Communications in
Analysis and Geometry
Volume 26, Number 6, 1223–1250, 2018

Intrinsic chirality of graphs in 3-manifolds

Erica Flapan and Hugh Howards

The main result of this paper is that for every closed, connected, orientable, irreducible 3-manifold M, there is an integer n_M such that if γ is a graph with no involution and a 3-connected minor λ with genus$(\lambda) > n_M$, then every embedding of γ in M is chiral. By contrast, the paper also proves that for every graph γ, there are infinitely many closed, connected, orientable, irreducible 3-manifolds M such that some embedding of γ in M is pointwise fixed by an orientation reversing involution of M.

1. Introduction

We say that a graph Γ embedded in a 3-manifold M is *achiral*, if there is an orientation reversing homeomorphism h of M leaving Γ setwise invariant. If such an h exists, we say that it is a homeomorphism of the pair (M, Γ). If no such homeomorphism exists, we say that the embedded graph Γ is *chiral*.

We can think of a knot as an embedding of a circular graph with some number of vertices in S^3. Independent of the number of vertices on the graph, some embeddings of it are chiral while others are achiral. By contrast, there exist graphs which have the property that all of their embeddings in S^3 are chiral. Such a graph is said to be *intrinsically chiral* in S^3 because its

The first author was supported in part by NSF Grant DMS-1607744.

chirality depends only on the intrinsic structure of the graph and not on the extrinsic topology of an embedding of the graph in S^3. The following theorem provides a method of constructing graphs that are intrinsically chiral in S^3.

Theorem 1. *[2] Every non-planar graph γ with no order 2 automorphism is intrinsically chiral in S^3.*

Since chirality is defined for graphs embedded in any 3-manifold, it is natural to ask whether a graph which is intrinsically chiral in S^3 would necessarily be intrinsically chiral in other 3-manifolds. Our first result shows that no graph can be intrinsically chiral in every 3-manifold, or even in every "nice" 3-manifold.

Theorem 2. *For every graph γ, there are infinitely many closed, connected, orientable, irreducible 3-manifolds M such that some embedding of γ in M is pointwise fixed by an orientation reversing involution of M.*

Theorem 2 can be thought of as a generalization of the fact that every planar graph has an embedding in S^3 which is pointwise fixed by a reflection of S^3.

On the other hand, our main result is a generalization of Theorem 1, which shows that for any given "nice" 3-manifold M, there are infinitely many graphs with no automorphism of order 2 which are intrinsically chiral in M. In particular, let M be a closed, connected, orientable, irreducible 3-manifold. Then by Kneser-Haken finiteness [12], M contains at most finitely many non-parallel disjoint incompressible tori. Let N_M denote either the maximal possible number of non-parallel disjoint incompressible tori in M or 2, whichever is larger. Now we define $n_M = \dim_{\mathbb{Z}_2}(H_1(M,\mathbb{Z}_2)) + N_M$. We will refer to the constants n_M and N_M in the statements and proofs of Theorem 3 and Proposition 2.

Finally, we make use of the following definitions from graph theory. A graph γ is said to be 3-*connected*, if at least three vertices together with the edges containing them must be removed to disconnect γ or reduce γ to a single vertex. A graph γ_1 is said to be a *minor* of a graph γ_2, if γ_1 can be obtained from γ_2 by deleting and/or contracting some number of edges. Our main result is as follows.

Theorem 3. *For every closed, connected, orientable, irreducible 3-manifold M, any graph with no order 2 automorphism which has a 3-connected minor λ with* $\mathrm{genus}(\lambda) > n_M$ *is intrinsically chiral in M.*

For example, let M be a closed, connected, orientable, irreducible Seifert fibered space whose base surface S has genus $g \geq 2$. Then by using a standard pants decomposition argument (see for example [5]), it can be shown that the closed surface S contains at most $3g-3$ disjoint, non-parallel essential circles. This implies that M contains at most $3g-3$ disjoint, non-parallel incompressible vertical tori. Since M is Seifert fibered and $g \geq 2$, all incompressible tori in M are vertical. Hence M has at most $3g-3$ disjoint, non-parallel incompressible tori. Now let $n_M = \dim_{\mathbb{Z}_2}(H_1(M, \mathbb{Z}_2)) + 3g - 3$. Then by Theorem 3, any graph with no order 2 automorphism which has a 3-connected minor λ with $\text{genus}(\lambda) > n_M$ is intrinsically chiral in M.

Note that by contrast with Theorem 3 which is for 3-manifolds without boundary, Ikeda [8] has shown in the theorem below that for "nice" 3-manifolds with aspherical boundary, any graph with large enough genus which has a certain type of automorphism of order 2 has an achiral hyperbolic embedding in the double of the manifold.

Ikeda's Theorem. [8] Let M be a compact, connected, orientable, 3-manifold with non-empty aspherical boundary. Then there is an integer n_M such that for any abstract graph λ with $\text{genus}(\lambda) > n_M$ and no vertices of valence 1, any automorphism of order 2 of λ that does not restrict to an orientation preserving automorphism of a cycle in λ can be induced by an orientation reversing involution of some hyperbolic embedding of λ in the double of M.

Since the word "graph" is used in different ways by different authors, we note here that in the current paper as well as in [2] a *graph* is defined to be a connected 1-complex consisting of a finite number of vertices and edges such that every edge has two distinct vertices and there is at most one edge with a given pair of vertices.

In Section 2, we prove Theorem 2. In Section 3, we prove Theorem 3 making use of a proposition, which we then prove in Section 4. Throughout the paper we work in the smooth category.

2. Achiral embeddings

The goal of this section is to prove Theorem 2. To that end, we prove the following proposition. Note that we use $\dim_{\mathbb{Z}}(H_1(M, \mathbb{Z}))$ to denote the rank of the first $\mathbb{Z}$-homology group of M and we use $\dim_{\mathbb{Z}_2}(H_1(M, \mathbb{Z}_2))$ to denote the rank of the first $\mathbb{Z}_2$-homology group of M.

Proposition 1. *Let S be a closed, orientable surface. Then for infinitely many closed, connected, orientable, irreducible 3-manifolds Q such that*

$\dim_{\mathbb{Z}_2}(H_1(Q, \mathbb{Z}_2)) = \text{genus}(S)$, *there is an embedding of S in Q which is pointwise fixed by an orientation reversing involution of Q.*

The proof of Proposition 1 will make use of the idea of a *disk-busting* curve, which is a simple closed curve in a handlebody that intersects every essential, properly embedded disk in the handlebody. For example, a core of a solid torus is disk busting in the solid torus. For a genus 2 handlebody whose fundamental group is generated by a and b, an example of a disk-busting curve is one that includes into the fundamental group of the handlebody as $abab^{-1}$ (see Figure 1). All handlebodies have disk-busting curves and Richard Stong [17] gives an algorithm to recognize them. For more on disk-busting curves see [7] and [17].

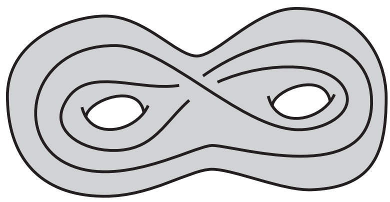

Figure 1: A disk-busting curve in a genus 2 handlebody.

Proof. Let g be the genus of a closed orientable surface S. Let M be the manifold obtained by gluing genus g handlebodies V_1 and V_2 together along S in such a way that there is an orientation reversing involution h interchanging V_1 and V_2 which pointwise fixes the surface S. Then M is a closed, connected, orientable 3-manifold M such that $\dim_{\mathbb{Z}_2}(H_1(M, \mathbb{Z}_2)) = \text{genus}(S)$. However, M is reducible.

In order to create an irreducible 3-manifold from M, we first remove neighborhoods N_1 and N_2 of identical disk busting curves C_1 and C_2 in the interiors of the handlebodies V_1 and V_2 such that N_1 and N_2 are interchanged by the involution h. Note that since $\text{cl}(M - (S \cup N_1 \cup N_2))$ consists of two identical handlebodies from which neighborhoods of disk busting curves have been removed, the inclusion of S into each component of $\text{cl}(M - (S \cup N_1 \cup N_2))$ is incompressible.

We prove by contradiction that $\text{cl}(V_i - N_i) \subset \text{cl}(M - (S \cup N_1 \cup N_2))$ is irreducible. Assume $\text{cl}(V_i - N_i)$ is reducible. Let F be a sphere in $\text{cl}(V_i - N_i)$ that doesn't bound a ball. F must separate S from ∂N_i since every sphere in the handlebody V_i bounds a ball in V_i. Now take any compressing disk D for $S = \partial V_i$ in the handlebody V_i chosen so that it intersects F minimally. Since

F is a sphere, an innermost loop argument implies $D \cap F = \emptyset$. This, however, implies that D is disjoint from ∂N_i and therefore that S is compressible in $\mathrm{cl}(V_i - N_i)$. This contradicts the conclusion above, so F must not exist and $\mathrm{cl}(V_i - N_i)$ is irreducible.

Now we sew in identical non-trivial knot complements Q_1 and Q_2 to $\mathrm{cl}(M - (N_1 \cup N_2))$ along ∂N_1 and ∂N_2 respectively so that the Seifert surfaces of the Q_i are glued in where the meridians of the N_i were.

Let Q denote the 3-manifold obtained in this way. Then the restriction $h|\mathrm{cl}(M - (N_1 \cup N_2))$ can be extended to an orientation reversing involution of Q that pointwise fixes S. Note that the components of $\mathrm{cl}(Q - (S \cup \partial Q_1 \cup \partial Q_2))$ that contain S are homeomorphic to the corresponding components of $\mathrm{cl}(M - (S \cup \partial N_1 \cup \partial N_2))$ that contain S.

Claim 1. The surfaces S, ∂Q_1, and ∂Q_2 are each incompressible in Q.

Proof of Claim 1. Assume one of S, ∂Q_1, or ∂Q_2 is compressible in Q. Then there is a compressing disk D for one of these surfaces whose interior meets $S \cup \partial Q_1 \cup \partial Q_2$ transversally in a minimal number of components. If the interior of D intersects one of the surfaces S, ∂Q_1, or ∂Q_2, then an innermost loop on D bounds a compressing disk Δ for that surface whose interior is disjoint from the other surfaces. Now since each Q_i is a knot complement, Δ cannot be contained in Q_i. Thus Δ must be contained in some component of $\mathrm{cl}(Q - (S \cup Q_1 \cup Q_2))$. Also, since the C_i are disk busting for the V_i, $\partial\Delta$ cannot be contained in S. So Δ is a compressing disk for one of the ∂N_i in $\mathrm{cl}(V_i - N_i)$. By compressing along Δ we obtain a sphere F in $\mathrm{cl}(V_i - N_i)$ separating S from ∂N_i. Then F is a reducing sphere for $\mathrm{cl}(V_i - N_i)$, contradicting the fact proven above that $\mathrm{cl}(V_i - N_i)$ is irreducible. $\square$

Claim 2. Q is irreducible.

Proof of Claim 2. Let F be a sphere in Q and assume without loss of generality that F intersects $S \cup \partial Q_1 \cup \partial Q_2$ transversally in a minimal number of components. If F intersects any of ∂Q_1, ∂Q_2, or S, then there is an innermost loop on F bounding a disk D that is a compressing disk for ∂Q_1, ∂Q_2, or S in $\mathrm{cl}(Q - (S \cup \partial Q_1 \cup \partial Q_2))$. But this violates Claim 1. Thus F is contained in some Q_i or $\mathrm{cl}(V_i - Q_i)$. However, since each Q_i is a knot complement it is irreducible, and we saw above that $\mathrm{cl}(V_i - N_i) = \mathrm{cl}(V_i - Q_i)$ is irreducible. Therefore F bounds a ball in Q. $\square$

Now, recall that $\dim_{\mathbb{Z}_2}(H_1(M, \mathbb{Z}_2)) = \mathrm{genus}(S)$. Also, it can be seen using a Mayer-Vietoris sequence that replacing the handlebodies N_1 and N_2 by the knot complements Q_1 and Q_2 does not change the first homology, since a meridional disk of N_i is replaced by a Seifert surface in Q_i. Thus $\dim_{\mathbb{Z}_2}(H_1(Q, \mathbb{Z}_2)) = \text{genus}(S)$, and hence Q has the properties required by the proposition.

In order to prove that we can find infinitely many such 3-manifolds Q, we begin by letting $X_1 = Q$. Now it follows from Kneser-Haken finiteness, that since X_1 is compact and orientable, there is a finite constant t_1 such that X_1 cannot contain more than t_1 disjoint closed, non-parallel, incompressible surfaces (for example, see Proposition 1.7 in [4] or see [12]).

Now let $n > t_1$, let $K_1, \dots, K_n$ be distinct non-trivial knots, and let D denote a disk with n holes. Then $D \times S^1$ has $n+1$ torus boundary components which we denote by $S_1, S_2, \dots, S_n$, and T_1. Now for each $j = 1, \dots n$, we glue the knot complement $R_j = \text{cl}(S^3 - N(K_j))$ to the torus S_j so that the boundary of a Seifert surface for R_j is attached to $S_j \cap D$. Then $Y_1 = (D \times S^1) \cup R_1 \cdots \cup R_n$ is the complement of the connected sum $K_1 \# K_2 \# \dots \# K_n$, the S_j are pairwise disjoint, non-parallel, incompressible tori in Y_1, and $\partial Y_1 = T_1$. Let Y_1' denote a copy of Y_1. Now we replace Q_1 by Y_1 and Q_2 by Y_1' glued in such a way that the closed manifold X_2 obtained in this way has an orientation reversing involution which interchanges Y_1 and Y_1' pointwise fixing S.

Thus X_2 contains two copies of the tori $S_1, S_2, \dots, S_n$. To see that these $2n$ tori are incompressible in X_2, suppose there is a compressing disk for one of the tori that intersects $\partial Y_1 \cup \partial Y_1' \cup S$ transversally in a minimal number of components. An innermost loop on the disk would be a compressing disk for ∂Y_1, $\partial Y_1'$, or S, again contradicting Claim 1. Thus the torus would have to be compressible in Y_1 or Y_1'. But this is impossible by our construction of Y_1. Thus the $2n$ tori together with ∂Y_1 and $\partial Y_1'$ must all be incompressible in X_2.

It follows that X_2 contains at least $2n + 2 > t_1$ disjoint, non-parallel, incompressible tori, and thus X_2 is distinct from X_1. By repeating this process, we can create an infinite sequence of such manifolds each containing more disjoint, non-parallel, incompressible tori than the previous manifold did. □

Remark 1. One of the referees observed that an alternate proof of Proposition 1 can be obtained using techniques of hyperbolic Dehn surgery together with results from [15]. We prefer the above proof because our construction is explicit.

Since every graph embeds in a closed orientable surface, the following theorem is an immediate consequence of Proposition 1.

Theorem 2. *For every graph γ, there are infinitely many closed, connected, orientable, irreducible 3-manifolds M such that some embedding of γ in M is pointwise fixed by an orientation reversing involution of M.*

3. Intrinsic chirality

The goal of this section is to prove our main result. We begin with the following definition.

Definition 3. We define the *genus of a closed surface S* by

$$\frac{2-\chi(S)}{2} = \frac{\dim_{\mathbb{Z}_2}(H_1(S,\mathbb{Z}_2))}{2}.$$

Then the *genus of a graph* is defined as the minimum genus of any closed (orientable or non-orientable) surface in which the graph embeds.

It is worth pointing out that the genus of a non-orientable surface is not consistently defined in the literature. Some papers use our definition, while others define the projective plane as having genus 1 instead of $\frac{1}{2}$.

Proposition 2. Let γ be a 3-connected graph with genus at least 2, and let Γ be an embedding of γ in a closed, connected, orientable, irreducible 3-manifold M such that (M,Γ) has an orientation reversing homeomorphism fixing every vertex of Γ.

Then there is an embedding Γ' of γ in a closed, connected, orientable 3-manifold M' such that (M',Γ') has an orientation reversing involution pointwise fixing Γ' and $\dim_{\mathbb{Z}_2}(H_1(M',\mathbb{Z}_2)) \leq n_M$.

The point of this proposition is that if we have an embedding Γ of a graph γ in a manifold M such that (M,Γ) has an orientation reversing *homeomorphism*, then we can find a (possibly different) manifold M' and an embedding Γ' of γ in M' such that (M',Γ') has an orientation reversing *involution*. Furthermore, even though M' might be homologically more complicated than M, there is a bound on $\dim_{\mathbb{Z}_2}(H_1(M',\mathbb{Z}_2))$ which depends only on M and not on the graph γ or the particular embedding of γ in M.

Note that the orientation reversing homeomorphism in the hypothesis of Proposition 2, leaves each edge of Γ setwise invariant since, according to our definition of a graph, there is at most one edge between any pair of vertices. While this homeomorphism does not necessarily pointwise fix Γ, it is isotopic to a homeomorphism which does fix Γ pointwise.

Proposition 2 will be proved in the next section. We now prove Theorem 3 (restated below) by making use of Proposition 2 together with the following result of Kobayashi.

Kobayashi's Theorem. [13] Let X be a closed, orientable, 3-manifold admitting an orientation reversing involution h. Then

$$\dim_{\mathbb{Z}_2}(H_1(\mathrm{fix}(h),\mathbb{Z}_2)) \leq \dim_{\mathbb{Z}_2}(H_1(X,\mathbb{Z}_2)) + \dim_{\mathbb{Z}}(H_1(X,\mathbb{Z})).$$

Theorem 3. *For every closed, connected, orientable, irreducible 3-manifold M, there is an integer n_M such that any abstract graph with no order 2 automorphism which has a 3-connected minor λ with $\mathrm{genus}(\lambda) > n_M$ is intrinsically chiral in M.*

Proof. Let γ be a graph with no order 2 automorphism. Suppose for the sake of contradiction that γ has an achiral embedding Γ in the manifold M. Let n_M and N_M be the constants associated with M that were defined in Section 1. Let λ be a 3-connected minor of γ. We will now show that λ satisfies the inequality

$$\mathrm{genus}(\lambda) \leq \dim_{\mathbb{Z}_2}(H_1(M,\mathbb{Z}_2)) + N_M = n_M.$$

First observe that if $\mathrm{genus}(\lambda) < 2$, then the above inequality is immediate since $N_M \geq 2$. Thus we assume that $\mathrm{genus}(\lambda) \geq 2$.

Since Γ is an achiral embedding of γ in M, there is an orientation reversing homeomorphism f of the pair (M,Γ). Let φ denote the automorphism that f induces on the graph Γ. Note that even though f does not necessarily have finite order, φ has finite order because Γ has a finite number of vertices. Hence we can express the order of φ as $2^r q$, where $r \geq 0$ and q is odd. Now since f is orientation reversing and q is odd, $g = f^q$ is an orientation reversing homeomorphism of (M,Γ). Also, g induces the automorphism φ^q on Γ and $\mathrm{order}(\varphi^q) = 2^r$. In particular, g^{2^r} fixes every vertex of Γ.

If $r \geq 1$, then $g^{2^{r-1}}$ would induce an order two automorphism on Γ. As we assumed that no such automorphism exists, we must have $r = 0$. Thus $g = g^{2^r}$ is an orientation reversing homeomorphism of (M,Γ) which fixes every vertex of Γ.

Since λ is a minor of the abstract graph γ, by deleting and/or contracting some edges of the embedding Γ of γ in M, we obtain an embedding Λ of λ in M. Furthermore, by composing the homeomorphism g with an isotopy in a neighborhood of each edge that was contracted, we obtain an orientation reversing homeomorphism of (M,Λ) which fixes every vertex of Λ.

Since λ is a 3-connected graph with genus$(\lambda) \geq 2$, we can now apply Proposition 2 to get an embedding Λ' of λ in a closed, connected, orientable, 3-manifold M' such that (M', Λ') has an orientation reversing involution h pointwise fixing Λ' and

$$\dim_{\mathbb{Z}_2}(H_1(M', \mathbb{Z}_2)) \leq n_M = \dim_{\mathbb{Z}_2}(H_1(M, \mathbb{Z}_2)) + N_M.$$

Let F be the component of the fixed point set $\mathrm{fix}(h)$ containing Λ', and let $x \in \Lambda'$. Since h is a smooth involution, we can now pick a neighborhood $N(x)$ which is homeomorphic to a ball and is invariant under h (see for example, Theorem 2.2 in [1]). Then by Smith Theory [16], since $h|N(x)$ is an orientation reversing involution of the ball $N(x)$, the fixed point set of $h|N(x)$ is either a single point or a properly embedded disk. Since $N(x) \cap \Lambda$ contains more than one point, $\mathrm{fix}(h|N(x))$ is a properly embedded disk, and hence F is a closed surface. Thus

$$\begin{aligned} \mathrm{genus}(\lambda) \leq \mathrm{genus}(F) &= \frac{2 - \chi(F)}{2} \\ &= \frac{\dim_{\mathbb{Z}_2}(H_1(F, \mathbb{Z}_2))}{2} \leq \frac{\dim_{\mathbb{Z}_2}(H_1(\mathrm{fix}(h), \mathbb{Z}_2))}{2}. \end{aligned}$$

Hence we have the inequality

$$2\mathrm{genus}(\lambda) \leq \dim_{\mathbb{Z}_2}(H_1(\mathrm{fix}(h), \mathbb{Z}_2)).$$

Also, since h is an orientation reversing involution and M' is a closed orientable manifold, we can apply Kobayashi's Theorem [13] to obtain the inequality

$$\dim_{\mathbb{Z}_2}(H_1(\mathrm{fix}(h), \mathbb{Z}_2)) \leq \dim_{\mathbb{Z}_2}(H_1(M', \mathbb{Z}_2)) + \dim_{\mathbb{Z}}(H_1(M', \mathbb{Z})).$$

It follows that

$$\dim_{\mathbb{Z}_2}(H_1(\mathrm{fix}(h), \mathbb{Z}_2)) \leq 2\dim_{\mathbb{Z}_2}(H_1(M', \mathbb{Z}_2)).$$

Combining the above inequalities, we now have

$$\mathrm{genus}(\lambda) \leq \dim_{\mathbb{Z}_2}(H_1(M', \mathbb{Z}_2)).$$

But M' was given by Proposition 2 such that

$$\dim_{\mathbb{Z}_2}(H_1(M', \mathbb{Z}_2)) \leq \dim_{\mathbb{Z}_2}(H_1(M, \mathbb{Z}_2)) + N_M.$$

Hence we obtain the required inequality

$$\text{genus}(\lambda) \leq \dim_{\mathbb{Z}_2}(H_1(M, \mathbb{Z}_2)) + N_M = n_M.$$

It now follows that if γ has a 3-connected minor whose genus is greater than n_M, then γ must be intrinsically chiral in M. □

4. Proof of Proposition 2

In the course of the proof of Proposition 2, we will use the well known "Half Lives, Half Dies" Theorem, which we state below. See [6] or [4] for a proof of this theorem.

Theorem 4. (Half Lives, Half Dies) *Let M be a compact orientable 3-manifold. Then the following equation holds with any field coefficients*

$$\dim\left(\text{Kernel}(H_1(\partial M) \to H_1(M))\right) = \frac{1}{2} \dim H_1(\partial M).$$

Corollary 1. *Let M be a manifold which has a torus boundary component T. Then for any pair of generators a and b of $H_1(T, \mathbb{Z}_2)$, at least one of a or b is non-trivial in $H_1(M, \mathbb{Z}_2)$.*

Proof. Suppose for the sake of contradiction that the generators a and b are both trivial in $H_1(M, \mathbb{Z}_2)$. Attach handlebodies to all boundary components of M except T to form a new manifold J with a single boundary component. Then a and b are both trivial in $H_1(J, \mathbb{Z}_2)$. Since a and b generate the homology of the only boundary component of J, we see that $\dim\left(\text{Kernel}(H_1(\partial J, \mathbb{Z}_2)) \to H_1(J, \mathbb{Z}_2)\right) = \dim_{\mathbb{Z}_2}(H_1(\partial J, \mathbb{Z}_2)) = 2$. But this contradicts the Half Lives, Half Dies Theorem. Thus at least one of the generators of $H_1(T, \mathbb{Z}_2)$ must have been non-trivial in M. □

Note that it is tempting to assume that the Half Lives, Half Dies Theorem implies that one of a or b must be trivial in J. But this is not always true. In particular, if M is the product of a torus and an interval, then no non-trivial curve in the boundary of M is in the kernel.

The proof of Proposition 2 will make use of the Characteristic Decomposition Theorem for Pared Manifolds. For readers who may not be familiar with pared manifolds, we include the following definitions and results.

Definition 5. Let M be an orientable 3-manifold and let P be a set of disjoint incompressible annuli and tori in ∂M. Then we say (M, P) is a *pared* 3*-manifold.*

Note that a pared manifold is a special type of 3-manifold with boundary pattern in the sense of Johannson [11] as well as a 3-manifold pair in the sense of Jaco-Shalen [9]. The following definitions are consistent with those of [11] and [9].

Definition 6. A pared manifold (M, P) is *simple* if it satisfies the following conditions:

1) M is irreducible and $\partial M - P$ is incompressible.
2) Every incompressible torus in M is parallel to a torus component of P.
3) Every annulus A in M such that $\partial A \subseteq \partial M - P$ is either compressible or parallel to an annulus A' in ∂M with $\partial A' = \partial A$ such that $A' \cap P$ consists of zero or one annular component of P.

Definition 7. A pared manifold (M, P) is *Seifert fibered* if there is a Seifert fibration of M for which P is a union of fibers. A pared manifold (M, P) is *I-fibered* if there is an I-bundle map of M over a surface B such that P is the preimage of ∂B.

Characteristic Decomposition Theorem for Pared Manifolds. [9], [11] Let (X, P) be a pared manifold where X is irreducible and $\partial X - P$ is incompressible. Then, up to an isotopy of (X, P), there is a unique set Ω of disjoint incompressible tori and annuli with $\partial\Omega \subseteq \partial X - P$ such that the following conditions hold:

1) If W is the closure of a component of $X - \Omega$, then the pared manifold $(W, W \cap (P \cup \Omega))$ is either simple, Seifert fibered, or I-fibered.
2) There is no set Ω' with fewer elements than Ω which satisfies the above.

Thurston's Theorem for Pared Manifolds. [18] Suppose that a pared manifold (M, P) is simple, M is connected, and ∂M is non-empty. Then either $M - P$ admits a finite volume complete hyperbolic metric with totally geodesic boundary, (M, P) is Seifert fibered, or (M, P) is I-fibered.

We are now ready to prove Proposition 2. Recall that the definition of the constant n_M is given in Section 1.

Proposition 2. *Let γ be a 3-connected graph with genus at least 2, and let Γ be an embedding of γ in a closed, connected, orientable, irreducible 3-manifold M such that (M, Γ) has an orientation reversing homeomorphism g fixing every vertex of Γ.*

Then there is an embedding Γ' *of* γ *in a closed, connected, orientable 3-manifold* M' *such that* (M', Γ') *has an orientation reversing involution pointwise fixing* Γ' *and* $\dim_{\mathbb{Z}_2}(H_1(M', \mathbb{Z}_2)) \leq n_M$.

Proof. Let Λ denote either Γ or Γ with one edge deleted, and suppose that Λ is contained in a ball B in M. Since $g(B)$ is isotopic to B in M, we can assume that g leaves B setwise invariant. Also, since g fixes every vertex of Λ and Λ has at most one edge between any pair of vertices, g takes each edge to itself. Thus $g|B$ is isotopic to a homeomorphism which pointwise fixes Λ. So we assume that g leaves B setwise invariant and pointwise fixes Λ. Let f be an embedding of (B, Λ) in S^3. Then $f \circ g \circ f^{-1}$ is an orientation reversing homeomorphism of $f(B)$ pointwise fixing $f(\Lambda)$. Now, $f \circ g \circ f^{-1}$ can be extended to an orientation reversing homeomorphism of S^3 pointwise fixing $f(\Lambda)$.

However, Jiang and Wang [10] showed that no graph containing $K_{3,3}$ or K_5 has an embedding in S^3 which is pointwise fixed by an orientation reversing homeomorphism of S^3. Thus Λ cannot contain $K_{3,3}$ or K_5, and hence is abstractly planar. But this implies that genus$(\Gamma) \leq 1$, which is contrary to our hypothesis. Thus neither Γ nor Γ with an edge deleted can be contained in a ball in M. We will use this result later in the proof.

Since the remainder of the proof is quite lengthy, we break it into steps.

Step 1: We define a neighborhood $N(\Gamma)$.

Let V and E be the sets of vertices and edges of Γ respectively. For each vertex $v \in V$, define $N(v)$ to be a ball around v in M (i.e., a 0-handle containing v), and let $N(V)$ denote the union of the balls around the vertices. Also, for each edge $e \in E$, let $N(e) = D \times I$ be a solid tube around $\mathrm{cl}(e - N(V))$ in M (i.e., a 1-handle containing the portion of e outside of the 0-handles). Then the intersection $N(V) \cap N(e)$ follows the standard convention for attaching 1-handles to 0-handles. In other words, in $N(e)$ this intersection consists of the disks $D \times \{0\}$ and $D \times \{1\}$ in $\partial N(V)$. Let $N(E)$ denote the union of the tubes around the edges. Then $N(\Gamma) = N(E) \cup N(V)$ is a neighborhood of Γ.

For convenience we introduce the following terminology. For each vertex v, we let $\partial' N(v)$ denote the sphere with holes $\partial N(v) \cap \partial N(\Gamma)$, and for each edge e we let $\partial' N(e)$ denote the annulus $\partial N(e) \cap \partial N(\Gamma)$. Thus $\partial N(\Gamma) = \partial' N(E) \cup \partial' N(V)$.

Now since $g(\Gamma) = \Gamma$ fixing each vertex of Γ, we know that $g(N(\Gamma))$ is isotopic to $N(\Gamma)$ setwise fixing Γ and fixing each vertex. Thus we can modify

g by an isotopy (and by an abuse of notation, still refer to the map as g) so that $g(\Gamma) = \Gamma$, and for each vertex v and edge e we have $g(N(v)) = N(v)$ and $g(N(e)) = N(e)$. Because this modification was by an isotopy, our new g is still orientation reversing.

Step 2: We split $\mathrm{cl}(M - N(\Gamma))$ along a family τ of JSJ tori and choose an invariant component X.

Since M is irreducible and we have assumed that Γ is not contained in a ball, $\mathrm{cl}(M - N(\Gamma))$ is irreducible. Thus we can apply the Characteristic Decomposition Theorem of Jaco-Shalen [9] and Johannson [11] to get a minimal family of incompressible tori τ for $\mathrm{cl}(M - N(\Gamma))$ such that each closed up component of $M - (N(\Gamma) \cup \tau)$ is either Seifert fibered or atoroidal. Since the characteristic family τ is unique up to isotopy, we can again modify g by an isotopy (and again by an abuse of notation still refer to the map as g) so that $g(\tau) = \tau$ and still have $g(\Gamma) = \Gamma$, $g(N(v)) = N(v)$ and $g(N(e)) = N(e)$ for each vertex v and edge e. Let X be the closed up component of $M - (N(\Gamma) \cup \tau)$ containing $\partial N(\Gamma)$ (see for example Figure 2), then $g(X) = X$.

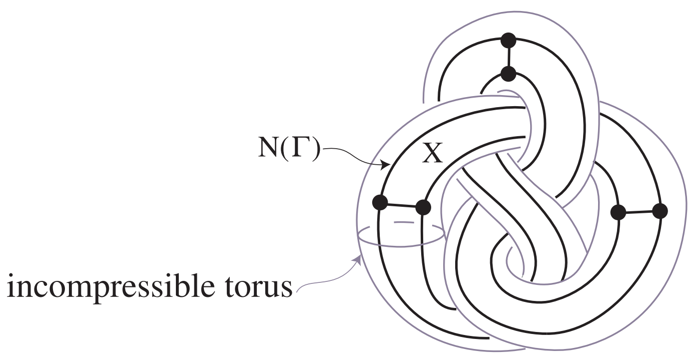

Figure 2: X is the closed up component of $M - (N(\Gamma) \cup \tau)$ between the grey incompressible torus and the black $\partial N(\Gamma)$.

Also, since Γ is 3-connected, $\mathrm{genus}(\partial N(\Gamma)) > 1$. Thus the component X is not Seifert fibered, and hence is atoroidal. Let P denote the set of torus boundary components of X together with the annuli that make up the components of $\partial' N(E)$. Since Γ is 3-connected, $\partial X - P = \partial N(\Gamma) - \partial' N(E) = \partial' N(V)$ is incompressible in $\mathrm{cl}(M - N(\Gamma))$. It follows that $\partial X - P$ is incompressible in X. Furthermore, X is irreducible since $\mathrm{cl}(M - N(\Gamma))$ is irreducible and X is a component of the JSJ decomposition of $\mathrm{cl}(M - N(\Gamma))$.

In the next step, we consider the pared manifold (X, P).

Step 3: We show that any sphere obtained by capping off an annulus A in the JSJ decomposition of (X, P) bounds a ball $B \subseteq M$ such that if the components of ∂A are in distinct components of $\partial N(V)$, then B meets $\Gamma - N(V)$ in a single edge, and if the components of ∂A are in the same component of $\partial N(V)$, then B is disjoint from $\Gamma - N(V)$.

We now apply the Characteristic Decomposition Theorem for Pared Manifolds [9, 11] to the pared manifold (X, P). Since X is atoroidal, this gives us a characteristic family σ of incompressible annuli in X with boundaries in $\partial X - P$ such that if W is the closure of any component of $X - \sigma$, then the pared manifold $(W, W \cap (P \cup \sigma))$ is either simple, Seifert fibered, or I- fibered. Once again, since the characteristic family σ is unique up to isotopy, we can modify g by an isotopy (and again by an abuse of notation, still refer to the map as g) so that $g(\sigma) = \sigma$.

Let A be an annulus component of $P \cup \sigma$, and let S denote the sphere obtained by capping off A by a pair of disjoint disks D_1 and D_2 in $\partial N(v_1)$ and $\partial N(v_2)$, where v_1 and v_2 may or may not be distinct vertices. Suppose that each component of $M - S$ intersects more than one edge of $\Gamma - N(V)$. Then by removing the vertices v_1 and v_2 and the edges that contain them we would obtain two non-empty subgraphs (see Figure 3). But this contradicts our hypothesis that Γ is 3-connected. Thus one of the components of $M - S$ meets $\Gamma - N(V)$ in at most one edge of Γ.

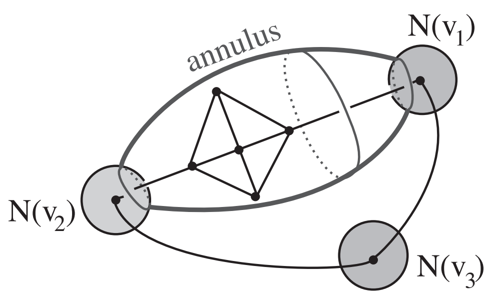

Figure 3: In this figure, there is more than one edge on each side of this capped off annulus, and hence Γ is not 3-connected.

Now, since M is irreducible, one of the closed up components of $M - S$ is a ball B. However, we showed at the beginning of our proof that neither Γ nor Γ with an edge removed can be contained in a ball in M. Thus B must be the closed up component of $M - S$ intersecting at most one edge of $\Gamma - N(V)$. Furthermore, since the annulus A is incompressible in X, if

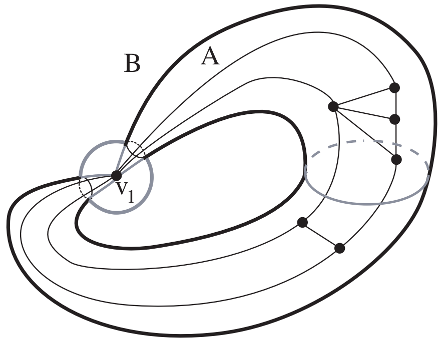

Figure 4: The component B of $M - S$ is disjoint from $\Gamma - N(V)$.

$v_1 \neq v_2$ then there is some edge e with vertices v_1 and v_2 such that B contains $\mathrm{cl}(e - (N(v_1) \cup N(v_2)))$. On the other hand, suppose that $v_1 = v_2$. Then by our definition of a graph, every edge in Γ has two distinct vertices. Thus the component of $M - S$ which meets $\Gamma - N(V)$ in at most one edge, is actually disjoint from $\Gamma - N(V)$. In this case, B must be disjoint from $\Gamma - N(V)$ as illustrated in Figure 4.

Note that since a ball cannot contain an incompressible torus, no torus boundary component of X can occur inside of B. It follows that every torus boundary component of X must also be a boundary component of $X - B$.

Step 4: We define a collection of balls $U_{e_1}, \ldots, U_{e_n}$, $V_{F_1}, \ldots, V_{F_m}$ in M such that the manifold $W = \mathrm{cl}(X - (U_{e_1} \cup \cdots \cup U_{e_n} \cup V_{F_1} \cup \cdots \cup V_{F_m}))$ is the closure of a single component of $X - (\sigma \cup P)$.

Let A be an annulus in $P \cup \sigma$ with one boundary in $\partial N(v_1)$ and the other boundary in $\partial N(v_2)$ where $v_1 \neq v_2$. By capping off A with disks in $\partial N(v_1)$ and $\partial N(v_2)$ we obtain a sphere which (as we saw in Step 3) bounds a ball B that meets Γ in $\mathrm{cl}(e - (N(v_1) \cup N(v_2)))$ for some edge e.

Note that by our definition of a graph, e is the only edge of Γ whose vertices are v_1 and v_2. Thus we can let $\mathcal{C}_e$ denote the collection of all annuli in $P \cup \sigma$ with one boundary in $\partial N(v_1)$ and the other boundary in $\partial N(v_2)$. Now we cap off the annuli in $\mathcal{C}_e$ with disks in $\partial N(v_1)$ and $\partial N(v_2)$ to obtain a collection of spheres which bound nested balls containing $\mathrm{cl}(e - (N(v_1) \cup N(v_2)))$. Let A_e denote the annulus in $\mathcal{C}_e$ which, when capped off in this way, is outermost with respect to the nested balls.

Observe that the boundaries of A_e bound disks $D_1 \subseteq \partial N(v_1)$ and $D_2 \subseteq \partial N(v_2)$ which each meet Γ in a single point of e, and the sphere $A_e \cup D_1 \cup D_2$

bounds a ball U_e which contains $\mathrm{cl}(e - (N(v_1) \cup N(v_2)))$ and all of the other annuli in $\mathcal{C}_e$. Furthermore, as we saw at the end of Step 3, no torus boundary component of X can be contained in U_e.

We see as follows that $X - U_e$ has a single component. If $A_e = \partial' N(e)$, then $U_e = N(e)$ and hence $\mathrm{cl}(X - U_e) = X$. Thus we assume that $A_e \neq \partial' N(e)$. In this case, $U_e \cap X$ is obtained by removing the interior of $N(e)$ from the ball U_e. Thus $\partial(X - U_e) \cap U_e = A_e$ is an annulus. Now let p and q be points in $X - U_e$. Since X is path connected, there are paths P and Q in $X - U_e$ from p and q respectively to $\partial(X - U_e) \cap U_e$. Now since $\partial(X - U_e) \cap U_e$ is an annulus, there is a path in $\partial(X - U_e) \cap U_e$ joining P and Q. Thus there is a path from p to q in $X - U_e$. We will use this at the end of this step.

We repeat the above construction for each annulus in $P \cup \sigma$ with boundaries in distinct components of $\partial N(V)$. This gives us a collection of pairwise disjoint balls $U_{e_1}, \dots, U_{e_n}$. Observe that for every edge e_j, the annulus $\partial' N(e_j)$ is in P and its boundaries are in distinct components of $\partial N(V)$. Thus every edge e_j in Γ is contained in some U_{e_j}. It follows that $U_{e_1} \cup \cdots \cup U_{e_n}$ contains $\mathrm{cl}(\Gamma - N(V))$ and contains every annulus in $P \cup \sigma$ with boundaries in distinct components of $\partial N(V)$.

Next we consider an annulus F in σ both of whose boundaries are in a single $\partial N(v)$. We saw in Step 3 that if we cap off F by a pair of disjoint disks in $N(v)$, we obtain a sphere which bounds a ball that is disjoint from $\Gamma - N(v)$. Thus the components of ∂F are parallel in $\partial' N(v)$. We now cap off every such annulus to obtain a collection of spheres which bound balls that are either nested or disjoint. Thus we can choose an annulus $F \in \sigma$ which, when capped off with disks in $N(v)$, bounds an outermost ball V_F with respect to this nesting (see Figure 5).

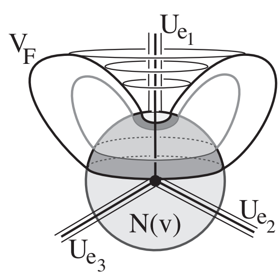

Figure 5: V_F is an outermost ball.

As we saw above for U_e, the ball V_F cannot contain any torus boundary components of X. In this case, $V_F \cap X$ is obtained by removing the interior of $V_F \cap N(v)$ from the ball V_F. Thus $\partial(X - V_F) \cap V_F = F$ is an annulus. Now by an argument analogous to the argument for $X - U_e$ we see that for any pair of points p and q in $X - V_F$ there is a path from p to q in $X - V_F$. Thus $X - V_F$ also has a single component.

Now for each annulus $F \in \sigma$ with both components of ∂F in a single $\partial N(v)$ such that F is not contained in one of the balls U_{e_j}, we use the above argument to define a ball V_F. In this way, we obtain a collection of pairwise disjoint balls $U_{e_1}, \ldots, U_{e_n}$, $V_{F_1}, \ldots, V_{F_m}$. Furthermore, each of the sets $X - U_{e_j}$ and $X - V_{F_i}$ has a single component, and each $\partial U_{e_j} - \partial X$ and $\partial V_{F_i} - \partial X$ is an annulus in $\sigma \cup P$. Now it follows that the manifold

$$W = \mathrm{cl}(X - (U_{e_1} \cup \cdots \cup U_{e_n} \cup V_{F_1} \cup \cdots \cup V_{F_m}))$$

is the closure of a single component of $X - (\sigma \cup P)$ (see Figure 6).

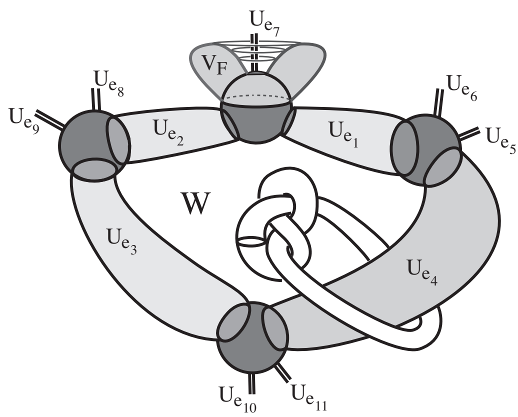

Figure 6: W is the closure of a single component of $X - \sigma$.

Step 5: We show that $g(W) = W$ and the pared manifold $(W, W \cap (P \cup \sigma))$ is simple.

Recall that g fixes each vertex and leaves each edge setwise invariant. Also, $g(N(\Gamma)) = N(\Gamma)$, $g(P) = P$, and $g(\sigma) = \sigma$. Now, since the sets of balls $\{U_{e_1}, \ldots, U_{e_n}\}$ and $\{V_{F_1}, \ldots, V_{F_m}\}$ were chosen to be outermost, each of these sets is also invariant under g. Furthermore, we know from Step 4 that each U_{e_j} intersects $\Gamma - N(V)$ only in e_j. Now since g leaves each e_j setwise invariant, g must also leave each U_{e_j} setwise invariant. It follows that each A_{e_j} must also be setwise invariant under g; and since each vertex is fixed by g, each boundary component of A_{e_j} is also setwise invariant under g.

Now let v be a vertex such that some F_i has both its boundary components in $\partial N(v)$. Since F_i is incompressible in X, the components of ∂F_i bound disjoint disks in $\partial N(v)$ which each intersect at least one edge of Γ. Since every vertex and edge of Γ is invariant under g, it follows that F_i and its boundary components are also setwise invariant under g. Thus V_{F_i} is setwise invariant under g. Finally, since all of the U_{e_j} and V_{F_i} are setwise invariant under g, we know that W must be setwise invariant under g as well.

To show that the pared manifold $(W, W \cap (P \cup \sigma))$ is simple, first recall that W is the closure of a single component of $X - \sigma$. Hence by JSJ for pared manifolds [9, 11], $(W, W \cap (P \cup \sigma))$ is either I-fibered, Seifert fibered, or simple as a pared manifold. We see as follows that $(W, W \cap (P \cup \sigma))$ cannot be Seifert fibered or I-fibered.

First observe that for every vertex v, there is some edge e such that the ball U_e meets $\partial N(v)$. It follows that ∂W meets every component of $\partial N(V)$. Furthermore, since every vertex v has valence at least three, $\partial' N(v)$ is a sphere with at least three holes. Also, each U_{e_j} contains at most one boundary component of $\partial' N(v)$, and each V_{F_i} contains no boundary components of $\partial' N(v)$. Note that if F_i has its boundaries in $\partial N(v)$, then V_{F_i} separates $\partial'(N(v))$ into two components.

This means that for each vertex v, $W \cap \partial N(v)$ is obtained from $\partial' N(v)$ by deleting some (possibly zero) number of annuli. Since $\partial' N(v)$ is a sphere with at least three holes, deleting a collection of annuli may create new components, but some component must still be a sphere with at least three holes. It follows that the component of ∂W meeting $\partial N(\Gamma)$ has genus more than one, and thus the pared manifold $(W, W \cap (P \cup \sigma))$ cannot be Seifert fibered.

Next, suppose for the sake of contradiction that the pared manifold $(W, W \cap (P \cup \sigma))$ is I-fibered. By definition of I-fibered for pared manifolds, this means that there is an I-bundle map of W over a base surface Y such that $W \cap (P \cup \sigma)$ is in the pre-image of ∂Y. It follows that Y must be homeomorphic to a component of $\partial' N(V) \cap W$, and hence must be a sphere with holes. In particular, the base surface Y is orientable.

Now since W is an orientable 3-manifold which is I-fibered over an orientable surface, it follows that W must actually be a product $Y \times I$. Thus $W \cap (P \cup \sigma) = \partial Y \times I$, and $Y_0 = Y \times \{0\}$ and $Y_1 = Y \times \{1\}$ are the only components of $\partial' N(V) \cap W$. However, since ∂W meets every component of $\partial N(V)$, this means that Γ contains at most two vertices. But this contradicts our hypothesis that Γ is 3-connected. Therefore, the pared $(W, W \cap (P \cup \sigma))$ is not I-fibered. Since it is also not Seifert fibered, it must be simple.

Step 6: We prove that $g|W$ is isotopic to an orientation reversing involution h of $(W, W \cap (P \cup \sigma))$.

Now it follows from Thurston's Hyperbolization Theorem for Pared Manifolds [18] applied to the simple pared manifold $(W, W \cap (P \cup \sigma))$ that $W - (W \cap (P \cup \sigma))$ admits a finite volume complete hyperbolic metric with totally geodesic boundary. Let D denote the double of $W - (W \cap (P \cup \sigma))$ along its boundary. Then D is a finite volume hyperbolic manifold, and $g|W$ can be doubled to obtain an orientation reversing homeomorphism of D (which we still call g) taking each copy of $W - (W \cap (P \cup \sigma))$ to itself. Now by Mostow's Rigidity Theorem [14] applied to D, the homeomorphism $g : D \to D$ is homotopic to an orientation reversing finite order isometry $h : D \to D$ that restricts to an isometry of $W - (W \cap (P \cup \sigma))$. By removing horocyclic neighborhoods of the cusps of $W - (W \cap (P \cup \sigma))$, we obtain a copy of the pair $(W, W \cap (P \cup \sigma))$ which is contained in $W - (W \cap (P \cup \sigma))$ and is setwise invariant under h. We abuse notation and now consider h to be an orientation reversing finite order isometry of $(W, W \cap (P \cup \sigma))$ instead of a copy of $(W, W \cap (P \cup \sigma))$. Furthermore, h induces isometries on the collection of tori and annuli in $W \cap (P \cup \sigma)$ with respect to a flat metric. Furthermore, the sets $\partial' N(V) \cap W$, $\partial' N(E) \cap W$, and $\tau \cap W$ are each setwise invariant under h. Finally, it follows from Waldhausen's Isotopy Theorem [20] that h is actually isotopic to $g|W$ by an isotopy leaving $W \cap (P \cup \sigma)$ setwise invariant.

Now, recall that the boundary components of W consist of tori in τ, and the union of spheres with holes in $\partial' N(V) \cap W$ together with annuli in $P \cup \sigma$. Recall from the first paragraph in Step 5 that g setwise fixes each annulus $A_{e_j} \subseteq \partial U_{e_j}$ with boundaries in distinct components of $\partial N(V)$, each annulus $F_i \subseteq \partial V_{F_i}$ with both boundaries in a single component of $\partial N(V)$, each component of ∂A_{e_j}, and each component of ∂F_i. Since h is isotopic to $g|W$ by an isotopy leaving $W \cap (P \cup \sigma)$ setwise invariant, h leaves invariant the same sets as g. It follows that for each vertex v, we have $h(\partial' N(v)) \cap W) = \partial' N(v) \cap W$, and h takes each component of $W \cap \partial N(v)$ to itself, leaving each boundary component setwise invariant.

Since h has finite order, h restricts to a finite order homeomorphism of every component of $W \cap \partial N(V)$. We saw in Step 5 that for every vertex v, at least one component C_v of $W \cap \partial N(v)$ is a sphere with at least three holes. Since h restricts to a finite order homeomorphism of C_v taking each boundary component of C_v to itself, h must be a reflection of C_v which also reflects each component of ∂C_v. Now h^2 is a finite order, orientation preserving isometry of W that pointwise fixes the surface C_v. It follows that h^2 is the identity, and hence h is an involution of W.

Step 7: We extend h to an orientation reversing involution of $X \cup N(\Gamma)$ which pointwise fixes an embedding Γ' of γ.

Observe that since every annulus in $P \cup \sigma$ is incompressible in W, no component of $W \cap \partial N(V)$ can be a disk. Thus every component of $W \cap \partial N(V)$ is a sphere with two or more holes.

As we saw in Step 6, for each vertex v, h reflects some component C_v of $W \cap \partial N(v)$ which is a sphere with at least three holes, and h reflects every component of ∂C_v. Let b_0 denote some boundary component of C_v. Then b_0 is also a boundary component of either an annulus A_{e_j} or an annulus F_i. Since h reflects b_0, we know that h must also reflect the annulus A_{e_j} or F_i, whichever contains b_0 in its boundary. Since the boundaries of the annulus are not interchanged, h must also reflect each boundary component of A_{e_j} or F_i. Below we extend h to U_{e_j} or V_{F_i}.

First we consider the case where b_0 is in the boundary of an annulus $A_{e_j} \subseteq \partial U_{e_j}$. Let D_j and D'_j denote the disks whose union is in $\mathrm{cl}(\partial U_{e_j} - A_{e_j})$. Then D_j and D'_j each meet Γ in a single point of e_j. Since h reflects the annulus A_{e_j} together with each boundary component of A_{e_j}, we can extend h radially to the disks D_j and D'_j to get a reflection of the sphere $A_{e_j} \cup D_j \cup D'_j$ pointwise fixing a circle containing the points $D_j \cap e_j$ and $D'_j \cap e_j$. Recall that the sphere $A_{e_j} \cup D_j \cup D'_j$ bounds the ball U_{e_j} in M. Now, we express U_{e_j} as a product $D_j \times I$ whose core $\overline{e_j}$ has endpoints $D_j \cap e_j$ and $D'_j \cap e_j$ (see Figure 7). Then we extend h from a reflection of the sphere $A_{e_j} \cup D_j \cup D'_j$ to a reflection of the product $D_j \times I$ which pointwise fixes the core $\overline{e_j}$.

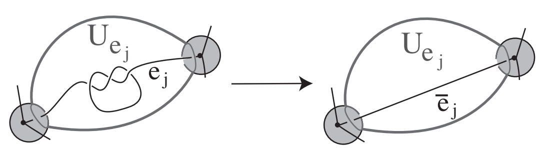

Figure 7: We can think of U_{e_j} as a product $D_j \times I$ with core $\overline{e_j}$.

Next we consider the case where b_0 is a boundary component of an annulus $F_i \subseteq \partial V_{F_i}$ which has both boundaries in a single $\partial N(v)$. Recall that $\mathrm{cl}(\partial V_{F_i} - F_i)$ consists of disks D_i and D'_i properly embedded in $N(v)$. Without loss of generality, D_i and D'_i each meet Γ at a single point on an edge. Since h reflects the annulus F_i together with each of its boundary components, we can extend h radially to the disks D_i and D'_i to get a reflection of the sphere $F_i \cup D_i \cup D'_i$ pointwise fixing a circle containing the points $\Gamma \cap D_i$ and $\Gamma \cap D'_i$. Thus we can extend h to a reflection of the ball V_{F_i} which pointwise fixes a disk containing the segment $V_{F_i} \cap \Gamma$ (see Figure 8).

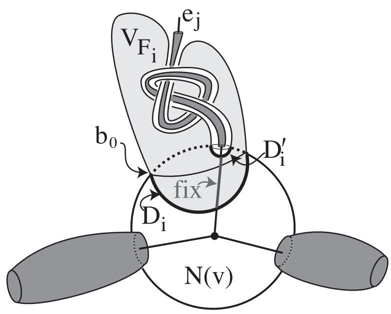

Figure 8: We extend h to a reflection of the ball V_{F_i} which pointwise fixes $\Gamma \cap V_{F_i}$.

In either case, the extension of h reflects the sphere with holes C_v, and one of the balls U_{e_j} or V_{F_i} depending on whether b_0 is a boundary component of A_{e_j} or F_i, respectively. Next we let S_1 denote the union of C_v together with the annulus A_{e_j} or F_i glued along b_0. Now h reflects S_1 taking every boundary component of S_1 to itself, and hence reflecting every boundary component of S_1. Let b_1 be a boundary component of S_1. If b_1 is not the other boundary of the annulus A_{e_j} or F_i, then we repeat the above argument with b_1 in place of b_0.

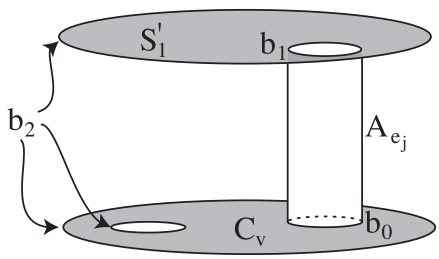

Figure 9: In this illustration, we have three choices for a boundary component b_2 of $S_2 = C_v \cup A_{e_j} \cup S'_1$.

If b_1 is the other boundary of the annulus A_{e_j} or F_i, then b_1 is also a boundary of some other component S'_1 of $W \cap \partial N(V)$, as illustrated in Figure 9. In this case, since b_1 is reflected by h and every boundary component of S'_1 is setwise invariant under h, we know that h must reflect $S_2 = S_1 \cup S'_1$. Now let b_2 denote a boundary component of S_2, and repeat the above argument with b_2 in place of b_0.

In general, for a given surface S_n obtained in this way, the surface S_{n+1} is the union of S_n together with either an annulus of the form A_{e_j} or F_i, or a sphere with at least two holes contained in $W \cap \partial N(V)$. Furthermore, S_{n+1} is reflected by h. This process will only stop when the surface that we obtain has no boundary components. Since ∂W has only one component which intersects $\partial N(V)$, the closed surface that we obtain in this way must be this component. Thus we have extended h to an orientation reversing involution of each of the balls $U_{e_1}, \dots, U_{e_n}$, $V_{F_1}, \dots, V_{F_m}$.

Since $N(E) \cup (X - W) \subseteq U_{e_1} \cup \dots U_{e_n}$, at this point we have defined h on X, and the only part of $N(\Gamma)$ on which we have not defined h is $N = N(V) - (V_{F_1} \cup \dots \cup V_{F_m})$. Then N is a collection of disjoint balls (for example in Figure 8, $N(v) - V_{F_i}$ is two balls one of which contains v). Also, h is a reflection of each component of ∂N that fixes each point in $\partial N \cap \Gamma$. Now we extend h radially to a reflection of each ball of N in such a way that h pointwise fixes each component of $N \cap \Gamma$. Thus h is a reflection of each component of $N(V)$ which pointwise fixes $N(V) \cap \Gamma$.

We have now extended h to an orientation reversing involution of the manifold

$$Y = W \cup V_{F_1} \cup \dots \cup V_{F_m} \cup U_{e_1} \cup \dots \cup U_{e_n} \cup N.$$

Recall from the end of Step 4 that ∂X and ∂W have the same collection of tori in their boundary components. Furthermore, we have filled in the boundary component of W meeting $\partial N(\Gamma)$ with a collection of balls in $X \cup N(\Gamma)$. Thus in fact $Y = X \cup N(\Gamma)$ and ∂Y is the collection of tori in ∂X.

Finally, we define a new embedding Γ' of γ in $X \cup N(\Gamma)$ as follows. Let $\Gamma' \cap N(V) = \Gamma \cap N(V)$. Then for each edge e_j define an embedding of $e_j - N(V)$ in Γ' as the core $\overline{e_j}$ of $U_{e_j} = D_j \times I$, which we know is pointwise fixed by h according to the way we extended h to U_{e_j} (recall Figure 7).

Step 8: We prove that if an essential curve in a component of $\partial(X \cup N(\Gamma))$ compresses in M, then it compresses in $X \cup N(\Gamma)$.

Let $T_1, \dots, T_r$ be the components of $\partial(X \cup N(\Gamma))$. Then each T_i is contained in the characteristic family τ, and hence is incompressible in $\mathrm{cl}(M - N(\Gamma))$.

Suppose that an essential curve λ_i on some T_i compresses in M. Let D_i be a compressing disk for λ_i whose intersection with the set of tori $\{T_1, \dots, T_r\}$ is minimal. Let $D = D_i$ if the interior of D_i is disjoint from T_i. Otherwise, there exists some D in the interior of D_i such that D is a compressing disk for T_i whose interior is disjoint from T_i. In either case, the intersection of D with $\{T_1, \dots, T_r\}$ is minimal.

Suppose that D contains at least one curve of intersection in its interior. Hence there is an innermost disk Δ on D which is a compressing disk for some T_j with $j \neq i$. Since T_j compresses in M but is incompressible in $\mathrm{cl}(M - N(\Gamma))$, we know that Δ intersects Γ.

Since M is irreducible, any compressible torus is separating in M. Thus we can let X_j denote the closed up component of $M - T_j$ containing X and let V_j denote the closed up component of $M - T_j$ whose interior is disjoint from X. Now let S denote the region of D which is adjacent to the innermost disk Δ. Then $S \subseteq V_j$, since $\Delta \subseteq X_j$. Also, $\partial D \subseteq T_i \subseteq X \subseteq X_j$ implies that S is adjacent to another region of D which is contained in X_j. In particular, there must be another circle of intersection α of $D \cap T_j$ which bounds a disk $\overline{D} \subseteq D$ such that $\overline{D}$ contains $\Delta \cup S$. We illustrate the abstract disk D and its intersections with T_j in Figure 10. The white regions in the figure are contained in V_j, and the grey regions are contained in X_j. Note we do not illustrate any circles of intersection of D with any T_k with $k \neq j$.

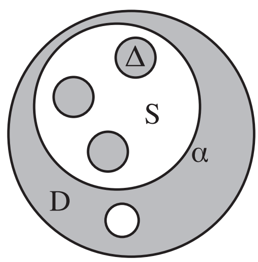

Figure 10: A picture of the abstract disk D and its circles of intersection with T_j.

Now, since the intersection of D with the tori $T_1, \ldots, T_r$ is minimal, all of the curves of intersection of $D \cap T_j$ must be essential on T_j. In particular, $\partial\Delta$ and α must both be essential on T_j. Since there cannot be two essential, disjoint, non-parallel curves on a torus, this means that α is parallel to $\partial\Delta$ on T_j. It follows that α must bound a disk $\overline{\Delta}$ which is parallel to Δ in M. In particular, since the interior of Δ is disjoint from $T_1, \ldots, T_r$, the interior of $\overline{\Delta}$ is as well. But now by replacing the disk $\overline{D}$ with the disk $\overline{\Delta}$ in the compressing disk D we obtain a new compressing disk D' for T_i which has fewer curves of intersection with $T_1, \ldots, T_r$ than D has. From this contradiction we conclude that the interior of D must be disjoint from $T_1 \cup \cdots \cup T_r$, and hence $D \subseteq X \cup N(\Gamma)$.

If $\partial D = \lambda_i$, then λ_i compresses in $X \cup N(\Gamma)$ as required. Otherwise, the compression disk D was contained in the interior of the original disk D_i

and $\partial D \subseteq T_i$. In this case, since the intersection of D_i with the set of tori $\{T_1, \dots, T_r\}$ was minimal, ∂D is essential in T_i. But now ∂D and λ_i are disjoint essential curves on T_i. Hence as we saw above, the disks D and D_i must be parallel in M. Now, since $D \subseteq X \cup N(\Gamma)$, it follows that λ_i must compress in $X \cup N(\Gamma)$ as well.

Step 9: We fill each component T_i of $\partial(X \cup N(\Gamma))$ with a solid torus such that h extends to an involution of the resulting manifold M', and M' satisfies the condition below.

Condition. *If T_i is compressible in M, then every element of $H_1(T_i, \mathbb{Z}_2)$ is trivial in $H_1(M', \mathbb{Z}_2)$, and if T_i is incompressible in M then at least one non-trivial element of $H_1(T_i, \mathbb{Z}_2)$ is trivial in $H_1(M', \mathbb{Z}_2)$.*

Let T_i be a component of $\partial(X \cup N(\Gamma))$. By Corollary 1 at the beginning of Section 4 there is a curve μ_i on T_i which is non-trivial in $H_1(X \cup N(\Gamma), \mathbb{Z}_2)$. Also, we know from Step 8 that if some essential curve λ_i on T_i compresses in M, then λ_i also compresses in $X \cup N(\Gamma)$. In particular, λ_i is not homologous in T_i to μ_i.

Now suppose that for some $j \neq i$, the involution h interchanges T_i and T_j. Since $h : X \cup N(\Gamma) \to X \cup N(\Gamma)$ is a homeomorphism and μ_i is a curve on T_i which is non-trivial in $H_1(X \cup N(\Gamma), \mathbb{Z}_2)$, we know that $h(\mu_i)$ is a curve on T_j which is also non-trivial in $H_1(X \cup N(\Gamma), \mathbb{Z}_2)$. Now we fill $X \cup N(\Gamma)$ along T_i by adding a solid torus V_i with its meridian attached to the non-trivial curve μ_i, and we fill along T_j by adding a solid torus V_j with its meridian attached to $h(\mu_i)$. Then we extend the involution h radially in $V_i \cup V_j$ (abusing notation and still calling the involution h). We repeat this process for every boundary component of $X \cup N(\Gamma)$ which is not setwise invariant under h.

Thus for every T_i along which we have glued a solid torus V_i, the curve μ_i on T_i is now trivial in $H_1(X \cup N(\Gamma) \cup V_i, \mathbb{Z}_2)$. Furthermore, if T_i is compressible in M, then there is an essential curve λ_i on T_i which compresses in $X \cup N(\Gamma)$. Hence, λ_i is not homologous to μ_i in $H_1(T_i, \mathbb{Z}_2)$, and together λ_i and μ_i generate $H_1(T_i, \mathbb{Z}_2)$. Furthermore, both λ_i and μ_i are trivial in $H_1(X \cup N(\Gamma) \cup V_i, \mathbb{Z}_2)$.

Let Z be the manifold that we have obtained by filling all of the boundary components of $X \cup N(\Gamma)$ which are not setwise fixed by h, and let T_i be a component of ∂Z. Recall from Step 7 that Γ' is an embedding of γ in $X \cup N(\Gamma)$ which is pointwise fixed by h. Now $h : Z \to Z$ is an orientation reversing involution pointwise fixing Γ', and T_i is setwise invariant under h. Furthermore, by the proof of Step 6, $h|T_i$ is an isometry with respect to a flat metric. Since h is an orientation reversing involution of Z and T_i is a boundary component of Z, $h|T_i$ is also an orientation reversing involution. Thus $h|T_i$ is either a reflection pointwise fixing two parallel circles on T_i, or a composition of a reflection and an order 2 rotation of T_i. In either case, for any non-trivial element $a_i \in H_1(T_i, \mathbb{Z})$, there is a non-trivial curve $b_i \in H_1(T_i, \mathbb{Z})$ such that a_i and b_i meet transversely in a single point and $h(b_i)$ is homologous to $\pm b_i$ in $H_1(T_i, \mathbb{Z})$.

Now suppose that some essential curve λ_i on T_i compresses in M. Then by Step 8, λ_i also compresses in $X \cup N(\Gamma)$, and hence in Z. Pick a non-trivial curve b_i on T_i that intersects λ_i in a single point such that $h(b_i)$ is homologous to $\pm b_i$ in $H_1(T_i, \mathbb{Z})$. Then $\langle \lambda_i, b_i \rangle = H_1(T_i, \mathbb{Z})$. Also, since λ_i is null homologous in $H_1(Z, \mathbb{Z}_2)$, by Corollary 1, b_i is non-trivial in $H_1(Z, \mathbb{Z}_2)$. Now we fill Z along T_i by adding a solid torus V_i with its meridian attached to the non-trivial curve b_i. Since $h(b_i)$ is homologous to $\pm b_i$ in $H_1(T_i, \mathbb{Z})$, we can extend h radially to an orientation reversing involution of the solid torus V_i. Then $h : Z \cup V_i \to Z \cup V_i$ is an orientation reversing involution, and both λ_i and b_i are trivial in $H_1(Z \cup V_i, \mathbb{Z}_2)$.

On the other hand, suppose that some T_i is incompressible in M. By Corollary 1, there is some curve b_i on T_i which is non-trivial in $H_1(Z, \mathbb{Z}_2)$, and b_i is homologous to $\pm h(b_i)$ in $H_1(T_i, \mathbb{Z})$. Now we fill T_i by adding a solid torus V_i with its meridian attached to the curve b_i, and then extend h to V_i. Then $h : Z \cup V_i \to Z \cup V_i$ is again an orientation reversing involution, and b_i is trivial in $H_1(Z \cup V_i, \mathbb{Z}_2)$.

In this way, we glue a solid torus to each of the T_i in $\partial(X \cup N(\Gamma))$ to obtain a closed manifold M' satisfying the required condition. Since $\Gamma' \subseteq X \cup N(\Gamma)$, this gives us an embedding Γ' of γ in M'. Furthermore, we have extended h to an orientation reversing involution of (M', Γ') which pointwise fixes Γ'.

To prove the proposition it remains to show that $\dim_{\mathbb{Z}_2}(H_1(M', \mathbb{Z}_2)) \leq n_M$. We do this in steps 10 and 11.

Step 10: We prove that at most N_M of the tori in $\partial(X \cup N(\Gamma))$ are incompressible in M.

Recall that N_M is either the number of non-parallel disjoint incompressible tori in M or 2, whichever is larger. If no pair of distinct tori of $\partial(X \cup N(\Gamma))$ are parallel in M, then at most N_M of the components of $\partial(X \cup N(\Gamma))$ are incompressible in M.

Thus we assume that some pair of distinct tori T_i and T_j of $\partial(X \cup N(\Gamma))$ are parallel in M. Then T_i and T_j co-bound a region R in M which is homeomorphic to a product of a torus and an interval. However, since T_i and T_j are in the characteristic family for $\mathrm{cl}(M - N(\Gamma))$, they cannot be parallel in $\mathrm{cl}(M - N(\Gamma))$. Thus R intersects Γ. But since $\partial R = T_i \cup T_j$ and Γ is disjoint from $T_i \cup T_j$, this implies that $\Gamma \subseteq R$. It now follows that $X \cup N(\Gamma) \subseteq R$.

Now suppose that some component T_k of $\partial(X \cup N(\Gamma))$ is incompressible in M and distinct from T_i and T_j. Then $T_k \subseteq R$. But since $R \cong T_j \times I$, either T_k is parallel in R to both T_i and T_j, or T_k is compressible in R. However, since $R \subseteq M$, the latter would imply that T_k is compressible in M.

Thus T_k must be parallel to both T_i and T_j in $R \subseteq M$. Since T_k does not intersect Γ, this means that T_k is parallel to one of T_i or T_j in $\mathrm{cl}(M - N(\Gamma))$. But this is impossible since T_i, T_j, and T_k are all in the characteristic family for $\mathrm{cl}(M - N(\Gamma))$. Thus the only components of $\partial(X \cup N(\Gamma))$ which could be incompressible in M are T_i and T_j. Hence $\partial(X \cup N(\Gamma))$ has at most $2 \leq N_M$ components which are incompressible in M.

Step 11: We prove the inequality $\dim_{\mathbb{Z}_2}(H_1(M', \mathbb{Z}_2)) \leq n_M$.

Recall that M' is obtained from $X \cup N(\Gamma)$ by adding a collection of solid tori $V_1, \ldots, V_r$ along the components $T_1, \ldots, T_r$ of $\partial(X \cup N(\Gamma))$. Also, by the condition in Step 9, we know that if T_i is compressible in M, then every element of $H_1(T_i, \mathbb{Z}_2)$ is trivial in $H_1(M', \mathbb{Z}_2)$; and if T_i is incompressible in M then there exists some curve β_i on T_i such that β_i is non-trivial in $H_1(T_i, \mathbb{Z}_2)$ but trivial in $H_1(M', \mathbb{Z}_2)$. It follows that for every T_i which is incompressible in M, there is at most one element of $H_1(T_i, \mathbb{Z}_2)$ which is non-trivial in $H_1(M', \mathbb{Z}_2)$. Let α_i be a representative of this homology class in $H_1(T_i, \mathbb{Z}_2)$. Thus every curve in $H_1(T_i, \mathbb{Z}_2)$ is either trivial in $H_1(M', \mathbb{Z}_2)$ or homologous to α_i in $H_1(M', \mathbb{Z}_2)$.

Now let α be a curve in M' which is non-trivial in $H_1(M', \mathbb{Z}_2)$, and for each solid torus V_i, let C_i denote its core. Then by general position, we can choose α to be disjoint from $C_1 \cup \cdots \cup C_r$, and hence disjoint from $V_1 \cup \cdots \cup V_r$. Now since $\alpha \subseteq X \cup N(\Gamma) \subseteq M$, we can also consider α in $H_1(M, \mathbb{Z}_2)$.

Suppose that α is trivial in $H_1(M, \mathbb{Z}_2)$. It follows that α is homologous in $X \cup N(\Gamma)$ to a collection of curves on $T_1 \cup \cdots \cup T_r$ which are trivial in $H_1(M, \mathbb{Z}_2)$ but non-trivial in $H_1(M', \mathbb{Z}_2)$. Since for any T_i which is compressible in M every element of $H_1(T_i, \mathbb{Z}_2)$ is trivial in $H_1(M', \mathbb{Z}_2)$, α must be homologous in $H_1(M', \mathbb{Z}_2)$ to a collection of curves on only those T_i which are incompressible in M. It now follows that α is homologous in $H_1(M', \mathbb{Z}_2)$ to a sum of α_i's on T_i's that are incompressible in M. But by Step 10, there are at most N_M such tori T_i. Hence there are at most N_M distinct α_i which are non-trivial in $H_1(M', \mathbb{Z}_2)$. It follows that these N_M curves generate every non-trivial element of $H_1(M', \mathbb{Z}_2)$ which is trivial in $H_1(M, \mathbb{Z}_2)$. This gives us the required inequality:

$$\dim_{\mathbb{Z}_2}(H_1(M', \mathbb{Z}_2)) \leq \dim_{\mathbb{Z}_2}(H_1(M, \mathbb{Z}_2)) + N_M = n_M.$$

Hence the proposition follows. □

References

[1] G. Bredon, Introduction to Compact Transformation Groups, Academic Press (1972).

[2] E. Flapan, *Rigidity of graph symmetries in the* 3*-sphere*, Journal of Knot Theory and Its Ramifications **4** (1995), 373–388.

[3] A. Hatcher, *The classification of* 3*-manifolds — a brief overview*, `http://www.math.cornell.edu/~hatcher/3M/3Mdownloads.html`.

[4] A. Hatcher, *Notes on basic* 3*-manifold topology*, `http://www.math.cornell.edu/~hatcher/3M/3Mdownloads.html`.

[5] A. Hatcher, P. Lochak, and L. Schneps, *On the Teichmüller tower of mapping class groups*, J. Reine Angew. Math. **521** (2000), 1–24.

[6] H. Howards, *Generating disjoint incompressible surfaces*, Topology and Its Applications **58** (2011), 325–343.

[7] H. Howards, *Surfaces and* 3*-manifolds*, Dissertation University of California San Diego, Mathematics (1997).

[8] T. Ikeda, *Rigidly achiral hyperbolic spatial graphs in* 3*-manifolds*, Journal of Knot Theory and Its Ramifications **22** (2013), 12 pages.

[9] W. Jaco and P. Shalen, *Seifert fibred spaces in* 3*-manifolds*, Memoirs Amer. Math. Soc. **220**, Amer. Math. Soc., Providence (1979).

[10] B. Jiang and S. Wang, *Achirality and planarity*, Communications in Contemporary Mathematics **2** (2000), 299–305.

[11] K. Johannson, Homotopy Equivalences of 3-Manifolds with Boundaries, Lecture Notes in Mathematics **761**, Springer-Verlag, New York, Berlin, Heidelberg (1979).

[12] Helmuth Kneser, *Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten*, Jahresbericht der Deutschen Mathematiker-Vereinigung **38** (1929), 248–259.

[13] M. Kobayashi, *Fixed point sets of orientation reversing involutions on* 3*-manifolds*, Osaka J. Math. **25** (1988), 877–879.

[14] G. Mostow, Strong Rigidity of Locally Symmetric Spaces, Annals of Mathematics Studies **78**, Princeton University Press, Princeton, NJ, (1973).

[15] R. Qiu and S. Wang, *Simple, small knots in handlebodies*, Topology Appl. **144** (2004), 211–227.

[16] P. A. Smith, *Transformations of finite period. II*, Ann. of Math. **40** (1939), 690–711.

[17] Richard Stong, *Diskbusting elements of the free group*, Mathematical Research Letters **4** (1997), 201–210.

[18] W. Thurston, *Three-dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. **6** (1982), 357–381.

[19] F. Waldhausen, *Eine Klasse von* 3*-dimensionaler Mannigfaltigkeiten I*, Inventiones **3** (1967), 308–333.

[20] F. Waldhausen, *On irreducible* 3*-manifolds which are sufficiently large*, Ann. Math. **87** (1968), 56–88.

Department of Mathematics, Pomona College
Claremont, CA 91711, USA
E-mail address: ELF04747@pomona.edu

Department of Mathematics, Wake Forest University
Winston-Salem, NC 27109, USA
E-mail address: howards@wfu.edu

Received December 14, 2015