AUTOMORPHISMS OF THE SPHERE COMPLEX OF AN INFINITE GRAPH

THOMAS HILL, MICHAEL C. KOPRESKI, REBECCA RECHKIN, GEORGE SHAJI, AND BRIAN UDALL

ABSTRACT. For a locally finite, connected graph Γ, let $\mathrm{Map}(\Gamma)$ denote the group of proper homotopy equivalences of Γ up to proper homotopy. Excluding sporadic cases, we show $\mathrm{Aut}(\mathcal{S}(M_\Gamma)) \cong \mathrm{Map}(\Gamma)$, where $\mathcal{S}(M_\Gamma)$ is the sphere complex of the doubled handlebody M_Γ associated to Γ. We also construct an exhaustion of $\mathcal{S}(M_\Gamma)$ by finite strongly rigid sets when Γ has finite rank and finitely many rays, and an appropriate generalization otherwise.

1. Introduction

Let Γ be a locally finite, connected graph, and let $\mathrm{Map}(\Gamma)$ denote its **mapping class group**, defined by Algom-Kfir–Bestvina to be the group of proper homotopy equivalences of Γ up to proper homotopy. Let M_Γ denote the doubled handlebody associated to Γ (see Definition 2.3). We prove in that the sphere complex $\mathcal{S}(M_\Gamma)$ (see Definition 2.4) satisfies an Ivanov-type theorem.

Theorem 1.1. *Let Γ, Γ' be two locally finite connected graphs. Suppose $f : \mathcal{S}(M_\Gamma) \to \mathcal{S}(M_{\Gamma'})$ is an isomorphism. Then f is induced by a diffeomorphism $h : M_\Gamma \to M_{\Gamma'}$. In particular, (i) Γ and Γ' are proper homotopy equivalent and (ii) when Γ is not a graph of rank r with s rays such that $2r + s < 4$ or $(r, s) \in \{(0, 4), (2, 0)\}$, $\mathrm{Aut}(\mathcal{S}(M_\Gamma)) \cong \mathrm{Map}(\Gamma)$ as topological groups.*

Our proof adapts an observation from Bavard–Dowdall–Rafi [BDR20] that the link of a sphere system $\sigma \subset \mathcal{S}(M_\Gamma)$ is isomorphic to the join of the $\mathcal{S}(M_i)$ for components $M_i \subset M_\Gamma \setminus \sigma$. For a maximal sphere system σ, we obtain an isomorphism of dual graphs $\Delta_\sigma \to \Delta_{f\sigma}$ from which we construct the diffeomorphism h.

In addition, we generalize the results of Bering–Leininger [BL24a] to $\mathcal{S}(M_\Gamma)$.

Theorem 1.2. *Suppose Γ has rank r and s rays such that $6 \le 2r + s < \infty$. Then there exists an exhaustion of $\mathcal{S}(M_\Gamma)$ by finite strongly rigid subcomplexes.*

From Theorem 1.2 and techniques from [BDR20], we obtain another proof of Theorem 1.1. Finally, we extend Theorem 1.2 to the infinite-type setting using techniques from the first proof of Theorem 1.1. We say a graph Γ is of **infinite-type** if it is not proper homotopy equivalent to a finite rank graph with finitely many ends.

Theorem 1.3. *For Γ infinite-type, there exists an exhaustion of $\mathcal{S}(M_\Gamma)$ by topologically locally finite subcomplexes that are strongly rigid over maximal maps.*

A subcomplex $X \subset \mathcal{S}(M_\Gamma)$ is **strongly rigid (*resp.* over maximal maps)** if any locally injective simplicial map $X \to \mathcal{S}(M_{n,s})$ (*resp.* preserving the maximality of sphere systems) extends uniquely to an automorphism of $\mathcal{S}(M_\Gamma)$. The subcomplex X is **topologically locally finite** if every compact set $K \subset M_\Gamma$ intersects finitely many vertices of X.

1.1. **Motivation.** Recall that $\mathrm{Out}(\mathbb{F}_n)$ is the group of outer automorphisms of the free group $\mathbb{F}_n$, defined as $\mathrm{Out}(\mathbb{F}_n) := \mathrm{Aut}(\mathbb{F}_n)/\,\mathrm{Inn}(\mathbb{F}_n)$. From a topological perspective, $\mathrm{Out}(\mathbb{F}_n)$ can be thought of as

$$\mathrm{Out}(\mathbb{F}_n) = \{\, \text{homotopy equivalences } \Gamma \to \Gamma \,\} \,/\, \text{homotopy}$$

where Γ is a finite graph of rank n (see [Hat02, Proposition 1B.9]). There is a rich dictionary between mapping class groups of surfaces, $\mathrm{Map}(S)$, and $\mathrm{Out}(\mathbb{F}_n)$ (see [BD19]). Analogies between the two include:

$$\begin{array}{rcl} \mathrm{Map}(S) & \longleftrightarrow & \mathrm{Out}(\mathbb{F}_n) \\ \mathrm{Teich}(S) & \longleftrightarrow & \text{Outer space, } \mathrm{CV}_n \\ \text{curve complex} & \longleftrightarrow & \begin{cases} \text{sphere complex} \\ \text{free factor complex} \\ \text{free splitting complex} \end{cases} \end{array}$$

Within the last decade, there has been a surge of interest in **big** mapping class groups of surfaces, that is, $\mathrm{Map}(S)$ for a surface S whose fundamental group is not finitely generated. Motivated by the parallels between graphs and surfaces, a natural question is: what is the big version of $\mathrm{Out}(\mathbb{F}_n)$? Generalizing the interpretation of $\mathrm{Out}(\mathbb{F}_n)$ as homotopy equivalences of a graph up to homotopy, [AB21] propose the definition of the mapping class group of a locally finite, infinite graph Γ to be:

$$\mathrm{Map}(\Gamma) := \{\text{proper homotopy equivalences } \Gamma \to \Gamma\}/\text{proper homotopy}.$$

The group $\mathrm{Map}(\Gamma)$ is sometimes referred to as "big $\mathrm{Out}(\mathbb{F}_n)$." When Γ is a finite graph of rank n, these definitions coincide: $\mathrm{Out}(\mathbb{F}_n) \cong \mathrm{Map}(\Gamma)$. [AB21, DHK23a, DHK23b, Uda24] have demonstrated that $\mathrm{Map}(\Gamma)$ exhibits many similarities with big mapping class groups of surfaces.

The curve complex $\mathcal{C}(S)$ is one of the most important tools for studying mapping class groups of surfaces. A celebrated theorem of Ivanov states that $\mathrm{Aut}(\mathcal{C}(S)) \cong \mathrm{Map}^{\pm}(S)$, illustrating an underlying connection between the curve complex and the surface mapping class group [Iva97, Luo99]. Ivanov's theorem has inspired a broader metaconjecture: "every object naturally associated to a surface S and possessing a sufficiently rich structure has $\mathrm{Map}^{\pm}(S)$ as its group of automorphisms" [Iva06]. There have been many subsequent results supporting this metaconjecture (see [BM19] for further discussion).

When S is an infinite-type surface, the curve complex is geometrically uninteresting (it has diameter 2); however, it is sufficiently rich *combinatorially* that Ivanov's theorem still holds [HMV18, BDR20]. The aim of this project is to introduce a complex that plays a similar role in the setting of big $\mathrm{Out}(\mathbb{F}_n)$. Specifically, we will show that the sphere complex of a 3-manifold associated to a locally finite graph Γ, denoted by $\mathcal{S}(M_\Gamma)$, satisfies an analog of Ivanov's theorem (see Theorem 1.1). When Γ has rank zero, $\mathrm{Map}(\Gamma) \cong \mathrm{Homeo}(\mathrm{Ends}(\Gamma))$ and Theorem 1.1 coincides with a recent result of Branman–Lyman, who show the homeomorphism group of a Stone space is isomorphic to the automorphism group of its complex of cuts [BL24b, Theorem A]. Our proofs provide an alternate perspective on this result.

Bavard–Dowdall–Rafi [BDR20] ultimately apply Ivanov's theorem to prove algebraic rigidity of the mapping class groups of infinite-type surfaces, showing that two infinite-type surfaces S and S' are homeomorphic if and only if $\mathrm{Map}^{\pm}(S) \cong \mathrm{Map}^{\pm}(S')$. In contrast, algebraic rigidity fails in general for locally finite infinite graphs, though all known examples are when both graphs are trees. For example, if Γ has end space a Cantor set, and Γ' has end space a Cantor set with one extra isolated point, then $\mathrm{Map}(\Gamma) \cong \mathrm{Map}(\Gamma')$ (see Theorem 2.1). A natural next question asks what conditions are needed to guarantee algebraic rigidity in the big $\mathrm{Out}(\mathbb{F}_n)$ setting.

1.2. **Outline of paper.** In Section 2, we define the sphere complex of the doubled handlebody associated to a locally finite infinite graph Γ, discuss some general tools for studying spheres in 3-manifolds, and recall relevant results from [BL24a]. In Section 3, we prove Theorem 1.1 by first showing that sphere graph isomorphisms induce isomorphisms of the dual graphs of pants decompositions of the doubled handlebody M_Γ. These graph isomorphisms define diffeomorphisms of M_Γ, which are compatible with pants decompositions that differ by a flip move. In Section 4 we generalize the results of [BL24a] to finite-type doubled handlebodies with S^2 boundaries by proving Theorem 1.2, with the aim of giving an alternative proof of Theorem 1.1 in Section 5 in the style of [BDR20]. In Section 6 we prove Theorem 1.3, generalizing a finite-type result of [BL24a]. Finally, in Section 7, we consider the existence of finite rigid sets in the low complexity cases not covered by Theorem 1.2.

Acknowledgements. The authors express their gratitude to Sanghoon Kwak, Vivian He, Priyam Patel, and Hannah Hoganson for their insightful conversations, which initially sparked our interest in this project. We also extend our appreciation to George Domat and Chris Leininger for valuable discussions, and special thanks to Mladen Bestvina, whose thoughtful input greatly contributed to the project's development. We would also like to thank the organizers of the 2023 Wasatch Topology Conference and 2024 Log Cabin Workshop for providing a collaborative environment that helped shape the direction of this work.

The first, second, and third authors acknowledge support from RTG DMS-1840190. The third author also acknowledges support from NSF grant DMS–2046889. The second and fourth authors acknowledge support from NSF grant DMS-2304774. The fifth author acknowledges support from NSF grant DMS-1745670.

2. Preliminaries

Let Γ be a locally finite graph. A proper map $f : \Gamma \to \Gamma$ is a **proper homotopy equivalence** if there exists a proper map $g : \Gamma \to \Gamma$ such that fg and gf are properly homotopic to the identity.

We denote by $\mathrm{Ends}(\Gamma)$ the (Freudenthal) **end space** of Γ. We recall that given a compact exhaustion $K_1 \subset K_2 \subset \dots$ of Γ, with maps $\pi_0(\Gamma \setminus K_i) \to \pi_0(\Gamma \setminus K_{i-1})$ induced by inclusion, $\mathrm{Ends}(\Gamma)$ is the inverse limit $\varprojlim_i \pi_0(\Gamma \setminus K_i)$. Like for surfaces, ends of graphs come in two flavors: those accumulated by loops (also called unstable) and those that are not accumulated by loops (also called stable ends). Heuristically, if you always see loops as you move out toward an end, then it is accumulated by loops. The (possibly empty) subset of all ends accumulated by loops is a closed subset of $\mathrm{Ends}(\Gamma)$, denoted $\mathrm{Ends}_\ell(\Gamma)$.

The rank and end space classify locally finite, infinite graphs up to proper homotopy equivalence.

Theorem 2.1 ([ADMQ90])**.** *Let Γ and Γ' be two locally finite, infinite graphs. Then Γ is* ***properly homotopy equivalent*** *to Γ' if and only if*

$$(\mathrm{rk}(\Gamma), \mathrm{Ends}(\Gamma), \mathrm{Ends}_\ell(\Gamma)) \cong (\mathrm{rk}(\Gamma'), \mathrm{Ends}(\Gamma'), \mathrm{Ends}_\ell(\Gamma')).$$

By "$\cong$" we mean that $\mathrm{rk}(\Gamma) = \mathrm{rk}(\Gamma')$ and there exists a homeomorphism of pairs $f : (\mathrm{Ends}(\Gamma), \mathrm{Ends}_\ell(\Gamma)) \to (\mathrm{Ends}(\Gamma), \mathrm{Ends}_\ell(\Gamma))$. The tuple $(\mathrm{rk}(\Gamma), \mathrm{Ends}(\Gamma), \mathrm{Ends}_\ell(\Gamma))$ is called the ***charateristic triple.***

Let $\mathrm{PHE}(\Gamma)$ denote the group of proper homotopy equivalences of Γ, equipped with the compact-open topology.

Definition 2.2. The **mapping class group** of Γ, denoted $\mathrm{Map}(\Gamma)$, is defined as

$$\mathrm{Map}(\Gamma) \coloneqq \mathrm{PHE}(\Gamma)/\text{proper homotopy}.$$

$\mathrm{Map}(\Gamma)$ is a topological group with the quotient topology. Not every homotopy equivalence of Γ that is proper is a proper homotopy equivalence as defined above (see [AB21, Example 4.1]).

We now construct the doubled handlebody associated to Γ, whose study will be the main focus of this paper.

Definition 2.3. Let N_Γ be the 3-manifold with 0-handles and 1-handles glued according to the vertices and edges, respectively, in Γ. The **doubled handlebody** M_Γ **associated to** Γ is the double of N_Γ, obtained by gluing two copies of N_Γ along ∂N_Γ.

One can think of N_Γ as a regular neighborhood of the image of a proper embedding of Γ in $\mathbb{R}^3$ (see Figure 1). When Γ is finite of rank n, $M_\Gamma \cong \#_n(S^2 \times S^1)$, the connect sum of n copies of $S^2 \times S^1$. Let $M_{n,s}$ denote $\#_n(S^2 \times S^1)$ with s disjoint open balls removed; note that $M_{n,s} \setminus \partial M_{n,s} \cong M_\Gamma$ for Γ rank n with s rays.

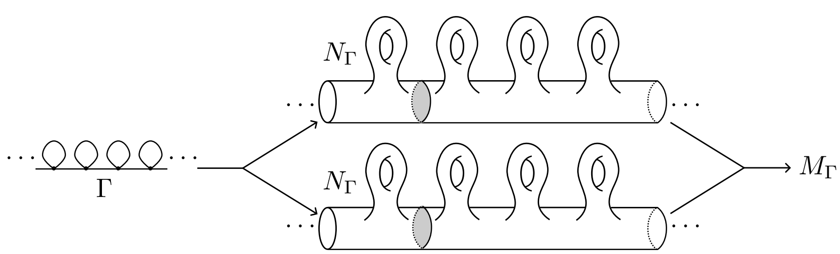

FIGURE 1. The construction of M_Γ for the ladder graph. Disk cross-sections in distinct copies of N_Γ form hemispheres of a 2-sphere in M_Γ.

As usual, we define the **mapping class group** of M_Γ to be

$$\mathrm{Map}(M_\Gamma) = \mathrm{Diff}^+(M_\Gamma)/\text{isotopy}$$

where $\mathrm{Diff}^+(M_\Gamma)$ denotes the orientation preserving diffeomorphisms of M_Γ.

Remark. We will generally work in the smooth category: submanifolds are assumed to be smoothly embedded, and homotopies and isotopies are smooth.

2.1. **Spheres in 3-manifolds.** The **sphere complex** on M_Γ is analogous to the curve complex on surfaces. An embedded sphere in a 3-manifold M is **essential** if it does not bound a ball and is not peripheral, *i.e.* not isotopic into a small neighborhood of a boundary sphere or puncture of M.

Definition 2.4. The **sphere complex** associated with a 3-manifold M is denoted by $\mathcal{S}(M)$. It is a simplicial complex with

- **vertices** corresponding to isotopy classes of essential embedded 2-spheres in M;
- k**-cells** spanned by vertices $a_0, \dots, a_k$ if these spheres can be isotoped to be pairwise disjoint.

We equip the automorphism group $\mathrm{Aut}(\mathcal{S}(M_\Gamma))$ of $\mathcal{S}(M_\Gamma)$ with the **permutation topology**. This topology is defined by the subbasis $\{U_a\}$ at the identity, where a ranges over the vertices of $\mathcal{S}(M_\Gamma)$ and

$$U_a = \{\phi \in \mathrm{Aut}(\mathcal{S}(M_\Gamma)) \mid \phi(a) = a\}.$$

Note that $\mathrm{Map}(M_\Gamma)$ acts naturally on $\mathcal{S}(M_\Gamma)$, as $\mathrm{Diff}^+(M_\Gamma)$ acts on the collection of essential embedded spheres of M_Γ, preserving disjointness between pairs of spheres. We have the following from [Uda24], which is the first step in proving Theorem 1.1:

Theorem 2.5 ([Uda24, Theorem 1.1])**.** *There is a surjective continuous map* $\Psi : \mathrm{Map}(M_\Gamma) \to \mathrm{Map}(\Gamma)$ *with* $\ker(\Psi)$ *contained in the kernel K of the action of* $\mathrm{Map}(M_\Gamma)$ *on* $\mathcal{S}(M_\Gamma)$*, hence inducing an action of* $\mathrm{Map}(\Gamma)$ *on* $\mathcal{S}(M_\Gamma)$*. If Γ is not a graph with rank r and s rays such that $2r + s < 4$ or $(r, s) \in \{(0, 4), (2, 0)\}$, then $K = \ker(\Psi)$ and the induced action is faithful.*

In the next two sections, we present tools used to characterize the links of simplices in $\mathcal{S}(M_\Gamma)$. In particular, we give a method for constructing intersecting spheres and describe certain full subcomplexes of links.

2.1.1. *Intersecting spheres.* Henceforth, let M be an oriented 3-manifold. A submanifold N is **essential** if it is not null-homotopic or peripheral, and we say two submanifolds **intersect essentially** if they cannot be made disjoint up to homotopy. Given transverse oriented submanifolds $S, T \subset M$ of complementary dimension, the **signed intersection number** of S and T is the sum

$$\hat{\imath}(S, T) \coloneqq \sum_{x \in S \cap T} \epsilon(x)$$

where $\epsilon(x) = \pm 1$ according to the orientation induced by S, T and M. The signed intersection number $\hat{\imath}$ is invariant up to homotopy in the following sense:

Lemma 2.6. *Let $S \subset M$ be an oriented submanifold and T an oriented manifold such that* $\dim S + \dim T = \dim M$*. Let $\Psi : I \times T \to M$ be a homotopy transverse to S such that $\Psi^{-1}(S)$ is compact and $\Psi(I \times \partial T) \cap S = \Psi(I \times T) \cap \partial S = \varnothing$. Let $\psi_t(x) = \Psi(t, x)$. If ψ_0, ψ_1 are embeddings transverse to S, then $\hat{\imath}(\psi_0(T), S) = \hat{\imath}(\psi_1(T), S)$.*

Proof. Fix the usual orientation for I and endow $I \times T$ with the product orientation. $Z = \Psi^{-1}(S)$ is then an oriented compact 1-manifold with (oriented) boundary $-\psi_0^{-1}(S) \sqcup \psi_1^{-1}(S)$. In particular, $\hat{\imath}(\psi_1^{-1}(S), Z) - \hat{\imath}(\psi_0^{-1}(S), Z) = \hat{\imath}(\partial Z, Z) = 0$, from which the claim follows. $\square$

By [Lau73], if two essential spheres in M are homotopic, then they are isotopic, and this isotopy extends to an ambient isotopy of $M \setminus \partial M$. In particular, if two spheres are disjoint up to homotopy, then there exists an isotopy of *one* sphere realizing their disjointness. Likewise, if a sphere and an arc are disjoint up to homotopy, then we may homotope the arc (rel ends) to be disjoint.

Corollary 2.7. *Let $a \subset M \setminus \partial M$ be an embedded non-peripheral sphere and $\gamma \subset M$ a transverse simple arc. If $\hat{\imath}(a, \gamma) \neq 0$, then a, γ are essential and intersect essentially.*

Proof. By the above, if a, γ do not intersect essentially, then there exists some arc γ' homotopic to γ (rel ends) and disjoint from a. In fact, γ' is then properly homotopic to γ, and up to a small homotopy leaving $\hat{\imath}$ unchanged we assume this homotopy is transverse to a. As a is compact and disjoint from ∂M, we satisfy the hypotheses of Lemma 2.6, and $\hat{\imath}(a, \gamma) = \hat{\imath}(a, \gamma') = 0$. □

Lemma 2.8. *Let $a \subset M \setminus \partial M$ be a non-peripheral sphere and let γ be a simple arc between distinct punctures. Let $b \cong S^2$ be the boundary of a regular neighborhood of γ. If γ intersects a essentially, then b must also intersect a essentially.*

Proof. We prove the contrapositive. Suppose that a, b do not intersect essentially. Then a is isotopic to a sphere a' disjoint from b. Let $M' \sqcup M'' = M \setminus b$; M' is disjoint from γ and M'' is a thrice-punctured 3-sphere. If $a' \subset M'$ then a does not intersect γ essentially. If $a' \subset M''$, then a' is peripheral in M''; since a is not peripheral in M, a' is homotopic to b which is disjoint from γ. □

2.1.2. *Full subcomplexes of $\mathcal{S}(M_\Gamma)$.* A **sphere system** is a (possibly infinite) collection of distinct pairwise disjoint essential spheres. We show that the sphere complex of a submanifold obtained by cutting along a sphere system is a full subcomplex of $\mathcal{S}(M_\Gamma)$:

Proposition 2.9. *Let $Z \subset M_\Gamma \setminus \sigma$ be a complementary component of a sphere system σ. Then $Z \hookrightarrow M_\Gamma$ induces an injective full simplicial map $\mathcal{S}(Z) \hookrightarrow \mathcal{S}(M_\Gamma)$.*

The proof of Proposition 2.9 will make use of the following lemma and two results of [Hat95].

Lemma 2.10. *$Z \hookrightarrow M_\Gamma$ is π_2-injective. In particular, any essential sphere in Z is essential in M_Γ.*

Proof. Let $p : \tilde{M}_\Gamma \to M_\Gamma$ be the universal covering of M_Γ and fix $\tilde{Z} \subset \tilde{M}_\Gamma$ a connected component of $p^{-1}(Z)$; observe $\tilde{Z}$ is a complementary component of the sphere system $\tilde{\sigma} = p^{-1}(\sigma)$. We show the inclusion $\tilde{Z} \hookrightarrow \tilde{M}_\Gamma$ is π_2-injective and preserves essentiality, which suffices: two spheres are homotopic if and only if they have a pair of homotopic lifts, and neighborhoods of punctures lift homeomorphically.

An application of van Kampen's theorem shows that Z is π_1-injective, hence $\tilde{Z}$ is simply connected. The pair $(\tilde{M}_\Gamma, \tilde{Z})$ induces the exact sequence

$$H_3(\tilde{M}_\Gamma, \tilde{Z}) \to H_2(\tilde{Z}) \xrightarrow{\iota_\natural} H_2(\tilde{M}_\Gamma)\,;$$

since each component of $\tilde{M}_\Gamma \setminus \tilde{Z}$ is non-compact, $H_3(\tilde{M}_\Gamma, \tilde{Z}) \cong H_3(\tilde{M}_\Gamma \setminus \tilde{Z}, \partial \tilde{Z}) = 0$. Hence $\iota_\natural$, which is induced by the inclusion $\iota : \tilde{Z} \hookrightarrow \tilde{M}_\Gamma$, is injective. Since $\tilde{Z}, \tilde{M}_\Gamma$ are simply connected, the Hurewicz (natural) isomorphism implies that $\iota_* : \pi_2(\tilde{Z}) \to \pi_2(\tilde{M}_\Gamma)$ is likewise injective. It follows that two spheres $a, a' \subset \tilde{Z}$ are homotopic in

$\tilde{M}_\Gamma$ only if they are likewise in $\tilde{Z}$, and in particular $a \subset \tilde{Z}$ is null-homotopic in $\tilde{M}_\Gamma$ only if it is likewise in $\tilde{Z}$. Extending along paths to a basepoint $z_0 \in \tilde{Z}$, we obtain homotopic based spheres $\bar{a}, \bar{a}' \subset \tilde{M}_\Gamma$ also homotopic to a, a' respectively. Since $\tilde{M}_\Gamma$ is simply connected, $\bar{a}, \bar{a}'$ can be chosen to be homotopic to a, a' via homotopies in $\tilde{Z}$, and π_2-injectivity implies that $\bar{a}, \bar{a}'$ are homotopic in $\tilde{Z}$ as well.

It remains to show that if $a \subset \tilde{Z}$ is peripheral in $\tilde{M}_\Gamma$, then it is null-homotopic or peripheral in $\tilde{Z}$. Suppose that a is peripheral to an (isolated) puncture $e \in \mathrm{Ends}(\tilde{M}_\Gamma)$. If $e \in \mathrm{Ends}(\tilde{Z})$, then a is peripheral in $\tilde{Z}$. Else, let M^+ be obtained by replacing a neighborhood of e disjoint from $\tilde{Z}$ with a ball and $\sigma^+ \subset M^+$ by removing components (in fact, at most one) of $\tilde{\sigma}$ peripheral in M^+. Then $\tilde{Z} \hookrightarrow M^+$ is isotopic to a component of $M^+ \setminus \sigma^+$, hence π_2-injective by the above: since a is null-homotopic in M^+, it is also null-homotopic in $\tilde{Z}$. □

We briefly review Hatcher normal form. Given a maximal sphere system $\sigma \subset M \cong M_{n,s}$ and a transverse sphere $a \subset M$, a is in **normal form** with σ if either a is equal to a component of σ or if (the closure of) each component of $a \setminus \sigma$ meets each sphere in σ at most once and no component is homotopic rel boundary to a disk in σ. The intersection of a with the closure of a component of $M \setminus \sigma$ is called a **piece** of a.

Proposition 2.11 ([Hat95, Proposition 1.1])**.** *Let $a \subset M_{n,s}$ be an essential sphere and $\sigma \subset M_{n,s}$ a maximal sphere system. Then a is homotopic to a sphere a' in normal position with σ, and if $a \cap \sigma_0 = \varnothing$ for some subsystem $\sigma_0 \subset \sigma$, then there exists a homotopy disjoint from σ_0.*

Remark. Our statement is slightly stronger than that in [Hat95], but immediate from the construction in its proof.

Proposition 2.12 ([Hat95, Proposition 1.2])**.** *Let $a, a' \subset M_{n,s}$ be essential homotopic spheres in normal form with a maximal sphere system $\sigma \subset M_{n,s}$. Then a, a' are homotopic via a homotopy which restricts to an isotopy on intersections with each sphere in σ.*

Proof of Proposition 2.9. The map is well defined by Lemma 2.10: a sphere is essential in Z only if it is essential in M_Γ. To show injectivity, we verify that two spheres in Z are isotopic in M_Γ only if they likewise are in Z (*n.b.* π_2-injectivity is insufficient since M_Γ is not simply connected). To show fullness, we prove that two spheres intersect essentially in Z only if they likewise do in M_Γ; the converse is immediate, implying that the map is simplicial.

First suppose that $a, b \subset M_\Gamma \setminus \sigma$ are essential spheres homotopic in M_Γ, and fix some $N \cong M_{n,s} \subset M_\Gamma$ containing the image of the homotopy such that the components of ∂N are essential spheres disjoint from σ. Let $\sigma' \supset \{b\} \cup \sigma$ be a maximal sphere system on N. By Proposition 2.11, a is homotopic to some a' in normal form with σ' disjointly from σ; since b is likewise in normal form with σ', disjoint from σ, and homotopic with a (hence a'), a', b are homotopic disjoint from σ by Proposition 2.12. Then, a, b are likewise homotopic disjointly from σ.

If $a, b \subset Z$ are essential spheres homotopic in M_Γ, then they are disjoint from σ and thus homotopic in Z by the above. If instead $a, b \subset Z$ are essential spheres that do not intersect essentially in M_Γ, then fix an embedded sphere a' homotopic to a in M_Γ and simultaneously disjoint from b, σ. Again by the above, a, a' are homotopic in Z and hence a, b do not intersect essentially in Z. □

Remark 2.13. The image of $\mathcal{S}(Z)$ in Proposition 2.9 is exactly the subcomplex in $\mathcal{S}(M_\Gamma)$ spanned by spheres in $\mathrm{link}(\sigma) \subset \mathcal{S}(M_\Gamma)$ contained in Z. We show that a sphere $a \in \mathrm{link}(\sigma)$ contained in Z is essential in Z, which suffices. If $a \subset Z$ is not essential in Z, then a bounds a ball in Z thus likewise in M_Γ, or a is peripheral in Z hence either $a \in \sigma$ or it is peripheral in M_Γ: in all cases, $a \notin \mathrm{link}(\sigma)$.

Corollary 2.14. *Let $Z_i \subset M_\Gamma \setminus \sigma$ denote the complementary components of a sphere system σ. Then* $\mathrm{link}(\sigma)$ *is isomorphic to the join $\ast_i \mathcal{S}(Z_i)$.*

Proof. By Remark 2.13, it suffices to show that the $\mathcal{S}(Z_i)$ have pairwise disjoint images in $\mathrm{link}(\sigma)$. Suppose that $a \in \mathrm{link}(\sigma)$ is in the image of $\mathcal{S}(Z_i)$ and $\mathcal{S}(Z_j)$, hence we may fix homotopic representatives $\alpha \subset Z_i$ and $\alpha' \subset Z_j$. As in the proof of Proposition 2.9 we obtain a homotopy between α, α' disjoint from σ, hence α, α' lie in the same complementary component of σ and $Z_i = Z_j$. □

2.1.3. *Pants decompositions and $M_{0,s}$.* We now shift our focus and introduce the tools we will need to prove Theorem 1.2. The proof relies on a generalization of the results of Bering–Leininger in [BL24a], which we will discuss in Section 4. We state here the relevant preliminary definitions and results from [BL24a].

Definition 2.15. Any manifold homeomorphic to $M_{0,3}$ is called a **pair of pants**. A maximal sphere system $P \subset \mathcal{S}(M_{n,s})^{(0)}$ is called a **pants decomposition**. Fixing an open regular neighborhood $\mathrm{nbd}(P) \supset P$, each component of $M_{n,s} \setminus \mathrm{nbd}(P)$ is homeomorphic to a pair of pants.

Definition 2.16. Suppose P is a pants decomposition of $M_{n,s}$. Two spheres $a, b \in P$ are **adjacent** in P if they are two of the boundary spheres of some pair of pants component of $M_{n,s} \setminus P$. A sphere $a \in P$ is **self-adjacent** in P if it bounds two cuffs of a single pair of pants in $M_{n,s} \setminus P$.

Definition 2.17. Let P_a be a pants decomposition containing a sphere a that is not self-adjacent. Consider the connected component of $M_{n,s} \setminus \mathrm{nbd}(P_a \setminus a)$ containing a. This is homeomorphic to $M_{0,4}$, and $\mathcal{S}(M_{0,4}) = \{\, a, a', a'' \,\}$. There are pants decompositions $P_{a'} = P_a \setminus \{\, a \,\} \cup a'$ and similarly $P_{a''}$. A change in pants decomposition from $P_a \mapsto P_{a'}$ or $P_a \mapsto P_{a''}$ is called a **flip move.** See Figure 2.

Definition 2.18 ([BL24a, Definition 2.2])**.** Let $X \subset \mathcal{S}(M_{n,s})$ be a subcomplex. Two spheres $a, a' \in X^{(0)}$ which intersect essentially have **X-detectable intersection** if there are pants decompositions $P_a, P_a' \subset X^{(0)}$ such that $P_a, P_{a'}$ differ by a **flip move** $a \to a'$, *i.e.* for which $a \in P_a$, $a' \in P_{a'}$ and $P_a \setminus \{a\} = P_{a'} \setminus \{a'\}$.

Since $P_a, P_{a'}$ differ by a flip move $a \to a'$, the component of $M_{n,s} \setminus \mathrm{nbd}(P_a \setminus \{a\})$ containing $a \cup a'$ deformation retracts to $a \cup a'$ and is homeomorphic to $M_{0,4}$ (see Figure 2). If a, a' has an intersection X-detectable, then they intersect essentially: we apply Proposition 2.9 to $M_{0,4} \setminus \partial M_{0,4} \hookrightarrow M_{n,s} \setminus \partial M_{n,s}$.

We will consider locally injective simplicial maps $f : X \to \mathcal{S}(M_{n,s})$, where X is a subcomplex of $\mathcal{S}(M_{n,s})$. The map f is **simplicial** if it sends vertices to vertices and edges to edges, and **locally injective** if it is injective when restricted to the stars of vertices in X.

Lemma 2.19 ([BL24a, Lemma 8])**.** *Let $X \subset \mathcal{S}(M_{n,s})$ be a subcomplex, and suppose $f : X \to \mathcal{S}(M_{n,s})$ is a locally injective simplicial map. If $a, a' \in X^{(0)}$ have X-detectable intersection, then $f(a)$ and $f(a')$ have $f(X)$-detectable intersection.*

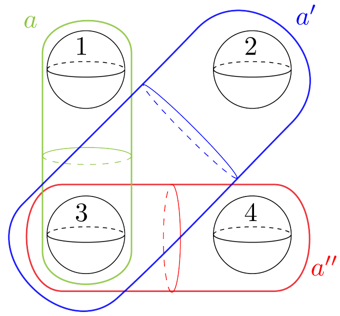

FIGURE 2. The three spheres a, a', and a'' in $M_{0,4}$ differ by a flip move. $\partial M_{0,4}$ is labeled 1–4.

The following two results discuss the sphere graph of $M_{0,s}$.

Lemma 2.20 ([BL24a, Lemma 9]). *A sphere $a \in \mathcal{S}(M_{0,s})^{(0)}$ is determined by the partition of $\partial M_{0,s}$ induced by the connected components of $M_{0,s} \setminus a$.*

Lemma 2.21 ([BL24a, Corollary 11]). *If $N \cong M_{0,s} \cong N'$ and $f : \mathcal{S}(N) \to \mathcal{S}(N')$ is a simplicial automorphism, then there is a homeomorphism $h : N \to N'$ so that $h(a) = f(a)$ for every $a \in \mathcal{S}(N)$.*

The following lemma provides a useful topological criterion to distinguish spheres combinatorially.

Lemma 2.22 ([BL24a, Lemma 12]). *Suppose $a, b, c \in \mathcal{S}(M_{n,s})$ are such that b and c are disjoint and a essentially intersects b and c, each in one circle, such that the boundary components of $\overline{\mathrm{nbd}}(a \cup b \cup c)$ are either essential or peripheral. Let $b' \in \mathcal{S}(\mathrm{nbd}(a \cup b))$ be the unique sphere in $\mathrm{nbd}(a \cup b)$ distinct from a and b and $c' \in \mathcal{S}(\mathrm{nbd}(a \cup c))$ the unique sphere in $\mathrm{nbd}(a \cup b)$ distinct from a and c. Then b' and c' intersect, and both intersect b and c.*

We include a useful technical result, which is equivalent to the connectedness of the pants graph for $M_{n,s}$. In particular, this proposition replaces the use of Outer space in the proof of Proposition 22 in [BL24a].

Proposition 2.23. *Any two pants decompositions in $M_{n,s}$ differ by a finite sequence of flip moves.*

Proof. Fix $M_{n,s}$ with minimal dimension $k = \dim \mathcal{S}(M_{n,s})$ such that the proposition does not hold: $k > 0$, else $\mathcal{S}(M_{n,s})$ is empty or $M_{n,s}$ is $M_{1,1}$ or $M_{0,4}$ and the claim is immediate. Let $\mathcal{S}'$ denote the barycentric subdivision of $\mathcal{S}(M_{n,s})$ and let $\mathcal{S}''$ denote the barycentric subdivision of $\mathcal{S}(M_{n,s})^{(k-2)}$. It is equivalent to show that the subcomblex $\mathcal{P} = \mathcal{S}' \setminus \mathcal{S}''$ is connected: in particular, any flip move corresponds to moving between (the barycenters of) two k-simplices in $\mathcal{S}(M_{n,s})$ via a $(k-1)$-simplex. We claim that if two k-simplices in $\mathcal{S}(M_{n,s})$ share a face then their barycenters are in the same component of $\mathcal{P}$, which suffices. In particular any simplex in $\mathcal{S}(M_{n,s})$ is the face of a k-simplex, hence any path in $\mathcal{S}'$ gives a sequence of k-simplices P_i with $P_i \cap P_{i+1} \neq \varnothing$. Since $k > 0$, $\mathcal{S}(M_{n,s})$ and $\mathcal{S}'$ are connected, hence by the claim $\mathcal{P}$ is connected.

Suppose that P, P' are k-simplices with non-empty intersection $\sigma = P \cap P'$. Let N_i denote the non-pants components of $M_{n,s} \setminus \sigma$. Since σ is non-empty and each $\mathcal{S}(N_i)$ is a full subcomplex of $\mathcal{S}(M_{n,s})$, $\dim(\mathcal{S}(N_i)) < k$. By minimality we may choose a sequence of flip moves in N_i between $P \cap N_i$ and $P' \cap N_i$ for each i; concatenating these sequences gives a path between P, P' in $\mathcal{P}$. □

2.2. **Edge isomorphisms and rigidity.** In this subsection, we state a version of Whitney's graph isomorphism theorem which will be useful in Section 3. Let Δ, Δ' denote graphs, possibly with loops and multiple edges between vertices.

Definition 2.24. A bijection $\psi : E(\Delta) \to E(\Delta')$ is an **edge isomorphism** if for all $e, e' \in E(\Delta)$, there is an isomorphism of subgraphs $e \cup e' \to \psi(e) \cup \psi(e')$ inducing $\psi|_{\{e,e'\}}$.

Definition 2.25. Given a map $\psi : E(\Delta) \to E(\Delta')$ and graphs G, G', ψ has a G, G'**-pair** if, for some $\eta \subset E(\Delta)$, the subgraphs $\bigcup_{e \in \eta} e, \bigcup_{e \in \eta} \psi(e)$ are isomorphic to G, G', in any order.

Let K_3 denote the 3-clique and $K_{1,3}$ the 3-star. The following theorem is immediate from [Gar84, *e.g.* Corollary 2.2].

Theorem 2.26 (Gardner)**.** *Suppose Δ, Δ' are finite. An edge isomorphism $\psi : E(\Delta) \to E(\Delta')$ is induced by an isomorphism $\Delta \to \Delta'$ if and only if ψ does not have a $K_3, K_{1,3}$-pair.*

Remark. If Δ is connected and $|\Delta| \neq 2$, then a unique graph isomorphism induces ψ. In particular, if ψ' were another such then $\psi^{-1}\psi'$ is the identity on edges, which implies identity unless Δ consists of two vertices with edges only between them.

Corollary 2.27. *Theorem 2.26 likewise holds if Δ, Δ' are infinite and Δ is connected with $|\Delta| \neq 2$.*

Proof. Fix a compact exhaustion of Δ by connected subgraphs Δ_i of order $n_i \neq 2$. Let Δ'_i denote the subgraph $\bigcup_{e \in E(\Delta_i)} \psi(e)$; note that Δ'_i is likewise a compact exhaustion of Δ', since every $e' \in E(\Delta')$ is the image of an edge $e \in E(\Delta_i)$ for some i. Then ψ restricts to edge isomorphisms $\psi_i : E(\Delta_i) \to E(\Delta'_i)$; if any has a $K_3, K_{1,3}$-pair then so does ψ, hence it follows that ψ_i is induced by a unique isomorphism $\tilde{\psi}_i : \Delta_i \to \Delta'_i$. By uniqueness, these maps form a direct system, the direct limit of which is an isomorphism $\tilde{\psi} : \Delta \to \Delta'$ inducing ψ. □

As in the remark, we observe that the $\tilde{\psi}$ above is likewise unique.

3. A proof of the main theorem

We prove Theorem 1.1. We first show that a sphere graph isomorphism induces an isomorphism between the dual graphs of a pants decomposition and its image, which we will use to define the desired diffeomorphism of M_Γ.

3.1. **Dual graphs of pants decompositions.** Henceforth, let $f : \mathcal{S}(M_\Gamma) \xrightarrow{\sim} \mathcal{S}(M_{\Gamma'})$ be an isomorphism of sphere graphs. Given a maximal sphere system $\sigma \subset \mathcal{S}(M_\Gamma)$, $f\sigma$ is likewise maximal and $\sigma, f\sigma$ define pants decompositions of $M_\Gamma, M_{\Gamma'}$ respectively. Let Δ_σ, $\Delta_{f\sigma}$ denote the dual graphs to these pants decompositions. For a sphere $a \in \sigma$, we denote by e_a the dual edge in Δ_σ. Thus $E(\Delta_\sigma) = \{e_t : t \in \sigma\}$ and the restriction $f|_\sigma$ defines a bijection $f|_\sigma : E(\Delta_\sigma) \to E(\Delta_{f\sigma})$.

Lemma 3.1. *$f|_\sigma$ is an edge isomorphism.*

Proof. For any $e_a, e_b \in E(\Delta_\sigma)$, we must show there is an isomorphism on subgraphs $e_a \cup e_b \to e_{fa} \cup e_{fb}$ for which $e_a \mapsto e_{fa}$ and $e_b \mapsto e_{fb}$. By examining complementary components and applying Corollary 2.14 we obtain the following:

(i) e_a is a loop in Δ_σ if and only if $\operatorname{link}(\sigma \setminus a) = \{a\} \cong \mathcal{S}(M_{1,1})$. Otherwise, $\operatorname{link}(\sigma \setminus a) = \{a, a', a''\} \cong \mathcal{S}(M_{0,4})$ for some distinct a', a''.

(ii) e_a, e_b have common incident vertices (*i.e.* a, b are adjacent) in Δ_σ if and only if $\operatorname{link}(\sigma \setminus \{a, b\}) \cong \mathcal{S}(M_{1,2})$ or $\mathcal{S}(M_{0,5})$.

(iii) If e_a, e_b are adjacent in Δ_σ and e_a or e_b is a loop, then they are incident on exactly one common vertex. If e_a, e_b are not loops and incident on exactly one common vertex, then $\operatorname{link}(\sigma \setminus \{a, b\}) = \mathcal{S}(M_{0,5})$. If they are incident on two common vertices (*i.e.* they form a bigon), then $\operatorname{link}(\sigma \setminus \{a, b\}) = \mathcal{S}(M_{1,2})$.

These properties are sufficient to determine $e_a \cup e_b$ up to order preserving isomorphism. Moreover, since (i)-(iii) are specified combinatorially (because links are preserved by isomorphisms), they are preserved by f: if any hold for e_a, e_b, then likewise do they for e_{fa}, e_{fb}, which suffices. □

Proposition 3.2. *$f|_\sigma$ is induced by an isomorphism $\Delta_\sigma \to \Delta_{f\sigma}$.*

Proof. By Theorem 2.26 and Corollary 2.27, it suffices that $f|_\sigma$ does not have a $K_3, K_{1,3}$-pair. Suppose that $f|_\sigma$ has such a pair on $e_a, e_b, e_c \in E(\Delta_\sigma)$. If $e_a \cup e_b \cup e_c \cong K_3$ then $e_{fa} \cup e_{fb} \cup e_{fc} \cong K_{1,3}$. It follows that $\operatorname{link}(\sigma \setminus \{a, b, c\}) \cong \mathcal{S}(M_{1,3})$ and $\operatorname{link}(f\sigma \setminus \{fa, fb, fc\}) \cong \mathcal{S}(M_{0,6}) \not\cong \mathcal{S}(M_{1,3})$, a contradiction. See Figure 3. An analogous argument applies exchanging K_3 and $K_{1,3}$. □

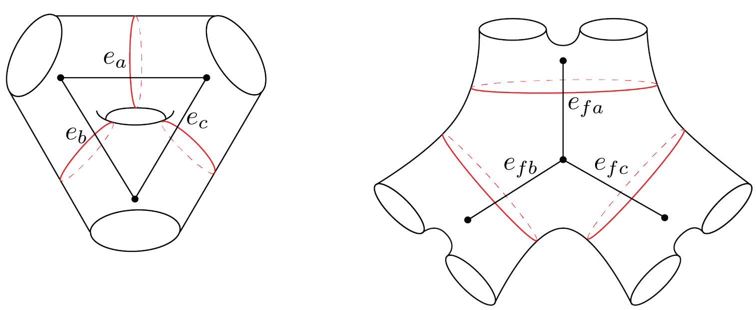

FIGURE 3. The $M_{1,3}$ and $M_{0,6}$ components in $M_\Gamma \setminus (\sigma \setminus \{a, b, c\})$ and $M_{\Gamma'} \setminus (f\sigma \setminus \{fa, fb, fc\})$ respectively. The figure shows one-half of the doubled handlebodies.

3.2. **Constructing a diffeomorphism.** We are interested in diffeomorphisms $M_\Gamma \to M_{\Gamma'}$ that agree with f locally about σ.

Definition 3.3. For a maximal sphere system $\sigma \subset \mathcal{S}(M_\Gamma)$, let

$$X_\sigma \coloneqq \bigcup_{a \in \sigma} \operatorname{link}(\sigma \setminus a)\,.$$

We note that $\sigma \subset X_\sigma$. Let $f_\sigma : \Delta_\sigma \to \Delta_{f\sigma}$ denote an isomorphism inducing $f|_\sigma$. Realizing $M_\Gamma, M_{\Gamma'}$ as the gluing of pairs of pants along spheres in $\sigma, f\sigma$ respectively, f_σ is induced by a diffeomorphism $h_0 : M_\Gamma \to M_{\Gamma'}$ defined up to isotopy and precomposition by diffeomorphisms fixing pointwise Δ_σ.

Definition 3.4. A separating sphere $a \subset M_\Gamma$ is **almost peripheral** if it bounds a pair of pants or a $M_{1,1}$.

For an almost peripheral sphere $a \in \sigma$, $\mathrm{link}(\sigma \setminus a) = \{a, a', a''\}$ and the half-twist fixes a and exchanges a', a'': obtain h from h_0 by precomposing by half-twists on almost peripheral $a \in \sigma$ such that h agrees with f on $\mathrm{link}(\sigma \setminus a)$.

Lemma 3.5. *h agrees with f on X_σ.*

Proof. By construction h_* agrees with f on σ and $\mathrm{link}(\sigma \setminus a)$ for $a \in \sigma$ almost peripheral. If a is a loop in Δ_σ, then $\mathrm{link}(\sigma \setminus a) = \{a\} \subset \sigma$. Thus it remains to show that h_* agrees with f on $\mathrm{link}(\sigma \setminus a)$ when a is not a loop or almost peripheral.

Let $\mathrm{link}(\sigma \setminus a) = \{a, a', a''\}$ and $\sigma' = (\sigma \setminus a) \cup a'$ and $\sigma'' = (\sigma \setminus a) \cup a''$. Let $M \cong M_{0,4} \setminus \partial M_{0,4}$ denote the complementary component of $\sigma \setminus a$ containing a, hence $M' = h(M)$ is the complementary component of $f\sigma \setminus fa$ containing fa. Since a is not almost peripheral, there exist $b, b' \in \sigma$ which are separated by a in M (possibly $b = b'$). Without loss of generality, a' separates b, b' in M and a'' does not, hence likewise ha and ha' separate hb, hb' in M' and ha'' does not. It follows that $\mathrm{link}(\sigma \setminus \{b, b'\}) \cong \mathrm{link}(\sigma' \setminus \{b, b'\}) \not\cong \mathrm{link}(\sigma'' \setminus \{b, b'\})$ and $\mathrm{link}(h\sigma \setminus \{hb, hb'\}) \cong \mathrm{link}(h\sigma' \setminus \{hb, hb'\}) \not\cong \mathrm{link}(h\sigma'' \setminus \{hb, hb'\})$. See Figure 4.

We have that $ha = fa$, $hb = fb$, and $hb' = fb'$. Suppose that $ha' = fa''$ and $ha'' = fa'$. Then $h\sigma'' = f\sigma'$ (and *vice versa*), hence $\mathrm{link}(h\sigma'' \setminus \{hb, hb'\}) = \mathrm{link}(f\sigma' \setminus \{fb, fb'\}) \cong \mathrm{link}(\sigma' \setminus \{b, b'\}) \cong \mathrm{link}(\sigma \setminus \{b, b'\}) \cong \mathrm{link}(h\sigma \setminus \{hb, hb'\})$, a contradiction with the above. Hence $ha' = fa'$ and $ha'' = fa''$ as required. □

Remark 3.6. The map f need not be an isomorphism for Lemma 3.5 to hold. Suppose that $f : Z \subset \mathcal{S}(M_\Gamma) \to \mathcal{S}(M_{\Gamma'})$ such that $X_\sigma \subset Z$, $f|_\sigma$ is the edge map of an isomorphism $\Delta_\sigma \xrightarrow{\sim} \Delta_{f\sigma}$, and $h : M_\Gamma \to M_{\Gamma'}$ is the diffeomorphism constructed as above. If f extends to an isomorphism $\mathrm{link}(\sigma \setminus \{a, b, c\}) \xrightarrow{\sim} \mathrm{link}(f\sigma \setminus \{fa, fb, fc\})$ for any $a, b, c \in \sigma$ such that $e_a \cup e_b \cup e_c$ is connected in Δ_σ, then h induces f on X_σ: in particular, in the proof above $\mathrm{link}(\sigma' \setminus \{b, b'\}) \subset \mathrm{link}(\sigma \setminus \{a, b, b'\})$ and $\mathrm{link}(f\sigma' \setminus \{fb, fb'\}) \cong \mathrm{link}(\sigma' \setminus \{b, b'\})$, hence the argument applies without modification. We will use this fact in Section 6.1.

Lemma 3.7. *Suppose that ρ, ρ' are maximal sphere systems that differ by a flip move. If a diffeomorphism $g : M_\Gamma \to M_{\Gamma'}$ agrees with f on X_ρ, then it likewise agrees with f on $X_{\rho'}$.*

Proof. Suppose that ρ' is obtained from ρ by a flip move $a \mapsto a'$, and let $\rho_0 \subset \rho \cap \rho'$ denote the set of spheres in ρ adjacent to a (equivalently, spheres in ρ' adjacent to a'). We note that $\mathrm{link}(\rho \setminus a) = \mathrm{link}(\rho' \setminus a')$, $\rho \triangle \rho' = \{a, a'\} \subset \mathrm{link}(\rho \setminus a)$, and for $\gamma \in (\rho \cap \rho') \setminus \rho_0$, $\mathrm{link}(\rho \setminus \gamma) = \mathrm{link}(\rho' \setminus \gamma)$.

It follows that $X_{\rho'} \setminus X_\rho \subset \bigcup_{b \in \rho_0} \mathrm{link}(\rho' \setminus b)$, thus it suffices that g_* agrees with f_* on $\mathrm{link}(\rho' \setminus b)$ for $b \in \rho_0$. If b is a loop in $\Delta_{\rho'}$, then $\mathrm{link}(\rho' \setminus b) = \{b\} \subset \rho$. Otherwise, let M be the complementary component of $\rho \setminus \{a, b\} = \rho' \setminus \{a', b\}$ containing a, a', and b; $M' = g(M)$ is then the complementary component of $f\rho \setminus \{fa, fb\}$ containing

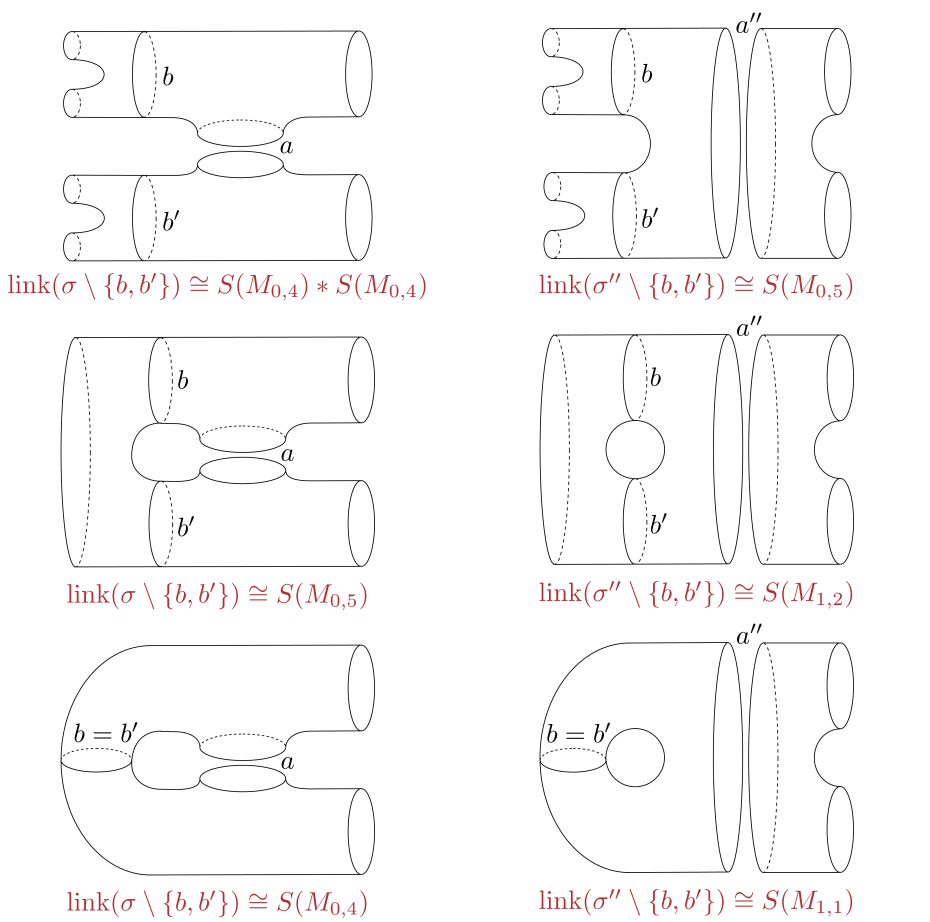

FIGURE 4. Various links in the proof of Lemma 3.5. Row 1 shows when e_a and e_b are not adjacent, Row 2 when e_a and e_b are incident on a single common vertex, and Row 3 when $b = b'$. Because a and a' both separate b and b' in M, the left column is identical replacing a with a', σ with σ', and b with b'.

fa, fa' and fb. Let $\mathrm{link}(\rho \setminus b) = \{b, b', b''\}$ and $\mathrm{link}(\rho' \setminus b) = \{b, b^\dagger, b^\ddagger\}$, all spheres of which are contained in M. We consider two cases (see Figure 5):

(i) *e_a and e_b are incident on two vertices* ($M \cong M_{1,2} \setminus \partial M_{1,2}$). Then b is non-separating in M, hence exactly one of $b^\dagger, b^\ddagger$ is separating in M, say $b^\dagger$. Thus $\mathrm{link}(b^\dagger \cup (\rho' \setminus \{a', b\})) \cong \mathcal{S}(M_{1,1})$ and $\mathrm{link}(b^\ddagger \cup (\rho' \setminus \{a', b\})) \cong \mathcal{S}(M_{0,4})$. The same holds for the f images of these sets, and since only $gb^\dagger$ is separating in M', likewise for their g_* images. Hence $g_* b^\dagger = fb^\dagger$ and $g_* b^\ddagger = fb^\ddagger$.

(ii) *e_a and e_b are incident on only one vertex* ($M \cong M_{0,5} \setminus \partial M_{0,5}$). Exactly one of $b^\dagger, b^\ddagger$ is disjoint from b', say $b^\dagger$. Then likewise $gb^\dagger$ is disjoint from gb' and $gb^\ddagger$ is not. Since f preserves disjointness and $g_* b' = fb'$, $g_* b^\dagger = fb^\dagger$ and $g_* b^\ddagger = fb^\ddagger$. $\square$

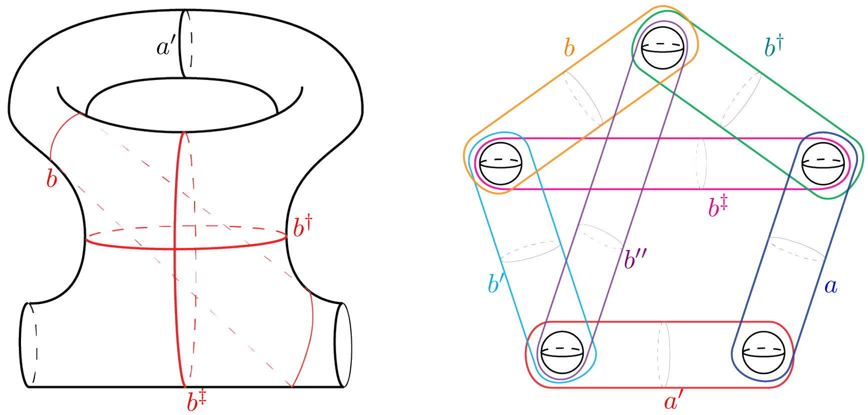

FIGURE 5. On the left is $M \cong M_{1,2} \setminus \partial M_{1,2}$ illustrating part (i). On the right is $M \cong M_{0,5}$ with the spheres in part (ii) drawn in.

Corollary 3.8. *If a diffeomorphism g agrees with f on X_σ, then $g_* = f$.*

Proof. Fix a compact exhaustion M_i of M_Γ such that $\partial M_i \subset \sigma$. For any $b \in \mathcal{S}(M_\Gamma)$, choose i such that $b \subset M_i$; let $\bar{\sigma} = \sigma \cap M_i \setminus \partial M_i$ and fix $\bar{\sigma}'$ a maximal sphere system in M_i containing b. By Proposition 2.23 there is a finite sequence of maximal sphere systems in M_i between $\bar{\sigma}$ and $\bar{\sigma}'$ by successive flip moves, which then extends to such a sequence between σ and $\sigma' = \bar{\sigma}' \cup (\sigma \setminus \bar{\sigma}) \ni b$ of maximal sphere systems in M_Γ. Inductively applying Lemma 3.7 shows that g agrees with f on $X_{\sigma'}$, and in particular $g_*(b) = fb$. As this is true for any b, it follows that $g_* = f$. □

Remark 3.9. It follows that g_* is uniquely determined by its restriction to X_σ among the class of isomorphisms induced by diffeomorphisms $M_\Gamma \to M_{\Gamma'}$. In particular, if g' is another diffeomorphism that agrees with g_* over X_σ, then $g'_* = g_*$.

We now restate and prove the main theorem of this paper:

Theorem 1.1. *Let Γ, Γ' be two locally finite connected graphs. Suppose $f : \mathcal{S}(M_\Gamma) \to \mathcal{S}(M_{\Gamma'})$ is an isomorphism. Then f is induced by a diffeomorphism $h : M_\Gamma \to M_{\Gamma'}$. In particular, (i) Γ and Γ' are proper homotopy equivalent and (ii) when Γ is not a graph of rank r with s rays such that $2r + s < 4$ or $(r, s) \in \{(0, 4), (2, 0)\}$, $\mathrm{Aut}(\mathcal{S}(M_\Gamma)) \cong \mathrm{Map}(\Gamma)$ as topological groups.*

Proof. From Lemma 3.5 and Corollary 3.8 we obtain that f is induced by a diffeomorphism, and in particular the action $\xi : \mathrm{Map}(M_\Gamma) \to \mathrm{Aut}(\mathcal{S}(M_\Gamma))$ is surjective. Thus M_Γ and $M_{\Gamma'}$ are diffeomorphic, hence $M_\Gamma, M_{\Gamma'}$ have the same characteristic triple and likewise do Γ, Γ'. By Theorem 2.1 Γ and Γ' are proper homotopy equivalent, which implies *(i)*.

Let $\rho : \mathrm{Map}(\Gamma) \to \mathrm{Aut}(\mathcal{S}(M_\Gamma))$ be the action induced by ξ in Theorem 2.5; note that ρ and ξ have the same image in $\mathrm{Aut}(\mathcal{S}(M_\Gamma))$, hence ρ is surjective. To show *(ii)*, suppose Γ is a graph satisfying the hypotheses of *(ii)*: by Theorem 2.5 ρ is injective, hence a group isomorphism. To obtain that ρ is a homeomorphism, we observe that the pullback of the permutation topology on $\mathrm{Aut}(\mathcal{S}(M_\Gamma))$ is compatible with the usual quotient topology on $\mathrm{Map}(\Gamma)$. In particular, by [Uda24, Prop. 7.1] the quotient topology is identical to the topology generated by the subbasis at identity $\{U'_a\}_{a\in\mathcal{S}(M_\Gamma)^{(0)}}$, where

$$U'_a = \rho^{-1}(U_a) = \{\phi \in \mathrm{Map}(\Gamma) : \rho(\phi)(a) = a\}$$

and $\{U_a\}$ is the subbasis for $\mathrm{Aut}(\mathcal{S}(M_\Gamma))$ given after Definition 2.4. □

Consequently, we obtain another proof of the following result, originally proven in Proposition 4.11 of [AB21].

Corollary 3.10. *For any locally finite connected graph* Γ*,* $\mathrm{Map}(\Gamma)$ *is Polish.*

Proof. It is a standard result of descriptive set theory that the automorphism group of a countable graph equipped with the permutation topology is Polish. Except for the finitely many cases excluded in Theorem 1.1(ii), $\mathrm{Map}(\Gamma)$ is topologically isomorphic to the automorphism group of a countable graph by Theorem 1.1(ii), so the result follows in these cases. For the excluded cases, $\mathrm{Map}(\Gamma)$ is countable and discrete, hence Polish. □

4. Rigidity of the sphere complex for finite-type doubled handlebodies with S^2-boundaries

The goal of this section is to establish Theorem 1.2. The setup and arguments closely follow [BL24a, Section 3], where analogous results are proven for doubled handlebodies with empty boundary. For completeness, we reintroduce their setup and highlight the adjustments needed to prove the results in the setting of doubled handlebodies with nonempty boundary. For clarity, we will match the notation conventions in [BL24a] throughout this section. We first construct a set X_0 that can be extended to a finite strongly rigid set X by adding a collection of "good spheres."

Throughout this section, we will implicitly use Proposition 2.9 to identify sphere complexes of submanifolds with subcomplexes of the sphere complex of a parent manifold. In particular, links of simplices will be identified with sphere complexes of submanifolds.

Given a subcomplex X of $\mathcal{S}(M_\Gamma)$, we say that X is **geometrically rigid** if for every simplicial locally injective map $f : X \to \mathcal{S}(M_\Gamma)$ there is a diffeomorphism h of M_Γ such that the restriction of h_* to X agrees with f. To prove Theorem 1.2, we first exhaust $\mathcal{S}(M_{n,s})$ by a sequence of geometrically rigid sets in Proposition 4.10, and then show that these sets are strongly rigid.

4.1. Constructing a geometrically rigid set X.

Let Y be a maximal collection of disjoint spheres $S_i \subset \mathcal{S}(M_{n,s})$ whose union is non-separating. Then, by removing a small regular neighborhood of Y from $M_{n,s}$, we obtain a manifold with no genus. Specifically,

$$N \coloneqq M_{n,s} \setminus \mathrm{nbd}(Y) \cong M_{0,2n+s}$$

Let Z be the collection of all spheres in $\mathcal{S}(M_{n,s})$ that are disjoint from Y. By construction of N, it follows $Z = \mathcal{S}(N)$.

The spheres in ∂N come in two types:

1) those coming from the original boundary of $M_{n,s}$;
2) pairs of spheres S_i^+, S_i^- coming from removing a sphere $S_i \in Y$.

We define a labeling map $\delta : \partial N \to Y \cup \partial M_{n,s}$ recording where boundary components of N came from. Specifically, for spheres $S_i^{\pm} \in \partial N$ coming from removing a sphere $S_i \in Y$, $\delta(S_i^{\pm}) = S_i$. On the other hand, for spheres originally in the boundary $S_j \subset \partial M_{n,s}$, we define $\delta(S_j) = S_j$. See Figure 6 for an illustration.

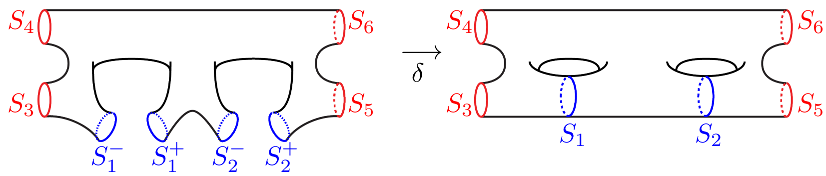

FIGURE 6. The handlebodies pictured above illustrate the process of removing a neighborhood of Y from $M_{n,s}$. The red spheres correspond to boundary components of the uncut manifold. The blue spheres play the role of Y.

We are interested in the subgraph of the 1-skeleton of $\mathcal{S}(M_{n,s})$ spanned by Y and Z, which we denote by

$$X_0 := \langle Y \cup Z \rangle.$$

Because X_0 is a join of Y and Z, and Y is complete, X_0 is not geometrically rigid. See Figure 7.

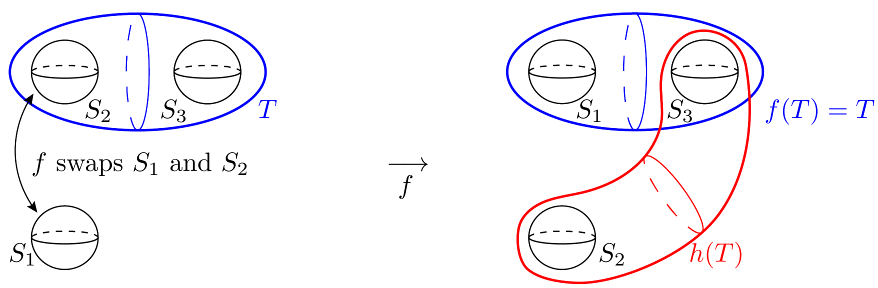

FIGURE 7. Let $S_1, S_2, S_3 \in Y$ and $T \in Z$ be as shown above. Consider the simplicial isomorphism $f\colon X \to X$ with $f|_Z = \mathrm{Id}_Z$, and $f|_Y$ permuting S_1 and S_2. Any homeomorphism inducing the permutation of S_1 and S_2 cannot fix T. So, f is not induced by any homeomorphism h.

Fix a locally injective simplicial map $f\colon X_0 \to \mathcal{S}(M_{n,s})$. Then intersecting pairs of spheres in Z have X_0-detectable intersection, so by Lemma 2.19 $f(Z)$ must fill a connected submanifold. As $f(Z)$ is disjoint from $f(Y)$, the submanifold filled by $f(Z)$ must lie in a component of $M_{n,s} \setminus \mathrm{nbd}(f(Y))$. By a complexity argument, this is only possible if $M_{n,s} \setminus \mathrm{nbd}(f(Y))$ is connected, and thus homeomorphic to $M_{0,2n+s}$. That is, if $M_{n,s} \setminus \mathrm{nbd}(f(Y))$ is disconnected, there is no way the manifold filled by $f(Z)$ could embed into any of the components of $M_{n,s} \setminus \mathrm{nbd}(f(Y))$, as can be seen by analyzing the possible components of $M_{n,s} \setminus \mathrm{nbd}(f(Y))$. By Lemma 2.21, there is a homeomorphism $h\colon N = M_{n,s} \setminus \mathrm{nbd}(Y) \to N' = M_{n,s} \setminus \mathrm{nbd}(f(Y))$ arising from $f|_{X_0}$, such that $h|_Z = f|_Z$.

To build a finite rigid set, we will add spheres to the set X_0, namely pairs of "good spheres." The addition of these good spheres will allow us to keep track of the pairs of boundary components of $M_{0,2n+s}$ that correspond to a component of Y. In particular, we will show that when X_0 is extended to a set X containing good pairs $h : N \to N'$ ascends to a homeomorphism $\hat{h}\colon M_{n,s} \to M_{n,s}$, which induces $f : X \to \mathcal{S}(M_{n,s})$.

We will now recall the definitions of a good sphere and a good pair.

Definition 4.1 ([BL24a, Section 3])**.** Given $A \subset Y$, let $A^+, A^- \subset \partial N$ be the boundary spheres obtained from removing a neighborhood of A from $M_{n,s}$; that is, $\delta(A^\pm) = A$. Let a be an essential sphere in $M_{n,s}$ that essentially intersects A in a single simple closed curve, and such that a is disjoint from all other spheres of Y.

When we descend to the cut manifold N, the sphere a decomposes as the union of the two disjoint disks a^+ and a^-, with $\partial a^+ \subset A^+$ and $\partial a^- \subset A^-$. The boundary of regular neighborhoods of $A^+ \cup a^+$ and $A^- \cup a^-$ are disjoint pairs of pants. See Figure 8 for an illustration of this setup.

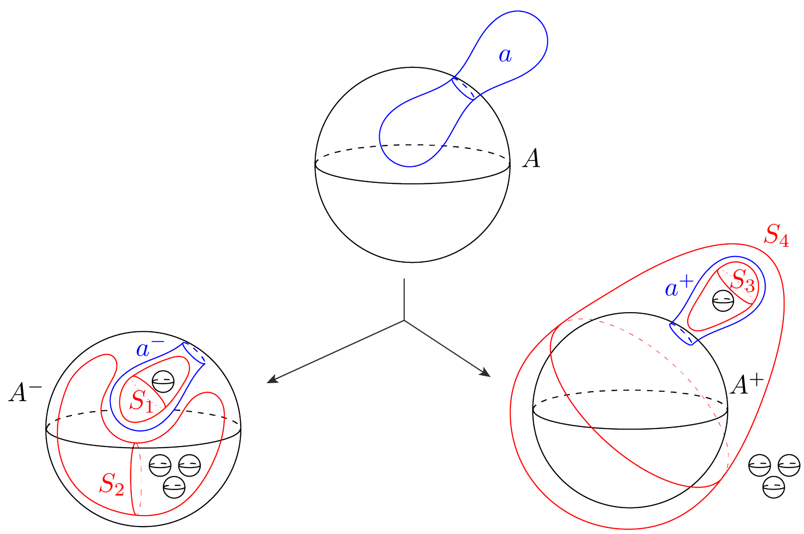

FIGURE 8. The sphere a is good for A. In the top picture, the spheres a and A are pictured in the uncut manifold $M_{n,s}$, and they intersect essentially. After cutting along Y, we obtain the picture in N below. The six spheres pictured inside of S_2 and outside of S_4 are meant to illustrate that these spheres are essential and not necessarily peripheral, unlike S_1 and S_3, which are peripheral since a is good.

We denote $\partial(\mathrm{nbd}(A^- \cup a^-)) = A^- \cup S_1 \cup S_2$ and $\partial(\mathrm{nbd}(A^+ \cup a^+)) = A^+ \cup S_3 \cup S_4$. Let $\partial(A, a) := A^+ \cup A^- \cup S_1 \cup S_2 \cup S_3 \cup S_4$. If S_1 and S_3 are peripheral in the cut manifold Y, we say that a is a **good sphere** for A. Suppose that a and a' are disjoint good spheres for A. Then if $\partial(A, a) \cap \partial(A, a') = A^+ \cup A^-$, then a and a' are said to be a **good pair** for A.

Good pairs can always be found when $2n + s \geq 6$, i.e., when ∂N has at least 6 components. This is because each sphere in a good pair for A requires the use of two boundary components of N other than $A^\pm$, and by assumption, these pairs of boundary components must be distinct for each sphere in the pair.

Let X be the subcomplex of $\mathcal{S}(M_{n,s})$ spanned by the vertices of X_0, together with a choice of a good pair for each sphere $A \subset Y$.

4.2. **Geometric rigidity of X for $M_{n,s}$.** Throughout the rest of Section 4, we fix a choice of $M_{n,s}$, with $2n + s \geq 6$.

Up to this point, the setup has been the same as in [BL24a, Section 3]. A key difference in the setting of $M_{n,s}$ when $s \neq 0$ is that the good pairs in X serve an additional role as discussed below in the proof of Proposition 4.2.

Proposition 4.2 ([BL24a, Lemma 13]). *Let $f \colon X \to \mathcal{S}(M_{n,s})$ be a locally injective, simplicial map. Let $h : N \to N'$ be the homeomorphism inducing $f|_Z$ as in Lemma 2.21. Then, for every sphere $A \in Y$ we have $\delta(h(A^{\pm})) = f(\delta(A^{\pm}))$ where A^+ and A^- denote the corresponding boundary spheres in N, i.e., $\delta(A^{\pm}) = A$.*

Proof. For all $A \in Y$, the spheres $A^{\pm}$ are not mapped by h to spheres $S \in \partial N'$ coming from the original boundary of $M_{n,s}$. This is true because the original boundary spheres of $M_{n,s}$ do not intersect any essential spheres in $M_{n,s}$. In particular, the boundary spheres do not admit any good spheres or good pairs. Thus, the good spheres in X allow us to differentiate between the boundary components of N that result from removing Y and the original boundary components of $M_{n,s}$. A priori, h could map a sphere coming from Y to an original boundary sphere of $M_{n,s}$. However, following the argument in [BL24a, Lemma 13] this cannot happen. Given this additional role of the good pairs, the proof of [BL24a, Lemma 13] applies in the more general setting of $M_{n,s}$ when $s \neq 0$. □

Proposition 4.3 ([BL24a, Proposition 14]). *The set X is geometrically rigid.*

Proof. Proposition 4.2 shows that $h : N \to N'$ ascends to a map $\hat{h} : M_{n,s} \to M_{n,s}$ such that $\hat{h}_*$ and f agree on X_0. It remains to verify that $\hat{h}_*$ and f agree on the good spheres in X. The argument proceeds exactly as in [BL24a, Proposition 14]. □

4.3. **Exhaustion by geometrically rigid sets.** In this section, we generalize Proposition 22 of [BL24a] and prove that we can find a nested family of geometrically rigid sets that exhaust $\mathcal{S}(M_{n,s})$ (see Proposition 4.10). The proof will primarily follow the argument used in [BL24a], with two differences that will be explicitly stated.

Definition 4.4. Suppose P is a pants decomposition of $M_{n,s}$ and $a \in P$. A sphere $b \in \mathcal{S}(M_{n,s})$ is a **split sphere for** (a, P) if a is the unique sphere in P intersecting b.

Definition 4.5. If $X \subseteq \mathcal{S}(M_{n,s})$ is a subcomplex, $P \subseteq X^{(0)}$ and $b \in X^{(0)}$ is a split sphere for (a, P), then we say that P is $X-$split at a by b. We say that P is $X-$split if it is $X-$split at a for some $a \in P$. If X contains every split sphere for P, then we say it is fully $X-$split.

Definition 4.6. Suppose $X \subseteq \mathcal{S}(M_{n,s})$ is a subcomplex and $a \in X^{(0)}$. A pair of distinct, disjoint spheres $(b_1.b_2)$ in $\mathcal{S}(M_{n,s})^{(0)}$ is a **split pair** for a relative to X if there exists pants decompositions $P_1, P_2 \subseteq X^{(0)}$, both containing a such that b_i is a split sphere for (a, P_i) for $i = 1, 2$. See Figure 9.

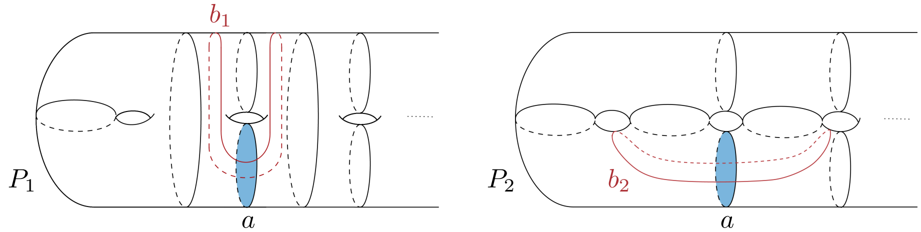

FIGURE 9. The spheres b_1 and b_2 give a split pair for a with respect to the pants decompositions P_1 and P_2.

Lemma 4.7 ([BL24a, Lemma 18])**.** *Suppose $X \subseteq \mathcal{S}(M_{n,s})$ is a geometrically rigid subcomplex and $a \in X^{(0)}$ has a split pair (b,c). Then the subcomplex $X_{b,c}$ induced by $X \cup \{b,c\}$ is geometrically rigid.*

The following is Lemma 20 of [BL24a]. The lemma as stated in [BL24a] claims that if a and c are two distinct adjacent spheres in a pants decomposition P of $M_{n,s}$, then the component of $M_{n,s} \setminus (P \setminus \{a,c\})$ is homeomorphic to $M_{0,5}$. However, consider Case 2 of Figure 10, where when a and c are the two nonperipheral spheres on the left, and the given component is homeomorphic to $M_{1,2}$.

Instead, we note the following lemma, which follows directly from their proof. We simply assume that a and c have the property that $M_{n,s} \setminus (P \setminus \{a,c\})$ is actually homeomorphic to $M_{0,5}$. Afterward, we note how to modify the proof of Lemma 21 of [BL24a], which is the only place where Lemma 20 is used in that paper.

Lemma 4.8 ([BL24a, Lemma 20])**.** *Suppose $X \subseteq \mathcal{S}(M_{n,s})$ is a subcomplex and $P \subseteq X^{(0)}$ is a pants decomposition that is $X-$split at a by $b \in X^{(0)}$. For every sphere $c \in P$ that is adjacent to a so that the component of $M_{n,s} \setminus (P \setminus \{a,c\})$ containing a and c is homeomorphic to $M_{0,5}$, there are spheres d and e such that (d,e) is a split pair for c.*

We now discuss Lemma 21 of [BL24a] and note the small difference that has to be made in the proof.

Lemma 4.9 ([BL24a, Lemma 21])**.** *Suppose $X \subseteq \mathcal{S}(M_{n,s})$ is a finite geometrically rigid set and $P \subseteq X^{(0)}$ a pants decomposition is $X-$split. Then there is a finite geometrically rigid set $X^P \supseteq X$ so that P is fully X^P-split, that is, X^P contains every split sphere for P.*

Proof. We inductively define the sets $P_i \subset P$ as follows. Suppose $a_0 \in P$ is X-split. Define $P_0 = \{a_0\}$, and for $i \geq 1$ let

$$P_i = \left\{ s \in P \;\middle|\; \begin{array}{l} s \text{ is adjacent to } a \in P_{i-1} \text{ and the component of } M_{n,s} \setminus (P \setminus \{a,s\}) \\ \text{containing } a, s \text{ is homeomorphic to } M_{0,5} \end{array} \right\}.$$

These sets differ from the sets denoted P_i in the proof of Lemma 21 of [BL24a], as there it is not assumed that the component of $M_{n,s} \setminus (P \setminus \{a,s\})$ is homeomorphic to $M_{0,5}$ (as they implicitly assume this). Even so, just as in their proof, there is a k so that $\cup_{i=1}^{k} P_i$ contains all the spheres in P which have a split sphere (i.e., are not self-adjacent). To see this, consider the dual graph to P, which has a vertex for every component of $M_{n,s} \setminus P$, and an edge connecting two vertices if they share a common component of P. This graph is connected, and the subgraph spanned by the edges whose corresponding spheres are not self-adjacent is connected (as such spheres correspond to the non-loop edges of the graph).

Then, starting from the edge corresponding to a_0, one can reach any other sphere by only going along edges corresponding to spheres a and s so that the component of $M_{n,s}\backslash(P\backslash\{a,s\})$ containing a and s is homeomorphic to $M_{0,5}$. To see this, note that every edge which is not a loop contains a vertex of valence 2, or a vertex of valence 3 that is not incident to a loop (potentially one of both). If a vertex v is valence 2 and s_1 and s_2 are the two spheres in P that correspond to the edges containing v, then it is easy to see that the component of $M_{n,s} \setminus (P \setminus \{s_1, s_2\})$ containing s_1 and s_2 is homeomorphic to $M_{0,5}$. Thus if $s_1 \in P_{i-1}$, then $s_2 \in P_i$. If v has valence 3 and is incident to no loops, with edges corresponding to spheres s_1, s_2, and s_3, then one can see that one of these spheres, say s_3, is such that, for $i = 1, 2$, the component of $M_{n,s} \setminus (P \setminus \{s_i, s_3\})$ containing s_i and s_3 is homeomorphic to $M_{0,5}$. In this case, one needs to explicitly use the fact that $M_{2,0}$ has been excluded, see Case 5 in Figure 10 below. Thus if s_1 is in P_{i-1}, then $s_3 \in P_i$ and $s_2 \in P_{i+1}$, and similarly if $s_2 \in P_{i-1}$. If $s_3 \in P_{i-1}$, both s_1 and s_2 are in P_i.

Figure 10 depicts all the cases for v and helps illustrate the above argument. The important cases for the choice of vertex v are Case 4, which illustrates the valence 2 case, and Cases 10 and 11, which illustrate the valence 3 case.

The proof then proceeds identically as in [BL24a], as the connectivity argument above shows that every non-self-adjacent sphere in P will eventually be contained in some P_i. □

We now note the main proposition, which is the key ingredient for proving Theorem 1.2. The proof is nearly identical, but we will modify the start of the argument slightly so that it makes sense in the generality we are working in.

Proposition 4.10 ([BL24a, Proposition 22])**.** *There exists a nested family of finite geometrically rigid sets $X_j \subseteq \mathcal{S}(M_{n,s})$ such that*

$$\mathcal{S}(M_{n,s}) = \bigcup_j X_j$$

Proof. Let X be the strongly rigid set constructed before Proposition 4.2. By construction, X contains a pants decomposition P_0 which is X-split.

We define a sequence of collections of pants decompositions of $M_{n,s}$ as follows. Begin with $\mathcal{P}_0 = \{P_0\}$, and inductively define

$$\mathcal{P}_i = \{P \text{ a pants decomposition} \mid \text{there is a } P' \in \mathcal{P}_{i-1} \text{ such that } |P \Delta P'| = 2\}.$$

For every $P \in \mathcal{P}_i$, there is a $P' \in \mathcal{P}_{i-1}$ so that P is obtained from P' by exchanging split spheres. By Proposition 2.23, every pants decomposition can be reached from P_0 by applying these exchanges. Thus, every pants decomposition of $M_{n,s}$ is contained in some $\mathcal{P}_i$. The proof then proceeds in exactly the same way as in Proposition 22 of [BL24a]. □

To finish the proof of Theorem 1.2, it suffices to show that each X_j in Proposition 4.10 is strongly rigid. To do this, we summarize the corresponding argument in [BL24a]. The proofs of the next lemma and its corollary follow exactly as in [BL24a].

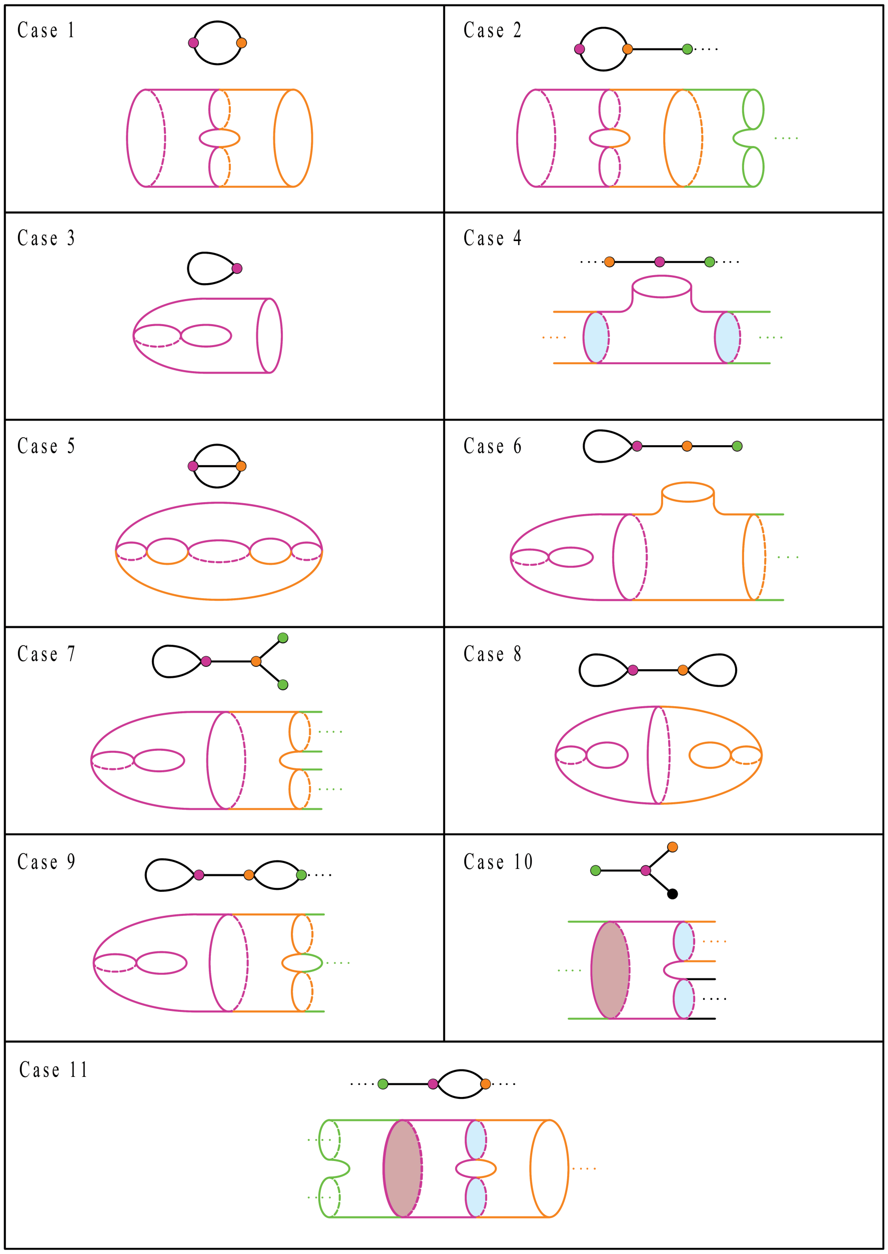

FIGURE 10. In all the cases, the vertex v is the pink vertex. In cases 1-4, v has valence 2, and the blue spheres denote s_1 and s_2. In cases 5-11, v has valence three, the blue spheres denote s_1 and s_2 while the red sphere denotes s_3. The cases without these spheres colored in don't apply, either because the pink pair of pants has a self-adjacent sphere (Cases 3, 6, 7, 8 and 9) or because we may assume the complexity is high enough (Cases 1, 3, 5 and 8).

Lemma 4.11 ([BL24a, Lemma 23]). *Suppose $h \in \mathrm{Map}(M_{n,s})$ is the point-wise stabilizer of a pants decomposition P so that every sphere is the boundary of two distinct complementary components of P. Then h induces the identity map on $\mathcal{S}(M_{n,s})$.*

Corollary 4.12 ([BL24a, Corollary 24]). *If $X \subset \mathcal{S}(M_{n,s})$ is geometrically rigid and contains a pants decomposition P so that every sphere is the boundary of two distinct complementary components of P, then it is uniquely geometrically rigid in the sense that any mapping class fixing X pointwise induces the identity on $\mathcal{S}(M_{n,s})$.*

Proof of Theorem 1.2. The proof of Theorem 25 of [BL24a] generalizes to imply 1.1(ii) a graph of rank n with s ends, where $2n+s \geq 6$ (using Theorem 2.5 in place of Laudenbach's result in [Lau73]).

In particular, if X_j is a set as in Proposition 4.10 and $f : X_j \to \mathcal{S}(M_{n,s})$ is a simplicial locally injective map, then by Proposition 4.10 there is a diffeomorphism h of $M_{n,s}$ inducing f. By Corollary 4.12, as each X_j contains a pants decomposition as required, any other diffeomorphism inducing f agrees with h on $\mathcal{S}(M_{n,s})$. In particular, h_* is an automorphism of $\mathcal{S}(M_{n,s})$ that agrees with f on X_j. It is the only such automorphism since there is some diffeomorphism inducing every automorphism, and we just showed that any other diffeomorphism inducing f agrees with h on $\mathcal{S}(M_{n,s})$. Thus, X_j is strongly rigid. □

5. Another proof of Theorem 1.1

In this section, we utilize the results in Section 4 along with an argument analogous to one that appears in [BDR20] to give another proof of Theorem 1.1.

We first show that the sphere complex uniquely determines finite-rank doubled handlebodies. This result follows from those in Section 3, but we give another proof of it here to show the proof of Theorem 1.1 that appears here is independent of the proofs in Section 3.

Proposition 5.1. *Suppose $\phi : \mathcal{S}(M_{n,s}) \to \mathcal{S}(M_{m,r})$ is an isomorphism between the sphere graphs of two handlebodies of finite type. Then $n = m$ and $s = r$.*

Proof. Note that the dimension of the largest simplex of the two sphere complexes must be the same, which puts a dimensional restriction on when the two graphs can be isomorphic. Thus, if $s, r \in \{0, 1, 2\}$, the result follows, as the maximal simplex dimension of $\mathcal{S}(M_{n,s})$ is $3n+s-3$, and similarly it is $3m+r-3$ for $\mathcal{S}(M_{m,r})$. For these to be equal, $s-r$ must be a multiple of 3, which means $s = r$ in this case. Thus, $n = m$ as well.

Now suppose we have an isomorphism when $m > n$ and $s \geq 3$. Then it follows that $s > r$. In $\mathcal{S}(M_{n,s})$ there is a sphere S cutting off a homeomorphic copy of $M_{0,s+1}$. In particular, by Proposition 5.3, as the equivalence classes of the link must be preserved by the isomorphism, $\phi(S)$ must also be a separating sphere, and one of its complementary components is homeomorphic to $M_{0,s+1}$. This is because the only doubled handlebodies with boundary that have a finite sphere graph are $M_{1,1}$ and $M_{0,p}$, and the only time when they have the same number of spheres as $M_{0,s+1}$ is when $p = s+1$ ($\mathcal{S}(M_{1,1})$ only has a single sphere). But this is impossible as $s > r$, so no such spheres exist in $M_{m,r}$. Hence, $m = n$, and thus $s = r$, finishing the proof. □

5.1. **Links of sphere systems.** In analogy with [BDR20], we first define an equivalence relation on $\mathrm{link}(\sigma)$ for σ a sphere system in $\mathcal{S}(M_\Gamma)$.

Definition 5.2. Let $a, b \in \mathrm{link}(\sigma)$. Then $a \sim b$ if and only if there exists $c \in \mathrm{link}(\sigma)$ non-adjacent to both a, b.

Remark. Since $\mathrm{link}(\sigma)$ is loop-free, in particular $a \sim a$. Likewise, if a, b are non-adjacent, then a is non-adjacent to both and $a \sim b$.

While defined combinatorially, $\sim$ may be characterized topologically. In particular, the following shows that $\sim$ is an equivalence relation.

Proposition 5.3. *Let $a, b \in \mathrm{link}(\sigma)$. Then $a \sim b$ if and only if a, b lie in the same complementary component of σ.*

Proof. To prove the forward direction, note that if $c \in \mathrm{link}(\sigma)$ is not adjacent to a, b, then $a \cup b \cup c$ is connected and disjoint from σ, hence is contained in a single connected component.

Conversely, suppose that $a, b \subset M$ for $M \in \pi_0(M_\Gamma \setminus \sigma)$. We assume a, b are disjoint and distinct, else $a \sim b$ by the remark above; by Remark 2.13, a, b are essential in M. Suppose first that M has at least two distinct punctures x, y. By Lemma 2.8, it suffices to find a simple arc $\gamma \subset M$ between x, y that intersects a, b essentially: the sphere c bounding a regular neighborhood of γ intersects a, b essentially in M, thus by Proposition 2.9 essentially in M_Γ. Thus c is essential in M_Γ and disjoint from σ, hence $c \in \mathrm{link}(\sigma)$ and $a \sim b$. We consider four cases:

(i) *$a \cup b$ is non-separating.* We may choose γ to intersect a, b each exactly once.

(ii) *$a \cup b$ is separating but a, b are non-separating.* We may choose γ intersecting a, b each exactly once if x, y are not separated by $a \cup b$; else, choose γ to intersect b once and a with signed intersection ± 2.

(iii) *Only a is separating.* If a separates some punctures of M, then replace x, y such that they are separated by a and fix γ to intersect b exactly once. Else, the complementary component M' of a not containing x, y must have genus; if $b \not\subset M'$, fix a non-separating sphere $d \subset M'$ and choose γ to intersect b and d (if defined) each exactly once.

(iv) *a, b are separating.* Let U, V, W be components of $M \setminus (a \cup b)$, where a separates U and $V \cup W$ and b separates $U \cup V$ and W. Let P denote the set of punctures of M. If U has no punctures in P, then it has genus, hence fix a non-separating sphere $X \subset U$, and likewise if W has no punctures in P fix a non-separating sphere $Y \subset W$. Replace $x \in P$ to be a puncture in U, if one exists, and $y \in P$ to be a puncture in W, if one exists, or else some puncture distinct from x. If defined, let γ intersect X, Y each exactly once.

M is obtained by removing σ from M_Γ, hence must have at least one puncture. If M has exactly one puncture, then we may assume M has genus at least 2; otherwise, M has at most one distinct essential sphere, and the statement follows. Hence there exists an essential sphere q distinct from a, b and non-separating in M; we note that $M' = M \setminus q$ has three punctures. Replacing σ with $\sigma' = \sigma \cup \{q\}$ and M with M', the argument above obtains $c \in \mathrm{link}(\sigma') \subset \mathrm{link}(\sigma)$ non-adjacent to a, b. ∎

Let $\mathrm{link}(\sigma)|_{[a]}$ denote the full subcomplex of $\mathrm{link}(\sigma)$ (or equivalently of $\mathcal{S}(M_\Gamma)$, since it is flag) induced by the equivalence class of a in $\mathrm{link}(\sigma)$. By Proposition 5.3,

$[a]$ is exactly the spheres in $\mathrm{link}(\sigma)$ in the same component $M \subset M_\Gamma \setminus \sigma$ as a. From Proposition 2.9 and Remark 2.13 we then have the following:

Corollary 5.4. *Let $\sigma \subset \mathcal{S}(M_\Gamma)$ and $a \in \mathrm{link}(\sigma)$. Let M be the component of $M_\Gamma \setminus \sigma$ containing a. Then the equivalence class $[a]$ is the set of all essential embedded 2-spheres in M, and $\mathrm{link}(\sigma)|_{[a]} \cong \mathcal{S}(M)$.* □

5.2. **Diffeomorphisms from automorphisms.** The following theorem will be the main result to proving Theorem 1.1 in the general case. It is analogous to the work in Section 3 to construct a diffeomorphism inducing an isomorphism between sphere graphs, but the method in which we build the diffeomorphism differs from what is in Section 3. The proof below is inspired by the proof of [BDR20, Theorem 1.3], and for the most part, follows it very closely.

Theorem 5.5. *Let Γ and Γ' be two locally finite, infinite graphs, with associated 3-manifolds M_Γ and $M_{\Gamma'}$. Every simplicial isomorphism of the sphere complexes $\mathcal{S}(M_\Gamma) \to \mathcal{S}(M'_\Gamma)$ is induced by a homeomorphism $M_\Gamma \to M_{\Gamma'}$.*

Proof. Fix a simplicial isomorphism $\Psi\colon \mathcal{S}(M_\Gamma) \to \mathcal{S}(M'_\Gamma)$. Fix a compact exhaustion $K_1 \subset K_2 \subset \dots$ of M_Γ such that all components $M_\Gamma \setminus K_i$ are infinite-type and K_i is homeomorphic to M_{n_i,s_i} for $2n_i + s_i \geq 6$. By enlarging each K_i if necessary, we may assume that the boundary spheres a_i of K_i are essential embedded spheres of K_{i+1}.

Let E_i denote the equivalence class of $\mathrm{link}(a_i)$ that contains essential embedded 2-spheres in K_i. The equivalence class E_i is distinguishable combinatorially from the other equivalence classes of $\mathrm{link}(a_i)$ since it is the unique class with a finite clique number. This is because K_i is finite-type and, therefore, the spheres in a pants decomposition of Y_i for a finite maximal clique. Since connected components of $M_\Gamma \setminus Y_i$ are of infinite type, their associated equivalence classes in $\mathrm{link}(a_i)$ have unbounded clique numbers.

In particular, as the restriction of Ψ from $\mathrm{link}(a_i)$ to $\mathrm{link}(\Psi(a_i))$ sends equivalence classes to equivalence classes (as the classes are defined combinatorially), it follows from the discussion in the previous paragraph that:

- $\mathrm{link}(\Psi(a_i))$ contains exactly one equivalence class E'_i which corresponds to a compact doubled handlebody $K'_i \subset M_{\Gamma'}$ whose boundary is $\Psi(a_i)$, and
- Ψ restricts to an isomorphism $\Psi_i : \mathcal{S}(K_i) \to \mathcal{S}(K'_i)$.

By Proposition 5.1, $K_i \cong K'_i$. Recall that the proof of Theorem 1.2 showed that every element of $\mathrm{Aut}(\mathcal{S}(M_{n_i,s_i}))$ is induced by a diffeomorphism. It follows that Ψ_i is induced by a diffeomorphism $\phi_i : K_i \to K'_i$. The ambiguity on the choice of ϕ_i is only up to a product of sphere twists. By possibly modifying each ϕ_i by sphere twists, we may thus assume that these diffeomorphisms are compatible in the sense that $\phi_{i+1}(K_i) = K'_i$ and the restriction of ϕ_{i+1} to K_i agrees with ϕ_i. In particular, the direct limit of this sequence of diffeomorphisms induces a diffeomorphism $\phi : M_\Gamma \to M_{\Gamma'}$ which, by construction, induces the isomorphism Ψ. □

To finish the proof of Theorem 1.1, one can follow the proof in Section 3, utilizing Theorem 5.5 in place of Lemma 3.5 and Corollary 3.8.

6. Locally finite strongly rigid sets

Let M_Γ be the doubled handlebody associated to a locally finite graph Γ. We would like to extend to $\mathcal{S}(M_\Gamma)$ the results of Section 4 and construct an exhaustion by (the appropriate generalization of) finite strongly rigid sets.

If Γ is infinite-type, then M_Γ does not admit a finite strongly rigid set. Indeed, a finite set $X_0 \subset \mathcal{S}(M_\Gamma)^{(0)}$ has compact union $K = \bigcup_{a \in X_0} a \subset M_\Gamma$. There exists a diffeomorphism $h : M_\Gamma \to M_\Gamma$ acting non-trivially on $\mathcal{S}(M_\Gamma)$ but supported disjointly from K, hence fixing X_0. The identity automorphism and $h_* \neq \mathrm{id}$ both restrict to the inclusion map of X_0 into $\mathcal{S}(M_\Gamma)$, hence $X_0 \hookrightarrow \mathcal{S}(M_\Gamma)$ does not extend to a unique automorphism of $\mathcal{S}(M_\Gamma)$. We instead consider the rigidity of subcomplexes that satisfy a local version of finiteness.

Definition 6.1. A subcomplex $X \subset \mathcal{S}(M_\Gamma)$ is **topologically locally finite** if every compact $K \subset M_\Gamma$ essentially intersects finitely many components of $X^{(0)}$.

In fact, for all Γ of infinite-type there exist *no* strongly rigid sets in $\mathcal{S}(M_\Gamma)$, as we describe in Section 6.2. Nonetheless, we will construct subcomplexes which always admit unique isomorphism extensions of locally injective maps, provided those maps also preserve maximal sphere systems.

Definition 6.2. A simplicial map $f : X \subset \mathcal{S}(M_\Gamma) \to \mathcal{S}(M_{\Gamma'})$ is **maximal** if for any sphere system $\sigma \subset X$ that is maximal in M_Γ, $f\sigma$ is maximal in $M_{\Gamma'}$.

This condition is necessary in the following sense: if $\sigma \subset \mathcal{S}(M_\Gamma)$ is a maximal sphere system σ and a map $f : \sigma \to \mathcal{S}(M_\Gamma)$ extends to an automorphism, then $f\sigma$ must also be maximal. If $X \subset \mathcal{S}(M_\Gamma)$ is rigid, then any locally injective map $X \to \mathcal{S}(M_\Gamma)$ extends to an automorphism, hence is maximal.

A subcomplex $X \subset \mathcal{S}(M_\Gamma)$ for which any maximal locally injective map $f : X \to \mathcal{S}(M_\Gamma)$ extends to a unique automorphism is **strongly rigid over maximal maps**. For brevity, by **locally finite strongly rigid set**, we will mean a topologically locally finite subcomplex that is strongly rigid over maximal maps. Note that when Γ has infinite type, such subcomplexes are not locally finite in the usual simplicial sense. In Section 6.1, we prove the following, which implies Theorem 1.3:

Theorem 6.3. *Suppose that Γ is connected and $\dim(\mathcal{S}(M_\Gamma)) \geq 4$, and in particular if Γ is infinite-type. Then $\mathcal{S}(M_\Gamma)$ is covered by nested locally finite strongly rigid sets Z_i.*

6.1. Constructing locally finite strongly rigid sets.

We will need the following simplicial non-embeddings:

Lemma 6.4.

(i) $\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})$ does not embed in $\mathcal{S}(M_{0,5})$ and vice versa.
(ii) $\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})$ and $\mathcal{S}(M_{0,5})$ do not embed in $\mathcal{S}(M_{1,2})$.
(iii) $\mathcal{S}(M_{0,6})$ does not embed in $\mathcal{S}(M_{1,3})$.

Proof.

(i) Suppose $\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4}) \cong K_{3,3}$ can be embedded via a map f into $\mathcal{S}(M_{0,5})$, which is isomorphic to the Petersen graph (see Figure 11). Consider the vertex $v = (1,2)$ in $K_{3,3}$. Without loss in generality, we may assume $f(v) = (1,4)$. Since both v and $f(v)$ have valence 3, we know that

$$\{(1',2'),(1',3'),(2',3')\} \overset{f}{\mapsto} \{(3,5),(3,2),(2,5)\}$$

Again, without loss in generality, we may assume that $f(1', 2') = (3, 5)$. Since $(1, 3)$ and $(2, 3)$ are adjacent to $(1', 2')$ in $K_{3,3}$,

$$\{(1, 3), (2, 3)\} \overset{f}{\mapsto} \{(2, 4), (1, 2)\}$$

This is not possible since neither $(2, 4)$ nor $(1, 2)$ are adjacent to $(3, 2)$ in the Petersen graph.

On the other hand, $\mathcal{S}(M_{0,5})$ has $\binom{5}{2} = 10$ vertices, while $\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})$ has only 6 vertices, so $\mathcal{S}(M_{0,5})$ cannot embed into $\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})$.

(ii) The graph $\mathcal{S}(M_{1,2})$ a tree (see Figure 16). As seen in Figure 11, both $\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})$ and $\mathcal{S}(M_{0,5})$ admit embedded loops, and thus cannot be embedded into a tree.

(iii) The complex $\mathcal{S}(M_{1,3})$ is contractible [HV04]. On the other hand, $\mathcal{S}(M_{0,6})$ is a two-dimensional simplicial complex with 25 vertices, 105 edges, and 105 faces, which means its Euler characteristic is 25. Since the Euler characteristic is greater than two, $H_2(\mathcal{S}(M_{0,6})) \neq 0$. Suppose that $\mathcal{S}(M_{0,6})$ embeds into $\mathcal{S}(M_{1,3})$. Then the pair $(\mathcal{S}(M_{1,3}), \mathcal{S}(M_{0,6}))$ gives the exact sequence

$$H_3(\mathcal{S}(M_{1,3}), \mathcal{S}(M_{0,6})) \to H_2(\mathcal{S}(M_{0,6})) \to H_2(\mathcal{S}(M_{1,3}))\,.$$

Since $\mathcal{S}(M_{1,3})$ is 2-dimensional the first group vanishes and $H_2(\mathcal{S}(M_{0,6})) \hookrightarrow H_2(\mathcal{S}(M_{1,3})) = 0$, a contradiction. $\square$

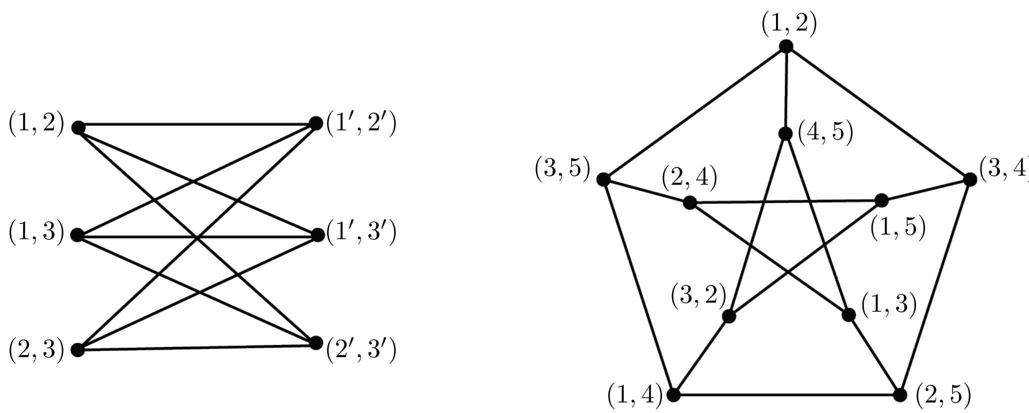

FIGURE 11. On the left is $\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4}) \cong K_{3,3}$. With the boundary components of each $\partial M_{0,4}$ labeled 1–4 and $1'$–$4'$, the pair (i, j) determines the sphere by Lemma 2.20. Similarly, for $\mathcal{S}(M_{0,5})$, which is pictured on the right and is isomorphic to the Petersen graph.

Recall that $X_\eta \coloneqq \bigcup_{a \in \eta} \operatorname{link}(\eta \setminus a)$ for $\eta \subset \mathcal{S}(M_\Gamma)$ a maximal sphere system.

Lemma 6.5. *Let Γ be connected and let $\eta \subset \mathcal{S}(M_\Gamma)$ be a maximal sphere system such that $|\eta| \geq 4$. Suppose $Y_\eta \subset \mathcal{S}(M_\Gamma)$ is a full subcomplex such that*

(1) $X_\eta \subset Y_\eta$,

(2) for $a \neq b \subset \eta$, $\operatorname{link}(\eta \setminus \{a, b\}) \subset Y_\eta$, unless $e_a \cup e_b$ has non-zero rank in Δ_η, in which case Y_η contains a subset of $\operatorname{link}(\eta \setminus \{a, b\})$ of size at least 11, and

(3) for $a, b, c \subset \eta$, if $e_a \cup e_b \cup e_c \cong K_{1,3}$ then $\operatorname{link}(\eta \setminus \{a, b, c\}) \subset Y_\eta$; if $e_a \cup e_b \cup e_c \cong K_3$ then Y_η contains a subset of $\operatorname{link}(\eta \setminus \{a, b, c\})$ of size at least 26.

If $f : Y_\eta \to \mathcal{S}(M_{\Gamma'})$ is a locally injective simplicial map for which $f\eta$ is maximal, then $M_\Gamma \cong M_{\Gamma'}$.

Proof. It suffices to find a maximal sphere system $\eta' \subset \mathcal{S}(M_{\Gamma'})$ such that $\Delta_\eta \cong \Delta_{\eta'}$. Since $|\eta| \geq 4$, f is injective on the sets $Y_\eta \cap \mathrm{link}(\eta\setminus a) = \mathrm{link}(\eta \setminus a)$, $Y_\eta \cap \mathrm{link}(\eta\setminus\{a,b\})$, and $Y_\eta \cap \mathrm{link}(\eta \setminus \{a,b,c\})$, as each of these sets is contained in the star of some element of η.

First, assume that Δ_η is loop-free and $a \in \eta$. If e_{fa} is a loop in $\Delta_{f\eta}$ then $\mathrm{link}(\eta \setminus a) \cong \mathcal{S}(M_{0,4})$ does not embed into $\mathrm{link}(f\eta \setminus fa) \cong \mathcal{S}(M_{1,1})$. Thus, $\Delta_{f\eta}$ is also loop-free. Now let $a, b \in \eta$. The edges e_a and e_b are disjoint in Δ_η if and only if $\mathrm{link}(\eta\setminus\{a,b\})$ is isomorphic to $\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})$ and e_a and e_b are incident on one common vertex or two common vertices if and only if $\mathrm{link}(\eta \setminus \{a,b\})$ is isomorphic to $\mathcal{S}(M_{0,5})$ or $\mathcal{S}(M_{1,2})$, respectively. An analogous statement holds for e_{fa} and e_{fb}.

The map f restricts to an embedding $Y_\eta \cap \mathrm{link}(\eta \setminus \{a,b\}) \hookrightarrow \mathrm{link}(f\eta \setminus \{fa, fb\})$. There are three possibilities for the domain of this embedding: $Y_\eta \cap \mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4}) = \mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})$, $Y_\eta \cap \mathcal{S}(M_{0,5}) = \mathcal{S}(M_{0,5})$, or $Y_\eta \cap \mathcal{S}(M_{1,2})$. We observe the following embeddings are not possible:

$$\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4}) \begin{array}{l} \not\hookrightarrow \mathcal{S}(M_{0,5}) \\ \not\hookrightarrow \mathcal{S}(M_{1,2}) \end{array} \quad (1) \qquad \mathcal{S}(M_{0,5}) \begin{array}{l} \not\hookrightarrow \mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4}) \\ \not\hookrightarrow \mathcal{S}(M_{1,2}) \end{array} \quad (2) \qquad Y_\eta \cap \mathcal{S}(M_{1,2}) \begin{array}{l} \not\hookrightarrow \mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4}) \\ \not\hookrightarrow \mathcal{S}(M_{0,5}) \end{array} \quad (3)$$

where (1) and (2) follow from Lemma 6.4, and (3) is true since $|Y_\eta \cap \mathcal{S}(M_{1,2})| \geq 11$, but $|\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})|, |\mathcal{S}(M_{0,5})| < 11$. The only remaining possibility is that all links (intersected with Y_η) embed into isomorphic copies of themselves. Hence it follows that $\mathrm{link}(\eta \setminus \{a,b\}) \cong \mathrm{link}(f\eta \setminus \{fa, fb\})$ and $e_a \cup e_b \cong e_{fa} \cup e_{fb}$, without loss of generality preserving order: *i.e.* via a graph isomorphism with $e_a \mapsto e_{fa}$ and $e_b \mapsto e_{fb}$.

Thus, $f|_\eta \colon \Delta_\eta \to \Delta_{f\eta}$ is an edge isomorphism. To obtain a graph isomorphism, we need only show that $f|_\eta$ does not have a $K_3, K_{1,3}$-pair. Now, $e_a \cup e_b \cup e_c \cong K_{1,3}$ if and only if $\mathrm{link}(\eta \setminus \{a,b,c\}) \cong \mathcal{S}(M_{0,6})$ and $e_a \cup e_b \cup e_c \cong K_3$ if and only if $\mathrm{link}(\eta \setminus \{a,b,c\}) \cong \mathcal{S}(M_{1,3})$ (see Figure 3) and likewise for $e_{fa} \cup e_{fb} \cup e_{fc}$ as well. The restriction $f : Y_\eta \cap \mathrm{link}(\eta \setminus \{a,b,c\}) \hookrightarrow \mathrm{link}(f\eta \setminus \{fa, fb, fc\})$, along with Lemma 6.4 and that $|\mathcal{S}(M_{0,6})^{(0)}| = 25$, implies $f|_\eta$ has no $K_3, K_{1,3}$-pairs. Therefore, $\Delta_\eta \cong \Delta_{f\eta}$ by Theorem 2.26 and Corollary 2.27.

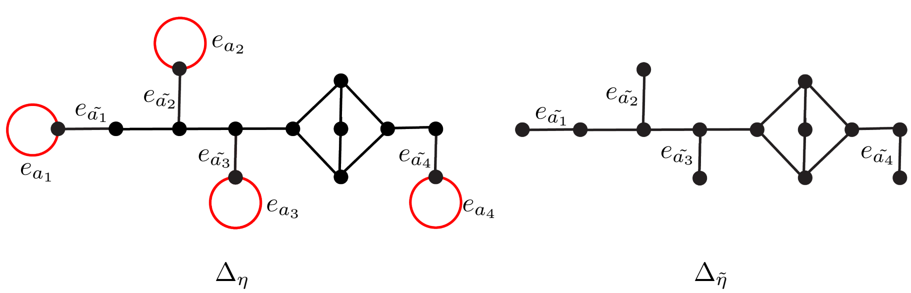

FIGURE 12. On the left the graph Δ_η has loops e_a incident to a unique edge $e_{\tilde{a}}$. On the right, $\Delta_{\tilde{\eta}}$ is the loop-free subgraph.

Now consider the case where Δ_η has loops. Since $|\eta| \geq 4$, no edge in Δ_η can be adjacent to two loops. Hence, if $a \in \eta$ and e_a is a loop, then we denote by $\tilde{a} \in \eta$ the unique sphere such that e_a is adjacent to $e_{\tilde{a}}$. Define $\tilde{\eta} = \eta \setminus \eta_0$, where $\eta_0 = \{a \in \eta \mid e_a \text{ is a loop}\}$. The subgraph $\Delta_{\tilde{\eta}} \subset \Delta_\eta$ is a graph without loops (see Figure 12) and $\tilde{\eta}$, $f\tilde{\eta}$ are maximal in $M_\Gamma \setminus \eta_0$ and $M_{\Gamma'} \setminus f\eta_0$ respectively. By the previous case, $\Delta_{\tilde{\eta}} \cong \Delta_{f\tilde{\eta}}$ via an isomorphism that induces $f|_{\tilde{\eta}}$ on the edges. Focusing on $a \in \eta_0$ (where e_a is a loop) we see that $e_{fa} \cup e_{f\tilde{a}}$ is isomorphic to either two disjoint edges, a loop disjoint from an edge, two edges incident on one common vertex, a bigon or $e_a \cup e_{\tilde{a}}$. In the first three cases, $\mathrm{link}(f\eta \setminus \{fa, f\tilde{a}\})$ is isomorphic to $\mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})$, $\mathcal{S}(M_{1,1}) * \mathcal{S}(M_{0,4})$ and $\mathcal{S}(M_{0,5})$ respectively. Since $|Y_\eta \cap \mathrm{link}(\eta \setminus \{a, \tilde{a}\})| \geq 11$, the map $f : Y_\eta \cap \mathrm{link}(\eta \setminus \{a, \tilde{a}\}) \hookrightarrow \mathrm{link}(f\eta \setminus \{fa, f\tilde{a}\})$ cannot be an embedding: $e_{fa} \cup e_{f\tilde{a}}$ is isomorphic to a bigon or $e_a \cup e_{\tilde{a}}$.

Let $\eta_1 \subset \eta_0$ be the subset of spheres $a \in \eta_0$ for which $e_{fa} \cup e_{f\tilde{a}}$ is a bigon. For each $a \in \eta_1$, we may perform a flip move $f\tilde{a} \mapsto \tilde{a}'$ such that $e_{fa} \cup e_{\tilde{a}'} \cong e_a \cup e_{\tilde{a}}$ (see Figure 13). Vertices in $\Delta_{f\eta}$ have valence at most 3, hence no two bigons are adjacent: we may perform these flips disjointly for each $a \in \eta_1$ to obtain a new pants decomposition η' for $M_{\Gamma'}$ with $\Delta_{\eta'} \cong \Delta_\eta$, as desired. ☐

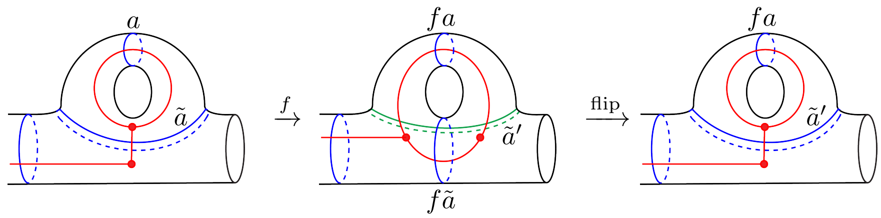

Figure 13. The flip move in Lemma 6.5.

Remark 6.6. If η is a finite maximal system and $|\eta| \geq 5$, then $M_\Gamma \cong M_{n,s} \setminus \partial M_{n,s}$ for some n, s and one verifies that $2n + s \geq 6$. By Theorem 1.2, there exists a finite strongly rigid set $Y_\eta \subset \mathcal{S}(M_\Gamma)$ satisfying the hypotheses of Lemma 6.5, in which case a locally injective map $Y_\eta \to \mathcal{S}(M_{\Gamma'})$ extends uniquely to an automorphism $\mathcal{S}(M_\Gamma) \to \mathcal{S}(M_{\Gamma'})$; by Theorem 1.1, this automorphism is induced by a diffeomorphism $M_\Gamma \to M_{\Gamma'}$.

Suppose that Γ is connected and of infinite or finite type such that $\dim(\mathcal{S}(M_\Gamma)) \geq 5$. Let $\sigma \subset \mathcal{S}(M_\Gamma)$ be a maximal sphere system and let Ω be the collection of sets of distinct spheres $\eta = \{a, b, c, d, e\} \subset \sigma$ such that $\bigcup_{t \in \eta} e_t \subset \Delta_\sigma$ is connected. For each $\eta \in \Omega$, denote by M_η the complementary component of $\sigma \setminus \eta$ that contains η (and in particular is not a pair of pants), and let $Y_\eta \subset \mathrm{link}(\sigma \setminus \eta) \cong \mathcal{S}(M_\eta)$ as in Remark 6.6. Define $Z_\sigma \subset \mathcal{S}(M_\Gamma)$ to be the full subcomplex induced by the Y_η:

$$Z_\sigma \coloneqq \langle Y_\eta \rangle_{\eta \in \Omega} .$$

Note that $\sigma \subset X_\sigma \subset Z_\sigma$, and that Z_σ is topologically locally finite by construction.

Proposition 6.7. *Suppose that $f : Z_\sigma \to \mathcal{S}(M_{\Gamma'})$ is a locally injective simplicial map such that $f\sigma$ is maximal. Then f extends to a unique automorphism $\mathcal{S}(M_\Gamma) \to \mathcal{S}(M_{\Gamma'})$.*

Proof. For $\eta \in \Omega$, let $M_{f\eta} \subset M_{\Gamma'}$ denote the complementary component of $f\sigma \setminus f\eta$ containing $f\eta$. By Remark 6.6, $f|_{Y_\eta}$ extends to a unique isomorphism $\mathrm{link}(\sigma \setminus \eta) \to$

$\mathrm{link}(f\sigma \setminus f\eta)$, which is induced by a diffeomorphism $h_\eta : M_\eta \to M_{f\eta}$ agreeing with f over $X_\eta \subset Y_\eta$.

We first prove that $f|_\sigma$ is an edge isomorphism $E(\Delta_\sigma) \to E(\Delta_{f\sigma})$ without a $K_3, K_{1,3}$-pair. If $a, b \in \sigma$ are such that e_a and e_b are adjacent in Δ_σ, fix $\eta \in \Omega$ such that $a, b \in \eta$. Then from the above we have $e_a \cup e_b \cong e_{h_\eta a} \cup e_{h_\eta b} = e_{fa} \cup e_{fb}$, preserving ordering. A similar argument implies that e_a is a loop if and only if e_{fa} is a loop, and prohibits a $K_3, K_{1,3}$-pair (*e.g.* fix $\eta \supset K_3$). It remains to show that when e_a, e_b are non-adjacent, so are e_{fa}, e_{fb}. Suppose that e_{fa}, e_{fb} are adjacent and e_a is a loop. As e_{fa} is a loop, e_{fb} is the unique sphere adjacent to e_{fa}. Let $c \in \sigma$ be a sphere so that e_c is adjacent to e_a in Δ_σ. Since e_c and e_a are adjacent, $e_c \cup e_a \cong e_{fc} \cup e_{fa}$ and e_{fc}, e_{fa} are adjacent, hence $fc = fb$; by injectivity $b = c$ and e_a, e_b are adjacent. Finally, if neither e_a, e_b are loops and e_a, e_b are non-adjacent, then, since $X_\sigma \subset Z_\sigma$, both $\mathrm{link}(\sigma \setminus a) \subset Z_\sigma$ and $\mathrm{link}(\sigma \setminus b) \subset Z_\sigma$. Consequently $\mathrm{link}(\sigma \setminus a) * \mathrm{link}(\sigma \setminus b) = \mathrm{link}(\sigma \setminus \{a, b\}) \subset Z_\sigma$. Now, as in the proof of Lemma 6.5, we know that $\mathrm{link}(\sigma \setminus b) \cong \mathcal{S}(M_{0,4}) * \mathcal{S}(M_{0,4})$ does not embed into $\mathrm{link}(f\sigma \setminus \{fa, fb\})$ if e_{fa} and e_{fb} are adjacent. Hence e_{fa} and e_{fb} must be non-adjacent.

By Theorem 2.26 and Corollary 2.27, $f|_\sigma$ is induced by an isomorphism $\Delta_\sigma \to \Delta_{f\sigma}$. Let h be a diffeomorphism constructed as in Section 3.2; we show that h_* extends f. For any $a, b, c \subset \sigma$ such that $e_a \cup e_b \cup e_c \subset \Delta_\sigma$ is connected, fix $\rho \in \Omega$ containing a, b, c. Then f extends to an isomorphism $\mathrm{link}(\sigma \setminus \rho) \to \mathrm{link}(f\sigma \setminus f\rho)$, which suffices to apply Remark 3.6: h_* agrees with f over X_σ. For $\eta \in \Omega$, h_* agrees with f on $X_\eta = X_\sigma \cap \mathrm{link}(\sigma \setminus \eta)$, hence by Remark 3.9 $(h|_{M_\eta})_* = (h_\eta)_*$. Since $(h_\eta)_*$ agrees with f over Y_η, we conclude. Finally, to show uniqueness, by Theorem 1.1 any isomorphism extending f is induced by a diffeomorphism g that agrees with h_* over X_σ: $g_* = h_*$ again by Remark 3.9. □

If σ, σ' are maximal sphere systems that differ by a flip move, then $Z_\sigma \cup Z_{\sigma'}$ exhibits the same rigidity. Suppose that $f : Z_\sigma \cup Z_{\sigma'} \to \mathcal{S}(M_{\Gamma'})$ is a locally injective simplicial map and $f\sigma, f\sigma'$ are maximal. Apply Lemma 3.7: the isomorphism extensions with respect to σ, σ' are identical. More generally, let $\mathcal{P}$ be a finite collection of maximal sphere systems for which any two $\rho, \rho' \in \mathcal{P}$ differ by a sequence of successive flip moves in $\mathcal{P}$. Then $Z_{\mathcal{P}} \coloneqq \bigcup_{\rho \in \mathcal{P}} Z_\rho$ is strongly rigid over maps f which preserve the maximality of all $\rho \in \mathcal{P}$. Let $\mathcal{P}_i$ be a nested family of such sets such that every sphere in $\mathcal{S}(M_\Gamma)$ is contained in some $\rho \in \mathcal{P}_i$ for some i. Then $Z_{\mathcal{P}_i}$ is a (nested) exhaustion of $\mathcal{S}(M_\Gamma)$ by topologically locally finite subcomplexes that are strongly rigid over maximal maps, proving Theorem 6.3.

6.2. **Non-existence of rigid sets.** We now show that for all infinite-type graphs Γ, there is no way to strengthen Theorem 6.3 to remove the assumption that the locally injective maps involved send maximal sphere systems to maximal sphere systems. We do this by producing nonsurjective embeddings of $\mathcal{S}(M_\Gamma)$ into itself so that the image does not contain any sphere systems which are maximal in M_Γ. It follows immediately from this that no subcomplex of $\mathcal{S}(M_\Gamma)$ can be strongly rigid.

The construction below is inspired by a similar construction for surfaces, showing that the curve graph of a hyperbolic surface embeds into the curve graph of the same surface but with one extra puncture: see Theorem 2.3 of [RS09]. We start by discussing two technical lemmas about pairs of spheres in normal form. The first is a consequence of Proposition 1.1 in [Hat02].

Lemma 6.8. *Suppose a and b are essential spheres in M_Γ whose homotopy classes can be realized disjointly. Fix a maximal sphere system σ. Then a and b contain disjoint homotopy representatives which are in normal form with respect to σ.*

Lemma 6.9. *Suppose a and a' are homotopic spheres in normal form with respect to a maximal sphere system σ. Let p be a point in a component of σ not contained in either a or a'. Then there is a homotopy from a to a' whose image does not intersect p.*

Proof. This follows directly from the proof of Proposition 1.2 in [Hat02]. Namely, we first take two homotopic lifts $\tilde{a}$ and $\tilde{a}'$ of a and a', respectively, as well as a lift $\widetilde{\sigma}$ of σ. Following the result in [Hat02], one can first homotope $\tilde{a}$ via a homotopy whose image is disjoint from all the lifts of p so that intersection of each piece of $\tilde{a}$ with $\widetilde{\sigma}$ agrees with the intersection of $\tilde{a}'$ with $\widetilde{\sigma}$ in the component of $\widetilde{M_\Gamma} \setminus \widetilde{\sigma}$ that they both lie in. This homotopy can be built by choosing a neighborhood of each element of $\widetilde{\sigma}$ so that the intersections of $\tilde{a}$ and $\tilde{a}'$ with each neighborhood are either both empty or a cylinder. Note that Proposition 1.2 in [Hat02] implies that such neighborhoods exist, as the intersections of $\tilde{a}$ and $\tilde{a}'$ with the components of $\widetilde{M_\Gamma} \setminus \widetilde{\sigma}$ agree combinatorially (that is, the pieces that show up in each component of $\widetilde{M_\Gamma} \setminus \widetilde{\sigma}$ are the same for both spheres). Now, in each such neighborhood with nonempty intersection with $\tilde{a}$ and $\tilde{a}'$, one can slide $\tilde{a}$ so that the desired intersection agreement holds. In each neighborhood containing a lift of p, such a homotopy can be chosen to avoid this lift as the complement of a point in S^2 is simply connected.

Now that the intersection circles of $\tilde{a}$ and $\tilde{a}'$ with $\widetilde{\sigma}$ agree, one can follow the argument of Proposition 1.2 in [Hat02] verbatim to homotope $\tilde{a}$ to $\tilde{a}'$ via a homotopy supported in the complement of $\widetilde{\sigma}$, finishing the proof. □

The desired embedding from $\mathcal{S}(M_\Gamma)$ into itself will be factored into two maps, the first given by the following lemma. Given a locally finite graph Γ, let M_Γ^* denote the doubled handlebody M_Γ with an extra puncture.

Lemma 6.10. *For any locally finite graph Γ, there is a simplicial embedding of $\mathcal{S}(M_\Gamma)$ into $\mathcal{S}(M_\Gamma^*)$.*

Proof. Put every element of $\mathcal{S}(M_\Gamma)$ in normal form with respect to some maximal system σ. Place the extra puncture of M_Γ^* in a component s of σ so that the puncture is not contained in any of the fixed normal form representatives of the elements of $\mathcal{S}(M_\Gamma)$. Such a choice for the puncture can be made since the intersections of all the spheres with s is a countable collection of embedded circles, which is measure 0 in s (one needs to isotope the normal form representative of s off itself by a small isotopy as well, so the puncture is disjoint from it too).

Then Lemma 6.8 implies that given any two spheres in normal form a and b which can be realized disjointly in M_Γ, there are spheres a', b' also in normal form homotopic to a, b, respectively, which disjoint from each other. From Lemma 6.9 we obtain homotopies (and thus isotopies by [Lau73]) in M_Γ^* from a to a' and from b to b'. In particular, it follows that two spheres which can be realized disjointly in M_Γ can also be realized disjointly in M_Γ^*, if the elements of $\mathcal{S}(M_\Gamma)$ are realized in normal form first, and then included into M_Γ^* as above. In particular, this induces a simplicial embedding of $\mathcal{S}(M_\Gamma)$ into $\mathcal{S}(M_\Gamma^*)$, as desired. □

Proposition 6.11. *Suppose Γ is an infinite-type graph. Then there is a simplicial embedding of $\mathcal{S}(M_\Gamma)$ into itself so that the image of every maximal system is not maximal.*

Proof. We can classify infinite-type graphs into three types.

(1) Graphs with infinite rank.
(2) Finite rank graphs with infinitely many isolated ends..
(3) Finite rank graphs whose space of ends is a Cantor set along with possibly finitely many isolated ends.

To see this, suppose Γ is not of the first two types. Then as we are assuming Γ is infinite-type, it must have an infinite space of ends. It follows by a theorem of Brouwer that, if E' is the space of ends of Γ without its isolated ends, then E' is homeomorphic to a Cantor set[Bro10]. This is because E' is compact, perfect, totally disconnected, and metrizable.

We choose a separating sphere a in each case. For type (1) graphs, let a cut off a copy of $M_{1,1}$. For type (2) graphs, let a cut off 2 isolated ends. Finally, for type (3) graphs, let a cut off a Cantor set of ends on one side. Then there is an embedding $\mathcal{S}(M_\Gamma^*)$ into $\mathcal{S}(M_\Gamma)$ which is given by removing a from M_Γ and choosing a diffeomorphism from M_Γ^* to a component of $M_\Gamma \setminus a$, sending the extra puncture of M_Γ^* to the puncture corresponding to a. Such a diffeomorphism exists, as in any case the characteristic triple of M_Γ^* is the same as one of the two components of $M_\Gamma \setminus a$.

The composition of the embedding from Lemma 6.10 and that from the previous paragraph sends $\mathcal{S}(M_\Gamma)$ to a subcomplex of itself so that in either case, a is not in the image of this map and the image of every sphere can be realized disjointly from a. In particular, every maximal sphere system is sent to something not maximal. □

7. Geometric rigidity in low complexity cases

In this section we consider the existence of finite rigid sets of $\mathcal{S}(M_{n,s})$ for n, s not covered by Theorem 1.2 The cases are $M_{0,4}, M_{0,5}, M_{1,0}, M_{1,1}, M_{1,2}, M_{1,3}, M_{2,0}$, and $M_{2,1}$. In the first four cases the complex is finite so the existence of finite rigid sets is trivial. It follows from Proposition 26 of [BL24a] that $\mathcal{S}(M_{2,0})$ has no finite rigid sets.

The remaining cases are then $M_{1,2}, M_{1,3}$, and $M_{2,1}$. We give a direct argument in the first case.

Lemma 7.1. *The graph $\mathcal{S}(M_{1,2})$ is geometrically rigid. On the other hand, $\mathcal{S}(M_{1,2})$ has no finite rigid sets.*

Proof. That $\mathcal{S}(M_{1,2})$ is geometrically rigid follows directly from Theorem 1.1. To show that it has no finite rigid sets, we will show that $\mathcal{S}(M_{1,2})$ is isomorphic to the graph in Figure 16, and deduce that this graph has no finite rigid sets.

Given a non-separating sphere a of $M_{1,2}$, there are exactly three spheres disjoint from a. To see why, note that $M_{1,2} \setminus \mathrm{nbd}(a) \cong M_{0,4}$, which has exactly three distinct spheres, say a_1, a_2, and b. See Figure 14. Observe that a_1, a_2 and b are disjoint from a, so in $\mathcal{S}(M_{1,2})$, a is a trivalent vertex. The spheres a_1 and a_2 are non-separating, and b is separating. The complementary components of b are a pair of pants and a one-holed torus. A pair of pants contains no essential spheres, and a one-holed torus contains a single homotopy class of spheres. Consequently, any essential sphere in $M_{1,2}$ disjoint from a must intersect b. Therefore, b has valence one in the sphere complex. Further, this separating sphere is only disjoint from S.

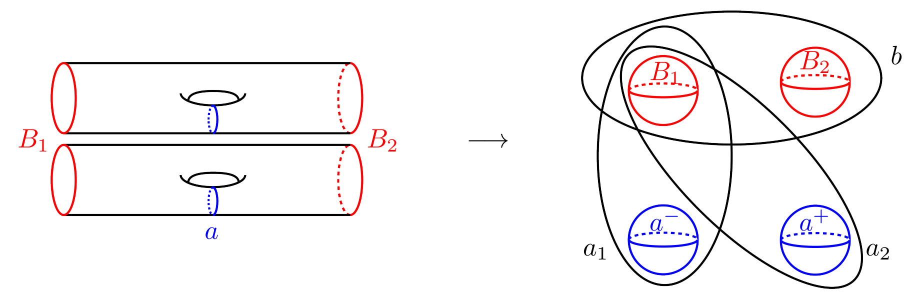

FIGURE 14. On the left, $M_{1,2}$ is pictured with boundary components B_1 and B_2 and a non-separating sphere a. To visualize the spheres a_1, a_2, and b, it is useful to cut along a, as seen on the right.

For completeness, we show explicitly that $\mathcal{S}(M_{1,2})$ is an infinite graph. Let h be a homeomorphism which pushes one of the boundary components of $M_{1,2}$ around the manifold once, and let a be a fixed non-separating sphere. Then $h^n(a)$ is not homotopic to a for any $n \geq 1$. If it were, they would have lifts in the universal cover $\widetilde{M}_{1,2}$ which are homotopic. But it is clear that the decompositions of the boundary components of $\widetilde{M}_{1,2}$ induced by any lifts of a and $h^n(a)$ must differ, and thus they cannot be homotopic. See Figure 15.

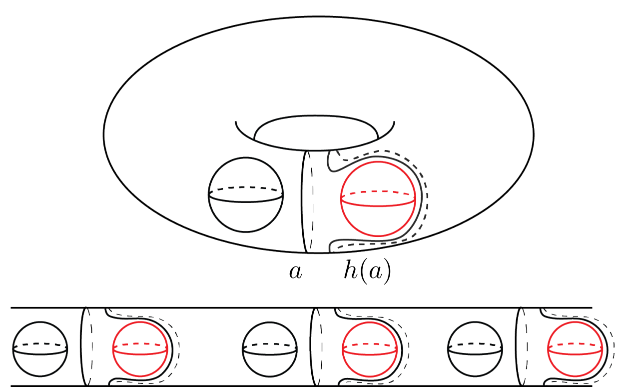

FIGURE 15. The spheres a and $h(a)$ cannot be homotopic as their none of the lifts in the universal cover are. A similar picture works for a and $h^n(a)$ for all n.

Thus, $\mathcal{S}(M_{1,2})$ is an infinite graph consisting of a collection of trivalent vertices, each of which is connected to two other trivalent vertices and one valence 1 vertex. Such a graph is isomorphic to the real line with a vertex at each integer point with a another edge attached at each vertex connecting to a valence 1 vertex, as in Figure 16.

Suppose X is a finite rigid subgraph for $\mathcal{S}(M_{1,2})$. Then X must be larger than a point, or else one could send a non-separating vertex to a separating vertex or vice versa. It must also be connected, or else one could embed X in many ways by fixing the image of one component and letting another vary, and all but finitely many of these maps cannot be induced by an automorphism of $\mathcal{S}(M_{1,2})$. Thus assume X is connected. By the structure of $\mathcal{S}(M_{1,2})$ there must be some vertex $v \in X^{(0)}$ which is a non-separating sphere which is connected by an edge to a vertex $w \in X^{(0)}$ which is valence 1 in X. Then we can embed X into $\mathcal{S}(M_{1,2})$ so that v is sent to itself, and w is sent to a separating sphere if it is non-separating, and to a non-separating sphere if it is separating. This is a contradiction, so $\mathcal{S}(M_{1,2})$ has no finite rigid sets. $\square$

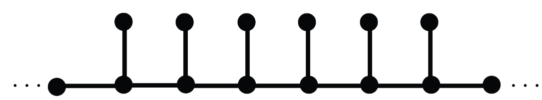

FIGURE 16. The graph $\mathcal{S}(M_{1,2})$.

We ask the following question for the final remaining cases.

Question 7.2. *Do $\mathcal{S}(M_{1,3})$ and $\mathcal{S}(M_{2,1})$ have finite rigid sets?*

References

[AB21] Yael Algom-Kfir and Mladen Bestvina. Groups of proper homotopy equivalences of graphs and Nielsen realization. *arXiv preprint arXiv:2109.06908*, 2021.

[ADMQ90] R. Ayala, E. Dominguez, A. Márquez, and A. Quintero. Proper homotopy classification of graphs. *Bull. London Math. Soc.*, 22(5):417–421, 1990.

[BD19] Mladen Bestvina and George Domat. Out($\mathbb{F}_n$) lecture notes, Fields Institute, Fall 2019.

[BDR20] Juliette Bavard, Spencer Dowdall, and Kasra Rafi. Isomorphisms between big mapping class groups. *International Mathematics Research Notices*, 2020(10):3084–3099, 2020.

[BL24a] Edgar A Bering IV and Christopher J Leininger. Finite rigid sets in sphere complexes. *Topology and its Applications*, page 108862, 2024.

[BL24b] Beth Branman and Rylee Alanza Lyman. The complex of cuts in a stone space. *arXiv preprint arXiv:2408.06994*, 2024.

[BM19] Tara Brendle and Dan Margalit. Normal subgroups of mapping class groups and the metaconjecture of ivanov. *Journal of the American Mathematical Society*, 32(4):1009–1070, 2019.

[Bro10] L. E. J. Brouwer. On the structure of perfect sets of points. *Proc. Akad. Amsterdam*, (12):785–794, 1910.

[DHK23a] George Domat, Hannah Hoganson, and Sanghoon Kwak. Coarse geometry of pure mapping class groups of infinite graphs. *Advances in Mathematics*, 413:108836, 2023.

[DHK23b] George Domat, Hannah Hoganson, and Sanghoon Kwak. Generating sets and algebraic properties of pure mapping class groups of infinite graphs. *arXiv preprint arXiv:2309.07885*, 2023.

[Gar84] M.L. Gardner. Hypergraphs and whitney's theorem on edge-isomorphisms of graphs. *Discrete Mathematics*, 51(1):1–9, 1984.

[Hat95] Allen Hatcher. Homological stability for automorphism groups of free groups. *Commentarii mathematici Helvetici*, 70(1):39–62, 1995.

[Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

[HMV18] Jesús Hernández Hernández, Israel Morales, and Ferrán Valdez. Isomorphisms between curve graphs of infinite-type surfaces are geometric. *The Rocky Mountain Journal of Mathematics*, 48(6):1887–1904, 2018.

[HV04] Allen Hatcher and Karen Vogtmann. Homology stability for outer automorphism groups of free groups. *Algebraic & Geometric Topology*, 4(2):1253–1272, 2004. Publisher: Mathematical Sciences Publishers.

[Iva97] Nikolai V Ivanov. Automorphisms of complexes of curves and of teichmuller spaces. *International Mathematics Research Notices*, 1997(14):651–666, 1997.

[Iva06] Nikolai V. Ivanov. Fifteen problems about the mapping class groups, 2006.

[Lau73] Franccois Laudenbach. Sur Les 2-Spheres D'une Variete de Dimension 3. *Annals of Mathematics*, 97(1):57–81, 1973.

[Luo99] Feng Luo. Automorphisms of the complex of curves. *arXiv preprint math/9904020*, 1999.

[RS09] Kasra Rafi and Saul Schleimer. Covers and the curve complex. *Geometry & Topology*, 13(4):2141–2162, May 2009.

[Uda24] Brian Udall. The sphere complex of a locally finite graph. *arXiv preprint arXiv:2407.07976*, 2024.