A PROOF OF THE GORDON CONJECTURE

RUIFENG QIU AND MARTIN SCHARLEMANN

Abstract. A combinatorial proof of the Gordon Conjecture: The sum of two Heegaard splittings is stabilized if and only if one of the two summands is stabilized.

1. Introduction and basic background

In 2004 the first author [Q] presented a proof of the Gordon Conjecture, that the sum of two Heegaard splittings is stabilized if and only if one of the two summands is stabilized. The same year, and a bit earlier, David Bachman [Ba] presented a proof of a somewhat weaker version, in which it is assumed that the summand manifolds are both irreducible. (A later version dropped that assumption.)

The proofs in [Q] and [Ba] are quite different. The former is heavily combinatorial, essentially presenting an algorithm that will create, from a pair of stabilizing disks for the connected sum Heegaard splitting, an explicit pair of stabilizing disks for one of the summands. (Earlier partial results towards the conjecture, e. g. [Ed] have been of this nature.) In contrast, the proof in [Ba] is a delicate existence proof, based on analyzing possible sequences of weak reductions of the connected sum splitting. Both proofs have been difficult for topologists to absorb. The present manuscript arose from the second author's efforts, following a visit to Dalian in 2007, to simplify and clarify the ideas in [Q]. (During that visit, MingXing Zhang was very helpful in providing the groundwork for this simplified version.)

The most important strategic change here is an emphasis on symmetry. In [Q] the roles of the two stabilizing disks on opposite sides of the summed Heegaard surface are quite different. Here symmetry between the sides is maintained for as long as possible. (In fact until Propositon 10.1.) This adds a bit of complexity to the argument, but also some major efficiencies.

The figures in this manuscript are meant to be viewed in color; readers confused by figures in a black-and-white version may find it helpful to look at an electronic version, e. g http://www.math.ucsb.edu/∼mgscharl/papers/Qiu3.pdf .

Since the argument easily extends to Heegaard splittings of bounded manifolds, for convenience we restrict to closed 3-manifolds.

A *Heegaard splitting* of a closed orientable 3-manifold M is a description of M as the union of two handlebodies along their homeomorphic boundary. That is $M = \mathcal{V} \cup_S \mathcal{W}$, where $\mathcal{V}$ and $\mathcal{W}$ are handlebodies and $S = \partial\mathcal{V} = \partial\mathcal{W}$. The splitting

Date: July 17, 2009.

Key words and phrases. Heegaard splitting, Heegaard stabilization.

First author supported by NSFC No. 10625102. Second author partially supported by NSF grant DMS-0706740.

is *stabilized* if there are properly embedded disks $V \subset \mathcal{V}$ and $W \subset \mathcal{W}$ so that $\partial V \cap \partial W$ is a single point in S.

Suppose $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ and $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ are two Heegaard split 3-manifolds. There is a natural way to obtain a Heegaard splitting $M = \mathcal{V} \cup_S \mathcal{W}$ for the connect sum $M = M_+ \# M_-$, where $S = S_+ \# S_-$: Remove a 3-ball $B^3_\pm$ from each of $M_\pm$, a ball that intersects $S_\pm$ in a single 2-disk $D_\pm$. Then attach ∂B^3_+ to ∂B^3_- so that the disk $B_{\mathcal{V}} = \partial B^3_+ \cap V_+$ is identified to the disk $\partial B^3_- \cap V_-$, the disk $B_{\mathcal{W}} = \partial B^3_+ \cap W_+$ is identified to the disk $\partial B^3_- \cap W_-$ and so ∂D_+ is identified to ∂D_- to create $S = S_+ \# S_-$. This gives a Heegaard splitting $M = \mathcal{V} \cup_S \mathcal{W}$ with $\mathcal{V} = \mathcal{V}_+ \natural_{B_{\mathcal{V}}} \mathcal{V}_-$ and $\mathcal{W} = \mathcal{W}_+ \natural_{B_{\mathcal{W}}} \mathcal{W}_-$. See Figure 1 (surfaces P and F to be explained later.)

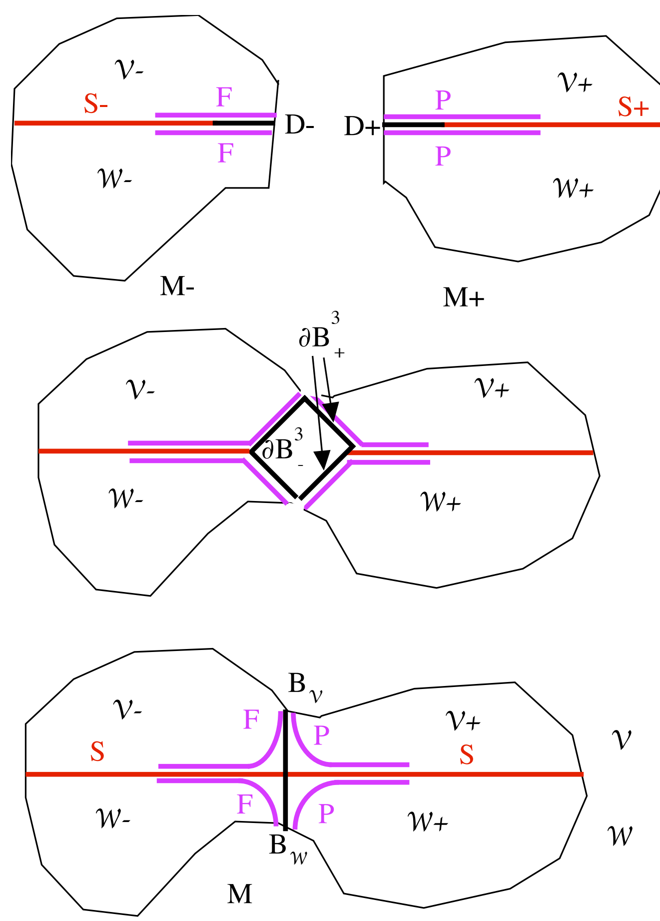

Figure 1

In problem 3.91 of the Kirby problem [Ki] list Gordon conjectured:

Conjecture 1.1. *$\mathcal{V} \cup_S \mathcal{W}$ is stabilized if and only if either $\mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ or $\mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ is stabilized.*

One direction of implication is obvious: a pair of stabilizing disks in $\mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ or $\mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ becomes a pair of stabilizing disks in $\mathcal{V} \cup_S \mathcal{W}$. The interest is in the opposite direction.

2. The framework part 1: Rooted forests of disks in handlebodies

Definition 2.1. *A* rooted tree *is a tree with a distinguished vertex called the root. A* coherent numbering *of the vertices of a rooted tree is a numerical labelling of the vertices $\alpha_i, i \in \mathbb{N} \cup \{0\}$ that increases along paths that move away from the root. That is, if the path in the tree from the root to the vertex α_i passes through the vertex α_k (or if α_k is the root) then $k < i$.*

A rooted forest *is a collection of rooted trees, one of which contains a distinguished root α_0. A* coherent numbering *of the vertices of a rooted forest is a numerical labelling of the vertices $\alpha_i, i \in \mathbb{N} \cup \{0\}$ which restricts to a coherent numbering in each of the rooted trees.*

Given an arbitrary forest with a distinguished root, it is easy to assign a coherent numbering: imagine the forest as a real forest in a hilly region with the distinguished root the lowest of all roots and the branches of all trees in the forest ascending upward. Take a generic height function on the forest and assign numbers to each vertex in order of their height. Feel free to skip some numbers; there is no requirement that the set of numbers assigned to vertices is contiguous in $\mathbb{N} \cup \{0\}$. Numbers that are assigned to vertices will be called *active* numbers.

Examples: A rooted tree with coherent numbering is clearly also a rooted forest with coherent numbering. Delete a vertex (other than the root) from a coherently numbered rooted tree and also delete all contiguous edges. The result is a coherently numbered rooted forest $\mathcal{F}$, with as many components as the valence of the vertex that is removed. The root of each component of $\mathcal{F}$ that does not contain the original root (now the distinguished root) is the vertex that was closest to the root in the original tree. More generally, if $\mathcal{F}$ is a coherently numbered rooted forest, and a vertex other than the distinguished root is removed, along with all contiguous edges, then the result is still a coherently numbered rooted forest, but with the number assigned to the vertex that has been removed now inactive.

Definition 2.2. *Suppose $\mathcal{V}_+, \mathcal{V}_-$ is a pair of disjoint handlebodies, and $P \subset \partial\mathcal{V}_+, F \subset \partial\mathcal{V}_-$ are subsurfaces of their respective boundaries. A* forest of disks *(modeled on the rooted forest $\mathcal{F}$) in the pair of handlebodies $\mathcal{V}_+, \mathcal{V}_-$ is a properly embedded collection of disks $\overline{V} = \{V_i\}$, one for each vertex α_i of $\mathcal{F}$ so that:*

1. *The disks alternate between lying in $\mathcal{V}_+$ and $\mathcal{V}_-$. That is, suppose vertices α_i, α_k are incident to the same edge in $\mathcal{F}$. Then $V_i \subset \mathcal{V}_+$ if and only if $V_k \subset \mathcal{V}_-$.*
2. *Suppose α_i is a vertex of $\mathcal{F}$ and $V_i \subset \mathcal{V}_+$ (resp $V_i \subset \mathcal{V}_-$) is the corresponding disk. If α_i is not a root, or is the distinguished root α_0, there is a one-to-one correspondence between the edges of $\mathcal{F}$ incident to α_i and arcs of $\partial V_i \cap P$ (resp $\partial V_i \cap F$). If α_i is a non-distinguished root then there is one extra arc of $\partial V_i \cap P$ (resp $\partial V_i \cap F$) called the* root arc.
3. *Corresponding to each root arc in $\partial V_i \cap P$ (resp $\partial V_i \cap F$) there is a normally oriented pair of properly embedded arcs in F (resp P) called* overpass arcs *(abbreviated* op-arcs*). The op-arcs are all disjoint, both from each other and from $\partial\overline{V}$. The collection of all op-arcs will be denoted ν.*

The pairs of op-arcs will be required to have certain properties, which will be discussed below (see the end of Section 3 and Section 5).

It will turn out that each surface P and F can be viewed as obtained from a disk by repeatedly plumbing it to itself. The role of the op-arcs will be to explicitly describe how the plumbing iss done. The reader is encouraged to jump ahead to Figure 9 to see how they play this role via a construction reminiscent of the building of a freeway overpass. That is, the op-arcs in P describe how to abstractly de-plumb P until it becomes a disk $\hat{P}$. The details of how this is accomplished will emerge as we proceed.

Seminal Example: Suppose the handlebody $\mathcal{V}$ is expressed as the ∂-connected sum of two handlebodies $\mathcal{V}_+$ and $\mathcal{V}_-$ along a disk D. That is $\mathcal{V} = \mathcal{V}_+ \natural_D \mathcal{V}_-$. Consider a ∂-reducing disk V in $\mathcal{V}$ and a distinguished point $x_0 \in \partial V$. It's easy to isotope V rel ∂V so it intersects D only in arcs. Then the components of $V - \eta(D)$ are disks.

Here is a natural description of a tree $\mathcal{T}$ embedded in the disk V: For vertices, choose a point in the interior of each disk component of $V - D$. For edges, choose, for each arc of $V \cap D$, an arc connecting the two vertices in the components of $V - D$ incident to that arc. Define the root of $\mathcal{T}$ to be the vertex α_0 that lies in the component of $V - D$ that has x_0 in its boundary. Then the components of $V - \eta(D)$ constitute a tree of disks in $\mathcal{V}_+ \cup \mathcal{V}_-$, modeled on $\mathcal{T}$, with P the copy of D in $\partial \mathcal{V}_+$ and F the copy of D in $\partial \mathcal{V}_-$. Since the only root is the distinguished root, there are no op-arcs.

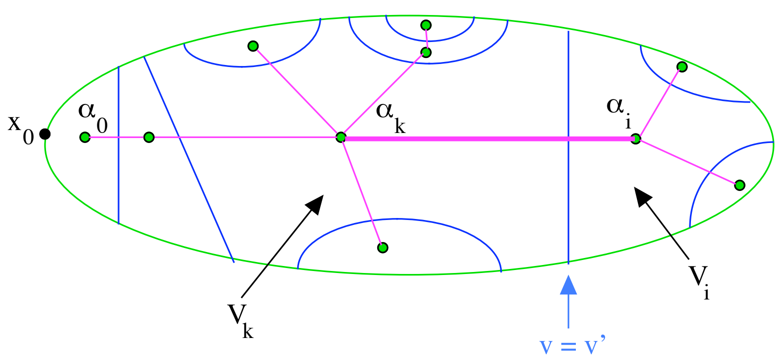

FIGURE 2

This example, though seminal to our discussion, is deceptive in two ways: First, in this example the surfaces $P \subset \partial \mathcal{V}_+$ and $F \subset \partial \mathcal{V}_-$ are simply two sides of the same surface (namely D) and so can be naturally identified. In general this will not be true. Second, and most deceptively, an edge in the tree $\mathcal{T}$ between two vertices, say α_i (representing $V_i \subset \mathcal{V}_-$) and α_k (representing $V_k \subset \mathcal{V}_+$) corresponds, as required, both to a component v of $\partial V_i \cap F$ and v' of $\partial V_k \cap P$. But what is true here and will not be true in general, is that both v and v' can be thought of as *the same* arc, namely a single component of $V \cap D$. In general (only in part because there will be no natural identification of P and F) the two arcs $v \subset \partial V_i$ and $v' \subset \partial V_k$

determined by a single edge in $\mathcal{T}$ will, at least *prima facie*, have nothing to do with each other.

Labeling convention: There is an efficient way to label the properly embedded arcs in F and P that come from a forest $\overline{V}$ of disks in $\mathcal{V}_+, \mathcal{V}_-$ that is modeled on a coherently numbered forest $\mathcal{F}$.

First note that there is a natural way to assign a unique label to each edge in the forest $\mathcal{F}$, namely give each edge the label of the vertex at its end that is most distant from the root. That is, if the edge in $\mathcal{F}$ has ends at vertices α_i and α_k, with α_k either closer to the root or perhaps the root itself, so $k < i$, then label the edge e_i.

As discussed in the example above, each edge e_i in $\mathcal{F}$ actually represents two arcs since e_i is incident to two vertices e_k and e_i in $\mathcal{F}$. One arc is in $\partial\overline{V} \cap P \subset \partial\mathcal{V}_+$ and the other is an arc in $\partial\overline{V} \cap F \subset \partial\mathcal{V}_-$. If, say, $V_i \subset \mathcal{V}_+$, so $V_k \subset \mathcal{V}_-$ then one end of e_i corresponds, under Definition 2.2, to an arc of $\partial V_i \cap P$, and the other end of e_i corresponds to an arc of $\partial V_k \cap F$. It is natural to call these arcs v_i^+ and v_i^- respectively, though it is perhaps counterintuitive that with this convention, $v_i^- \subset \partial V_k$. Symmetrically, if $V_i \subset \mathcal{V}_-$, so $V_k \subset \mathcal{V}_+$ then the arc of $\partial V_i \cap F$ corresponding to the end of e_i at α_i is called v_i^-, and the arc of $\partial V_k \cap P$ corresponding to the end of e_i at α_k is called v_i^+.

Now extend this labeling in the natural way to the root arcs and op-arcs: If $V_i \subset \mathcal{V}_+$ (resp $\mathcal{V}_-$) is a non-distinguished root, label the root arc $v_i^+ \subset P$ (resp $v_i^- \subset F$). Label the corresponding *pair* of op-arcs in F (resp P) by v_i^- (resp v_i^+). See Figure 3 for how this labels arcs in the Seminal Example and Figure 4 for how the labelling may appear on $\partial\mathcal{V}_+ \cup \partial\mathcal{V}_-$ in the more general case.

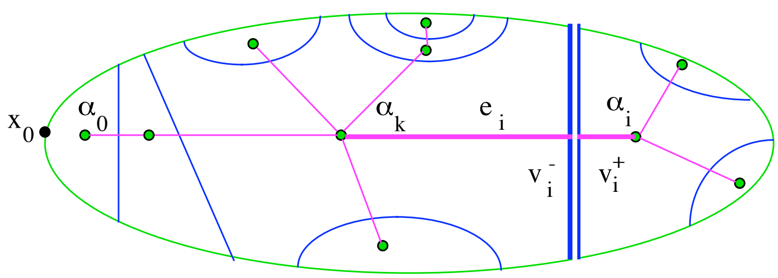

FIGURE 3

3. THE FRAMEWORK PART 2: STABILIZING FORESTS OF DISKS IN A HEEGAARD SPLITTING

We now extend this construction to a pair of Heegaard split 3-manifolds:

Definition 3.1. *Suppose $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$, $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ are two closed orientable 3-manifolds, and $P \subset S_+$ and $F \subset S_-$ are subsurfaces. Suppose $\overline{V} = \{V_i\}$ and $\overline{W} = \{W_j\}$ together with associated op-arcs ν, ω are forests of disks in the pairs $\mathcal{V}_+, \mathcal{V}_-$ and $\mathcal{W}_+, \mathcal{W}_-$ respectively. Let $\{v_i^\pm\}, \{w_j^\pm\}$ be the collections of arcs labeled as described above. (Some of these are root arcs, some of them pairs of op-arcs; typically they are arcs in $(\partial\overline{V} \cup \partial\overline{W}) \cap P$ and $(\partial\overline{V} \cup \partial\overline{W}) \cap F$).*

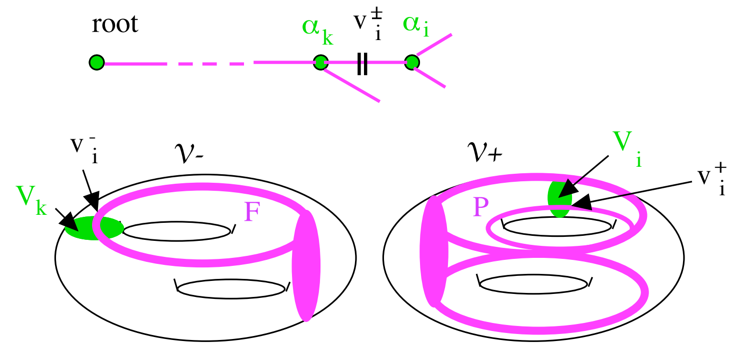

FIGURE 4

Suppose that only a single point $x_0 \in \partial\overline{V} \cap \partial\overline{W} \subset S_+ \cup S_-$ *lies outside of* $F \cup P$*, a point that lies in* $\partial V_0 \cap \partial W_0$*. Suppose further that all op-arcs among the* $\{v_i^{\pm}\}$ *(denoted* ν*) are disjoint from all op-arcs among the* $\{w_j^{\pm}\}$ *(denoted* ω*). Then* $\overline{V}, \overline{W}$ *together with associated op-arcs* ν, ω *is a* stabilizing pair *of forests of disks for* $\mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ *and* $\mathcal{V}_- \cup_{S_-} \mathcal{W}_-$*.*

Seminal example: Suppose $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ and $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ are two Heegaard split 3-manifolds and $M = \mathcal{V} \cup_S \mathcal{W}$ is the connected sum splitting on $M = M_+ \# M_-$, where $S = S_+ \# S_-$. Suppose that $\mathcal{V} \cup_S \mathcal{W}$ is a stabilized splitting. Then there are disks $V \subset \mathcal{V}$ and $W \subset \mathcal{W}$ so that $\partial V \cap \partial W$ is a single point $x_0 \in S$. Following the Seminal Example for handlebodies above, V and W give rise to rooted trees of disks $\overline{V}$ and $\overline{W}$ in the pairs $\mathcal{V}_+, \mathcal{V}_-$ and $\mathcal{W}_+, \mathcal{W}_-$ respectively. These rooted trees, having no non-distinguished roots, also have no op-arcs ν or ω. The original Heegaard splittings for $M_\pm$ are obtained from this picture of $\mathcal{V}_- \cup_{(S_- - B_-)} \mathcal{W}_-$ and $\mathcal{V}_+ \cup_{(S_+ - B_+)} \mathcal{W}_+$ by identifying the disks $B_{\mathcal{V}}$ with $B_{\mathcal{W}}$ in both manifolds. The resulting disk in S_- we regard as F and the resulting disk in S_+ we regard as P. Except for x_0, all intersections between $\overline{V}$ and $\overline{W}$ lie where the disks $B_{\mathcal{V}}$ and $B_{\mathcal{W}}$ have been identified, namely in the disk $F \subset S_-$ and the disk $P \subset S_+$. Hence $\overline{V}$ and $\overline{W}$ constitute a stabilizing pair of forests for the pair of Heegaard split manifolds $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ and $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$.

In this example, there is a clear connection between how the boundaries of the disks $\overline{V}$ and $\overline{W}$ intersect in S_+ and how they intersect in S_-. Consider a pair of arcs v_i^+ and w_j^+ in the disk P, arcs isotoped rel their boundary points in ∂P to intersect minimally. Then the arcs intersect (in precisely one point) if and only if the pair of points ∂v_i^+ separate the pair of points ∂w_j^+ in the circle ∂P. Since, in this example, v_i^+ and v_i^- are copies of the same arc of $V \cap B_{\mathcal{V}}$ (and, symmetrically, w_j^+ and w_j^- are copies of the same arc of $W \cap B_{\mathcal{W}}$), the pair of points ∂v_i^+ separate the pair of points ∂w_j^+ in the circle $\partial P = \partial D_+$ if and only if the pair of points ∂v_i^- separate the pair of points ∂w_j^- in the circle $\partial F = \partial D_-$. To summarize, $|v_i^+ \cap w_j^+| = |v_i^- \cap w_j^-| \leq 1$.

The relation between $|v_i^+ \cap w_j^+|$ and $|v_i^- \cap w_j^-|$ is more complicated in the general case. To begin with, as mentioned above, the arcs $v_+ \subset P$ and $v_- \subset F$ may not have anything to do with each other. Moreover, since P (resp F) is an arbitrary subsurface of S_+ (resp S_-), two proper arcs, even when isotoped rel boundary to intersect minimally, may still intersect in a large number of points.

Complicating things further, one of $v_i^\pm$ (or $w_j^\pm$) may represent a pair of op-arcs, about which so far we've said only this: Each pair of op-arcs, say $v_i^+ = v_i^a \cup v_i^b \subset P$, is normally oriented, disjoint from all other arcs $v_k^+ \subset P$ and also disjoint from all pairs of op-arcs $w_j^+ \in \omega \subset P$. We now introduce two properties which describe how such a pair of op-arcs v_i^+ is assumed to intersect the remaining arcs $\partial\overline{W} \cap P$, the arcs w_j^+ that are not themselves op-arcs. Symmetric statements apply to pairs of op-arcs $v_k^- \subset F$ and pairs of op-arcs $w_k^+ \subset P$ and $w_k^- \subset F$.

Near $v_i^+ = v_i^a \cup v_i^b$ in P, call the side of v_i^+ towards which the normal orientation of v_i^+ points the *inside of* v_i^+ and the other side the *outside* of v_i^+.

Separation Property of op-arcs: Suppose $v_i^+ = v_i^a \cup v_i^b$ is a pair of op-arcs in P. Then for any arc w_j^+, each component of $w_j^+ - v_i^+$ has ends

1. both incident to the outside of v_i^+ or
2. both incident to ∂P (when v_i^+ and w_j^+ are disjoint) or
3. one incident to ∂P and one incident to the outside of v_i^+ or
4. one incident to the inside of v_i^a and one incident to the inside of v_i^b

Subarcs of w_j^+ of the last type are said to be *on the overpass* associated to v_i^+ and, in analogy to railroad ties, are called *op-ties* for the overpass. See Figure 5. Components of $w_j^+ - v_j^+$ of the first three types are said to be *off the overpass* associated to v_i^+.

Parallelism Property of op-ties: Suppose $v_i^+ = v_i^a \cup v_i^b$ is a pair of op-arcs in P. Then all op-ties for the overpass associated to v_i^+ are parallel. To be explicit: suppose α and α' are two components of $\partial\overline{W} - v_i^+$ and each has one end incident to the inside of v_i^a and the other incident to the inside of v_i^b. Then the rectangle in P formed by the union of α, α' and subarcs of v_i^a and v_i^b bounds a disk in P. See Figure 5.

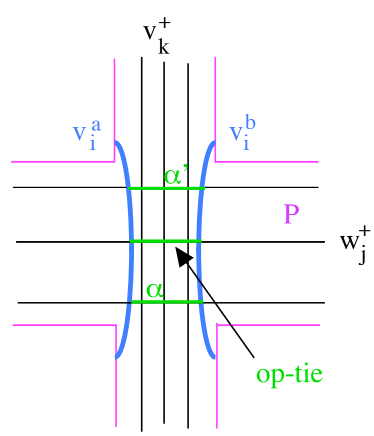

FIGURE 5

Two further properties of op-arcs that we assume will be given later. (See Section 5.) For now, we only introduce a useful definition:

Definition 3.2. *A component of $v_i - \omega$ that lies off of every ω overpass is called a* ground arc *in v_i. Symmetrically, a component of $w_j - \nu$ that lies off of every ν overpass is called a* ground arc *in w_j.*

Note that since the op-arcs ν and ω are disjoint, each overpass arc is, perhaps counterintuitively, a ground arc.

4. The first pairings $\rho_\pm$

We have seen that the Separation Property guarantees that for each arc w_j^+ and pair of op-arcs $v_i^+ = v_i^a \cup v_i^b$, $|v_i^+ \cap w_j^+| = 2|v_i^a \cap w_j^+| = 2|v_i^b \cap w_j^+|$. With this in mind, the following is a natural definition.

Definition 4.1. *Let $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$, $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$, $P \subset S_+$ and $F \subset S_-$ be as above. Suppose the families of disks $\overline{V}, \overline{W}$ and associated op-arcs ν, ω is a stabilizing pair of coherently numbered forests for the pair of Heegaard splittings. Define two pairings $\rho_\pm : \mathbb{N} \times \mathbb{N} \to \mathbb{N} \cup \{0\}$ by*

- *$\rho_\pm(i,j) = |v_i^\pm \cap w_j^\pm|$ when neither $v_i^\pm$ nor $w_j^\pm$ is a pair of op-arcs or*
- *$\rho_\pm(i,j) = |v_i^\pm \cap w_j^\pm|/2$ when either $v_i^\pm$ or $w_j^\pm$ is a pair of op-arcs*
- *$\rho_\pm(i,j) = 0$ if $v_i^\pm$ or $w_j^\pm$ is not defined, e. g. if i (resp j) is not among the indices of the disks in the rooted forest $\overline{V}$ (resp. $\overline{W}$). That is, when i (resp j) is an inactive index.*

Explanatory notes: Here $|v_i^+ \cap w_j^+|$ (resp $|v_i^- \cap w_j^-|$) means the number of intersection points, minimized by isotopy rel boundary, of the two arcs v_i^+ and w_j^+ in P (resp $v_i^- \cap w_j^-$ in F). See Figure 6. If v_i^+ is a pair of op-arcs in P then we have seen that $\rho_+(i,j)$ is the number of intersections of w_j^+ with either one v_i^a or v_i^b of the op-arc pair v_i^+ (and symmetrically for a pair of op-arcs v_i^- in F or a pair of op-arcs $w_j^+ \subset P$ or $w_j^- \subset F$).

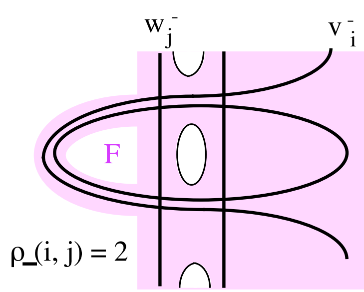

Figure 6

Definition 4.2. *Suppose $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ and $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$; $P \subset S_+$ and $F \subset S_-$; and forests of disks $\overline{V}, \overline{W}$ and associated op-arcs ν, ω are all given as above. A pair (i,j) is* peripheral *if for all $(i', j') \neq (i,j)$ with $i' \geq i$ and $j' \geq j$, $\rho_+(i', j') = \rho_-(i', j') = 0$.*

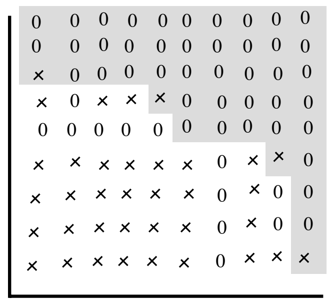

FIGURE 7. Grey shows peripheral lattice points (i, j); $\times$ means a non-zero entry

Peripheral pairs clearly exist: merely choose i and j larger than any index that appears among those for disks in the forest. Slightly less obvious is this: unless the pairings $\rho_\pm$ are both identically zero there is a peripheral pair (i_0, j_0) so that at least one of $\rho_\pm(i_0, j_0)$ is non-trivial. Simply choose j_0 to be the largest value for which there is an i with $\rho_\pm(i, j_0) \neq 0$. Then, among all such i, let i_0 be the largest.

Lemma 4.3. *Suppose* (i, j) *is peripheral and*

- $\rho_+(i, j) = 1$
- $V_i \subset \mathcal{V}_+$
- $W_j \subset \mathcal{W}_+$

Then $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ *is a stabilized splitting.*

Symmetrically, if $\rho_-(i, j) = 1$, $V_i \subset \mathcal{V}_-$ *and* $W_j \subset \mathcal{W}_-$, *then* $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ *is a stabilized splitting.*

Proof. $V_i \subset \mathcal{V}_+$ and $W_j \subset \mathcal{W}_+$ will be the stabilizing disks. By the labeling convention, $\partial V_i \cap P$ consists of arcs v_i^+ and (possibly) other arcs $v_{i'}^+, i' > i$. Similarly, $\partial W_j \cap P$ consists of the arc w_j^+ and (possibly) other arcs $w_{j'}^+, j' > j$. Since (i, j) is peripheral, each $v_{i'}^+$ with $i' > i$ is disjoint from ∂W_j and each $w_{j'}^+$ with $j' > j$ is disjoint from ∂V_i. Hence the only points in $\partial V_i \cap \partial W_j$ are those in $v_i^+ \cap w_j^+$. Since $\rho_+(i, j) = 1$, there is exactly one such point. Hence $\partial V_i \cap \partial W_j$ is a single point, so $\mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ is a stabilized splitting. $\square$

5. Further properties of the op-arcs

We now introduce two further properties which pairs of op-arcs are assumed to satisfy. Since there are no op-arcs in the Seminal Example above, these properties are vacuously satisfied in that example. Part of the argument will be to show that the fundamental construction described below preserves all these properties of pairs of op-arcs. This section describes properties of op-arcs v_i^+ in P; symmetric statements are true for op-arcs w_j^+ in P and op-arcs v_i^- and w_j^- in F. It may be helpful, when v_i^+ is specifically meant to be a pair of op-arcs, to denote it ν_i and when w_j^+ is meant to be a pair of op-arcs, denote it by ω_j.

Ordering Property for op-ties: Suppose $\alpha \subset \partial\overline{W}$ is an op-tie for the pair of op-arcs ν_i. For any $k \geq i$, v_k^+ is disjoint from the interior of α.

In particular, suppose $\nu_k = v_k^+, k > i$ is also a pair of op-arcs ν_k, and α' is an op-tie for the overpass associated with ν_k with $\alpha \cap \alpha' \neq \emptyset$. Then neither end of the arc $\alpha \cap \alpha'$ can lie on ν_k, so both ends lie on ν_i and $\alpha \subset \alpha'$. See Figure 8.

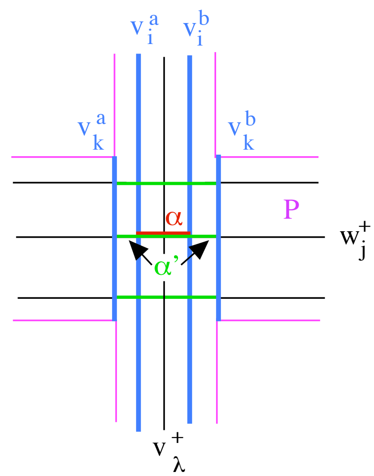

FIGURE 8. Here the Ordering Property implies $\lambda < i < k$.

The last property of op-arc pairs, stated below, requires some background: Here is a way of using the pairs of op-arcs in P to construct a new surface $\hat{P}$. This construction will be called *building the overpasses*. It is parallel to the idea of an "abstract tree" for F and P, found in [Q]. It's important to understand that, as with the abstract tree, the building of overpasses is done in the abstract, as a way to express a property of the op-arcs, and is *not* a construction actually performed inside of M_+ or M_-. Here we describe how to build a single overpass, one associated to the pair of op-arcs $\nu_i = \nu_i^a \cup \nu_i^b$. First cut P along ν_i. The resulting surface P' has two copies of ν_i^a and two copies of ν_i^b in its boundary. One copy of ν_i^a in $\partial P'$ is incident to the outside of ν_i^a in P and one copy of ν_i^b is incident to the outside of ν_i^b in P. Identify these two arcs in $\partial P'$ and call the resulting arc $v_i''^+$ and surface P''. The other copies of ν_i^a and ν_i^b remain in $\partial P''$.

Building the overpass as described does nothing particularly interesting to the other arcs v_k, since these are disjoint from ν_i. The arcs $\{w_j^+\}$ that intersect ν_i are cut up when the overpass is built: Each arc in $\{w_j^+\}$ naturally gives rise to perhaps many properly embedded arcs in P' (each w_j^+ is cut into pieces by ν_i) and, less obviously, to a single special arc in P'': By the Parallelism Property, the ends of $w_j^+ - \nu_i$ lying just outside ν_i^a match naturally with the ends of $w_j^+ - \nu_i$ lying just outside ν_i^b and so can be attached in P'' to become a proper arc $w_j''^+$ in P''. A simple picture of the special arc $w_j''^+$ is that it is the arc obtained from w_j^+ by collapsing all the op-ties of w_j^+ that lie on the overpass associated to ν_i. The upshot is that, in P'', w_j^+ is fractured into a collection of op-ties, each now a proper arc in P'' and no longer indexed, plus a single arc $w_j''^+$ that is the end-point union of all subarcs of w_j^- that lie off the overpass.

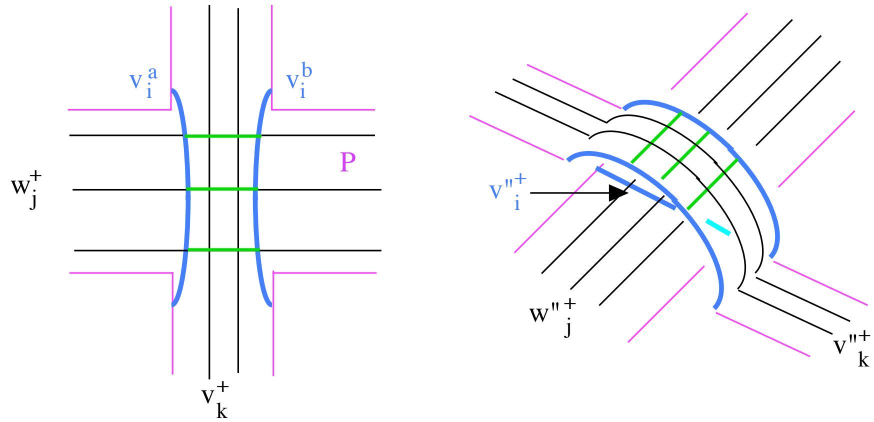

FIGURE 9. Building an overpass. Green arcs on right are unindexed

Define $\hat{P}$ to be the surface obtained from P by building all the overpasses at once. That is, perform the operation just described on all pairs of op-arcs ν and ω simultaneously. There is no ambiguity in the construction, since ν and ω are assumed to be disjoint. It may be worth noting (but is not important to the argument) that, following the Ordering Property above, when an op-tie α for ν_i overlaps with an op-tie α' for $\nu_k, k > i$ then $\alpha \subset \alpha'$. So in the construction of $\hat{P}$, α' is fractured into pieces by ν_i and the proper arc in $\hat{P}$ corresponding to α' is obtained by assembling those pieces that lie off the overpasses ν_i (and whatever other overpasses may pass through α').

Our final and most delicate assumption on op-arcs is:

Disk Property: The surface $\hat{P}$ is a disk.

In the disk $\hat{P}$ each curve v_i^+ and each curve w_j^+ may be fractured into many pieces. One component constructed out of w_j^+, for example, will be the end point union of all subarcs of w_j^+ that lie off of every overpass, i. e. the ground arcs of w_j^+. (See Definition 3.2.) This arc in $\hat{P}$ corresponding to w_j^+ is denoted $\hat{w}_j^+$ (and appears as $w"_j^+$ in Figure 9, where only one overpass is built). Another component constructed out of w_j^+ might be the end point union of all subarcs of w_j^+ that lie on the overpass determined by ν_k and off all overpasses determined by any $\nu_i, i < k$. We have no notation for such arcs, since arcs in $\hat{P}$ coming from op-ties will play no role in the argument. Among op-arcs, the pair of arcs ν_i (resp ω_j) in P becomes a single proper subarc of $\hat{P}$ which we denote $\hat{v}_i^+$ (resp. $\hat{w}_j^+$). (The curve $\hat{v}_i^+$ appears as $v"_i^+$ in Figure 9.) The union of all such curves (coming from ν and ω) in $\hat{P}$ will be denoted $\hat{\nu}$ and $\hat{\omega}$ respectively. Of course they are no longer op-arcs in $\hat{P}$ because all overpasses have been built there.

6. The pairing σ of arcs in $\hat{P}$

Lemma 6.1. *Any two arcs $\hat{v}_i^+$ and $\hat{w}_j^+$ (resp. $\hat{v}_i^-$ and $\hat{w}_j^-$) intersect efficiently in $\hat{P}$ (resp $\hat{F}$). That is, $|\hat{v}_i^+ \cap \hat{w}_j^+|$ cannot be reduced by isotopies of $\hat{v}_i^+$ and $\hat{w}_j^+$ in $\hat{P}$ rel $\partial \hat{P}$.*

Proof. We must show that no complementary component of the two curves in $\hat{P}$ is a bigon, that is, a disk bounded by the union of a subarc of $\hat{v}_i^+$ and a subarc of $\hat{w}_j^+$. Suppose, towards a contradiction, that there were such a bigon B. Since ν and ω, hence $\hat{\nu}$ and $\hat{\omega}$, are disjoint, at least one side of the bigon, say the side on $\hat{v}_i^+$ does not come from an op-arc.

Consider first the case in which the interior of B is disjoint from all curves $\hat{\nu} \cup \hat{\omega}$ coming from op-arcs. Then B would also lie in P since the interior of B is disjoint from the curves $\hat{\nu} \cup \hat{\omega}$ along which P was cut and glued. But the conclusion $B \subset P$ would violate our initial assumption that the curves $v_i^+ \subset P$ and $w_j^+ \subset P$ intersect efficiently in P.

Now suppose that the interior of B is not disjoint from $\hat{\nu} \cup \hat{\omega}$. Since $\hat{v}_i^+$ intersects only op-arcs in $\hat{\omega}$ and $\hat{w}_j^+$ intersects only op-arcs in $\hat{\nu}$, any component of $\hat{\omega} \cap B$ (resp $\hat{\nu} \cap B$) would have both ends on $\hat{v}_i^+$ (resp $\hat{w}_j^+$). Hence there is a bigon $B' \subset B$ between a subarc of $\hat{v}_i^+$ (say) and a subarc of some $\hat{\omega}_k^+ \in \hat{\omega}$. Moreover, if B' is chosen to be an innermost such example, then the interior of B' would be disjoint from $\hat{\nu} \cup \hat{\omega}$. Then, just as above, B' would lie entirely in P and this would violate the initial assumption that the arcs v_i^+ and w_k^+ intersect efficiently in P. □

Following Lemma 6.1, the Disk Property leads naturally to a new pairing:

Definition 6.2. *Analogous to the intersection pairings $\rho_\pm$ in P and F define intersection pairings $\sigma_\pm : \mathbb{N} \times \mathbb{N} \to \mathbb{N} \cup \{0\}$ in the disks $\hat{P}$ and $\hat{F}$ by*

- *$\sigma_\pm(i,j) = |\hat{v}_i^\pm \cap \hat{w}_j^\pm|$ or*
- *$\sigma_\pm(i,j) = 0$ if $\hat{v}_i^\pm$ or $\hat{w}_j^\pm$ is not defined. (That is, if i or j is an inactive index.)*

Lemma 6.3. *For each (i,j)*

(1) *$\sigma_\pm(i,j) \leq \rho_\pm(i,j)$ and*
(2) *$\sigma_\pm(i,j) \leq 1$*

Proof. As usual, we focus on σ_+ defined on arcs in P; the case for σ_- defined on arcs in F is symmetric.

For the first claim, note that any intersection point of $\hat{v}_i^+$ with $\hat{w}_j^+$ in P is merely a particular type of intersection point of v_i^+ with w_j^+, namely one which is not on any overpass.

The second claim follows immediately from the fact that $\hat{P}$ is a disk and, following Lemma 6.1, the arcs $\hat{v}_i^+$ and $\hat{w}_j^+$ intersect efficiently in $\hat{P}$. □

Corollary 6.4. *If (i,j) is peripheral, then $\rho_\pm(i,j) = \sigma_\pm(i,j) \leq 1$.*

Proof. Following Lemma 6.3, the statement is obvious if $\rho_\pm(i,j) = 0$. Suppose, say, $\rho_+(i,j) > 0$, and $x \in v_i^+ \cap w_j^+$. If x were in the interior of any op-tie in w_j^+, coming from a pair of op-arcs ν_k, say, then it would follow from the Ordering Property that $k > i$. Then the ends of the op-tie would be points in $v_k^+ \cap w_j^+$, contradicting the fact that (i,j) is peripheral. Hence x lies on no overpass associated with any of the ν. Symmetrically, it's on no overpass associated with any of the ω. Hence $x \in \hat{v}_i^+ \cap \hat{w}_j^+ \subset \hat{P}$.

Summarizing, this shows that for any peripheral (i,j), $\sigma(i,j)_+ \geq \rho(i,j)_+$. The result then follows from Lemma 6.3. $\square$

Definition 6.5. *Suppose* $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$*,* $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ *and surfaces* $P \subset S_+$ *and* $F \subset S_-$ *are given as above and disks* $\overline{V}, \overline{W}$ *and associated op-arcs* ν, ω *is a stabilizing pair of coherently numbered forests for the pair of Heegaard splittings. Then the forests are* coordinated *if for all* $(i,j) \in \mathbb{N} \times \mathbb{N}$*,* $\sigma_+(i,j) = \sigma_-(i,j)$*.*

Seminal Example: For the Seminal Example, it was observed that for all $(i,j) \in \mathbb{N} \times \mathbb{N}$, $\rho_+(i,j) = \rho_-(i,j)$. But in that example there are no op-edges, so $\hat{P} = P, \hat{F} = F$. Then for all (i,j), $\sigma_+(i,j) = \rho_+(i,j) = \rho_-(i,j) = \sigma_-(i,j)$. Hence the forests of disks in the Seminal Example are coordinated.

7. A digression on some operations on curves and surfaces

Suppose A is an annulus containing a core circle c and two spanning arcs e and w. Suppose λ_w is a proper arc in A that intersects w once and is disjoint from c and e. Then there is an arc λ_e in A, unique up to isotopy rel ∂, that has the same ends as λ_w but is disjoint from c and w and intersects e. One way of describing how λ_e is derived from λ_w is to band-sum λ_w to c along w. The same is true if λ_w consists of a disjoint family of arcs in A, each component of which intersects w in a single point and is disjoint from w and e. The change could be described as band-summing λ_w along w to c; as many copies of c are band-summed as there are components of λ_w. See Figure 10.

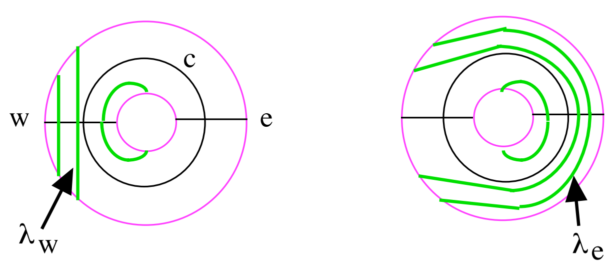

Figure 10

More generally, suppose that c is a simple closed curve in a surface P and w is a properly embedded arc in P that intersects c once. Suppose λ is a properly embedded 1-manifold in P that is disjoint from c and intersects w transversally. Then a small regular neighborhood $\eta(c \cup w) \subset P$ can be viewed as an annulus A in which w is a spanning arc, λ intersects A in proper arcs, each of which intersects w once, and each of which is disjoint from c and from a distant fiber of $\eta(c) \subset \eta(c \cup w)$. Performing the operation above to $\lambda \cap A$ will be called *band-summing* λ *to* c *along* w. See Figure 11.

Here is an additional feature of this band-sum operation. Suppose M is a 3-manifold and $P \subset \partial M$. Suppose there are proper disks C and D in M so that $\partial C = c$ and $\partial D = \lambda$. Then after the operation, λ still bounds a disk, one obtained

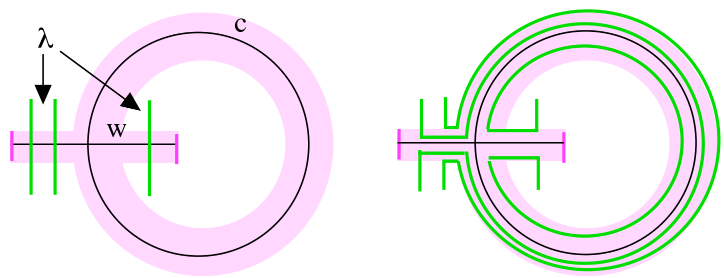

FIGURE 11

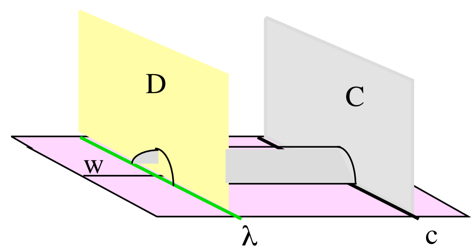

FIGURE 12

by boundary-summing ∂D to one copy of C for each point in $\lambda \cap w$. This operation will be called *tube-summing D to C along w.* See Figure 12.

Now suppose P is a compact orientable surface and $v, w \subset P$ are properly embedded arcs in P that meet at a single point. Define a new orientable surface P_{v-w} by the following operation: add a band to P with its ends attached at the pair of points $\partial v \subset \partial P$. Then remove a neighborhood of w.

P and P_{v-w} have the same Euler characteristic; whether they are homeomorphic or not then depends only on whether the operation changes the number of boundary components. In any case, we have:

Lemma 7.1. *There is a homeomorphism $\phi_{v,w} : P_{w-v} \to P_{v-w}$ that is the identity away from $\eta(v \cup w)$.*

Proof. The proof is illustrated in Figure 13. □

Suppose λ is a properly embedded curve in P, in general position with respect to w and disjoint from v. Then λ is unaffected by the operation that creates P_{w-v}. This observation then provides a natural embedding $\lambda \subset P_{w-v}$.

Lemma 7.2. *Let P_+ be the surface obtained from P by adding a band to P with its ends attached at the pair of points $\partial v \subset \partial P$. Let v_+ be the circle in P_+ which is the union of v and the core of the band. Then $\phi_{v,w}(\lambda) \subset P_{v-w} \subset P_+$ is the curve obtained from λ by band-summing λ along w to v_+.*

Proof. The proof is illustrated in Figure 14. □

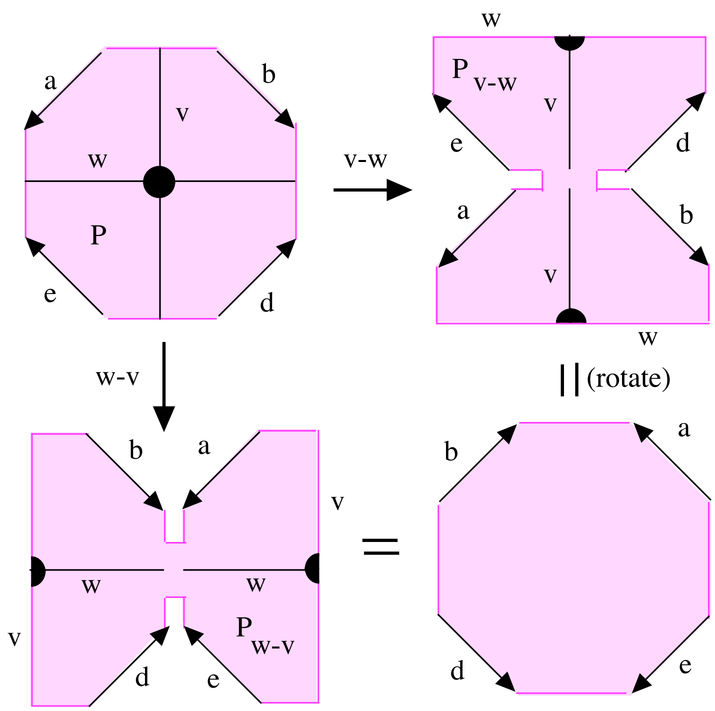

FIGURE 13. The arrows denote points of contact with the rest of the surface P.

8. The fundamental construction

Proposition 8.1. *Suppose $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$, $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ are Heegaard splittings. Suppose $P \subset S_+$ and $F \subset S_-$ are surfaces with respect to which a collection of disks $\overline{V} \cup \overline{W}$ and associated op-arcs $\nu \cup \omega$ is a stabilizing pair of coherently numbered coordinated forests of disks. Suppose further that for some peripheral (i, j) with $\rho_\pm(i, j) \neq 0$, $V_i \subset \mathcal{V}_+$ and $W_j \subset \mathcal{W}_-$ (or vice versa) and*

1. *$\partial V_i \cap P - v_i^+$ is disjoint from all op-arcs*
2. *$\partial W_j \cap F - w_j^-$ is disjoint from all op-arcs*
3. *either v_i^+ or $w_j^+ \subset P$ is disjoint from all op-arcs and*
4. *either v_i^- or $w_j^- \subset F$ is disjoint from all op-arcs.*

Then there are surfaces $P' \subset S_+$ and $F' \subset S_-$, with respect to which a collection of disks $\overline{V}' \cup \overline{W}'$ and associated op-arcs $\nu' \cup \omega'$ is a stabilizing coordinated pair of coherently numbered forests of disks. Moreover, there are fewer disks in $\overline{V}'$ than in $\overline{V}$ and fewer disks in $\overline{W}'$ than in $\overline{W}$.

Proof. We construct another stabilizing coordinated pair of coherently numbered forests of disks. We describe the construction in M_+ and later note the effect of the symmetric construction in M_-.

Start with the surface $P'' \supset P$ that is the union of P with a collar neighborhood $Y = \eta(\partial V_i)$ of ∂V_i in S_+. Since part of ∂V_i already lies in P, another way to view

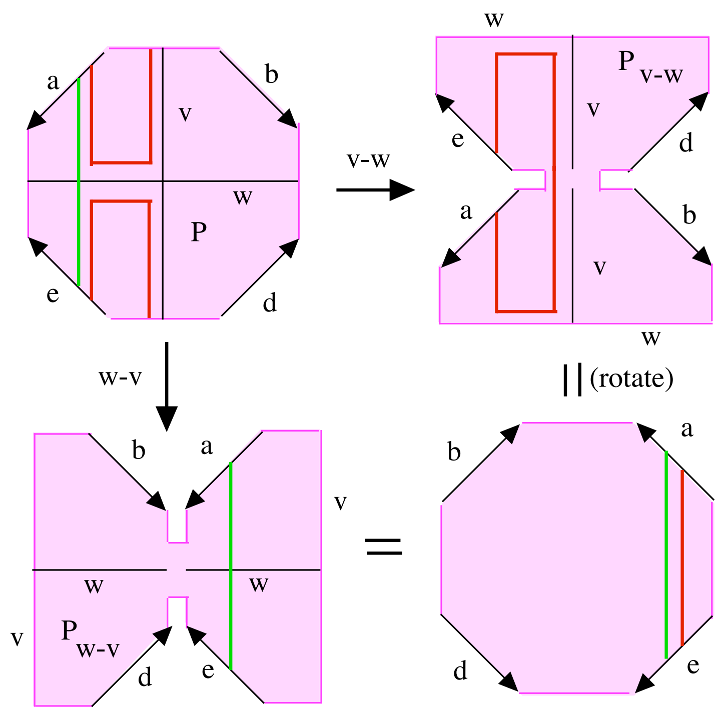

FIGURE 14

the construction of P'' from P is to add to P a band in $S_+ - P$ along each arc of $\partial V_i - P \subset S_+$.

We initially assume that w_j^+ is not a pair of op-arcs, but, like v_i^+, just a single proper arc in P. The arcs v_i^+, w_j^+ intersect in a single point, since by Corollary 6.4, $\rho_\pm(i,j) = 1$. Since (i,j) is peripheral, the arc w_j^+ may intersect other arcs v_ℓ^+ but only if $\ell < i$. Band-sum all such v_ℓ^+ along w_j to ∂V_i and call the result $v_\ell'^+ \subset P''$. If v_ℓ^+ was on the boundary of a disk in $\overline{V}$, tube-sum the disk to (copies of) V_i to obtain a corresponding disk in $\overline{V}'$. If v_ℓ^+ was a pair of op-arcs (so, by assumptions (1) and (3), ∂V_i is disjoint from all op-arcs ω) then $v_\ell'^+$ is a pair of op-arcs in P''. Although after this step $v_\ell'^+$ may not intersect all w_j^+ efficiently, it is straightforward to see that, when the pair $v_\ell'^+$ is isotoped in P'' to make all intersections efficient, the Separation, Parallel and Ordering Properties on the pair v_ℓ in P induce the same properties on the pair $v_\ell'^+$ in P''. New op-ties may have been introduced, each corresponding to an intersection point of some w_k^+ with an arc of $\partial V_i \cap P$. Now remove the original V_i from the collection of disks and call the result $\overline{V}'$.

After the operation described above, $w_j \subset P''$ is disjoint from all disks in $\overline{V}'$ and from all op-arcs in ν'. Let $P' = P'' - \eta(w_j)$. Augment the set of op-arcs ν' by adding the pairs of arcs $\partial Y \cap P$, one pair $v_k'^+$ for each arc v_k^+ in $\partial V_i \cap P - v_i^+$, and normally orient each $v_k'^+$ into Y. The assumptions of the proposition guarantee that the new pair of op-arcs $v_k'^+$ is disjoint from all other op-arcs and it is easy to see from the construction that it satisfies the Separation and Parallel Properties. See Figure 15.

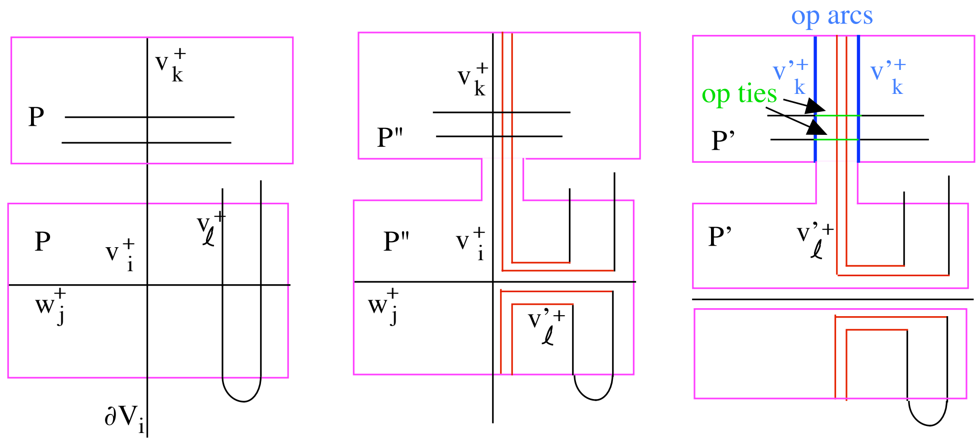

FIGURE 15

That such a pair of new op-arcs $v_k'^+$ satisfies the Order Property is only a little more difficult to show: By the coherence of the numbering, $v_k^+ \subset \partial V_i \cap P - v_i^+$ guarantees that $k > i$. The interior of each op-tie of the new pair of op-arcs $v_k'^+$ (corresponding to a point of $v_k^+ \cap \overline{W}$) intersects only those $v_\ell'^+$ that have been band-summed to v_i^+ along w_j^+, that is only those for which $v_\ell^+ \cap w_j^+ \neq 0$. Since (i, j) is peripheral, this implies $\ell < i$, hence $\ell < k$, as required.

If $w_j^+ = w_j^a \cup w_j^b$ is a pair of op-arcs, the construction is only slightly different. By the Parallelism Property, points of intersection of w_j^a with any v_ℓ^+ are paired to points of intersection of w_j^b by op-ties. So the band summing described above, using say the component w_j^a, in fact removes (when the intersections are made efficient) all points of intersection between $\partial\overline{V}'$ and w_j^b as well. So then *both* op-arcs w_j^a and w_j^b end up disjoint from $\partial\overline{V}'$ and neighborhoods of both should be removed. See Figure 16.

It is much more difficult to show that the new collection ν' of op-arcs still satisfies the Disk Property; that piece of the argument is postponed until later. (See Corollary 9.2.)

What is the effect of the construction described above on the forest of trees? Is the result a new pair of forests? First of all, w_j^+ disappears, so, if w_j^+ is not a pair of op-arcs, and so lies on ∂W_h for some disk $W_h \subset \mathcal{W}_+$ with $h < j$, then ∂W_h has one less arc of intersection with P'. Also, the disk $W_j \subset \mathcal{W}_-$ becomes the root of a tree with root arc w_j^-. Secondly the entire disk V_i^+ disappears, so each disk (in $\mathcal{V}_-$) whose vertex, in the forest, was adjacent to α_i away from v_i, becomes a root in the resulting forest, a root associated to the new pairs of op-arcs that we have created. But there are two immediately apparent defects: The arc or pair of op-arcs $v_i^- \subset F$ no longer has a matching arc $v_i^+ \subset P$, since V_i has been removed. Also, w_j^+ has been removed, whereas $w_j^- \subset \partial W_j$ remains as a root arc, violating the condition that each root arc in F is coordinated with a pair of op-arcs in P.

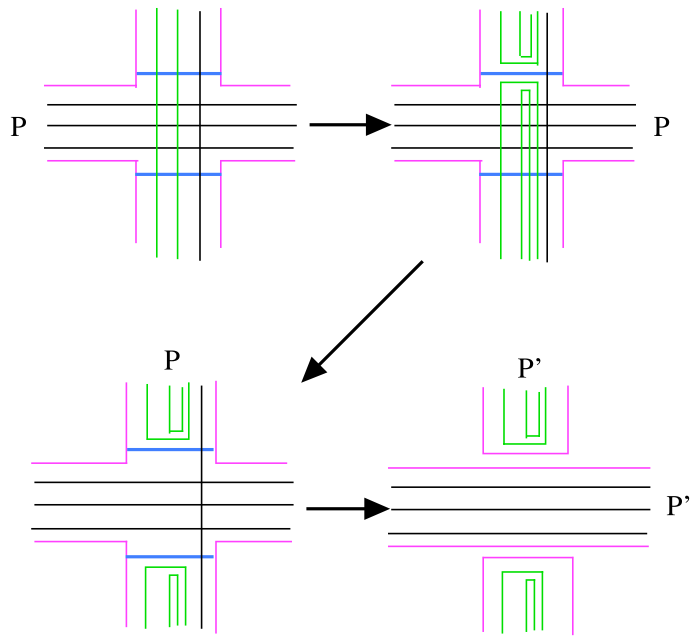

FIGURE 16

Both defects are overcome by doing the symmetric operation in M_- using now $\rho_-(i,j) = 1$. That is, tube-sum disks in $\mathcal{W}_-$ along v_i^- to W_j, alter F by removing a neighborhood of $v_i^- \subset F$ (thereby fixing the first defect) and add to F a neighborhood of the arcs $\partial W_j^- - F$. Then delete the disk W_j, fixing the second defect. See Figure 17.

We have shown that the new surfaces P' and F' and the new forests of disks satisfy all of the properties (except perhaps the Disk Property) of a coherently numbered stabilizing pair of forests of disks. The new forests have fewer disks since the disks V_i, W_j (corresponding to vertices α_i and β_j in the two forests) have been removed.

We now assume that the new framework also satisfies the Disk Property (we will show this later) and verify that then the forests are coordinated. That is,

Lemma 8.2. *For $\sigma'_\pm$ the new pairings in $\hat{P}'$ and $\hat{F}'$ constructed as above and for each $(\ell, k) \in \mathbb{N} \times \mathbb{N}$, $\sigma'_+(\ell, k) = \sigma'_-(\ell, k)$.*

Proof. Since the initial forests are coordinated, the statement is true before the construction. So the proof consists of showing that the construction process does not alter the relationship.

Whether or not any of the arcs $v_i^\pm$ or $w_j^\pm$ are op-arcs, all disappear from our accounting by the end of the construction, so they are irrelevant to the question. The focus is on other arcs, which may change during the construction. The curves that are altered (as P also is altered) in S_+ are the curves v_ℓ^+ which intersect w_j^+; those altered (as F also is altered) in S_- are the curves w_k^- which intersect v_i^-.

By Lemma 6.3 each number is either 0 or 1, so it suffices to prove

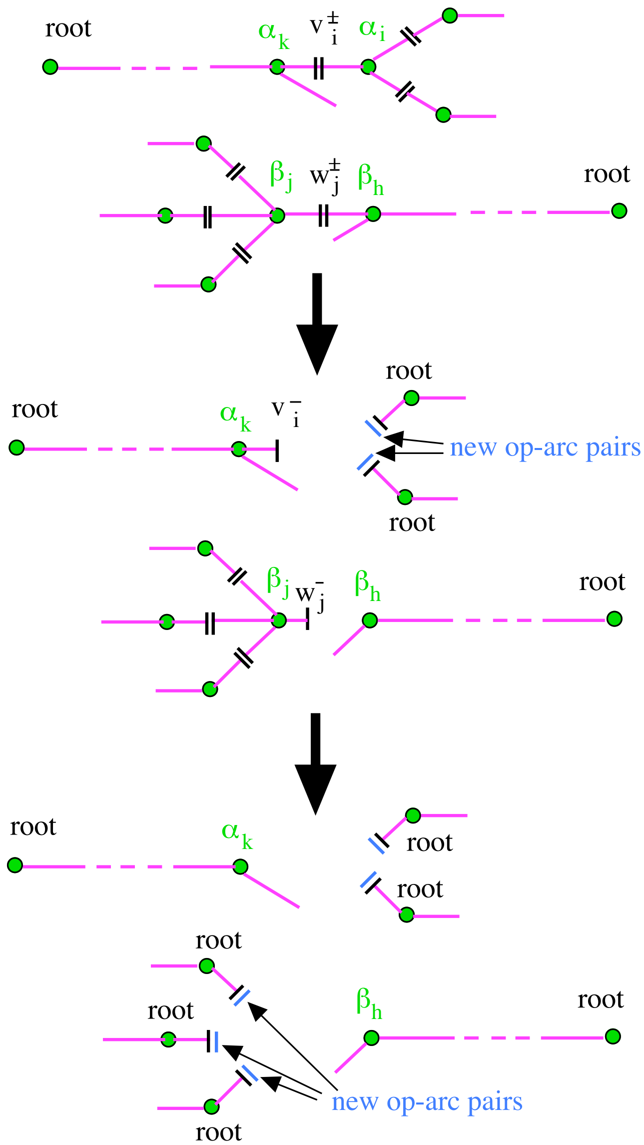

FIGURE 17

Claim: For all $\ell \neq i, k \neq j$ in $\mathbb{N}$,

$$\sigma'_{\pm}(\ell, k) \cong \sigma_{\pm}(\ell, k) + \sigma_{\pm}(\ell, j) \cdot \sigma_{\pm}(i, k) \; mod \; 2.$$

Proof of Claim: By symmetry, it suffices to show that this is true in S_+, that is for the intersection pairing σ'_+ on arcs in $\hat{P}'$.

Following Lemma 6.1 there is a way to accurately calculate $\sigma_+(\ell, k)$ in P. An intersection point of v_ℓ^+ with w_k^+ counts as a point in $\sigma_+(\ell, k)$ if and only if the point is not on any overpass from either ν or ω, that is the intersection point lies on a ground arc of both v_ℓ^+ and w_k^+. Similarly, *once* $v_\ell'^+$ *and* $w_k'^+$ *are isotoped rel* *boundary to intersect efficiently*, an intersection point of v'^+_ℓ with w'^+_k counts as a point in $\sigma'_+(\ell,k)$ if and only if the point is on ground arcs of both v'^+_ℓ and w'^+_k.

Since, to prove the claim, we only have to determine the parity of $\sigma'_\pm(\ell,k)$, the requirement that the arcs v'^+_ℓ and w'^+_k first be isotoped to intersect efficiently turns out to be irrelevant, as we now demonstrate. Two proper arcs in a surface can be isotoped to intersect efficiently by a sequence of isotopies, each removing a bigon of intersection. So, to demonstrate that this process does not change the parity of intersection points between ground arcs in v'^+_ℓ and ground arcs of w'^+_k, it suffices to show that for any bigon B in P' between a subarc α of v'^+_ℓ and a subarc β of w'^+_k, either both end points of α lie in ground arcs of v'^+_ℓ or neither does (and symmetrically for w'^+_k). It follows from the Separation Property that any subarc of v'^+_ℓ that has one end off an overpass and one end on must intersect the associated pair of op-arcs an odd number of times. On the other hand, since the op-arcs ω are disjoint from $\beta \subset w'^+_k$, any subarc of ω that lies in B must have both ends on α. That is, α intersects each op-arc in ω an even number of times. Hence if one end of α is on any overpass so is the other.

Since we do not have to make v'^+_ℓ and w'^+_k intersect efficiently, we need only count (parity of) points of intersection as they are originally constructed. We have already seen that the Ordering Property and the fact that (i,j) is peripheral guarantees that the point $x = v^+_i \cap w^+_j$ is a ground arc in both v^+_i and w^+_j.

Case 1: v_i is disjoint from all pairs of op-arcs in ω.

In this case, since x is on a ground arc of v_i, all of v_i is a ground arc. Suppose an intersection point $y \in v^+_\ell \cap w^+_j$ is on an overpass associated to a pair of op-arcs v^+_s. Then each of the pair of op-arcs v^+_s also intersects w^+_j, namely at the ends of the op-tie on which y lies. When these three arcs $(v^+_\ell \cup v^+_s)$ are band-summed to v^+_i the resulting subarc of v'^+_ℓ still lies entirely on the overpass associated to the new pair v'^+_s. So whatever intersections of v'^+_ℓ with w_k are created by this tube-summing do not count in $\sigma'(\ell,k)$. Hence in calculating how $\sigma'_+(\ell,k)$ differs from $\sigma_+(\ell,k)$ we can focus only on points of $v^+_\ell \cap w^+_j$ that lie in ground arcs of w^+_j. Similarly, we can focus only on points that lie in ground arcs of v^+_ℓ since if y is on an op-tie for some pair of op-arcs w^+_t, band-summing v^+_ℓ near y to v^+_i only creates a much longer tie, since v_i is disjoint from the op-arcs w^+_t; new points of intersection don't count in $\sigma'_+(\ell,k)$.

So the only relevant change caused by the construction could come from band-summing v^+_ℓ near a point z (unique, if it exists, by Lemma 6.3) in $v^+_\ell \cap w^+_j$ that lies on the ground arc of both curves. If no such point exists, then $\sigma_+(\ell,j) = 0$ and the number of intersection points of $v'^+_\ell \cap w^+_k$ that lie in ground arcs of each is unchanged. That is,

$$\sigma'_+(\ell,k) = \sigma_+(\ell,k) = \sigma_+(\ell,k) + 0 \cdot \sigma_+(i,k) = \sigma_+(\ell,k) + \sigma_+(\ell,j) \cdot \sigma_+(i,k)$$

as required. If $z \in v^+_\ell \cap w^+_j$ does lie on a ground arc of each, then the construction band-sums the ground arc of v^+_ℓ to v^+_i at x. It follows that the number of intersection points of ground arcs of v'^+_ℓ with ground arcs of w^+_k is increased by $\sigma_+(i,k)$ (before v'^+_ℓ is isotoped to have efficient intersection with w^+_k). That is,

$$\sigma'_+(\ell,k) = \sigma_+(\ell,k) + 1 \cdot \sigma_+(i,k) = \sigma_+(\ell,k) + \sigma_+(\ell,j) \cdot \sigma_+(i,k) \ mod\ 2$$

as required.

Case 2: w_j is disjoint from all pairs of op-arcs in ν

The proof is quite analogous to Case 1. Here, since x is on a ground arc of w_j, all of w_j is a ground arc. If $y \in v_\ell^+ \cap w_j^+$ is not on a ground arc of v_ℓ^+, consider the op-tie in v_ℓ^+ on which y lies, say for a pair of op-arcs w_t^+. Observe first the subtle fact that w_t^+ must be disjoint from v_i^+. For if it weren't, there would be an op-tie for w_t^+ contained in v_i^+, and by the Parallelism Property that op-tie also must intersect w_j^+ and so $\rho_+(i,j) \geq 2$, contradicting Corollary 6.4. It follows then that when the op-tie in v_ℓ^+ containing y is band-summed to v_i^+, the resulting arc becomes an op-tie in $v_\ell'^+$ for the pair of op arcs w_t^+. Thus none of the new points introduced affects $\sigma'(\ell,k)$. So, as in Case 1, we need only focus on the point z (unique, if it exists) at which a ground arc of v_ℓ^+ intersects w_j^+. The rest of the argument is essentially the same as in Case 1. This proves the Claim, and so (assuming the Disk Property is preserved by the construction) Lemma 8.2 and with it Proposition 8.1. □

□

9. The Disk Property is preserved

We want to understand how the fundamental construction, described in the proof of Proposition 8.1 above, that changes P to P' affects the topology of the surfaces $\hat{P}$ and $\hat{P}'$ obtained by building all overpasses in P and P'. The operation $P \to P_{v-w}$ described in Section 7 plays a role:

Lemma 9.1. $\hat{P}' = \hat{P}_{\hat{v}_i^+ - \hat{w}_j^+}$.

Proof. The first observation is this: The (abstract) surface obtained from P' by building exactly those overpasses that are newly created in P' is simply $P_{v_i^+ - w_j^+}$. This is immediate from the description: when the overpass is built for the overpass corresponding to an arc $v_{i'}^+$ of $\partial V_i \cap P - v_i^+$, the effect on the topology of P' is as if $v_{i'}^+ \subset \partial V_i$ were simply disjoint from P. Apply that logic to every component of $\partial V_i \cap P - v_i^+$, and so to every newly created overpass, and the effect is as if the entire arc $\partial V_i - v_i^+$ were disjoint from P. That is, once all the new overpasses are built, it is as if a single band were attached to P with core the arc $\partial V_i - v_i^+$, and then the arc w is deleted. This is the same description as the surface $P_{v_i^+ - w_j^+}$.

Case 1: w_j^+ is disjoint from all op-arcs ν and is a simple arc (not a pair of op-arcs).

$\hat{P}'$ is obtained from P' by building all overpasses. Build the new overpasses first, changing P' to $P_{v_i^+ - w_j^+}$. Since all the remaining op-arcs are unaffected by removing w_j^+ they persist into $P_{v_i^+ - w_j^+}$ and $\hat{P}'$ can be viewed as the result of building the overpasses in $P_{v_i^+ - w_j^+}$. By the hypothesis of this case, none of the old overpasses goes through the band attached at the ends of v_i^+ so we may as well attach it, and remove w_j^+ *after* building the old overpasses. But this is equivalent to first creating $\hat{P}$ (by building the old overpasses) then attaching the band to the ends of what was v_i^+ and is now $\hat{v}_i^+$ and then removing $w_j^+ = \hat{w}_j^+$.

Case 2: w_j^+ is a pair of op-arcs and so is disjoint from ν.

The argument is much the same as Case 1, but requires a preliminary move: before launching the argument above, first build the overpass corresponding to w_j^+, creating a surface P_j that plays the role of P in Case 1. w_j^+ becomes a single arc w_j in P_j intersecting v_i^+ in a single point and removing w_j from P_j gives the same surface as removing both of the original op-arcs w_j^+ from P.

Case 3: v_i^+ is disjoint from all op-arcs ω.

The important difference from Case 1 is that here the op-arcs ν may intersect w_j^+ in P; during the construction of P' they are rerouted. Begin the construction the same as in Case 1: build all new overpasses, so that P' becomes $P_{v_i^+ - w_j^+}$. The old op-arcs that previously intersected w_j^+ are rerouted through the new band via the same operation that is described in Lemma 7.2. So, according to that Lemma, an equivalent way of viewing the surface at this point would have been to construct $P_{w_j^+ - v_i^+}$, leaving the op-arcs where they are, disjoint from v_i^+ and then apply $\phi_{v_i^+, w_j^-}$. Then the argument of Case 1 applied to $P_{w_j^+ - v_i^+}$ shows that $\hat{P}' = \hat{P}_{\hat{w}_j^+ - \hat{v}_i^+}$ and Lemma 7.1 shows that $\hat{P}_{\hat{w}_j^+ - \hat{v}_i^+} \cong \hat{P}_{\hat{v}_i^+ - \hat{w}_j^+}$. □

Corollary 9.2. *$\hat{P}'$ is a disk.*

Proof. We are given before the construction that $\hat{P}$ is a disk, and for v, w any two properly embedded arcs in a disk D that intersect in a point, D_{v-w} is a disk. □

10. Dropping symmetry: a combinatorial proof of the Gordon Conjecture

Proposition 10.1. *Suppose $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ and $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ are Heegaard splittings. Suppose collections of disks $\overline{V}, \overline{W}$ and associated op-arcs ν, ω is a stabilizing pair of coherently numbered coordinated forests of disks for surfaces $P \subset S_+$ and $F \subset S_-$. Suppose further that there is a peripheral (i, j) with $\rho_\pm(i, j) \neq 0$ and that all op-arcs ν are disjoint from all arcs $\{w_k^\pm\}$ in both F and P. (Note: but* not *symmetrically: That is, op-arcs in ω may intersect $\{v_i^\pm\}$.)*

If neither $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ nor $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ is stabilized then there are surfaces $P' \subset S_+$, $F' \subset S_-$, and collections of disks $\overline{V}', \overline{W}'$ and associated op-arcs ν', ω' so that

- *$\overline{V}', \overline{W}'$ and associated op-arcs ν', ω' is a stabilizing pair of coherently numbered coordinated forests of disks with respect to P' and F' and*
- *there are fewer disks in $\overline{V}'$ than in $\overline{V}$ and fewer disks in $\overline{W}'$ than in $\overline{W}$ and*
- *all op-arcs ν' are disjoint from all arcs $\{w_\pm'^+\}$ in both F' and P'.*

Proof. Among all (i, j) with $\rho_\pm(i, j) \neq 0$ choose that in which i is maximal. If V_i and W_j both lie in the same manifold, say M_+ then Lemma 4.3 and Corollary 6.4 show that the splitting of M_+ is stabilized. So henceforth we assume, with no loss of generality, that $V_i \subset \mathcal{V}_+$ and $W_j \subset \mathcal{W}_-$.

Since i was chosen to be maximal among non-trivial peripheral pairs (i, j), each arc $v_{i'}^+ \subset \partial V_i \cap P - v_i^+$ is disjoint from all arcs $\{w_k^+\}$, else a maximal k with non-trivial intersection would be a peripheral pair with $i' > i$. It follows that the first requirement of Proposition 8.1, namely that $\partial V_i \cap P - v_i^+$ is disjoint from all op-arcs, is satisfied. All three other requirements are trivially satisfied, since any ν is disjoint from all all arcs $\{w_k'^+\}$. Hence we can apply the fundamental construction to the pair of disks V_i and W_j as was done in the proof of Proposition 8.1. New op-arcs are created in ν', one pair for each arc $v_{i'}^+$ of $\partial V_i \cap P - v_i^+$. But we have observed above that our choice of i guarantees that each of these will be disjoint from all arcs $\{w_k^+\}$, as required. □

Theorem 10.2. *If $\mathcal{V} \cup_S \mathcal{W}$ is stabilized either $\mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ or $\mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ is stabilized.*

Proof. Begin with the Seminal Example of a stabilizing pair of coherently numbered coordinated forests of disks. Since there are no op-arcs in this example, it clearly satisfies the hypotheses of Proposition 10.1. Repeatedly apply Proposition 10.1, stopping if some iteration shows that one of $M_+ = \mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$ or $M_- = \mathcal{V}_- \cup_{S_-} \mathcal{W}_-$ is stabilized. Since each application decreases the number of indices represented by disks in $\overline{V}$ and $\overline{W}$, Proposition 10.1 can only be applied a finite number of times. If the process does not stop because it detects a stabilized splitting, it must stop because there are no longer any peripheral (i, j) for which $\rho_\pm(i, j) \neq 0$. As noted in the comments following Definition 4.2, this implies that $\rho_\pm = 0$.

In this case, consider the disks V_0, W_0 that define the distinguished roots. They are both contained in M_+ or both in M_-, since their boundaries have the common intersection point x_0 outside of P or F. Say both are contained in M_+. The arcs $\partial V_0 \cap P$ are disjoint from the arcs $\partial W_0 \cap P$ since $\rho_+ = 0$. Hence the only intersection point in $\partial V_0 \cap \partial W_0$ is x_0. Thus V_0 and W_0 are a stabilizing pair of disks for the splitting $\mathcal{V}_+ \cup_{S_+} \mathcal{W}_+$. □

References

[Ba] David Bachman, Connected sums of unstabilized Heegaard splittings are unstabilized, ArXiv preprint math.GT/0404058

[Ed] Karen Edwards, Stabilizations of Heegaard Splittings with Respect to Connect-Sums of 3-Manifolds, PhD thesis, UC Berkeley, 2001.

[Ki] R. Kirby, Problems in Low-Dimensional Topology, in Geometric Topology, Edited by H. Kazez, AMS/IP Vol. 2, International Press, 1997.

[Q] Ruifeng Qiu, Stabilizations of Reducible Heegaard Splittings, ArXiv preprint math.GT/0409497

Ruifeng Qiu, Department of Mathematics, East China Normal University, 200062, China

E-mail address: `Email: qiurf@dlut.edu.cn`

Martin Scharlemann, Mathematics Department, University of California, Santa Barbara, Santa Barbara, CA 93117, USA

E-mail address: `mgscharl@math.ucsb.edu`