UNIFORM FELLOW TRAVELING BETWEEN SURGERY PATHS IN THE SPHERE GRAPH

MATT CLAY, YULAN QING, AND KASRA RAFI

ABSTRACT. We show that the Hausdorff distance between any forward and any backward surgery paths in the sphere graph is at most 2. From this it follows that the Hausdorff distance between any two surgery paths with the same initial sphere system and same target sphere system is at most 4. Our proof relies on understanding how surgeries affect the Guirardel core associated to sphere systems. We show that applying a surgery is equivalent to performing a Rips move on the Guirardel core.

1. INTRODUCTION

In this paper, we study the surgery paths in the sphere graph. Let $\underline{M}$ be the connected sum of n copies of $S^1 \times S^2$ (we reserve the notation M for the universal cover of $\underline{M}$ which is used more frequently in the body of the paper). The vertices of the sphere graph are essential sphere systems in $\underline{M}$ and edges encode containment (see Section 2 for precise definitions). We denote the sphere graph by $\mathcal{S}$ and the associated metric with $d_{\mathcal{S}}$. It is known that the sphere graph $(\mathcal{S}, d_{\mathcal{S}})$ is hyperbolic in the sense of Gromov [HM13, HH15]. The relationship between the optimal hyperbolicity constant and the rank of the fundamental group of $\underline{M}$ (which is isomorphic to $\mathbb{F}_n$, the free group of rank n) is unknown.

Given a pair of (filling) sphere systems $\underline{S}$ and $\underline{\Sigma}$, there is a natural family of paths, called surgery paths, connecting them. They are obtained by replacing larger and larger portions of spheres in $\underline{S}$ with pieces of spheres in $\underline{\Sigma}$. This process is not unique. Also, families of paths that start from $\underline{S}$ with target $\underline{\Sigma}$ are different from those starting from $\underline{\Sigma}$ with target $\underline{S}$. It follows from [HH15] that surgery paths are quasi-geodesics. Together with the hyperbolicity of the sphere graph, this implies that different surgery paths starting with $\underline{S}$ and with target $\underline{\Sigma}$ have bounded Hausdorff distance in the sphere graph. The bound depends on the optimal hyperbolicity constant, which as stated above, does not have a good qualitative estimate.

However, in this paper we show that, in any rank, any two surgery paths are within Hausdorff distance at most 4 of each other. This follows by comparing a surgery path that starts from $\underline{S}$ with target $\underline{\Sigma}$ to a surgery path starting from $\underline{\Sigma}$ with target $\underline{S}$.

Theorem 1.1. *Let $\underline{S}$ and $\underline{\Sigma}$ be two filling sphere systems and let*

$$\underline{S} = \underline{S}_1, \underline{S}_2, \ldots, \underline{S}_m, \qquad d_{\mathcal{S}}(\underline{S}_m, \underline{\Sigma}) \leq 1$$

be a surgery sequence starting from $\underline{S}$ towards $\underline{\Sigma}$ and

$$\underline{\Sigma} = \underline{\Sigma}_1, \underline{\Sigma}_2, \ldots, \underline{\Sigma}_\mu, \qquad d_{\mathcal{S}}(\underline{\Sigma}_\mu, \underline{S}) \leq 1$$

be a surgery sequence in the opposite direction. Then, for every $\underline{S}_i$ there is a $\underline{\Sigma}_j$ so that $d_{\mathcal{S}}(\underline{S}_i, \underline{\Sigma}_j) \leq 2$.

1

Using this, we get the bound of 4 between paths with the same initial sphere system and same target sphere system.

Theorem 1.2. *Let $\underline{S}$ and $\underline{\Sigma}$ be two filling sphere systems and let*

$$\underline{S} = \underline{S}_1, \underline{S}_2, \ldots, \underline{S}_m, \qquad d_{\mathcal{S}}(\underline{S}_m, \underline{\Sigma}) \leq 1$$

and

$$\underline{S} = \underline{S}'_1, \underline{S}'_2, \ldots, \underline{S}'_n, \qquad d_{\mathcal{S}}(\underline{S}'_n, \underline{\Sigma}) \leq 1$$

be two surgery sequences starting from $\underline{S}$ towards $\underline{\Sigma}$. Then, for every $\underline{S}_i$ there is a $\underline{S}'_j$ so that $d_{\mathcal{S}}(\underline{S}_i, \underline{S}'_j) \leq 4$.

Proof. Fix two filling sphere systems $\underline{S}$ and $\underline{\Sigma}$ and surgery paths as in the statement of the theorem. Let

$$\underline{\Sigma} = \underline{\Sigma}_1, \underline{\Sigma}_2, \ldots, \underline{\Sigma}_\mu \qquad d_{\mathcal{S}}(\underline{\Sigma}_\mu, \underline{S}) \leq 1$$

be a surgery sequence starting at $\underline{\Sigma}$ towards $\underline{S}$. Given $\underline{S}_i$, by Theorem 1.1, there is a $\underline{\Sigma}_k$ such that $d_{\mathcal{S}}(\underline{S}_i, \underline{\Sigma}_k) \leq 2$. Applying Theorem 1.1 again, there is a $\underline{S}'_j$ such that $d_{\mathcal{S}}(\underline{\Sigma}_k, \underline{S}_j) \leq 2$. Thus $d_{\mathcal{S}}(\underline{S}_i, \underline{S}'_j) \leq 4$ as desired. $\square$

The sphere graph is a direct analogue of the graph of arcs on a surface with boundary. In fact, there is an embedding of the arc graph into the sphere graph. The arc graph is known to be uniformly hyperbolic [Aou13, Bow14, CRS14, HPW15, PS15]. Since the solution of many algorithmic problems for mapping class groups or hyperbolic 3–manifold that fibers over a circle rely on the action of mapping class group on various curve and arc complexes, the uniform hyperbolicity clarifies which constant depend on the genus and which ones are genus independent. The uniform hyperbolicity of the sphere graph (or one of the other combinatorial complexes associated to $\mathrm{Out}(\mathbb{F}_n)$) is a central open question in the study of the group $\mathrm{Out}(\mathbb{F}_n)$, the group of outer automorphisms of the free group. Note that Theorem 1.2 is not sufficient to prove that the sphere graph is uniformly hyperbolic.

Summary of other results. The Guirardel core [Gui05] is a square complex associated to two trees equipped with isometric actions by a group, in our case $\mathbb{F}_n$. This is an analogue of a quadratic differential in the surface case; the area of the core is the intersection number between the two associated sphere systems. Following [BBC10], in Section 3, we describe how to compute the core using the change of marking map between the two trees. Lemma 3.7 gives a simple condition when a product of two edges is in the core, which will be used in future work to study the core. Also, in Section 4 we define the core for two sphere systems, $\mathrm{Core}(\underline{S}, \underline{\Sigma})$ directly using the intersection pattern of the spheres and show this this object is isomorphic to the Guirardel core for the associated tree (Theorem 4.9). Much of what is contained in these two sections is known to the experts, however, we include a self-contained exposition of the material since it is not written in an easily accessible way in the literature.

Applying a surgery to a sphere system amounts to applying a splitting move to the dual tree (see Example 5.6), however, not all splittings towards a given tree come from surgeries. In general, applying a splitting move could change the associated core in unpredictable ways potentially increasing the volume of the core. We will show that (Theorem 5.5) applying a surgery is equivalent to performing a Rips move on the Guirardel core. That is, there is a subset of all splitting paths between two trees that is natural from the point of view of the Guirardel core and it matches exactly with the set of splitting sequences that are associated to surgery paths.

Outline of the proof. Our proof of Theorem 1.1 analyses Guirardel core $\text{Core}(\underline{S}_i, \underline{\Sigma}_j)$. Generally, this does not have to be related to $\text{Core}(\underline{S}, \underline{\Sigma})$. However, we show that, for small values of i and j, the spheres $\underline{S}_i$ and $\underline{\Sigma}_j$ are still in normal form and $\text{Core}(\underline{S}_i, \underline{\Sigma}_j)$ can be obtained from $\text{Core}(\underline{S}, \underline{\Sigma})$ via a sequence of vertical and horizontal Rips moves (Proposition 5.8). For every i, there is the smallest j where this breaks down which is exactly the moment the surgery path from $\underline{\Sigma}$ to $\underline{S}$ passes near $\underline{S}_i$. The proof of Theorem 1.1 is completed in Section 6: for every $\underline{S}_i$, apply enough surgery on $\underline{\Sigma}$ until $\text{Core}(\underline{S}_i, \underline{\Sigma}_j)$ has a free edge, which implies $d_{\mathcal{S}}(\underline{S}_i, \underline{\Sigma}_j) \leq 2$.

Acknowledgements. We would like to thank Mladen Bestvina for helpful conversations, particularly, for giving us the idea to examine how the core changes along splitting sequence.

2. Sphere systems and free splittings

Let $\underline{M}$ be the connected sum of n copies of $S^1 \times S^2$ and fix an identification of $\pi_1(\underline{M})$ with $\mathbb{F}_n$. There is a well-known correspondence between spheres in $\underline{M}$ and graph of group decompositions of $\mathbb{F}_n$ with trivial edge groups. We explain this correspondence now.

Definition 2.1. A *sphere system* $\underline{S} \subset \underline{M}$ is a finite union of disjoint essential (does not bound a 3–ball) embedded 2–spheres in $\underline{M}$. We specifically allow for the possibility that a sphere system contains parallel, i.e., isotopic, spheres. A sphere system is *filling* if each of the complementary regions $\underline{M} - \underline{S}$ are simply-connected.

We define a preorder on the set of sphere systems by $\underline{S} \preceq \underline{\Sigma}$ if every sphere in $\underline{S}$ is isotopic to a sphere in $\underline{\Sigma}$. This induces an equivalence relation: $\underline{S} \sim \underline{\Sigma}$ if $\underline{S} \preceq \underline{\Sigma}$ and $\underline{\Sigma} \preceq \underline{S}$. The set of equivalence classes of sphere systems in $\underline{M}$ is denoted by $\mathcal{S}$; the subset of equivalence classes of filling sphere systems is denoted by $\mathcal{S}_{\text{fill}}$. When there can be no confusion, we denote the equivalence class of $\underline{S}$ again by $\underline{S}$.

The preorder induces a partial order on $\mathcal{S}$ that we continue to denote by $\preceq$. The *sphere graph* is the simplicial graph with vertex set $\mathcal{S}$ and edges corresponding to domination $\underline{S} \preceq \underline{\Sigma}$. For $\underline{S}, \underline{\Sigma} \in \mathcal{S}$, we denote by $d_{\mathcal{S}}(\underline{S}, \underline{\Sigma})$ the distance between $\underline{S}$ and $\underline{\Sigma}$ in the sphere graph. This the fewest number of edges in an edge path between the two vertices.

We denote the universal cover of $\underline{M}$ by M and the lift of the sphere system $\underline{S}$ to M by S; we will refer to S as *a sphere system in M*. To simplify notation, we use $\mathcal{S}$ and respectively $\mathcal{S}_{\text{fill}}$ to denote (equivalence classes of) sphere systems, respectively filling sphere systems, in M. Let $\text{Map}(\underline{M}) = \text{Homeo}(\underline{M})/_{\text{homotopy}}$. The natural map:

$$\text{Map}(\underline{M}) \to \text{Out}(\mathbb{F}_n)$$

is surjective and has finite kernel generated by Dehn twists about embedded 2–spheres in $\underline{M}$ [Lau74]. Such homomorphisms act trivially on spheres systems and hence there is a left action of $\text{Out}(\mathbb{F}_n)$ by automorphisms on the sphere graph. Specifically, realize the given outer automorphism by a homeomorphism of $\underline{M}$ and apply this homeomorphism to the members of a given equivalence class of sphere systems.

Definition 2.2. A *free splitting* G is a simplicial tree equipped with a cocompact action of $\mathbb{F}_n$ by automorphisms (without inversions) such that the stabilizer of every edge is trivial. We specifically allow for the possibility that a free splitting contains vertices of valence two. A free splitting is *filling* if the stabilizer of every vertex is trivial.

We define a preorder on the set of free splittings by $G \preceq \Gamma$ if there is an $\mathbb{F}_n$–equivariant cellular map $\Gamma \to G$ with connected point pre-images. This induces an equivalence relation: $G \sim \Gamma$ if $G \preceq \Gamma$ and $\Gamma \preceq G$. The set of equivalence classes of free splittings is denoted by

$\mathcal{X}$; the subset of equivalence classes of filling free splittings is denoted $\mathcal{X}_{\mathrm{fill}}$[1]. When there can be no confusion, we denote the equivalence class of G again by G.

The preorder induces a partial order on $\mathcal{X}$ that we continue to denote by $\preceq$. The *free splitting graph* is the simplicial graph with vertex set $\mathcal{X}$ and edges corresponding to domination $G \preceq \Gamma$. For $G, \Gamma \in \mathcal{X}$, we denote by $d_{\mathcal{X}}(G, \Gamma)$ the distance between G and Γ in the free splitting graph. This the fewest number of edges in an edge path between the two vertices.

Suppose that G is a free splitting and let $\rho \colon \mathbb{F}_n \to \mathrm{Aut}(G)$ denote the action homomorphism. Given $\Phi \in \mathrm{Aut}(\mathbb{F}_n)$, the homomorphism $\rho \circ \Phi \colon \mathbb{F}_n \to \mathrm{Aut}(G)$ defines a new free splitting we denote by $G \cdot \Phi$. This defines a right action by $\mathrm{Aut}(\mathbb{F}_n)$ on the free splitting graph. As $\mathrm{Inn}(\mathbb{F}_n)$ acts trivially, this induces an action of $\mathrm{Out}(\mathbb{F}_n)$ by automorphisms on the free splitting graph.

There is a natural $\mathrm{Out}(\mathbb{F}_n)$–equivariant map from the sphere graph to the free splitting graph. Given a sphere system $S \subset M$, we define a tree G with vertex set consisting of the components of $M - S$ and edges corresponding to non-empty intersection between the closures of the components. The action of $\mathbb{F}_n$ on M induces a cocompact action of $\mathbb{F}_n$ on G by automorphisms such that the stabilizer of every edge is trivial, i.e., G is a free splitting. This map is a simplicial isomorphism [AS11, Lemma 2].

3. THE GUIRARDEL CORE

In this section we give the definition of Guirardel core of two trees as it is presented in [Gui05] specialized to the case of trees in $\mathcal{X}_{\mathrm{fill}}$.

3.1. **A core for a pair of tree actions.** A *ray* in $G \in \mathcal{X}_{\mathrm{fill}}$ is an isometric embedding $\vec{r} \colon \mathbb{R}_+ \to G$. An *end* of G is an equivalence class of rays under the equivalence relation of having finite Hausdorff distance. The set of all ends is called the *boundary* of G and is denoted by ∂G.

A *direction* is a connected component of $G - \{x\}$, where x is a point in G. A direction $\delta \subset G$ determines a subset $\partial \delta \subset \partial G$ consisting of all ends for which every representative ray intersects δ in a non-empty (equivalently unbounded) subset. Given an edge $e \subset G$ we denote by $\vec{e}$ the edge with a specific orientation. This determines a direction $\delta_{\vec{e}} \subset G$ by taking the component of $G - \{x\}$ that contains e, where x is the initial vertex of $\vec{e}$. We will denote by $\vec{e}_\infty \subset \partial_\infty G$ the set of ends with a representative that crosses $\vec{e}$ with the specified orientation, i.e., $\vec{e}_\infty = \partial \delta_{\vec{e}}$.

A *quadrant* in $G \times \Gamma$ is the product $\delta_1 \times \delta_2$ of two directions $\delta_1 \subset G$ and $\delta_2 \subset \Gamma$.

Definition 3.1. Fix a basepoint $(*_1, *_2) \in G \times \Gamma$ and consider a quadrant $Q = \delta_1 \times \delta_2 \subset G \times \Gamma$. We say that Q is *heavy* if there exists a sequence $g_k \in \mathbb{F}_n$ so that

(1) $(g_k *_1, g_k *_2) \in Q$
(2) $d_G(g_k *_1, *_1) \xrightarrow{k \to \infty} \infty$ and $d_\Gamma(g_k *_2, *_2) \xrightarrow{k \to \infty} \infty$

Otherwise, we say that Q is *light*.

The core of $G \times \Gamma$ is what remains when one has removed the light quadrants.

Definition 3.2 (The Guirardel core)**.** Suppose that $G, \Gamma \in \mathcal{X}_{\mathrm{fill}}$ and let $\mathcal{L}(G, \Gamma)$ be the collection of light quadrants of $G \times \Gamma$. The (*Guirardel*) *core* of G and Γ is the subset

$$\mathrm{Core}(G, \Gamma) = (G \times \Gamma) - \left(\bigcup_{Q \in \mathcal{L}(G, \Gamma)} Q \right).$$

---

[1]Experts may recognize $\mathcal{X}_{\mathrm{fill}}$ as the vertices in the spine of the Culler–Vogtmann outer space [CV86].

It follows from the definition that $\mathrm{Core}(G, \Gamma)$ is isomorphic to $\mathrm{Core}(\Gamma, G)$ via the swap $(x, y) \mapsto (y, x)$. For more details and examples see [Gui05, BBC10].

3.2. **Computing the core.** There is an algorithm to compute the core for trees $G, \Gamma \in \mathcal{X}_{\mathrm{fill}}$. This suffices to compute the core for any free splittings $G_0, \Gamma_0 \in \mathcal{X}$. Indeed, if the given trees are not filling, they can be "blown-up" to filling trees $G, \Gamma \in \mathcal{X}_{\mathrm{fill}}$ by replacing vertices with non-trivial stabilizer in the quotient graph of groups $G_0/\mathbb{F}_n$ and $\Gamma_0/\mathbb{F}_n$ with roses of of the appropriate rank. There are domination maps $p\colon G \to G_0$ and $\pi\colon \Gamma \to \Gamma_0$ and we have that $\mathrm{Core}(G_0, \Gamma_0) = (p \times \pi)(\mathrm{Core}(G, \Gamma))$. This material appears in [BBC10, Section 2] with slightly different terminology and notation. We provide proofs of the most relevant parts necessary for the sequel.

Definition 3.3. Suppose that $G, \Gamma \in \mathcal{X}_{\mathrm{fill}}$. An $\mathbb{F}_n$–equivariant map $f\colon G \to \Gamma$ is called a *morphism* if:
  (1) f linearly expands every edge across a tight edge path; and
  (2) at each vertex of G there are adjacent edges e, e' such that $f(e) \cap f(e')$ is trivial, i.e., there is more than one gate at each vertex.

Such a map f induces an $\mathbb{F}_n$–equivariant homeomorphism $f_\infty\colon \partial G \to \partial \Gamma$. Indeed, this follows by bounded cancellation, for instance see [Coo87]. Next, we state the criterion provided in [BBC10] regarding the existence of squares in the core.

Lemma 3.4 ([BBC10, Lemma 2.3]). *Let $f\colon G \to \Gamma$ be a morphism between $G, \Gamma \in \mathcal{X}_{\mathrm{fill}}$. Given two edges $e \subset G$ and $\eta \subset \Gamma$, the square $e \times \eta$ is in the core $\mathrm{Core}(G, \Gamma)$ if and only if for every choice of orientations of the edges e and η the subset $f_\infty(\vec{e}_\infty) \cap \vec{\eta}_\infty$ is non-empty.*

This condition is very natural in the following way. Given a curve $\underline{\alpha}$ on a closed surface $\underline{X}$, each lift α of $\underline{\alpha}$ to the universal cover X determines a partition of ∂X (which is homeomorphic to S^1) into two subsets α_+ and α_-. For two curves $\underline{\alpha}, \underline{\beta}$ on $\underline{X}$ that intersect minimally, lifts α, β to X intersect if and only if for every choice of $*, *' \in \{+, -\}$ the set $\alpha_* \cap \beta_{*'}$ is non-empty.

Using f, it is a simple matter to determine when this condition is met for a given pair of edges. We discuss this now. By the *interior* of a simplicial subtree we mean all non-extremal edges.

Definition 3.5. Suppose that $G, \Gamma \in \mathcal{X}_{\mathrm{fill}}$, $f\colon G \to \Gamma$ is a morphism and $\eta \in \Gamma$ is an edge. We let $\mathcal{P}_\eta^f$ be the set of edges in G whose image under f traverses η. In other words, $\mathcal{P}_\eta^f$ is the set of edges containing $f^{-1}(\eta)$. Since f is a morphism, by bounded cancellation, the set $\mathcal{P}_\eta^f$ is finite.

Let $\mathcal{H}_\eta^f$ be the interior of the convex hull of $\mathcal{P}_\eta^f$ and let $\widehat{\mathcal{P}}_\eta^f = \mathcal{P}_\eta^f - \mathcal{H}_\eta^f$. Notice that the interior of the convex hull of $\widehat{\mathcal{P}}_\eta^f$ is also $\mathcal{H}_\eta^f$. Suppose $e \in \mathcal{H}_\eta^f$ and $\vec{e}$ is an orientation of e. We say $\vec{e}$ *can escape* $\mathcal{P}_\eta^f$ if there is an embedded ray of the form $\vec{e} \cdot \vec{r}$ such that $\vec{r}$ does not cross any edge of $\widehat{\mathcal{P}}_\eta^f$. Define the *consolidated convex hull* $\mathcal{CH}_\eta^f$ of $\mathcal{P}_\eta^f$ to be the set of edges in $e \in \mathcal{H}_\eta^f$ such that both orientations of e can escape $\mathcal{P}_\eta^f$.

Lemma 3.6. *For every vertex $v \in G$ there is a ray $\vec{r}$ originating at v that is disjoint from $\mathcal{P}_\eta^f$.*

Proof. If the lemma were false, then for every edge e adjacent to v the image $f(e)$ would contain the initial edge in the path connecting $f(v)$ to η. This violates condition (2) in Definition 3.3. $\square$
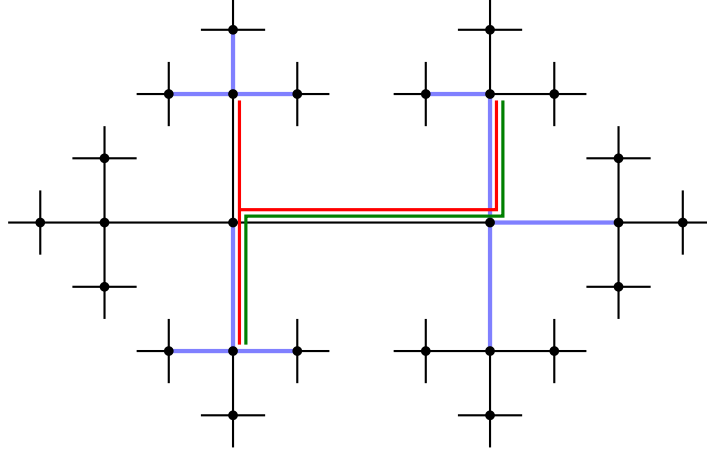
FIGURE 1. A schematic of the sets $\mathcal{P}_\eta^f$ (blue), $\widehat{\mathcal{P}}_\eta^f$, $\mathcal{H}_\eta^f$ (red) and $\mathcal{CH}_\eta^f$ (green).

The following simple condition tells exactly when a square is in the core.

Lemma 3.7. *Let $G, \Gamma \in \mathcal{X}_{\text{fill}}$, fix morphisms $f\colon G \to \Gamma$ and $\phi\colon \Gamma \to G$ and consider a pair of edges $e \subset G$ and $\eta \subset \Gamma$. The square $e \times \eta$ is in $\operatorname{Core}(G, \Gamma)$ if and only if one of the two following equivalent conditions holds:*

- *$e \subseteq \mathcal{CH}_\eta^f$.*
- *$\eta \subseteq \mathcal{CH}_e^\phi$.*

Proof. We prove the first of the two equivalent statements; the fact that they are equivalent follows from the symmetry of the construction of the core. For simplicity, we omit the superscript f on the various subsets from Definition 3.5 during the proof of this lemma.

By Lemma 3.4, what needs to be shown is that $e \subseteq \mathcal{CH}_\eta$ if and only if for each orientation $\vec{e}$ of e and orientation $\vec{\eta}$ of η there is a ray $\vec{r}$ crossing $\vec{e}$ with the specified orientation such that $f_\infty(\vec{r}) \in \vec{\eta}_\infty$.

First suppose that $e \subseteq \mathcal{CH}_\eta$ and fix an orientation $\vec{e}$ on e. As $e \subseteq \mathcal{CH}_\eta$, there is a ray $\vec{r}_0 = \vec{e} \cdot \vec{u}$ such that $\vec{u}$ is disjoint from $\widehat{\mathcal{P}}_\eta$. Let e_0 be the last edge on $\vec{r}_0$ that is in $\mathcal{CH}_\eta$ and decompose $\vec{r}_0 = \vec{u}_0 \cdot \vec{e}_0 \cdot \vec{u}_1$ where $\vec{u}_0$ may be trivial. It is easy to verify that the ray $\vec{u}_1$ is disjoint from $\mathcal{P}_\eta$. As $e_0 \subseteq \mathcal{H}_\eta$, there is a ray of the form $\vec{e}_0 \cdot \vec{u}_2$ where $\vec{u}_2$ is not disjoint from $\mathcal{P}_\eta$. (It may be that $\vec{u}_1$ and $\vec{u}_2$ have non-trivial intersection.) Let e_1 be the first edge on $\vec{u}_2$ that is contained in $\mathcal{P}_\eta$ and $\vec{p}$ the oriented edge path from $\vec{e}$ to $\vec{e}_1$. By Lemma 3.6, there is a ray $\vec{v}_1$ originating at the terminal vertex of $\vec{p}$ that is disjoint from $\mathcal{P}_\eta$. Let $\vec{r}_1 = \vec{p} \cdot \vec{v}_1$. We now see that:

$$\#|\vec{r}_1 \cap \mathcal{P}_\eta| = \#|\vec{r}_0 \cap \mathcal{P}_\eta| + 1$$

Since $\vec{r}_0$ and $\vec{r}_1$ originate from the same vertex, their f_∞–images lie in $\vec{\eta}_\infty$ for opposite choices of orientation of η. By Lemma 3.4, this shows that $e \times \eta \subseteq \operatorname{Core}(G, \Gamma)$.

For the converse we suppose that $e \nsubseteq \mathcal{CH}_\eta$. If further, $e \nsubseteq \mathcal{H}_\eta$, then there is a choice of orientation $\vec{e}$ such that for every ray of the form $\vec{e} \cdot \vec{r}$, the ray $\vec{r}$ misses $\mathcal{P}_\eta$. Therefore, there is an orientation on η, say $\vec{\eta}$, such that $f_\infty(\vec{e}_\infty) \cap \vec{\eta}_\infty = \emptyset$. By Lemma 3.4, $e \times \eta \nsubseteq \operatorname{Core}(G, \Gamma)$.

Thus we can assume that $e \subseteq \mathcal{H}_\eta - \mathcal{CH}_\eta$. Hence, there is a choice of orientation $\vec{e}$ that cannot escape, i.e., for every ray form $\vec{e} \cdot \vec{r}$, the ray $\vec{r}$ must contain some edge in $\widehat{\mathcal{P}}_\eta$. By Lemma 3.6, we see that each such ray $\vec{r}$ can only contain a single edge of $\mathcal{P}_\eta$. Again, there is an orientation on η, say $\vec{\eta}$, such that $f_\infty(\vec{e}_\infty) \cap \vec{\eta}_\infty = \emptyset$. By Lemma 3.4, $e \times \eta \nsubseteq \operatorname{Core}(G, \Gamma)$. $\qquad\square$
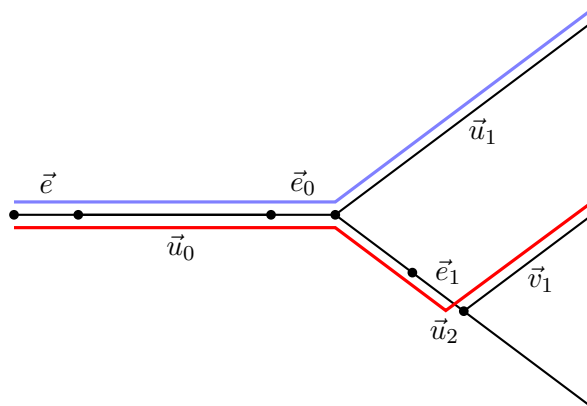
FIGURE 2. Rays $\vec{r}_0$ (blue) and $\vec{r}_1$ (red) witnessing $e \subseteq \mathcal{CH}_\eta$ in Lemma 3.7.

Since the Core(G, Γ) is defined without reference to the morphism $f \colon G \to \Gamma$, Lemma 3.7 shows that $\mathcal{CH}_\eta^f$ and $\mathcal{CH}_e^\phi$ do not depend on the actual morphism used to compute them. As such, we will drop the superscripts from these sets for the remainder.

4. SPHERE SYSTEMS AND THE CORE

In Section 2 we described an $\text{Out}(\mathbb{F}_n)$–equivariant association between sphere systems and free splittings respecting the notion of filling: $(\mathcal{S}, \mathcal{S}_{\text{fill}}) \leftrightarrow (\mathcal{X}, \mathcal{X}_{\text{fill}})$. In Section 3, given a pair of free splittings $G, \Gamma \in \mathcal{X}_{\text{fill}}$, we described how to construct their Guirardel core Core(G, Γ). The goal of this section is given a pair of filling sphere systems $S, \Sigma \in \mathcal{S}_{\text{fill}}$ to construct a 2–dimensional square complex Core(S, Σ). We then show that when the pair of sphere systems (S, Σ) is associated to the pair of free splittings (G, Γ) there is a $\mathbb{F}_n$–equivariant isomorphism from Core$(S, \Sigma) \to$ Core(G, Γ) of square complexes. Moreover, this association is $\text{Out}(\mathbb{F}_n)$–equivariant with respect to the actions on $\mathcal{X}_{\text{fill}}$ and $\mathcal{S}_{\text{fill}}$. This association is implicit in the proof of Proposition 2.1 in [Hor12]. We explain the connection in more details here and provide an alternative proof. The in depth description is necessary for understanding the effect of surgery on the core that we describe in Section 5.

4.1. **Hatcher's normal form.** Central to the understanding of sphere systems in M is Hatcher's notion of normal form. He originally defined normal form only with respect to a maximal sphere systems Σ [Hat95] and extended this to filling sphere systems in subsequent work with Vogtmann [HV96]. We recall this definition now. The sphere system S is said to be in *normal form* with respect to Σ if every sphere $s \in S$ either belongs to Σ, or intersects Σ transversely in a collection of circles that split s into components called *pieces* such that for each component $\Pi \subset M - \Sigma$ one has:

(1) each piece in Π meets each boundary sphere in $\partial \Pi$ in at most one circle, and
(2) no piece in Π is a disk that is isotopic relative to its boundary to a disk in $\partial \Pi$.

Hatcher proved that a sphere system S can always be homotoped into normal form with respect to the maximal sphere system Σ and that such a form is unique up to homotopy [Hat95, HV96]. Hensel–Osajda–Przytycki generalized Hatcher's definition of normal form to non-filling sphere systems and in a way that is obviously symmetric with respect to the two sphere systems [HOP14]. With their notion, two sphere systems S and Σ are in *normal form* if for all $s \in S$ and $\sigma \in \Sigma$ one has:

(1) s and σ intersect transversely in at most one circle, and
(2) none of the disks in $s - \sigma$ is isotopic relative to its boundary to a disk in σ.

These notions are equivalent when Σ is filling [HOP14, Section 7.1].

4.2. **A core for a pair of sphere systems.** Suppose that S and Σ are filling sphere systems in M and that they are in normal form. An *S–piece* is the closure of a component of $S - \Sigma$. Likewise, a *Σ–piece* is the closure of a component of $\Sigma - S$. By *piece*, we mean either an S–piece or a Σ–piece (this agrees with the use of piece in Section 4.1).

Lemma 4.1. *Suppose that X is the intersection of a component of $M - S$ and a component of $M - \Sigma$. Then the following statements are true.*

(1) *X is connected.*
(2) *∂X is the union of S–pieces and Σ–pieces and moreover, different pieces are subsets of different spheres.*
(3) *If Y is also the intersection of a component of $M - S$ and a component of $M - \Sigma$, then either $X = Y$, their closures are disjoint, or $\partial X \cap \partial Y$ is a piece.*

Proof. This follows from the description of normal form, details are left to the reader. $\square$

The first item in Lemma 4.1 implies that the intersection of a component of $M - S$ and a component of $M - \Sigma$ is either empty or a component of $M - (S \cup \Sigma)$.

Definition 4.2. Suppose that S and Σ are filling sphere systems in M and that they are normal form. The *core* of S and Σ, denoted $\text{Core}(S, \Sigma)$, is the square complex defined as follows.

- Vertices correspond to components of $M - (S \cup \Sigma)$. Such a region corresponds to the intersection of a component $P \subset M - S$ and a component $\Pi \subset M - \Sigma$. We denote the vertex by (P, Π).
- There is an edge between two vertices when the closures of the corresponding components of $M - (S \cup \Sigma)$ have non-trivial intersection. By Lemma 4.1, each edge corresponds to a piece. If it is an S–piece, then it is the closure of $s \cap \Pi$ for some sphere $s \in S$ and component $\Pi \subset M - \Sigma$. We denote the edge by (s, Π). Likewise, if it is an Σ–piece, then it is the closure of $P \cap \sigma$ for some component $P \subset M - S$ and sphere $\sigma \in \Sigma$. In this case, we denote the edge by (P, σ).
- Suppose that $s \in S$ and $\sigma \in \Sigma$ have non-empty intersection. Let $P_1, P_2 \subset M - S$ be the components whose boundary contains s and let $\Pi_1, \Pi_2 \subset M - \Sigma$ be the components whose boundary contains σ. Then four edges (s, Π_1), (P_1, σ), (s, Π_2) and (P_2, σ) form the boundary of a square with vertices (P_1, Π_1), (P_2, Π_1), (P_2, Π_2) and (P_1, Π_2) which is then filled in. The square is denoted by $s \times \sigma$.

Remark 4.3. We always assume that S and Σ do not share a sphere. Otherwise Theorem 1.1 is trivial. However the core in this case would be disconnected and make the exposition more complicated. There is a procedure to add diagonal edges resulting in the *augmented core*, which is connected. See [Gui05] for details.

Let G, $\Gamma \in \mathcal{X}_{\text{fill}}$ be the free splittings corresponding to S and Σ respectively. We will show that the two notions of the core, $\text{Core}(S, \Sigma)$ and $\text{Core}(G, \Gamma)$, are isomorphic as $\mathbb{F}_n$–square complexes. We will do so by showing that their horizontal hyperplanes agree. To this end we make the following definition.

Definition 4.4. The *shadow of $\sigma \in \Sigma$* is the union of the edges $e \subset G$ whose associated sphere in S intersects σ. We denote the shadow by $\text{Shadow}(\sigma) \subset G$.

Observe that the shadow of σ is isomorphic to the tree in σ that is dual to the intersection circles between σ and S. Now will show how to relate the two definitions of the core. We will make use of the following notion.
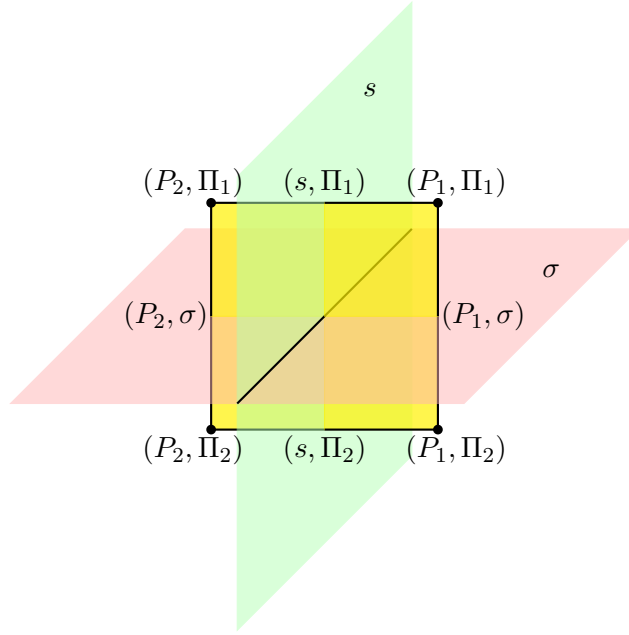
FIGURE 3. Edges (s, Π_1), (P_1, σ), (s, Π_2) and (P_2, σ) form the boundary of a square $s \times \sigma$.

Definition 4.5. If $G \in \mathcal{X}$ corresponds to a sphere system $S \in \mathcal{S}$, and $\iota \colon G \hookrightarrow M$ is an $\mathbb{F}_n$–equivariant embedding, we say $\iota(G)$ is *dual to S* if each sphere $s \in S$ intersects exactly one edge of $\iota(G)$, namely the image of the corresponding edge, and this intersection is transverse and a single point. We say that ι is a *dual embedding (for S)*.

It is a routine matter to construct a dual embedding for a given free splitting. We need to show that we can make it in some sense normal to Σ.

Lemma 4.6. *There exists a dual embedding $\iota \colon G \hookrightarrow M$ for S so that for each edge $e \in G$ and sphere $\sigma \in \Sigma$, $\iota(e)$ and σ are either disjoint or intersect transversely at a single point in the interior of $\iota(e)$.*

Proof. Let $\iota_0 \colon G \hookrightarrow M$ be a dual embedding. By general position, we can assume that $\iota_0(G) \cap S \cap \Sigma = \emptyset$ and that Σ is disjoint from the vertices of $\iota_0(G)$.

Suppose that for some edge $e \subset G$, the image $\iota_0(e)$ intersects some sphere in Σ in more than one point. Let $s \in S$ be the sphere corresponding to e. Fix some innermost pair of intersection points $x, y \in \iota_0(e)$ and let $\sigma \in \Sigma$ be the corresponding sphere. Let I be the subsegment of $\iota_0(e)$ with endpoints x and y.

Notice that any circle of intersection of $S \cap \sigma$ that separates x and y in σ must correspond to a sphere $s' \in S$ such that $s' \cap I \neq \emptyset$. Indeed, if not, since S and Σ are in normal form, there would a loop consisting of I and an arc in σ that intersects some sphere in S exactly once. This is a contradiction as spheres in S are separating.

Therefore, there is an arc $J \subset \sigma$ that intersects exactly the same set of spheres of S as I, which is either s or the empty set. We can then homotope I to J and continue pushing in this direction to reduce the number of intersection points between $\iota_0(e)$ and Σ by two. Equivariantly perform this process to obtain a new dual embedding $\iota_1 \colon G \to M$ that has fewer F_n–orbits of intersect between the image of G and Σ.

Iterating this procedure we arrive at $\iota \colon G \to M$ as in the statement of the lemma. $\square$

If $\iota\colon G \hookrightarrow M$ is an $\mathbb{F}_n$–equivariant embedding so that $\iota(G)$ is transverse to Σ we can create a map $k_\iota\colon G \to \Gamma$ by sending an edge e to the edge path in Γ corresponding to the spheres in Σ crossed by $\iota(e)$.

Lemma 4.7. *There exists a dual embedding $\iota\colon G \hookrightarrow M$ for S so that the associated map $k_\iota\colon G \to \Gamma$ is a morphism.*

Proof. Whenever a dual embedding $\iota_0\colon G \hookrightarrow M$ satisfies Lemma 4.6 the image of each edge $e \subset G$ is a tight edge path in Γ. Thus the only way that such a dual embedding $\iota_0\colon G \to M$ fails to produce a morphism is if there is some vertex $v \in G$ with adjacent edges $e_1, \ldots, e_\ell$ (oriented to have v as their initial vertex) and sphere $\sigma \in \Sigma$ so that the first intersection point of $\iota_0(e_i) \cap \Sigma$ lies in σ for each $i = 1, \ldots, \ell$. Arguing as in Lemma 4.6 we can equivariantly homotope ι_0 to locally reduce the number of intersections between the image of G and Σ by pushing the image of v across σ and pushing subarcs of edges with both endpoints on σ across σ as well.

Iterating this procedure we arrive at $\iota\colon G \to M$ as in the statement of the lemma. $\square$

A dual embedding $\iota\colon G \to M$ satisfying the conclusions of of Lemma 4.6 and 4.7 is said to be *normal to* Σ.

Proposition 4.8. *Suppose that S and Σ are filling sphere systems in M and that G and Γ are the associated trees. Fix an edge $\eta \in \Gamma$ and let $\sigma \in \Sigma$ be the associated sphere. If $\iota\colon G \hookrightarrow M$ is a dual embedding that is normal to Σ then $\mathrm{Shadow}(\sigma) = \mathcal{CH}_\eta = \mathcal{H}_\eta^{k_\iota}$.*

Proof. Suppose that $\iota\colon G \hookrightarrow M$ is a dual embedding that is normal to Σ. Let $e \subset G$ be an edge and $s \in S$ the sphere corresponding to e.

First suppose that $e \subseteq \mathrm{Shadow}(\sigma)$. Thus $s \cap \sigma$ is non-empty and as the sphere systems are in normal form, this intersection is a single circle. Let X be one of the four components of $M - (s \cup \sigma)$. Decompose $\partial X = d \cup \delta$ where d is a subdisk of s and δ is a subdisk of σ.

We claim that $\iota^{-1}(X) \subseteq G$ contains an infinite subtree. Suppose otherwise, thus $\iota^{-1}(X)$ is a finite sub-forest T. At most one extremal vertex of T corresponds to an intersection of $\iota(e)$ and s (which is in d), the remaining extremal vertices correspond to intersections of δ with edges in $\iota(G)$.

If T is empty or has some component contained in an edge of G then an innermost disk of δ (with respect to the intersection circles $\delta \cap S$) is homotopic relative to its boundary to a disk in S, violating the assumption that S and Σ are in normal form. Else, if for some component $T_0 \subseteq T$, we have that $\iota(T_0)$ does not intersect s, then for any interior vertex of T, as we saw in the proof of Lemma 4.7, the map k_ι only has one gate, violating the assumption that ι is not normal.

Thus we may assume that T is connected and has some interior vertex v, that we assume is adjacent to some extremal edge of T that is not contained in e. We label the edges $e_0, e_1, \ldots, e_\ell$ adjacent to v where $\iota(e_i)$ intersects σ for $i = 1, \ldots \ell$. Let s_i be the spheres of S corresponding to e_i for $i = 0, \ldots \ell$. Then σ must be disjoint from s_i for $i = 1, \ldots, \ell$ for otherwise there is a component of $M - (s_i \cup \sigma)$ whose pre-image in G contains a component that contained in a single edge, which we already ruled out. But in this case we have that $\sigma - s_0$ contains a disk isotopic relative to its boundary to a disk in s_0, which again violates the assumption that S and Σ are in normal form.

Hence we can find a ray $\vec{r}$ starting with e so that $\iota(\vec{r})$ is eventually contained in X. Since X was arbitrary, this shows that for each orientation of $\vec{e}$ for e and $\vec{eta}$ for η we can find a ray $\vec{r}$, crossing $\vec{e}$ with the specified orientation so that $k_\iota(\vec{r}) \in \vec{\eta}_\infty$. By Lemma 3.4, we have that $e \times \eta \subseteq \mathrm{Core}(G, \Gamma)$ and so $e \subseteq \mathcal{CH}_\eta$ by Lemma 3.7. Hence $\mathrm{Shadow}(\sigma) \subseteq \mathcal{CH}_\eta \subseteq \mathcal{H}_\eta^{k_\iota}$.

Now suppose that $e \subseteq \mathcal{H}_\eta^{k_\iota}$. Then for each orientation $\vec{e}$, there is a ray of the form $\vec{e} \cdot \vec{r}$ so that $k_\iota(\vec{r})$ intersects η and hence $\iota(\vec{r})$ intersects σ. Since s separates M and σ is connected, this shows that s intersects σ, i.e., $e \subseteq \mathrm{Shadow}(\sigma)$. Hence $\mathcal{H}_\eta^{k_\iota} \subseteq \mathrm{Shadow}(\sigma)$ completing the proof. □

In other words, Proposition 4.8 states that $\mathrm{Shadow}(\sigma) \subset G$ is the interior of the convex hull of $k_\iota^{-1}(\sigma)$.

Recall the relation between a sphere system S and the corresponding free splitting G mentioned in Section 2: vertices of G correspond to connected components $M - S$ and edges corresponding to non-empty intersection between the closures of the components, i.e., spheres in S. We can define a map $\mathrm{Core}(S, \Sigma) \to G \times \Gamma$ as follows:

- The image of a vertex (P, Π) is the vertex $(v, \nu) \in G \times \Gamma$ where v is the vertex corresponding to $P \subset M - S$ and ν is the vertex corresponding to $\Pi \subset M - \Sigma$.
- The image of an edge (s, Π) is the edge $(e, \nu) \subset G \times \Gamma$ where e is the edge corresponding to $s \in S$ and ν is the vertex corresponding to $\Pi \subset M - \Sigma$. Likewise, the image of an edge (P, σ) is the edge $(v, \eta) \subset G \times \Gamma$ where v is the vertex corresponding to $P \subset M - S$ and η is the edge corresponding to $\sigma \in \Sigma$.
- The image of the square $s \times \sigma$ is $e \times \eta \subset G \times \Gamma$ where e is the edge corresponding to $s \in S$ and η is the edge corresponding to $\sigma \in \Sigma$.

The following theorem is implicit in the proof of [Hor12, Proposition 2.1]. There, Horbez uses a characterization by Guirardel of the core as the minimal closed, connect, $\mathbb{F}_n$–invariant subset of $G \times \Gamma$ that have connected fibers [Gui05, Proposition 5.1]. We avoid using this characterization by using Lemma 3.7 and Proposition 4.8.

Theorem 4.9. *If $G, \Gamma \in \mathcal{X}_{\mathrm{fill}}$ correspond to $S, \Sigma \in \mathcal{S}_{\mathrm{fill}}$, which do not share a sphere, then the map $\mathrm{Core}(S, \Sigma) \to G \times \Gamma$ induces an $\mathbb{F}_n$–equivariant isomorphism of square complexes $\mathrm{Core}(S, \Sigma) \to \mathrm{Core}(G, \Gamma)$.*

Proof. It is clear that the map is injective, $\mathbb{F}_n$–equivariant and preserves the square structure. We just need to show that the image is $\mathrm{Core}(G, \Gamma)$. For each $\sigma \in \Sigma$, let $S_\sigma = \{s \in S \mid s \cap \sigma \neq \emptyset\}$. Notice that the edges in G corresponding to S_σ is $\mathrm{Shadow}(\sigma)$ by definition. We can decompose the core $\mathrm{Core}(S, \Sigma)$ vertically into horizontal slices $C_\sigma = \{s \times \sigma \mid s \in S_\sigma\}$. Now fix an σ and let η by the corresponding edges of Γ. Then image of the strip C_σ is exactly the set of squares $\{e \times \eta \mid e \subseteq \mathrm{Shadow}(\sigma)\}$. By Proposition 4.8 we can also write this as $\{e \times \eta \mid e \subseteq \mathcal{CH}_\eta\}$. By Lemma 3.7 we can further write this as $\{e \times \eta \mid e \times \eta \subseteq \mathrm{Core}(G, \Gamma)\}$. Hence the image of the map is as claimed. □

5. Surgery and the core

The purpose of this section is to show how the core changes along a surgery path in the sphere graph.

5.1. **Surgery Sequences.** Suppose that $S, \Sigma \in \mathcal{S}$ and assume that they are in normal form. We now describe a path from S to Σ in $\mathcal{S}$ using a surgery procedure introduced by Hatcher [Hat95]. It is exactly these paths that appear in the the main theorem of this paper.

Fix a sphere $\sigma \in \Sigma$ that intersects some spheres of S. The intersection circles define a pattern of disjoint circles on σ, each of which bounds two disks on σ. Choose an innermost disk δ in this collection, i.e., a disk that contains no other disk from this collection, and let α be its boundary circle. The sphere $s \in S$ containing α is the union of two disks d_+ and d_- that share the boundary circle α. Briefly, surgery replaces the sphere σ with new spheres $d_+ \cup \delta$ and $d_- \cup \delta$. One problem that arises is that the new sphere system and S

are not in normal form. This happens when some innermost disk δ' in a sphere $\sigma' \in \Sigma$ is parallel rel s to δ. To address this, we remove all such disks at once so that the resulting sphere system and S are in normal form (Lemma 5.1).

Let $\{\alpha_i\}_{i=k}^{\ell}$ be the maximal family of intersection circles in $s \cap \Sigma$ such that:

(1) $k \leq 0 \leq \ell$,
(2) $\alpha_i \subset d_-$ for $i \leq 0$ and $\alpha_i \subset d_+$ for $i \geq 0$ (this implies that $\alpha_0 = \alpha$), and
(3) for $k \leq i < \ell$, the circles α_i and α_{i+1} co-bound an annulus $A_i \subset s$ whose interior is disjoint from Σ.

Related to these circles, we let $\{\delta_i\}_{i=\kappa}^{\lambda}$ be the maximal family of innermost disks in Σ such that:

(1) $\kappa \leq 0 \leq \lambda$,
(2) $\partial \delta_i = \alpha_i$, and
(3) for $\kappa \leq i < \lambda$, the sphere $\delta_i \cup A_i \cup \delta_{i+1}$ bounds an embedded 3–ball, i.e., δ_i and δ_{i+1} are parallel rel s.

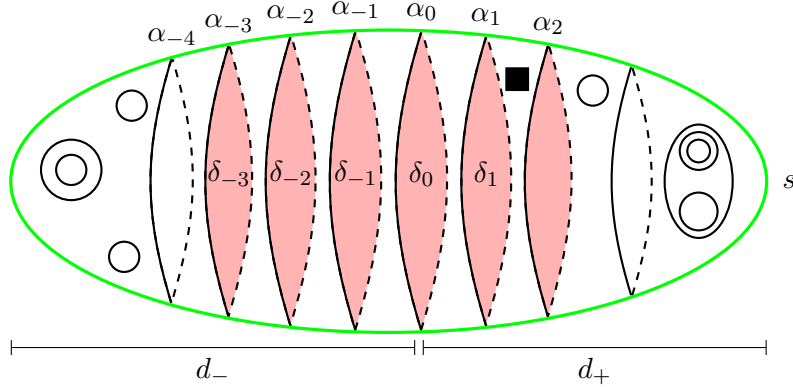See Figure 4 for an example illustrating this set-up and notation.

FIGURE 4. An example illustrating the curves $\{\alpha_i\}$ and the disks $\{\delta_i\}$. The green sphere is s, its intersection with Σ is in black and the red disks are the innermost disks in Σ. The small black box represents an obstruction to isotoping the disk bounded by α_2 to δ_1 relative to s. In this example $\kappa = -3$ and $\lambda = 1$.

Using this set-up we can now describe a surgery of S. Let δ_- be a parallel copy of δ_κ rel s such that $\partial \delta_-$ and α_κ co-bound an annulus whose interior is disjoint from Σ and A_κ. Similarly let δ_+ be a parallel copy of δ_λ rel s such that $\partial \delta_+$ and α_λ co-bound an annulus whose interior is disjoint from Σ and A_λ. Set $\hat{d}_-$ to be the subdisk of d_- with boundary $\partial \delta_-$ and set $\hat{d}_+$ to be the subdisk of d_+ with boundary $\partial \delta_+$ We get two new spheres $s_- = \hat{d}_- \cup \delta_-$ and $s_+ = \hat{d}_+ \cup \delta_+$. We say that $\hat{S} = (S - \mathbb{F}_n\{s\}) \cup \mathbb{F}_n\{s_+, s_-\}$ is obtained from S by performing a *surgery on S with respect to Σ*.

In what follows, it is important to record the history of the portions of the new spheres and so we introduce notion to this effect. Suppose that $\hat{S}$ is the result of a surgery of S with respect to Σ and that $\hat{s} \in \hat{S}$ is (a translate of) one of the newly created spheres $s_* = \hat{d}_* \cup \delta_*$ for $* \in \{+, -\}$. We call d_* the *portion of $\hat{s}$ from S*, denote it by $\hat{s}^S$. Similarly, we call δ_* the *portion of $\hat{s}$ from Σ*, denote it by $\hat{s}^\Sigma$. Thus $\hat{s} = \hat{s}^S \cup \hat{s}^\Sigma$. Notice that $\hat{s}^S \subseteq S$ and also that $\hat{s}^\Sigma$ is parallel rel s to a disk in Σ. For all other spheres $s \in \hat{S}$ we set $s^S = s$ and $s^\Sigma = \emptyset$.

Our definition of surgery differs slightly from the standard in three ways: one, we do not remove parallel spheres in $\hat{S}$, two, we perform surgery along parallel innermost disks in a single step, and three, we do not isotope S' to be in normal form with respect to Σ. That we do not remove parallel spheres is in keeping with our definition of sphere systems from Section 2. Justification of the latter two differences is the following lemma that shows that by performing surgery along the parallel innermost disks we can eliminate the need to perform a subsequent isotopy.

Lemma 5.1. *Let $\hat{S}$ be the result of a surgery on S with respect to Σ. Then $\hat{S}$ and Σ are in normal form.*

Proof. Suppose otherwise. As S and Σ are in normal form by assumption and normal form is a local condition, it must be that one of the newly created spheres is not in normal form with respect to Σ. Denote this sphere by $\hat{s} = \hat{d} \cup \delta$ where $\hat{d}$ is a subdisk of the surgered sphere $s \in S$ and δ is a disk parallel rel s to a disk in Σ. Any intersection between $\hat{s}$ and some sphere of Σ must lie in $\hat{d} \subset s$ and hence $\hat{s}$ and a given sphere in Σ intersect transversely in at most one circle as the same held for $s \in S$.

Therefore, if $\hat{s}$ is not in normal form with respect to Σ, then there is a sphere $\sigma \in \Sigma$ such that one of the disks in $\hat{s} - \sigma$, denote it d, is isotopic relative to its boundary to a disk in σ, denote it δ'. Without loss of generality, we can assume that this disk is innermost on $\hat{s}$, i.e., no subdisk of d is isotopic relative to its boundary to a disk in some sphere of Σ. The disk d cannot lie entirely in $\hat{d}$ since s and Σ are in normal form by assumption. Hence d contains δ. Let A be the annulus such that $d = A \cup \delta$. Since $d \cup \delta'$ bounds a 3–ball, the assumptions that S and Σ are in normal form and that d is innermost implies that Σ is disjoint from the interior of A. This contradicts the maximality assumption on the family of disks $\{\delta_i\}_{i=\kappa}^{\lambda}$. Indeed, without loss of generality we can assume that $\delta = \delta_+$. Then $A_\lambda \cup A$ is an annulus in s whose interior is disjoint from Σ and so $\partial \delta' = \alpha_{\lambda+1}$ and further $\delta_\lambda \cup (A_\lambda \cup A) \cup \delta'$ bounds an embedded 3–ball. Hence $\hat{S}$ and Σ are in normal form. $\square$

Definition 5.2. A *surgery sequence from S to Σ* is a finite sequence of sphere systems:

$$S = S_1, \ldots, S_m$$

such that S_{i+1} is the result of a surgery of S_i with respect to Σ and $d_{\mathcal{S}}(\Sigma, S_m) \leq 1$.

It is a standard fact that if $d_{\mathcal{S}}(S, \Sigma) \geq 2$, then there is a surgery sequence from S, see for instance [HH15, Lemma 2.2]. Further $d_{\mathcal{S}}(S_i, S_{i+1}) \leq 2$ as both S_i and S_{i+1} are dominated by $S_i \cup S_{i+1}$.

The discussion and notion regarding portions from S and from Σ make sense for surgery sequences as well by induction. Indeed, suppose that S_{i+1} is obtained from S_i by a surgery with respect for Σ, specifically, assume that the (orbit of the) sphere $s \in S_i$ is split into (the orbit of) two spheres $s_- = \hat{d}_- \cup \delta_-$ and $s_+ = \hat{d}_+ \cup \delta_+$ in S_{i+1}. Then we have $s = \hat{d}_- \cup A \cup \hat{d}_+$ for some annulus A, the boundary curves of which are parallel to circles in Σ rel s. By choosing A sufficiently narrow enough, we can assume that the annuli of s witnessing the isotopy are contained in s^S. We set $s_*^S = \hat{d}_* \cap s^S$ and $s_*^\Sigma = (\hat{d}_* \cap s^\Sigma) \cup \delta_*$ for $* \in \{+, -\}$. All other spheres in S_{i+1} are also in $S_i - \{s\}$ and as such the portions from S and Σ remain unchanged. See Figure 5.

Lemma 5.3. *Suppose that $S = S_1, \ldots, S_m$ is a surgery sequence from S to Σ. Then for every $s \in S_i$, the subset s^S is connected.*

Proof. Using induction, we can conclude that the subset s^Σ is a union of disks, each parallel rel s to a disk in Σ. Hence, s^S is the complement of finitely many disks in s and therefore connected. $\square$
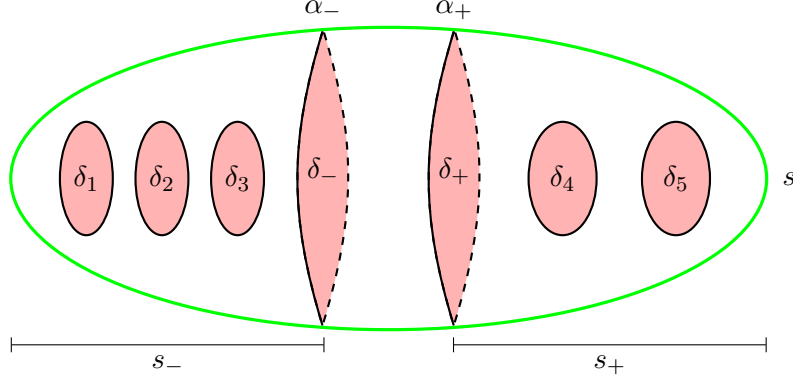
FIGURE 5. An illustration showing a decomposition of $s \in S_i$ and the resulting spheres $s_-, s_+ \in S_{i+1}$ into their portions from S and Σ. In this example $s^\Sigma = \{\delta_1, \delta_2, \delta_3, \delta_4, \delta_5\}$, $s^\Sigma_- = \{\delta_1, \delta_2, \delta_3, \delta_-\}$ and $s^\Sigma_- = \{\delta_4, \delta_5, \delta_+\}$. The portions from S are the complements in the respective spheres.

We remark that parallel spheres in S_i may have different histories, that is, $s_1, s_2 \in S_i$ may be parallel even though s_1^S and s_2^S are not parallel. For a surgery sequence $S = S_1, \ldots, S_m$ from S to Σ we set

$$S_i^S = \bigcup_{s \in S_i} s_i^S \text{ and } S_i^\Sigma = \bigcup_{s \in S_i} s_i^\Sigma.$$

5.2. **Rips moves and surgery steps.** Suppose that S and Σ are filling sphere systems and assume that they are in normal form. Let $S = S_1, S_2, \ldots, S_m$ be a surgery sequence from S to Σ and let $\Sigma = \Sigma_1, \Sigma_2, \ldots, \Sigma_\mu$ be a surgery sequence from Σ to S. We describe $\mathrm{Core}(S_i, \Sigma_j)$ as (in some appropriate sense) an intersection of $\mathrm{Core}(S_i, \Sigma)$ and $\mathrm{Core}(S, \Sigma_j)$. We start by giving an embedding of $\mathrm{Core}(S_i, \Sigma)$ and $\mathrm{Core}(S, \Sigma_j)$ into $\mathrm{Core}(S, \Sigma)$. This embedding is constructed inductively. A single step in the construction is reminiscent to one of the elementary moves of the Rips machine [BF95, CH14]. We start with a definition.

Definition 5.4. Suppose that S and Σ are filling sphere systems in M. We define ∂_S, the S–boundary of $\mathrm{Core}(S, \Sigma)$, to be the subset of $\mathrm{Core}(S, \Sigma)$ consisting of the (open) edges (P, σ) that are the face of exactly one square or vertices (P, Π) that are the vertex of exactly 3 edges of the form (P, σ), (P, σ') and (s, Π). A connected component of ∂_S is called an S–*sides* of $\mathrm{Core}(S, \Sigma)$. Similarly, we define ∂_Σ, the Σ–boundary of $\mathrm{Core}(S, \Sigma)$, to be the subset of $\mathrm{Core}(S, \Sigma)$ consisting of the (open) edges (s, Π) that are the face of exactly one square or vertices (P, Π) that are the vertex of exactly 3 edges of the form (s, Π), (s', Π) and (P, σ). A connected component of ∂_Σ is called an Σ–*sides* of $\mathrm{Core}(S, \Sigma)$.

The union of an S–side with the set of the (open) squares that have a face contained in that side is called a *maximal S–boundary rectangle*. That is, in a S–maximal boundary rectangle, all of the squares are of the form $s_0 \times \sigma$ for some fixed $s_0 \in S$. A Σ–*maximal boundary rectangle* is similarly defined from a connected component of the Σ–side. A *Rips move* on (S, Σ) is the removal of the $\mathbb{F}_n$–orbit of a (S– or Σ–)maximal boundary rectangle.

If R is a maximal boundary rectangle in $\mathrm{Core}(S, \Sigma)$, we let $\mathrm{Core}(S, \Sigma)_R$ denote the result of the associated Rips move. We like to think of the removal of the maximal boundary rectangle as collapsing the rectangle by pushing across the adjacent squares.

We postpone presenting an example until after the following theorem.

Theorem 5.5. *Suppose that S and Σ are filling sphere systems in M and let $\hat{S}$ be the result of a surgery on S with respect to Σ. There is a S–maximal boundary rectangle $R \subseteq \mathrm{Core}(S, \Sigma)$ so that $\mathrm{Core}(S, \Sigma)_R$ is isomorphic to $\mathrm{Core}(\hat{S}, \Sigma)$. Moreover, for each S–maximal boundary rectangle, R, there is a sphere $\sigma \in \Sigma$ and innermost disk on σ that defines a surgery $S \mapsto \hat{S}$ so that $\mathrm{Core}(\hat{S}, \Sigma)$ is isomorphic to $\mathrm{Core}(S, \Sigma)_R$.*

Proof. Assume $\hat{S}$ is obtained from S by a surgery on a sphere $s_0 \in S$ and a disk δ that is part of the sphere system Σ, whose boundary α lies on s and is otherwise disjoint from S. By Lemma 5.1 $\hat{S}$ and Σ are in normal form and so we can use the combinatorics of $\hat{S}$ and Σ to build $\mathrm{Core}(\hat{S}, \Sigma)$.

We make use of the notation introduced in Section 5.1. Let $\{\delta_i\}_{i=\kappa}^{\lambda}$ be the maximal family of disks in Σ parallel rel s where $\delta_0 = \delta$. Let A be the union of the annuli $A_i \subset s_0$ and d_+, d_- the components of $s_0 - A$. Thus $\delta_\kappa \cup A \cup \delta_\lambda$ bounds a 3–ball B. The two spheres obtained by surgery of s using this family, s_+ and s_-, are parallel to $d_+ \cup \delta_\lambda$ and $d_- \cup \delta_\kappa$ respectively.

Let $P_+ \in M - S$ be the component that contains the interior of B and let P_- be the other component with s as a boundary. Each disk δ_i is contained in some sphere $\sigma_i \in \Sigma$. For each $\kappa \leq i < \lambda$, there are components $\Pi_i \subset M - \Sigma$ such that both $\sigma_i, \sigma_{i+1} \subset \partial\Pi_i$. We claim that the collections of edges and vertices:

$$(P_+, \sigma_\kappa), (P_+, \Pi_\kappa), (P_+, \sigma_{\kappa+1}), \ldots, (P_+, \sigma_\lambda)$$

is a side. Indeed, each edge (P_+, σ_i) is the face of only $s \times \sigma_i$ and each vertex (P_+, Π_i) is only also adjacent to (s_0, Π_i). The first of these observations is due to the fact that $P_+ \cap \sigma_i = \delta_i$ is a disk; the second observation due to the fact that $P_+ \cap \Pi_i$ is bounded by $\delta_i \cup A_i \cup \delta_{i+1}$. Maximality of this collection follows from maximality of the collection $\{\delta_i\}_{i=\kappa}^{\lambda}$.

Let R be the corresponding maximal boundary rectangle of $\mathrm{Core}(S, \Sigma)$. We will show that $\mathrm{Core}(\hat{S}, \Sigma)$ is isomorphic to $\mathrm{Core}(S, \Sigma)_R$. To do so, we will construct an injection of square complexes $\mathrm{Core}(\hat{S}, \Sigma) \hookrightarrow \mathrm{Core}(S, \Sigma)$ whose image is $\mathrm{Core}(S, \Sigma)_R$.

Components in $M - S$ that are not in the orbit of P_+ and P_- are also components of $M - \hat{S}$. But $M - \hat{S}$ has 3 other components; $\hat{P}_-$ which is obtained from P_- by adding a neighborhood of s and a neighborhood of the 3–ball B bounded by $\delta_\kappa \cup A \cup \delta_\lambda$, P_+^+ and P_+^- which are contained in the two components of $P_+ - B$. In other words, we have:

$$M - \hat{S} = \big((M - S) - \mathbb{F}_n\{P_+, P_-\}\big) \cup \mathbb{F}_n\{\hat{P}_-, P_+^+, P_+^-\}.$$

There is an $\mathbb{F}_n$–equivariant map $\iota\colon M - \hat{S} \to M - S$ defined by $P_+^+, P_+^- \mapsto P_+$, $\hat{P}_- \mapsto P_-$ and the identity on the other orbits. Also, there is a $\mathbb{F}_n$–equivariant map $\epsilon\colon \hat{S} \to S$ defined by $s_+, s_- \mapsto s_0$ and the identity on the other orbits.

Using ι, we get a map on the 0–skeleton of $\mathrm{Core}(\hat{S}, \Sigma)$ defined by $(P, \Pi) \mapsto (\iota(P), \Pi)$. In order for this to be well-defined, we need to know that if $P \cap \Pi \neq \emptyset$, then $\iota(P) \cap \Pi \neq \emptyset$ also. If P is not in the orbit of $\hat{P}_-$, then this follows as $P \cap \Pi \subseteq \iota(P) \cap \Pi$. Finally, since $\hat{P}_- = P_- \cup B$ and no component of $M - \Sigma$ is contained in B, any component of $M - \Sigma$ that intersects $\hat{P}_-$ necessarily intersects P_- as well.

We extend over the 1–skeleton using $\epsilon\colon (s, \Pi) \mapsto (\epsilon(s), \Pi)$. This map is well-defined since any intersection between s_+, or s_-, with a component of $M - \Sigma$ is contained in the portion of s_+, or s_- respectively, from S, i.e., d_+, or d_- respectively. Notice that this is consistent with the mapping on the 0–skeleton. The edge (s_+, Π) in $\mathrm{Core}(\hat{S}, \Sigma)$ is sent to (s_0, Π). The vertices of (s_0, Π) are (P_+^+, Π) and $(\hat{P}_-, \Pi)$, which are the images of (P_+^+, Π) and $(\hat{P}_-, \Pi)$. Other verifications are similar.

Finally, we extend over the 2–skeleton: $s \times \sigma \mapsto \epsilon(s) \times \sigma$. Since any intersection of s_+ or s_- with Σ is contained in the portion from S, this map is well-defined. Again, the map on the 2–skeleton is consistent with the maps on 1–skeleton and 0–skeleton by construction.

The map $\mathrm{Core}(\hat{S}, \Sigma) \to \mathrm{Core}(S, \Sigma)$ is $\mathbb{F}_n$–equivariant and preserves the square structure. The map is not surjective as no 2–cell is mapped to the squares associated with $s \times \sigma_i$, i.e., the image of the map is exactly $\mathrm{Core}(S, \Sigma)_R$.

The converse is similar, if $s_0 \times \sigma_1, \ldots, s_0 \times \sigma_\lambda$ forms an S–maximal boundary rectangle R, then one shows that there are disks $\delta_i \subset \sigma_i$ that are parallel rel s_0 and that surgery using the family $\{\delta_i\}_{i=1}^\lambda$ results in the sphere system $\hat{S}$ where $\mathrm{Core}(\hat{S}, \Sigma)$ is isomorphic to $\mathrm{Core}(S, \Sigma)_R$.   □

Example 5.6. Here we describe a Rips move and the corresponding surgery explicitly in an example. Consider a sphere $s \in S$ associated to the edge b in the dual graph G. In the example depicted in Figure 6, s intersects 7 spheres in Σ; spheres $\sigma_0, \ldots, \sigma_6$ associated to edges $\eta_0, \ldots, \eta_6$ in Γ. We denote the intersection circle between s and σ_i by α_i. The slice over s in $\mathrm{Core}(S, \Sigma)$ consists of squares associated to intersection circles between s and Σ that is,

$$C_s = \{s \times \sigma_i \mid i = 0, \ldots, 6\}.$$

In the language of trees, C_s is associated to the slice over b, which is

$$b \times \mathrm{Shadow}(s) = \{b \times \eta_i \mid i = 0, \ldots, 6\} \subset G \times \Gamma.$$

There are two components of $M - S$ that have s as their boundary sphere. In this example, the component P_l, which we call left, has 3 other boundary spheres (associated to edges a_1, a_2 and a_3) and the component P_r on the right has two other boundary spheres (associated to edges c_1 and c_2).

Note that Figure 6 indicates that the sphere σ_1 intersects spheres in S associated to edges a_1, b and c_2 since the core contains squares $a_1 \times \eta_1$, $b \times \eta_1$ and $c_2 \times \eta_1$. However, the sphere σ_2 does not intersect spheres associated to edges a_1, a_2 and a_3. But, σ_2 intersects s, hence, the circle α_2 must bound a disk δ_2 that is the intersection of σ_2 with P_l. Similarly, circles α_3, α_4 and α_5 bound disks δ_3, δ_4 and δ_5 that are, respectively, intersections of spheres σ_3, σ_4 and σ_5 with P_l, (thus the squares $b \times \eta_i$, $i = 2, \ldots, 5$, have boundary edges on their left side). The circle α_2 also bounds a disk in σ_2 in P_r (thus the square $b \times \eta_2$ has a boundary edge on its right side).

The disks δ_2 and δ_3 are parallel and the disks δ_4 and δ_5 are also parallel, however, the two sets of disks are not parallel to each other (see Figure 7). Thus, there are two maximal boundary rectangles from the left; $R = b \times (\eta_2 \cup \eta_3)$ and $R' = b \times (\eta_4 \cup \eta_5)$. More precisely, let Π be the component of $M - \Sigma$ with σ_0, σ_3 and σ_4 as its boundary spheres. Then, referring to Definition 5.4, we see that the vertex (P_l, Π) is not in the S–boundary of $\mathrm{Core}(S, \Sigma)$ because it is the vertex of 5 different edges. Hence, the union of R and R' is not a boundary rectangle.

Define S' to be the sphere system obtained from S by applying the surgery on the set of parallel disks $\{\delta_4, \delta_5\}$ (and their $\mathbb{F}_n$–orbits). The surgery results in two spheres s_1 and s_2 associated to edges b_1 and b_2 and the removal of the maximal boundary rectangle R. It appears that removal of this rectangle makes the slice over b disconnected. However, the two components are slices over the edges b_1 and b_2.

To summarize, the surgery along the disks $\{\delta_4, \delta_5\}$, changes G by splitting the edge b and changes $\mathrm{Core}(S, \Sigma)$ by removing the maximal boundary rectangle $R = b \times (\eta_4 \cup \eta_5)$, resulting in $\mathrm{Core}(S, \Sigma)_R \cong \mathrm{Core}(S', \Sigma)$.

A different splitting of b into b_1 and b_2 partitioning the a edges into $\{a_1, a_3\}$ and $\{a_2\}$ does not arise as a surgery and could potentially increase the volume of the core.
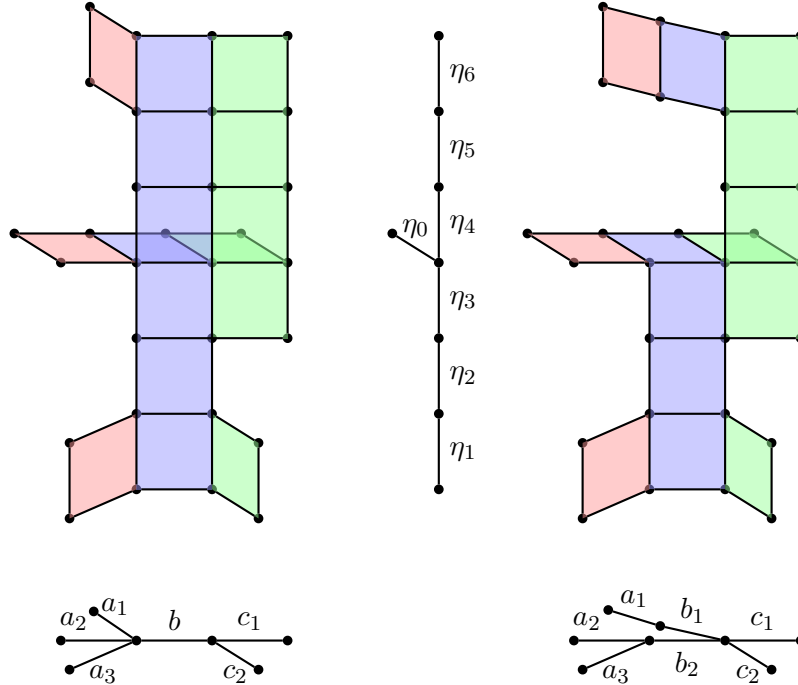
FIGURE 6. The left hand side depicts the slice C_s and squares attached to it in $\text{Core}(S, \Sigma)$. Consider the maximal S–boundary rectangle $R = b \times (\eta_4 \cup \eta_5)$. A Rips move along R is associated to a surgery on the sphere s or a splitting of the edge b in the graph G. The right side depicts the associated portion of $\text{Core}(S, \Sigma)_R = \text{Core}(S', R)$.
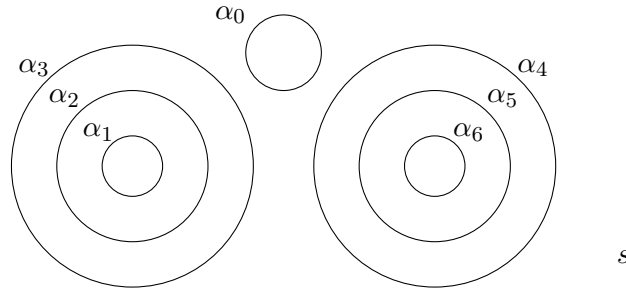

FIGURE 7. The curves $\alpha_0, \ldots, \alpha_6$ are intersection circles between the sphere s and the spheres $\sigma_0, \ldots, \sigma_6$ respectively. The circles α_2, α_3, α_4 and α_5 bound a disk in P_l. However, the circle α_3 and α_4 are not parallel.

5.3. The intersection of cores.

Applying Theorem 5.5 to the surgery sequence $S = S_1, S_2, \ldots, S_m$ we obtain maps for $i = 1, \ldots, m - 1$:

$$k_{i,i+1} : \text{Core}(S_{i+1}, \Sigma) \to \text{Core}(S_i, \Sigma).$$

that are the composition of the isomorphism $\text{Core}(S_{i+1}, \Sigma) \cong \text{Core}(S_i, \Sigma)_R$ for the corresponding maximal boundary rectangle and the natural inclusion $\text{Core}(S_i, \Sigma)_R \hookrightarrow \text{Core}(S_i, \Sigma)$.

By symmetry there are also maps for $j = 1, \mu - 1$:

$$\kappa_{j,j+1} \colon \ \mathrm{Core}(S, \Sigma_{j+1}) \to \mathrm{Core}(S, \Sigma_j).$$

Since these map exist for all $1 \le i \le m - 1$, $1 \le j \le \mu - 1$, we can define "inclusions" alluded to at the beginning of this section:

$$k_i = k_{1,2} k_{2,3} ... k_{i-2,i-1} k_{i-1,i} \colon \ \mathrm{Core}(S_i, \Sigma) \to \mathrm{Core}(S, \Sigma) \tag{5.1}$$

$$\kappa_j = \kappa_{1,2} \kappa_{2,3} ... \kappa_{j-2,j-1} \kappa_{j-1,j} \colon \ \mathrm{Core}(S, \Sigma_j) \to \mathrm{Core}(S, \Sigma) \tag{5.2}$$

Remark 5.7. On the level of squares, the map $k_i \colon \ \mathrm{Core}(S_i, \Sigma) \to \mathrm{Core}(S, \Sigma)$ is easy to describe. For each $\hat{s} \in S_i$, we have that $\hat{s}^S \subseteq s$ for a unique $s \in S$. The map is defined by $\hat{s} \times \sigma \to s \times \sigma$.

The following is the fundamental concept essential to the proof of the main theorem.

Proposition 5.8. *With the above set-up, assume*

$$k_i\big(\mathrm{Core}(S_i, \Sigma)\big) \cup \kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big) = \mathrm{Core}(S, \Sigma), \tag{5.3}$$

then S_i and Σ_j are in normal form. Furthermore, there exists an isomorphism

$$\Phi \colon \ \mathrm{Core}(S_i, \Sigma_j) \to k_i\big(\mathrm{Core}(S_i, \Sigma)\big) \cap \kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big).$$

Proof. First, we show that every intersection circle between S_i and Σ_j is in fact in $S_i^S \cap \Sigma_j^\Sigma$. This is because a square in $\mathrm{Core}(S, \Sigma)$ associated to an intersection circle in $S_i^\Sigma \cap \Sigma_j^S$ is neither in $k_i\big(\mathrm{Core}(S_i, \Sigma)\big)$ (S_i^Σ does not intersect Σ) nor in $\kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big)$ (Σ_j^S does not intersect S) and by the assumption (5.3) every square in $\mathrm{Core}(S, \Sigma)$ is in the image of one of these two maps.

This observation implies that S_i and Σ_j are in fact in normal form. In fact, pick spheres $s_i \in S_i$ and $\sigma_j \in \Sigma_j$. We will show that s_i and σ_j intersect at most once. Otherwise, s_i^S and σ_j^Σ intersect more than once. But, by Lemma 5.3, s_i^S and σ_j^Σ are connected, which means there is a sphere $s \in S$ that contains s_i^S and a sphere $\sigma \in \Sigma$ that contains σ_j^Σ. Hence, s and σ intersect more than once. This contradicts the fact that S and Σ are in normal form.

Now consider a square $s_i \times \sigma_j$ in $\mathrm{Core}(S_i, \Sigma_j)$ associated to an intersection circle α. Then α is an intersection circle in $S_i^S \cap \Sigma_j^\Sigma$. Which mean it is an intersection circle in both $S \cap \Sigma_j$ and $S_i \cap \Sigma$ and thus there are spheres $s \in S$, $\sigma \in \Sigma$ for which $s \cap \sigma = \alpha$ and $s_i^S \subseteq s$, $\sigma_j^\Sigma \subseteq \sigma$. Hence, $s \times \sigma$ is contained in both $k_i\big(\mathrm{Core}(S_i, \Sigma)\big)$ and $\kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big)$ and so we define $\Phi(s_i \times \sigma_j) = s \times \sigma$. Normal form implies that the map is injective.

To prove that Φ is surjective, suppose $s \times \sigma$ is in $k_i\big(\mathrm{Core}(S_i, \Sigma)\big)$. Then the associated intersection circle in S_i^S. Similarly, assuming $s \times \sigma$ is in $\kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big)$ implies that the associated intersection circle in Σ_j^Σ. Therefore, it also lies in $S_i \cap \Sigma_j$. Hence there are spheres $s_i \in S_i$ and $\sigma_j \in \Sigma_j$ such that $\Phi(s_i \times \sigma_j) = s \times \sigma$. $\qquad\square$

For future reference, we record the following corollary:

Corollary 5.9. *As long as* (5.3) *is satisfied,* $\mathrm{Core}(S_i, \Sigma_{j+1})$ *can obtained from* $\mathrm{Core}(S_i, \Sigma_j)$ *by a Rips move.*

Proof. Let R be a maximal Σ_j–boundary rectangle in $\mathrm{Core}(S, \Sigma_j)$ such that $\mathrm{Core}(S, \Sigma_j)_R \cong \mathrm{Core}(S, \Sigma_{j+1})$. Thus R consists of squares $s_1 \times \hat{\sigma}, \ldots, s_\ell \times \hat{\sigma}$ for some $\hat{\sigma} \in \Sigma_j$ and $s_1, \ldots, s_\ell \in S$. Let $\sigma \in \Sigma$ be such that $\hat{\sigma}^\Sigma \subseteq \sigma$ and consider the set $C_{\sigma,i}$ of squares of the form $\hat{s} \times \sigma$ in $\mathrm{Core}(S_i, \Sigma)$. Then $k_i(C_{\sigma,i}) \cap \kappa_j(R)$ corresponds via the isomorphism in Proposition 5.8 to a maximal Σ_j–boundary rectangle in $\mathrm{Core}(S_i, \Sigma_j)$ whose collapse results in $\mathrm{Core}(S_i, \Sigma_{j+1})$. $\qquad\square$

6. Proof of Theorem 1.1

To finish the proof of the main theorem, we proceed as follows using the set-up from the previous section. We start with a lemma giving a necessary condition for two sphere systems to be at a bounded distance. A *free edge* is an edge that does not bound any squares.

Lemma 6.1. *If* $\mathrm{Core}(S_i, \Sigma_j)$ *contains a free edge, then the two sphere systems are of distance at most 2 in the sphere graph.*

Proof. Edges in the core are associated to spheres in either sphere system S_i or Σ_j and squares are associated to intersection circles between sphere systems. Hence, a free edge in the core is associated to a sphere in either S_i or Σ_j that does not intersect any other spheres from the other system. Thus this sphere can be added to both sphere systems. That is, $\underline{S}_i$ and $\underline{\Sigma}_j$ have distance 2 in the sphere graph. $\qquad\square$

We now prove Theorem 1.1. We restate it for convenience.

Theorem 1.1. *Let $\underline{S}$ and $\underline{\Sigma}$ be two filling sphere systems and let*

$$\underline{S} = \underline{S}_1, \underline{S}_2, \ldots, \underline{S}_m, \qquad d_{\mathcal{S}}(\underline{S}_m, \underline{\Sigma}) \le 1$$

be a surgery sequence starting from $\underline{S}$ towards $\underline{\Sigma}$ and

$$\underline{\Sigma} = \underline{\Sigma}_1, \underline{\Sigma}_2, \ldots, \underline{\Sigma}_\mu, \qquad d_{\mathcal{S}}(\underline{\Sigma}_\mu, \underline{S}) \le 1$$

be a surgery sequence in the opposite direction. Then, for every $\underline{S}_i$ there is a $\underline{\Sigma}_j$ so that $d_{\mathcal{S}}(\underline{S}_i, \underline{\Sigma}_j) \le 2$.

Proof. Fix two filling sphere systems $\underline{S}$ and $\underline{\Sigma}$ and surgery paths as in the statement of the theorem. For every $\underline{S}_i$ we need to find $\underline{\Sigma}_j$ with $d_{\mathcal{S}}(\underline{S}_i, \underline{\Sigma}_j) \le 2$. Fix an $i = 1, \ldots, m$ and let j be the largest index where the equality

$$k_i\big(\mathrm{Core}(S_i, \Sigma)\big) \cup \kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big) = \mathrm{Core}(S, \Sigma) \tag{6.1}$$

still holds. Note that the equation holds when $j = 1$. But, since $\kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big)$ eventually contains no squares (for instance, when $j = \mu$) and $k_i\big(\mathrm{Core}(S_i, \Sigma)\big)$ is a proper subset of $\mathrm{Core}(S, \Sigma)$ for each $i > 1$, there exists an index $j + 1$ for which (6.1) does not hold.

We will show that $\mathrm{Core}(S_i, \Sigma_j)$ contains a free edge. By Lemma 6.1, this will complete the proof. Let $s \times \sigma$ be a square in $\mathrm{Core}(S, \Sigma)$ that is not contained in $k_i\big(\mathrm{Core}(S_i, \Sigma)\big) \cup \kappa_{j+1}\big(\mathrm{Core}(S, \Sigma_{j+1})\big)$. By (6.1), $s \times \sigma$ is contained in $\kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big)$. Thus a surgery on Σ_j has deleted the intersection circle associated to this square. By Corollary 5.9, $s \times \sigma$ is part of a maximal Σ_j–boundary rectangle. That is, there is a component $\Pi \subseteq M - \Sigma$ for which $\sigma \in \partial\Pi$ so that the edge (s, Π) is a boundary edge of $s \times \sigma$ but not the boundary edge of any other square in $\kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big)$.

We also know that $s \times \sigma$ is not contained in $k_i\big(\mathrm{Core}(S_i, \Sigma)\big)$. Thus, if (s, Π) is an edge in $k_i\big(\mathrm{Core}(S_i, \Sigma)\big)$ then we have that (s, Π) is a free edge in $k_i\big(\mathrm{Core}(S_i, \Sigma)\big) \cap \kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big) \cong \mathrm{Core}(S_i, \Sigma_j)$ (Proposition 5.8). If this is not the case, then there is some $i_0 < i$ such that (s, Π) lies between two squares $s \times \sigma'$ and $s \times \sigma''$ that are part of a maximal S_{i_0}–boundary rectangle in $\mathrm{Core}(S_{i_0}, \Sigma)$ that is collapsed in the formation of $\mathrm{Core}(S_{i_0+1}, \Sigma)$. Then neither of these squares are in $k_i\big(\mathrm{Core}(S_i, \Sigma)\big)$ at least one of these squares is not in $\kappa_j\big(\mathrm{Core}(S, \Sigma_j)\big)$. However, this would contradict (6.1). Therefore, (s, Π) is a free edge in $\mathrm{Core}(S_i, \Sigma_j)$. $\qquad\square$

References

[Aou13]   T. Aougab. Uniform hyperbolicity of the graphs of curves. *Geom. Topol.*, 17(5):2855–2875, 2013.

[AS11]   Javier Aramayona and Juan Souto. Automorphisms of the graph of free splittings. *Michigan Math. J.*, 60(3):483–493, 2011.

[BBC10]   Jason Behrstock, Mladen Bestvina, and Matt Clay. Growth of intersection numbers for free group automorphisms. *J. Topol.*, 3(2):280–310, 2010.

[BF95]   Mladen Bestvina and Mark Feighn. Stable actions of groups on real trees. *Invent. Math.*, 121(2):287–321, 1995.

[Bow14]   B.H. Bowditch. Uniform hyperbolicity of the curve graphs. *Pacific J. Math.*, 269(2):269–280, 2014.

[CH14]   Thierry Coulbois and Arnaud Hilion. Rips induction: index of the dual lamination of an $\mathbb{R}$-tree. *Groups Geom. Dyn.*, 8(1):97–134, 2014.

[Coo87]   Daryl Cooper. Automorphisms of free groups have finitely generated fixed point sets. *J. Algebra*, 111(2):453–456, 1987.

[CRS14]   M. Clay, K. Rafi, and S. Schleimer. Uniform hyperbolicity of the curve graph via surgery sequences. *Algebr. Geom. Topol.*, 14(6):3325–3344, 2014.

[CV86]   Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.*, 84(1):91–119, 1986.

[Gui05]   Vincent Guirardel. Cœur et nombre d'intersection pour les actions de groupes sur les arbres. *Ann. Sci. École Norm. Sup. (4)*, 38(6):847–888, 2005.

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

[HH15]   A. Hilion and C. Horbez. The hyperbolicity of the sphere complex via surgery paths. *J. reine angew. Math., Ahead of Print*, 2015. DOI 10.1515/crelle-2014-0128.

[HM13]   Michael Handel and Lee Mosher. The free splitting complex of a free group, I: hyperbolicity. *Geom. Topol.*, 17(3):1581–1672, 2013.

[HOP14]   Sebastian Hensel, Damian Osajda, and Piotr Przytycki. Realisation and dismantlability. *Geom. Topol.*, 18(4):2079–2126, 2014.

[Hor12]   Camille Horbez. Sphere paths in outer space. *Algebr. Geom. Topol.*, 12(4):2493–2517, 2012.

[HPW15] S. Hensel, P. Przytycki, and R. Webb. 1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs. *J. Eur. Math. Soc. (JEMS)*, 17(4):755–762, 2015.

[HV96]   Allen Hatcher and Karen Vogtmann. Isoperimetric inequalities for automorphism groups of free groups. *Pacific J. Math.*, 173(2):425–441, 1996.

[Lau74]   François Laudenbach. *Topologie de la dimension trois: homotopie et isotopie*. Société Mathématique de France, Paris, 1974. With an English summary and table of contents, Astérisque, No. 12.

[PS15]   P. Przytycki and A. Sisto. A note on acylindrical hyperbolicity of mapping class groups. arXiv:1502.02176, 2015.

Dept. of Mathematical Sciences, University of Arkansas, Fayetteville, Arkansas, USA 72701
  E-mail address: mattclay@uark.edu

Dept. of Mathematics, University of Toronto, Toronto, Ontario, Canada M5S 2E4
  E-mail address: yulan.qing@utoronto.ca

Dept. of Mathematics, University of Toronto, Toronto, Ontario, Canada M5S 2E4
  E-mail address: rafi@math.toronto.edu