathcal{Q}}_{k,d}$ restricts to a surjection $\text{Inn}(F_k) \rightarrow \text{Inn}(H_{k,d}) = \text{Ker}(\Psi)$. Thus we have a short exact sequence

$$1 \rightarrow \text{Ker}(\widehat{\mathcal{Q}}_{k,d}) \rightarrow \text{Ker}(\Psi \circ \widehat{\mathcal{Q}}_{k,d}) \rightarrow \text{Ker}(\Psi) \rightarrow 1.$$

Since $\text{Ker}(\Psi) = \text{Inn}(H_{k,d})$ is finitely generated, it must be that $\text{Ker}(\widehat{\mathcal{Q}}_{k,d})$ is not finitely generated.

Thus, it remains to prove that the group $\text{Ker}(\mathcal{Q}_{k,d})$ is not finitely generated for $k \geq 3$. As discussed in Section 8.2, Theorem A implies that $\text{Ker}(\mathcal{L}_p)$ is not finitely generated for $k \geq 3$. Thus Theorem D is a consequence of the following proposition.

Proposition 8.2. *The kernels of the maps*

$$\mathcal{L}_p : \text{Mod}^+(S^3, C_k) \rightarrow \text{Mod}_G(M)/G$$

and

$$\mathcal{Q}_{k,d} : \text{SymOut}(F_k) \rightarrow \text{SymOut}(H_{k,d})$$

are virtually isomorphic.

Proof. Let $\overline{R} : \text{SymOut}^+(H_{k,d}) \rightarrow \text{Out}_{G_*}(\pi_1(M))/G_*$ be the restriction map from Section 8.4. Observe that if f is a homeomorphism of (S^3, C_k) representing an element of $\text{Mod}^+(S^3, C_k)$, then the from the diagram (8.1), the action of a lift \tilde{f} on $\pi_1(M)$ is obtained by restricting the action of f on $\pi_1^{\text{orb}}(\mathcal{O}_p)$. Therefore we have the following commutative diagram:

$$\begin{array}{ccc} \text{Mod}^+(S^3, C_k) & \xrightarrow{\cong} & \text{SymOut}^+(F_k) \\ \downarrow \mathcal{L}_p & & \downarrow \mathcal{Q}_{k,d}^+ \\ & & \text{SymOut}^+(H_{k,d}) \\ & & \downarrow \overline{R} \\ \text{Mod}_G(M)/G & \xrightarrow{\Phi} & \text{Out}_{G_*}(\pi_1(M))/G_* \end{array}$$

Here Φ is induced by the action of $\text{Mod}(M)$ on $\pi_1(M)$ by outer automorphisms. To prove the proposition, it's enough to show that the maps Φ and \bar{R} have finite kernels. We know that \bar{R} has a finite kernel by Lemma 8.1. We also know that Φ has a finite kernel since the kernel of the map $\text{Mod}(M) \rightarrow \text{Out}(\pi_1(M))$ is the subgroup generated by sphere twists (see e.g. [HW10, Prop 2.1]), and this subgroup is finite [McC90, Prop 1.2]. \square

REFERENCES

- [ALS09] Javier Aramayona, Christopher J. Leininger, and Juan Souto. “Injections of mapping class groups”. In: *Geom. Topol.* 13.5 (2009), pp. 2523–2541. DOI: 10.2140/gt.2009.13.2523.
- [BF81] E. Binz and H. R. Fischer. “The manifold of embeddings of a closed manifold”. In: *Differential Geometric Methods in Mathematical Physics*. Ed. by Heinz-Dietrich Doebner. Berlin, Heidelberg: Springer Berlin Heidelberg, 1981, pp. 310–325. ISBN: 978-3-540-38573-8.
- [BH73] Joan S. Birman and Hugh M. Hilden. “On isotopies of homeomorphisms of Riemann surfaces”. In: *Ann. of Math. (2)* 97 (1973), pp. 424–439. DOI: 10.2307/1970830.
- [Bre72] Glen E Bredon. *Introduction to Compact Transformation Groups*. Vol. 46. Pure and Applied Mathematics. Academic Press, 1972. ISBN: 0121288501.
- [BS84] Ronald Brown and Abdul Razak Salleh. “A van Kampen theorem for unions of non-connected spaces”. In: *Archiv der Mathematik* 42.1 (Jan. 1984), pp. 85–88. ISSN: 1420-8938. DOI: 10.1007/bf01198133.
- [Cer61] Jean Cerf. “Topologie de certains espaces de plongements”. In: *Bulletin de la Société Mathématique de France* 79 (1961), pp. 227–380. ISSN: 2102-622X. DOI: 10.24033/bsmf.1567.
- [Cog11] José Ignacio Cogolludo-Agustín. “Braid Monodromy of Algebraic Curves”. In: *Annales mathématiques Blaise Pascal* 18.1 (2011), pp. 141–209. ISSN: 1259-1734. DOI: 10.5802/ambp.295.
- [Dah62] D. M. Dahm. “A Generalization of Braid Theory”. PhD thesis. Princeton University, 1962.
- [GMT03] David Gabai, G. Meyerhoff, and Nathaniel Thurston. “Homotopy hyperbolic 3-manifolds are hyperbolic”. In: *Annals of Mathematics* 157.2 (2003), pp. 335–431. ISSN: 0003-486X. DOI: 10.4007/annals.2003.157.335.
- [Gol81] Deborah L. Goldsmith. “The theory of motion groups”. In: *Michigan Math. J.* 28.1 (1981), pp. 3–17.
- [HD84] A. Haefliger and Quach Ngoc Du. “Appendice: une présentation du groupe fondamental d’une orbifold”. In: 116. Transversal structure of foliations (Toulouse, 1982). 1984, pp. 98–107.
- [Hig76] P. J. Higgins. “The Fundamental Groupoid of a Graph of Groups”. In: *Journal of the London Mathematical Society* s2-13.1 (May 1976), pp. 145–149. ISSN: 0024-6107. DOI: 10.1112/jlms/s2-13.1.145.
- [HW10] Allen Hatcher and Nathalie Wahl. “Stabilization for mapping class groups of 3-manifolds”. In: *Duke Math. J.* 155.2 (2010), pp. 205–269. DOI: 10.1215/00127094-2010-055.
- [Ill00] S. Illman. “Existence and uniqueness of equivariant triangulations of smooth proper G-manifolds with some applications to equivariant Whitehead torsion”. In: *Journal für die reine und angewandte Mathematik (Crelles Journal)* 2000.524 (2000). ISSN: 1435-5345. DOI: 10.1515/crll.2000.054.
- [KM63] Michel A. Kervaire and John W. Milnor. “Groups of Homotopy Spheres: I”. In: *The Annals of Mathematics* 77.3 (May 1963), p. 504. ISSN: 0003-486X. DOI: 10.2307/1970128.
- [Kob72] Shoshichi Kobayashi. *Transformation Groups in Differential Geometry*. Springer Berlin Heidelberg, 1972. ISBN: 9783642619816. DOI: 10.1007/978-3-642-61981-6.

- [Luc25] Trent Lucas. “Birman-Hilden theory for 3-manifolds”. In: *Advances in Mathematics* 468 (2025), p. 110204. ISSN: 0001-8708. DOI: 10.1016/j.aim.2025.110204.
- [Mag80] Wilhelm Magnus. “Rings of Fricke characters and automorphism groups of free groups”. In: *Math. Z.* 170.1 (1980), pp. 91–103. DOI: 10.1007/BF01214715.
- [McC86] J. McCool. “On basis-conjugating automorphisms of free groups”. In: *Canad. J. Math.* 38.6 (1986), pp. 1525–1529. DOI: 10.4153/CJM-1986-073-3.
- [McC90] Darryl McCullough. “Topological and algebraic automorphisms of 3-manifolds”. In: *Groups of self-equivalences and related topics (Montreal, PQ, 1988)*. Vol. 1425. Lecture Notes in Math. Springer, Berlin, 1990, pp. 102–113. DOI: 10.1007/BFb0083835.
- [MH75] C. Maclachlan and W. J. Harvey. “On mapping-class groups and Teichmüller spaces”. In: *Proc. London Math. Soc. (3)* 30 (1975), pp. 496–512. DOI: 10.1112/plms/s3-30.4.496.
- [MM96] Darryl McCullough and Andy Miller. “Symmetric automorphisms of free products”. In: *Mem. Amer. Math. Soc.* 122.582 (1996). DOI: 10.1090/memo/0582.
- [MW21] Dan Margalit and Rebecca R. Winarski. “Braids groups and mapping class groups: the Birman-Hilden theory”. In: *Bull. Lond. Math. Soc.* 53.3 (2021), pp. 643–659. DOI: 10.1112/blms.12456.
- [Pal60] Richard S. Palais. “Extending diffeomorphisms”. In: *Proceedings of the American Mathematical Society* 11.2 (1960), pp. 274–277. ISSN: 1088-6826. DOI: 10.1090/s0002-9939-1960-0117741-0.
- [Ram25] Sidhanth Raman. “A smooth Birman–Hilden theory for hyper-Kähler manifolds”. In: *Journal für die reine und angewandte Mathematik (Crelles Journal)* (2025). ISSN: 1435-5345. DOI: 10.1515/crelle-2025-0007.
- [Ran98] Andrew Ranicki. *High-dimensional Knot Theory*. Springer Berlin Heidelberg, 1998. ISBN: 9783662120118. DOI: 10.1007/978-3-662-12011-8.
- [Ser80] Jean-Pierre Serre. *Trees*. Springer Berlin Heidelberg, 1980. ISBN: 9783642618567. DOI: 10.1007/978-3-642-61856-7.
- [Spa81] Edwin H. Spanier. *Algebraic Topology*. Springer New York, 1981. ISBN: 9781468493221. DOI: 10.1007/978-1-4684-9322-1.
- [Vog77] Elmar Vogt. “Projecting isotopies of sufficiently large P^2 -irreducible 3-manifolds”. In: *Arch. Math. (Basel)* 29.6 (1977), pp. 635–642. DOI: 10.1007/BF01220467.
- [Wat72] Frank Wattenberg. “Differentiable Motions of Unknotted, Unlinked Circles in 3-Space”. In: *Mathematica Scandinavica* 30.1 (1972), pp. 107–135.