

REFILLING MERIDIANS IN A GENUS 2 HANDLEBODY COMPLEMENT

MARTIN SCHARLEMANN

Dedicated to the memory of Heiner Zieschang, who first noticed that genus two handlebodies could be interesting

ABSTRACT. Suppose a genus two handlebody is removed from a 3-manifold M and then a single meridian of the handlebody is restored. The result is a knot or link complement in M and it is natural to ask whether geometric properties of the link complement say something about the meridian that was restored. Here we consider what the relation must be between two not necessarily disjoint meridians so that restoring each of them gives a trivial knot or a split link.

1. BACKGROUND

For a knot or link in a 3-manifold, here are some natural geometric questions that arise, in roughly ascending order of geometric sophistication: is the knot the unknot? is the link split? is the link or knot a connected sum? are there companion tori? beyond connected sums, are there essential annuli in the link complement? beyond connected sums, are there essential meridional planar surfaces? One well-established context for such questions is that of Dehn surgery (cf [Go]) where one imagines filling in the knot or link complement with solid tori via different meridian slopes and then asks under what conditions two of the fillings have geometric features such as those listed above.

Another natural context is this: Suppose W is a genus 2 handlebody embedded in a compact orientable 3-manifold M . Suppose α, β are not necessarily disjoint essential properly embedded disks in W (called therefore meridian disks). Then $W - \eta(\alpha)$ (resp $W - \eta(\beta)$) is a regular neighborhood of a knot or link $L[\alpha]$ (resp $L[\beta]$) in M . Under what circumstances do $L[\alpha]$ and $L[\beta]$ have geometric features like those outlined above? At the most primitive level (and so presumably the easiest level) one can ask when both $L[\alpha]$ and $L[\beta]$ are split (if a link) or trivial (if a knot). Put another way,

Date: March 21, 2007.

Research partially supported by an NSF grant. Thanks also to Catalonia's Centre Recerca Matemàtica for their extraordinary hospitality while this work was underway, and to both the referee and Scott Taylor for very helpful comments.

suppose $M[\alpha], M[\beta]$ are the manifolds obtained from $M - W$ by restoring neighborhoods of the meridians α and β . Under what circumstances are both $M[\alpha]$ and $M[\beta]$ reducible and/or ∂ -reducible? (In the absence of Lens space or $S^1 \times S^2$ summands in a closed 3-manifold M , a ∂ -reducing disk for a knot complement is equivalent to an unknotting disk for the knot.)

Not only is this a natural question, the solution to it in specific cases has been of significant interest in knot theory over the past two decades. Here is a probably partial list of such results, for $M = S^3$:

- [Sc3] (see also [Ga]) considers the case in which one meridian disk is separating, the other is non-separating, and the two meridians intersect in a single arc.
- [BS1] and [BS2] (see also [ST2]) consider the case (among others) in which both meridian disks are non-separating and the two intersect in a single arc.
- [EM3] extends these earlier results and unifies them within the more general setting in which there are non-isotopic non-separating meridian disks μ, μ' for W which are disjoint from both α and β .

If we extend the question to whether one of the links is a connected sum, the literature becomes even more extensive, including [EM1], [EM2], [EM3], [Sc2], [ST1].

We briefly describe a typical conclusion in the arguments above. First some terminology: Say that the handlebody $W \subset M$ is *unknotted* if it is isotopic to the regular neighborhood of a figure 8 graph Σ that lies on a sphere in M . If W is unknotted and Σ is such a spine and if μ and μ' are the pair of meridian disks for W that are dual to the two edges of Σ , then μ and μ' are called an *unknotting pair of meridians for W* . Put another way, the meridians μ and μ' for W are called unknotting meridians, and W is said to be unknotted, if there are properly embedded disks $\lambda, \lambda' \subset M - \text{interior}(W)$ such that $|\mu \cap \lambda| = |\mu' \cap \lambda'| = 1$ and $\mu \cap \lambda' = \mu' \cap \lambda = \emptyset$. What is typically proven (most generally by Eudave-Muñoz in [EM3]) is this: We are given specific conditions on the filling meridians α and β , including that they are both disjoint from the same pair of non-isotopic meridians μ, μ' for W . We suppose further that the manifolds $M[\alpha]$ and $M[\beta]$ are both either reducible or ∂ -reducible, whereas $M - W$ is irreducible. The conclusion is that W is unknotted in M , and the meridians μ and μ' are an unknotting pair of meridians for W .

Put in this way, one wonders if the various conditions on α and β can be dropped to give a global theorem on the unknottedness of W . For example:

Naive Conjecture: Suppose α and β are meridian disks of a genus two handlebody $W \subset M$. Suppose further that $\overline{M - W}$ is irreducible and ∂M is

incompressible in $\overline{M - W}$. Then either W is unknotted or at least one of $M[\alpha], M[\beta]$ is both irreducible and ∂ -irreducible.

This naive conjecture is plainly false. Most obviously, take M to be merely a regular neighborhood of W ; no matter how α and β are chosen, both $M[\alpha], M[\beta]$ are ∂ -reducible. But there are easy counterexamples even for $M = S^3$. For example, attach an arc e to the unknot $U \subset S^3$, an arc chosen to be so complicated that the closed complement $S^3 - \eta(U \cup e)$ is both irreducible and ∂ -irreducible. Let $W = \eta(U \cup e)$ and choose both α and β to be copies of the meridian of W that is dual to the arc e . Then both $L[\alpha], L[\beta]$ are the unknot. Of course, taking α and β parallel like this might be regarded as cheating. Figure 1 (due to Kinoshita [Ki]) is a more subtle counterexample in which α and β aren't parallel. The complement of Kinoshita's graph is also called the Thurston wye manifold and is known to be ∂ -irreducible, so W is knotted.




FIGURE 1

In view of these counterexamples to the naive conjecture, are there simple conditions that ensure the conclusion of the conjecture, that either W is unknotted or at least one of $M[\alpha]$ or $M[\beta]$ is both irreducible and ∂ -irreducible? With reasonable conditions on the original pair (M, W) , it appears that there are such conditions on filling meridians α and β . These are outlined in the next section.

2. A CONJECTURE AND A WEAK CONVERSE

Suppose W is a genus two handlebody properly embedded in a compact orientable 3-manifold M . The pair is called *admissible* if these conditions are satisfied:

- any sphere in M is separating
- M contains no Lens space summands
- any pair of curves in ∂M that compress in M are isotopic in ∂M
- $M - W$ is irreducible

- ∂M is incompressible in the complement of W

These are reasonable conditions to assume in our context: The first two guarantee that the complement of a knot in a closed M is ∂ -reducible only if the knot is the unknot. The third condition (which is the most technical) removes the first counterexample above, in which M is merely a regular neighborhood of W . The last two conditions remove obvious counterexamples in which reducing spheres or ∂ -reducing disks exist even before filling meridians are added.

The precise conditions that we propose on the pair of filling meridians $\alpha, \beta \subset W$ depend on whether α and β are separating or non-separating. Suppose that α and β have been properly isotoped in W to minimize $|\alpha \cap \beta|$. In particular $\alpha \cap \beta$ consists of a possibly empty collection of arcs.

Definition 2.1. *If α and β are both non-separating then they are **aligned** if*

- *there is a non-separating meridian disk for W that is disjoint from both α and β and*
- *all arcs of $\alpha \cap \beta$ are parallel in both disks.*

Definition 2.2. *If α and β are both separating then they are aligned if*

- *there is a non-separating meridian disk μ for W that is disjoint from both α and β and*
- *There is a longitude in the boundary of the solid torus $W - \mu$ that is disjoint from α and β .*

Here a *longitude* of a solid torus means any simple closed curve in the boundary that intersects a meridian disk in a single point. If α and β are aligned and separating, then they both lie in the 4-punctured sphere $\partial W - (\partial\mu \cup \text{longitude})$ and they separate the same pairs of punctures there.

Finally

Definition 2.3. *If α is non-separating and β is separating, then the disks are aligned if they are disjoint or, when they are not disjoint,*

- *one solid torus component of $W - \beta$ has a meridian μ disjoint from both α and β and*
- *the other solid torus component of $W - \beta$ has a meridian β' that is disjoint from β and furthermore*
- *β' is maximally aligned with α . That is, α and β' are aligned (as defined in Definition 2.1) and $|\beta' \cap \alpha| = |\beta \cap \alpha| - 1$. See Figure 2.*

For a pair of non-separating disks, the condition that all arcs of $\alpha \cap \beta$ are parallel in both disks means that either α and β are disjoint, or exactly two components of $\alpha - \beta$ (and two components of $\beta - \alpha$) are outermost disks, that is disks incident to a single arc component of $\alpha \cap \beta$. The condition is




FIGURE 2

clearly satisfied whenever there are at most two arcs of intersection, ie if $|\partial\alpha \cap \partial\beta| \leq 4$.

Here is the main conjecture:

Conjecture 1. *If (M, W) is an admissible pair then either*

- *W is unknotted and $M = S^3$*
- *at least one of $M[\alpha]$ or $M[\beta]$ is both irreducible and ∂ -irreducible,*
or
- *α and β are aligned in M .*

In many appearances of this general problem cited in the literature above, $\partial\alpha$ and $\partial\beta$ intersect in few points, so they are automatically aligned. In any case, the conjecture would only partially recapture the known results but would extend them to a more general setting. Conjecture 1 appears to be true, at least with one additional technical hypothesis: When $|\partial\alpha \cap \partial\beta| \leq 4$, assume further that M contains no proper summand that is a rational homology sphere. A complete proof is not yet written down; even for a weaker result, in which ∂ -irreducible is removed from the second conclusion, the combinatorial argument is extremely complicated. The intention here is to offer the more straightforward proofs in these three important special cases:

- $M - W$ is ∂ -reducible.
- $|\partial\alpha \cap \partial\beta| \leq 4$. (This requires the additional technical condition.)
- Both α and β are separating.

In addition, we explain why the combinatorics becomes so difficult once non-separating meridians are considered.

Before starting to verify the conjecture in these special cases, here is a sort of weak converse to Conjecture 1. The assumption in Theorem 2.4 that α and β are aligned is meant to be the weakest natural condition for which the theorem is likely to be true.

Theorem 2.4. *Suppose α and β are aligned meridians in W . Then there is an unknotted embedding of W in S^3 (hence in any 3-manifold) so that each of $L[\alpha]$ and $L[\beta]$ is either the unknot or a split link.*

Proof. Suppose first that α and β are both separating. Embed W as the regular neighborhood of an eyeglass graph (ie two circles σ_1, σ_2 connected by an edge e) in $S^2 \subset S^3$. Since α is separating, such an embedding of W can be found so that α is the meridian dual to the edge e . Then $L[\alpha] = \sigma_1 \cup \sigma_2$ is split. Further choose the embedding W so that in the framing of the tori $W - \alpha$, the longitude λ of the solid torus $W - \mu$ that is disjoint from $\alpha \cup \beta$ (cf Definition 2.2) is one of the three curves $\partial W \cap S^2$. Then λ bounds a disk in $S^2 - W$, so one component of $L[\beta]$ bounds a disk in the complement of the other, showing that the link $L[\beta]$ is also split.

In case α and β are non-separating, a more subtle construction is required. Begin with an annulus A in $S^2 \subset S^3$ and draw a pair P of disjoint spanning arcs. Let μ_{\pm} be points in $\partial A - \partial P$, one in each component of ∂A but both lying on the boundary of the same rectangle component of $A - P$. A product neighborhood $A \times [-1, 1] \subset S^2 \times [-1, 1] \subset S^3$ is a solid unknotted torus in which $P \times [-1, 1]$ can be thought of as the union of two disjoint meridians α and β . Connect disk neighborhoods of $\mu_{\pm} \times \{1/2\}$ in $\partial A \times [-1, 1]$ by adding an unknotted 1-handle on the outside of $A \times [-1, 1]$ that lies above $S^2 \times \{0\}$. The result is an unknotted embedding of W , with the meridian μ of the 1-handle also a meridian of W that is disjoint from α and β . Now repeat the construction after first altering exactly one of the original spanning arcs in P by $n \geq 0$ Dehn twists around the core of the annulus A . The construction gives aligned meridians with exactly $n - 1$ arcs of intersection. It is not hard to show that, given a pair of aligned non-separating meridians in W which have $n - 1$ arcs of intersection, there is an automorphism of W that carries the pair to α and β . The corresponding knots $L[\alpha]$ and $L[\beta]$ are easily seen to be unknotted. (In fact they are isotopic: the isotopy merely undoes the Dehn twists used in the construction by adding twists around the meridian μ of the 1-handle.) See Figure 3.

In case β is separating and α is not, let μ and β' be meridians of the solid tori $M - \beta$ as given in Definition 2.3, so that

- both μ and β' are disjoint from β ,
- μ is also disjoint from α and
- β' is aligned with α .

Apply the previous construction to the aligned meridians α and β' . The definition requires that the aligned β' intersect α in almost as many components as β does. Viewed dually, this implies that β intersects α only once more than β' does. This constrains β to be a regular neighborhood of the arc β_- connecting $\mu_+ \times \{1/2\}$ to $\mu_- \times \{1/2\}$ in the disk $\partial(A \times I) - (A \times$




FIGURE 3

$\{-1\} \cup \partial\beta'$). As before, $W - \alpha$ is the unknot. $W - \beta$ is a regular neighborhood of the trivial link, with one component parallel to the core of $A \times \{0\}$ and the other to the union of β_- and the core of the 1-handle. See Figure 4. \square

A first step towards the proof of Conjecture 1 is to note that we can restrict to the case in which $M - W$ is ∂ -irreducible.

Proposition 2.5. *Conjecture 1 is true when $M - W$ is ∂ -reducible.*

Proof. Case 1: There is a ∂ -reducing disk P for $M - W$ such that $|\partial P \cap \partial\alpha| = 1$ (or, symmetrically, $|\partial P \cap \partial\beta| = 1$).

In this case, $W - \eta(\alpha) \cong W \cup \eta(P)$ so any reducing sphere for $M[\alpha]$ would be a reducing sphere for $M - W \supset M - (W \cup \eta(P))$, contradicting the assumption that (M, W) is an admissible pair. Similarly, if $M[\alpha]$ is ∂ -reducible via a disk whose boundary lies on ∂M , the disk would be a ∂ -reducing disk for $M - W$, contradicting the assumption that (M, W) is an admissible pair. Hence either $M[\alpha]$ is both irreducible and ∂ -irreducible (the second conclusion) or there is a ∂ -reducing disk for $M[\alpha]$ whose boundary lies on $\partial(W - \eta(\alpha))$. Since M contains no Lens space summands, such a disk is incident to the solid torus $W - \eta(\alpha)$ in a longitude and the union of the disk with $W - \eta(\alpha)$ is a ball. Equivalently, there is a ∂ -reducing disk D for $M - (W \cup \eta(P))$ so that $W \cup \eta(P) \cup \eta(D)$ is a ball. W is clearly a genus 2 Heegaard surface for the ball, so it follows from Waldhausen's theorem




FIGURE 4

[Wa] that W is unknotted, giving the first conclusion and completing the proof in this case.

Case 2: There is a ∂ -reducing disk $(P, \partial P) \subset (M - W, \partial W - \alpha)$ for $M[\alpha]$ (or symmetrically for β and $M[\beta]$).

In this case first note that $W - \eta(\alpha)$ consists of one or two solid tori (depending on whether α separates W) and ∂P lies on one of them. Since M contains no Lens space or non-separating 2-spheres, ∂P is in fact a longitude of the solid torus on whose boundary it lies. So ∂P is non-separating.

Suppose first that ∂P is disjoint from $\partial\beta$ as well as $\partial\alpha$; we show that α and β are then aligned. Let μ be a meridian for W which is disjoint from α and whose boundary is disjoint from ∂P . In particular, if α is non-separating, take a parallel copy of α for μ . If β intersects μ then an outermost disk of β cut off by μ would be a meridian of the solid torus $W - \mu$ for which ∂P is a longitude, a contradiction. Hence β lies entirely in the solid torus $W - \mu$. β can't be essential in that solid torus for the same reason. Hence β is inessential in $W - \mu$ and so is either parallel to μ (hence disjoint

from and therefore aligned with α) or separating. If α is non-separating then it is parallel to μ and so disjoint from and therefore aligned with β . If α , like β , is separating, then ∂P is the longitude required by Definition 2.2 to show that α and β are aligned.

So henceforth assume that ∂P and $\partial\beta$ are not disjoint. Apply the Jaco handle-addition lemma [Ja] (as generalized by [CG] to the reducible case) to the 2-handle $\eta(\beta)$ attached to the ∂ -reducible (via P) manifold $M - W$: If $M[\beta]$ is reducible or ∂ -reducible then there is a ∂ -reducing disk J for $M - W$ whose boundary is disjoint from β . Since ∂M is incompressible in $M - W$, ∂J lies on ∂W . Assume that such a disk J has been chosen to minimize $|J \cap P|$.

Suppose first that $J \cap P \neq \emptyset$ and consider an outermost disk E cut off from J by P . We can assume the arc $\partial E - \partial P$ is essential in the surface $\partial W - \partial P$, else an isotopy of that arc across ∂P would reduce $|J \cap P|$. It follows that ∂E is essential on the torus $\partial(W \cup \eta(P))$, for the alternative is that it cuts off a disk containing all of one copy of P in the torus but only part of the other, and so, absurdly, ∂J intersects one side of ∂P more often than the other. Then $W \cup \eta(P) \cup \eta(E)$ is the union of a solid torus and an essential 2-handle. Since M contains neither non-separating spheres nor Lens spaces, $W \cup \eta(P) \cup \eta(E)$ is a 3-ball. Cap the 3-ball off with a 3-handle and the result is a genus two Heegaard splitting of S^3 , in which W is one of the two handlebodies. But Waldhausen's theorem [Wa] states that any Heegaard splittings of S^3 is isotopic to a standard one. It follows that W may be isotoped in the 3-sphere (hence in a neighborhood of the 3-ball $W \cup \eta(P) \cup \eta(E)$) so that it is standard. Thus W is unknotted in M and we are done.

The remaining case is that J is disjoint from P . If ∂J is essential on the torus $\partial(W \cup \eta(P))$ then use J for E in the previous argument and again we are done. Suppose then that ∂J is inessential on $\partial(W \cup \eta(P))$. The disk it bounds can't contain just one copy of P , since ∂P intersects β and ∂J does not. So the disk contains both copies of P , hence ∂J is separating in ∂W . Since ∂J is disjoint from the meridian β , ∂J then also bounds a separating meridian $J' \subset W$. Since $\partial\beta$ and ∂P intersect, β is a meridian disk for the solid torus component of $W - J'$ for which ∂P is the longitude. Hence $|\partial\beta \cap \partial P| = 1$ and we are done by the first case.

Now consider the general case:

Apply the Jaco handle-addition lemma to $\eta(\alpha)$ and conclude that either $M[\alpha]$ is irreducible and ∂ -irreducible (the second conclusion) or a ∂ -reducing disk $(P, \partial P)$ can be found whose boundary is disjoint from α . Since ∂M is incompressible in $M - W$, $\partial P \subset \partial W - \alpha$. If ∂P is essential on $\partial M[\alpha]$ the proof follows from Case 2. If ∂P is inessential on $\partial M[\alpha]$ then

it is coplanar (possibly parallel with) α , so ∂P bounds a meridian P' for W . In particular, ∂P is separating, since M contains no non-separating spheres. Similarly there is a separating disk $(Q, \partial Q) \subset (M - W, \partial W - \beta)$. If ∂P and ∂Q are parallel in ∂W then it follows easily that α and β are disjoint from each other and hence aligned. If ∂P and ∂Q are not isotopic in ∂W then they must intersect, and an outermost arc of intersection in Q cuts off a disk $(F, \partial F) \subset (M - W, \partial W - \partial P)$ whose boundary is essential on, and hence a longitude of, one of the solid tori $W - P'$. If α is a meridian of that solid torus, use F as P in Case 1. If α is the meridian of the other solid torus, use F as P in Case 2. \square

Following Proposition 2.5, Conjecture 1 is equivalent to the apparently weaker

Conjecture 2. *Suppose (M, W) is an admissible pair and $M - W$ is ∂ -irreducible. Then either*

- *at least one of $M[\alpha]$ or $M[\beta]$ is both irreducible and ∂ -irreducible,*
- or*
- *α and β are aligned in M .*

Proposition 2.6. *When $|\partial\alpha \cap \partial\beta| \leq 4$, Conjecture 2 is true. If, furthermore, α is separating and M contains no proper summand that is a rational homology sphere then one of $M[\alpha]$ or $M[\beta]$ is irreducible and ∂ -irreducible.*

Proof. When $|\partial\alpha \cap \partial\beta| \leq 4$, there are at most two arcs of intersection, so each arc of intersection is incident to an outermost disk in both α and β and, moreover, all arcs of intersection are parallel in both disks. Gluing an outermost disk of one to an outermost disk of the other gives a meridian disk for W that is disjoint from both α and β . It follows that if α and β are both non-separating, then they are aligned.

So it suffices to consider the case in which at least one of them, say α , is separating. We aim to show that the disks are aligned and, assuming M contains no proper summand that is a rational homology sphere, if $M[\alpha]$ is either reducible or ∂ -reducible, then $M[\beta]$ is neither reducible nor ∂ -reducible. There are two cases to consider:

Case 1: $|\partial\alpha \cap \partial\beta| = 2$.

In this case β must be non-separating and it is immediate from Definition 2.3 that α and β are aligned. The argument that one of $M[\alpha]$ or $M[\beta]$ is irreducible and ∂ -irreducible now mimics in this more general setting the central argument of [Ga], which in fact provides the complete argument when M is the 3-sphere.

Let $L \subset M - W$ be the regular neighborhood of a circle parallel to $\partial\alpha$, so a longitude of ∂L bounds a disk D in $M - L$ that intersects W in a single copy of α . Since $\alpha \cap \beta$ is a single arc, D intersects the solid torus $T = W - \beta$ in two oppositely-oriented meridians. (For example, in Gabai's setting, $W - \alpha$ is a split link and $W - \beta$ is obtained from the split link by a band sum. ∂D encircles the band.)

First observe that $M[\beta] - L$ is irreducible and ∂ -irreducible. For suppose first that a ∂ -reducing disk or reducing sphere \hat{Q} were disjoint from D . Since $M[\beta] \cup \eta(D) \cong W$, \hat{Q} would also be a ∂ -reducing disk or reducing sphere in $M - W$, contradicting the hypothesis. Similarly, if \hat{Q} intersects D an innermost disk E in $\hat{Q} - D$ provides a ∂ -reducing disk for ∂W in the complement of $M[\beta] \cup \eta(D) \cong W$.

Let $M[\beta]_0$ be the manifold obtained from $M[\beta]$ by 0-framed surgery on L , that is, surgery with slope ∂D .

Claim: If $M[\alpha]$ is reducible or ∂ -reducible, then so is $M[\beta]_0$

To prove the claim, suppose \hat{P} is a reducing sphere or ∂ -reducing disk for $M[\alpha]$. Then $\text{int}(\hat{P}) \cap W$ is a collection of parallel copies of α . Let $A \subset \partial W$ be an annulus containing all their boundary components $\text{int}(\hat{P}) \cap \partial W$. Let A' be a copy of A pushed into W rel ∂A . Then $W - \eta(A')$ is isotopic to $W \cup L$ and under this isotopy the boundary components of the planar surface $P = \hat{P} - W$ that previously were on A become 0-framed curves on $\partial(L)$. After 0-framed surgery on L to get $M[\beta]_0$, these components on ∂L can be capped off to give a copy $\hat{P}' \subset M[\beta]_0$ of \hat{P} that either ∂ -reduces or reduces $M[\beta]_0$. See Figure 5. That \hat{P}' does not bound a ball is obvious if \hat{P}' is non-separating (eg m is odd). If \hat{P}' is separating, note that such a ball must have come from the component of $M - (W \cup L \cup \eta(P))$ not adjacent to W . But this component is completed by attaching 2-handles in the same way in both $M[\beta]_0$ and $M[\alpha]$, and we know that \hat{P} is a reducing sphere in $M[\alpha]$ and so does not bound a ball. This completes the proof of the Claim.

Following the Claim, apply the central theorem of [Sc4] to the link L in the manifold $M[\beta]$ and rule out possible conclusions a)-c): $L \subset M[\beta]$ is not homeomorphic to a braid in $D^2 \times S^1$ since, for one thing, L is null-homologous. Nor does $M[\beta]$ contain a Lens space summand, since M doesn't. A longitude of ∂L bounds a disk in $M - L$ so ∂L can't be cabled in M with the same boundary slope. Hence conclusion d) of the main theorem of [Sc4] applies to $M[\beta]$ and $M[\beta]_0$. We now examine the consequences.

Suppose that $M[\alpha]$ is reducible. Then $M[\beta]_0$ is reducible, so by [Sc4], $M[\beta]$ is irreducible and ∂ -irreducible.

Suppose that $M[\alpha]$ is ∂ -reducible via the disk \hat{P} . We note that if $\partial\hat{P}$ lies on the boundary of one of the solid tori components of $W - \alpha$ then that solid




FIGURE 5

torus lies in a 3-ball in M , so $M[\alpha]$ is reducible. Then, via the previous comment, $M[\beta]$ is irreducible and ∂ -irreducible. So we can assume that $\partial\hat{P} \subset \partial M$. Then $\partial\hat{P}$ is disjoint from any ∂ -reducing disk for $M[\beta]$ since, by hypothesis, any essential simple closed curves in ∂M that compress in M are isotopic. Hence $M[\beta]_0$ is ∂ -reducible via a disk whose boundary is disjoint from any ∂ -reducing disk for $M[\beta]$. So by [Sc4] $M[\beta]$ is irreducible and ∂ -irreducible. To summarize: If $M[\alpha]$ is reducible or ∂ -reducible, then $M[\beta]$ is neither reducible nor ∂ -reducible. Thus we have the first conclusion of Conjecture 2, completing the proof in this case.

Case 2: $|\partial\alpha \cap \partial\beta| = 4$.

In this case $\beta \cap \alpha$ consists of two arcs; call the rectangles that lie between them in α and β respectively R_α and R_β . The outermost disks of β cut off by α are both meridians of the same solid torus component W_1 of $W - \alpha$; in fact they are parallel copies of the same disk meridian $D_\beta \subset W_1$ whose boundary consists of an arc in R_α parallel to the arcs $\beta \cap \alpha$ and an arc on $\partial W - \alpha$. See Figure 6. The meridian D_β is properly isotopic in W to a meridian that is disjoint from both α and β ; a meridian for the other solid torus component of $W - \alpha$ intersects β in a single arc. It follows immediately that α and β are aligned. It remains to show that if M contains no

proper summand that is a rational homology sphere and if $M[\alpha]$ is reducible or ∂ -reducible, then $M[\beta]$ is neither reducible nor ∂ -reducible.




FIGURE 6

The union $A = R_\alpha \cup R_\beta$ is either an annulus or Möbius band in W (depending on the orientations of the arcs $\beta \cap \alpha$ in α) which, when ∂ -compressed to ∂W along D_β becomes β . Put another way, $W_- = W - A$ is a handlebody or pair of handlebodies (depending on whether A is a Möbius band or an annulus) which, when ∂ -reduced along D_β becomes $W - \beta$. A copy α' of α pushed into W_1 intersects D_β in a single arc. Essentially, we intend to settle Case 2 by applying Case 1 to the disks D_β and α' , two meridians of W_- that intersect in a single arc.

If A is an annulus, then W_- is the union of a solid torus W_2 (essentially the component of $W - \alpha$ that is not W_1) and a genus two component. In this case, let W' be the genus 2-component and $M' = M - W_2$. If A is a Möbius band let $W' = W_-$ and $M' = M$.

Claim 1: If $M[\alpha]$ is reducible or ∂ -reducible, so is $M'[\alpha]$.

Proof of Claim 1: $M'[\alpha]$ is obtained from $M[\alpha]$ by attaching A to $\partial M[\alpha]$ via ∂A and thickening A . This operation can't turn a complementary component of a sphere or properly embedded disk into a 3-ball, so it can't make a reducible manifold irreducible or a ∂ -reducible manifold ∂ -irreducible.

Claim 2: If $M[\beta]$ is reducible or ∂ -reducible, so is $M'[D_\beta]$.

Proof of Claim 2: By construction, ∂ -compressing A to ∂W via D_β gives β , so $M[\beta]$ is homeomorphic to $M'[D_\beta]$

Claim 3: $M' - W'$ is irreducible and ∂ -irreducible.

Proof of Claim 3: Suppose S is a reducing sphere or ∂ -reducing disk for $M' - W'$ and A' is the annulus or Möbius band dual to A in W . That is, $W' \cup \eta(A') \cong W$ if A is a Möbius band and $W' \cup \eta(A') \cup W_2 \cong W$ if A is an annulus. If S is disjoint from A' then S is a reducing sphere or ∂ -reducing disk for $M - W$, contradicting hypothesis.

On the other hand, suppose S and A' are not disjoint. Closed components of $S \cap A'$ that are inessential in A' can be removed by a standard innermost disk argument; if any essential closed components of $S \cap A'$ remain, then an innermost disk of $S - A'$ would be a compressing disk for ∂W in $M - W$, again contradicting hypothesis. Similarly, all arc components of $S \cap A'$ (which only can arise if S is a ∂ -reducing disk, not a reducing sphere) that are inessential in A' (ie non-spanning) can be removed by a standard outermost arc argument. On the other hand, a spanning arc in $S \cap A'$ is clearly impossible if A' is an annulus, since the ends of A' are on different components of $\partial(M' - W')$ whereas ∂S can only be incident to one. If A' is a Möbius band then an outermost disk cut off from S by A' will be a ∂ -reducing disk for $M - W$, again contradicting hypothesis.

Following Claim 3, Case 1) may be applied to the disks $\alpha, D_\beta \subset W'$ in M' : If $M'[\alpha]$ is reducible or ∂ -reducible then $M'[D_\beta]$ is irreducible and ∂ -irreducible. Hence, following Claims 1 and 2, if $M[\alpha]$ is reducible or ∂ -reducible, $M[\beta]$ is irreducible and ∂ -irreducible. \square

Proposition 2.7. *If both α and β are separating Conjecture 2 is true. Furthermore, if, in addition to the assumptions of Conjecture 2, $|\partial\alpha \cap \partial\beta| \geq 6$, then one of $M[\alpha]$ or $M[\beta]$ is irreducible and ∂ -irreducible.*

The rest of the section is devoted to the proof of Proposition 2.7. The proof requires a few internal lemmas.

Following Proposition 2.6 we may as well assume $|\partial\alpha \cap \partial\beta| \geq 6$.

Since both α and β are separating, $W - \eta(\alpha)$ and $W - \eta(\beta)$ are each a pair of solid tori. If $M[\alpha]$, say, is ∂ -reducible and the ∂ -reducing disk P is incident to a solid torus component W_1 of $\partial(W - \eta(\alpha))$ then $\partial(\eta(W_1 \cup P))$ is a sphere in M that separates the two solid tori, so $M[\alpha]$ is reducible as well as ∂ -reducible. On the other hand, if $M[\alpha]$ is ∂ -reducible and the ∂ -reducing disk P is incident to ∂M then P is a ∂ -reducing disk for M as well. In that case ∂P lies in the unique isotopy class of ∂M which compresses in M , following the assumption that (M, W) is admissible. We deduce that

if neither $M[\alpha]$ nor $M[\beta]$ are reducible but both are ∂ -reducible then the ∂ -reducing disks have disjoint boundaries, lying in ∂M .

The proof proceeds by assuming that both $M[\alpha]$ and $M[\beta]$ are ∂ -reducible or reducible and derives a contradiction by examining components of intersection between a reducing sphere or ∂ -reducing disk in $M[\alpha]$ and a reducing sphere or ∂ -reducing disk in $M[\beta]$. Following the comments above it would be natural to view the case of reducing spheres as entirely representative, since when there are no reducing spheres the ∂ -reducing disks that can be used necessarily have disjoint boundaries, so the combinatorics is unaffected by the boundary. In fact we will ignore ∂ -reducing disks whose boundaries lie on ∂M , because the combinatorics involved is identical to that for reducing spheres. But for somewhat technical reasons, it is best to allow another possibility, even when there are reducing spheres. Say that a ∂ -reducing disk for $M[\alpha]$ is *special* if its boundary p_0 lies in ∂W and p_0 is parallel in ∂W to the end-point union of a subarc of $\partial\alpha$ and a component β_0 of $\partial\beta - \partial\alpha$. The curve p_0 is necessarily a longitude of the solid torus component W_1 of $W - \alpha$ on which it lies. Furthermore, since each arc of $\partial\beta - \alpha$ in the punctured torus $\partial W_1 - \alpha$ is essential and lies in the annulus $\partial W_1 - (\alpha \cup \beta_0)$, it follows that each such arc component intersects p_0 at most once. There is a symmetric definition for a special ∂ -reducing disk for $M[\beta]$.

To begin the combinatorial argument, let \hat{P} (resp \hat{Q}) be either a reducing sphere or a special ∂ -reducing disk for $M[\alpha]$ (resp $M[\beta]$). Since $M - W$ is irreducible and ∂ -irreducible, \hat{P} must pass through the handle $\eta(\alpha)$ $m \geq 1$ times. Among all reducing spheres and special ∂ -reducing disks for $M[\alpha]$, choose \hat{P} to minimize m . Let $P = \hat{P} - W$ be the associated properly embedded planar surface in $M - W$. Then ∂P has components $\alpha_1, \dots, \alpha_m, m \geq 1$, each of them parallel on ∂W to $\partial\alpha$ and, if \hat{P} was a special ∂ -reducing disk, a component $p_0 = \partial\hat{P}$. Label the α_i in the order they appear in an annular neighborhood of $\partial\alpha$ in ∂W . There are two choices for the direction of the ordering. If \hat{P} is a ∂ -reducing disk, choose the ordering which puts α_1 adjacent to p_0 ; otherwise choose either ordering. For \hat{Q} there is a similar planar surface $(Q, \partial Q) \subset (M - W, \partial W)$ whose boundary components are similarly labeled $\beta_1, \dots, \beta_n, n \geq 1$ and possibly q_0 . Isotope P and Q so as to minimize $|P \cap Q|$. $|P \cap Q| > 0$ since α and β intersect. Let $\hat{\alpha}_i, \hat{\beta}_j \subset W$ be meridian disks in W bounded by α_i, β_j respectively, for each $1 \leq i \leq m$ and $1 \leq j \leq n$.

As is now a classical strategy, view $P \cap Q$ as giving rise to graphs Σ and Υ respectively in \hat{P} and \hat{Q} . The vertices of the graphs correspond respectively to the disks $\hat{\alpha}_i \subset \hat{P}$ and $\hat{\beta}_j \subset \hat{Q}$. Edges of each graph correspond to the arcs

of intersection in $P \cap Q$. Circles of intersection are ignored. The valence of each vertex in Σ is $|\partial\alpha \cap \partial\beta| \cdot n + |\partial\alpha \cap q_0|$ and in Υ is $|\partial\alpha \cap \partial\beta| \cdot m + |\partial\beta \cap p_0|$. When \hat{P} is a special ∂ -reducing disk, the number of ends of edges that are incident to $p_0 = \partial\hat{P}$ is less than half the valence of each vertex of Σ , since each arc of $\partial Q - \alpha$ lying on the torus component of $\partial(W - \alpha)$ containing p_0 contributes 2 to the valence of α_1 but intersects p_0 at most once, and at least one such component is disjoint from p_0 . Similarly, the number of ends of edges that are incident to $q_0 = \partial\hat{Q}$ is less than half the valence of each vertex of Υ .

Consider a point x in $\partial P \cap \partial Q$, say $x \in \alpha_i \cap \beta_j$. To x assign the ordered pair (i, j) , $1 \leq i \leq m, 1 \leq j \leq n$. When viewed in the graphs Σ and Υ the point x appears as an end of a unique edge. Assign to the end of the edge in Σ the label j . Similarly assign to the end of the edge in Υ the label i . For each $1 \leq j \leq n$, of the ends of edges incident to any vertex in Σ , exactly $|\alpha \cap \beta|$ will have label j . A similar remark holds for labeling around a vertex of Υ . Points (if any) in $p_0 \cap \beta_j$, $\alpha_i \cap q_0$ and $p_0 \cap q_0$ similarly are assigned the ordered pairs $(0, j)$, $(i, 0)$ and $(0, 0)$ respectively, giving rise to ends of edges labeled 0 in respectively Υ , Σ , and both Σ and Υ .

An important difference between the topology exploited in this proof and that used in the analysis of Dehn surgery (cf [CGLS]) is that in the latter, any two components α_i, β_j of the boundaries of P and Q respectively always intersect in the torus boundary of the 3-manifold with the same orientation. In the present situation, each α_i will intersect each β_j with both orientations. Indeed, since both bound meridian disks of W , the sum of the orientations of all points of intersection between the two will be trivial.

Let $A \subset \partial W$ be the annulus between α_1 and α_m whose core is $\partial\alpha$. To any proper arc in A that spans A and is properly isotoped to minimally intersect the α_i , assign the orientation that points from α_1 to α_m . This is called the *spanning* orientation of the arc. Similarly let $B \subset \partial W$ be the annulus between β_1 and β_n and define the spanning orientation for spanning arcs of B to be the orientation that points from β_1 to β_n . The spanning orientation for A gives what is called the *spanning* normal orientation in ∂W to each component $\alpha_1, \dots, \alpha_m$ of ∂P , namely the normal orientation that points from α_i to α_{i+1} . Similarly define the spanning normal orientation for each $\beta_j \subset \partial Q$.

Fix normal orientations for P and Q . Following such a choice, an edge in Σ whose corresponding edge in Υ connects two vertices (so neither end in Σ is labelled 0) is called *incoherent* if the normal orientation of Q along the arc of intersection agrees with the spanning orientation at one end and disagrees at the other. Otherwise (if it agrees at both ends or disagrees at both ends) it is called *coherent*. Whether the edge is coherent or incoherent

is independent of the original choice of orderings of the α_i and the β_j or the choice of normal orientation for Q . Note that by definition an incoherent edge cannot have ends with the same label j . If two coherent edges are parallel and adjacent, one with ends labeled j_1, j'_1 and the other with corresponding ends labeled j_2, j'_2 then $j_1 - j'_1 \equiv j_2 - j'_2 \pmod n$. Similarly if two incoherent edges are parallel and adjacent, then $j_1 + j'_1 \equiv j_2 + j'_2 \pmod n$. Similar remarks hold for labels of ends of edges in Υ . Note that if two edges are parallel and if one is coherent and the other is incoherent, then the spanning orientations agree on one end (i.e. induce the same orientation on that component of ∂P) and disagree on the other. The only way that spanning orientations can disagree at adjacent labels is if the interval between them does not lie on B , so one pair of adjacent ends are either both labeled 1 or both labeled n .

Lemma 2.8. *There are no trivial loops in either Υ or Σ .*

Proof. We show that an innermost trivial loop in Υ can be used to reduce m . $W - \hat{P}$ consists of $m - 1$ copies of $D^2 \times I$ (labelled W_1, \dots, W_{m-1}) and two solid tori W_0 and W_m . Each 1-handle W_i , $1 \leq i \leq m - 1$ lies between copies $\hat{\alpha}_i$ and $\hat{\alpha}_{i+1}$ of α ; the solid torus W_0 is incident to \hat{P} in $\hat{\alpha}_1$ (and possibly p_0) and the solid torus W_m is incident to \hat{P} in $\hat{\alpha}_m$.

An innermost trivial loop in Υ cuts off a disk D from Q whose boundary consists of an arc q on $\partial Q - P \subset \partial W - P$ (in the component of ∂Q on which the loop is based) and an arc γ in $P \cap Q$. Suppose first that, q lies in the annulus lying between some α_i and α_{i+1} , $1 \leq i < m$. q can't be inessential in the annulus, since $|\partial P \cap \partial Q|$ has been minimized up to isotopy in ∂W . So q spans the annulus. The 1-handle W_i can be viewed as a regular neighborhood of the arc q . Then D can be used to isotope W_i through $\gamma \subset P$, removing both $\hat{\alpha}_i$ and $\hat{\alpha}_{i+1}$ and reducing m by 2. A similar argument, reducing m by 1, applies if q has one end on α_1 and its other end on p_0 .

The only remaining possibility is that q is an essential arc in $\partial W - \eta(\alpha)$ with both its ends on one of the components α_1 or α_m , say α_m . (Both ends have to be on the same curve α_1 or α_m , since α is separating.) That is, q is an essential arc on the punctured torus $\partial W_m - \hat{\alpha}_m$.

Claim: q passes through a meridian of W_m exactly once.

Proof of Claim: The arc γ together with a subarc p of α_m form a closed curve in P , and so a closed curve that bounds a disk in $M[\alpha]$. If this disk is attached to D along γ the result is a disk D_+ in $M[\alpha]$ whose boundary $p \cup q$ is an essential closed curve on the boundary of W_m . Since M contains no Lens space summands or non-separating spheres, $p \cup q$ is a longitude of W_m , and so intersects a meridian of W_m in a single point. By isotoping p to be very short, the intersection point lies in q , establishing the claim.

Following the claim, W_m can be viewed as a regular neighborhood of q , attached to $\hat{\alpha}_m$. The disk D_+ appearing in the proof of the claim is a special ∂ -reducing disk for W . The number of times D_+ passes through α is the number of vertices lying between γ and α_m in \hat{P} and so in particular is at most $m - 1$. This contradicts the original choice of \hat{P} . The symmetric argument eliminates trivial loops in Σ . \square

Following Lemma 2.8 we may immediately assume that $m, n \geq 2$ since, for example, if $m = 1$ the fact that fewer ends of edges of Σ lie on p_0 (if it exists) than on α_1 guarantees that there will be a trivial loop at α_1 . In the absence of trivial loops in the graphs, our analysis will focus on edges that are parallel in the graph. Parallel edges cut off faces that are bigons. We will be interested in large families of parallel edges. So consider a collection of parallel edges e_1, \dots, e_t in Y , with ends on vertices v, w . Number them in order around v , making an arbitrary choice between the two possible ways of doing this. Arbitrarily call v the source vertex. Then the family of edges defines a function ϕ from a sequence of labels around v (namely the labels of the ends of e_1, \dots, e_t at v) to a sequence of labels around w (namely the labels of the ends of e_1, \dots, e_t at w). In a typical setting we will know or assume a lot about the label sequence at v (called the *source sequence*) and a little about the label sequence at w . Typically we will only know that the label sequence at w lies as a contiguous subsequence of a much larger sequence called the *target sequence*.

In our setting, and momentarily assuming no labels 0 appear, the labels around any vertex in Y appear in (circular) order

$$1, \dots, m, m, \dots, 1, 1, \dots, m, m, \dots, 1 \quad \dots \quad 1, \dots, m, m, \dots, 1$$

with $|\partial\alpha \cap \partial\beta| \geq 6$ determining the total number of sequences $1, \dots, m$ and $m, \dots, 1$ that appear. When labels 0 do appear, i.e. when \hat{P} is a special ∂ -reducing disk and p_0 intersects $\partial\beta$, then a label 0 appears between some (but not necessarily all) successive labels $1, 1$. If the source sequence of a set of parallel edges (that is, the sequence of labels at the vertex v) is of length $t \leq m + 1$ then the above long sequence is a natural target sequence. That is, we know that the ends of the edges e_1, \dots, e_t at w have labels some ordered contiguous subsequence in the long sequence above (with perhaps some 0's inserted). But we could equally well have used the shorter target sequences $1, \dots, m, m, \dots, 1, 1, \dots, m$ or $1, \dots, m, m, \dots, 1, 0, 1, \dots, m$ since any ordered sequence of $t \leq m + 1$ contiguous labels in the long sequence is also an ordered contiguous sequence in $1, \dots, m, m, \dots, 1, 1, \dots, m$ or $1, \dots, m, m, \dots, 1, 0, 1, \dots, m$.

The following lemmas are classical, going back at least to [GL], [Sc1]:

Lemma 2.9. *Suppose ϕ is a function as described above determined by a set of parallel edges in Υ , with source sequence $0, 1, \dots, m$. For all $0 \leq i \leq m-1$, if $\phi(i) = i+1$, then $\phi(i+1) \neq i$.*

Proof. The case $1 \leq i \leq m-1$ is representative. The bigon between the edges represents a rectangle embedded in M with one pair of opposite sides lying in \hat{P} . The other pair of opposite sides are parallel spanning arcs on the annulus in ∂W between α_i and α_{i+1} . Moreover these spanning arcs are oriented in such a way that if the rectangle between them in the annulus is added, the result is a Möbius band A with its boundary on \hat{P} . The union of the Möbius band A and a disk component of $\hat{P} - \partial A$ is a copy of \mathbb{RP}^2 in M , whose regular neighborhood is then a punctured \mathbb{RP}^3 in M , contradicting the hypothesis that M contains no Lens space summands. \square

Lemma 2.10. *Suppose ϕ is a function as described above determined by a set of parallel edges in Υ , with source sequence contained in $0, 1, \dots, m$. Then:*

- *For all $1 \leq i \leq m-1$, if $\phi(i) = i$, then $\phi(i+1) \neq i-1$.*
- *If $\phi(1) = 1$, then $\phi(2) \neq 1$.*
- *If $\phi(m) = m$, then $\phi(m-1) \neq m$.*

Proof. In each case, the hypothesis implies that an edge in Υ with the same label on each end is incoherent. \square

Jointly call these lemmas *the standard \mathbb{RP}^3 contradiction*, and note that they are typically applied to show that a set of mutually parallel edges can't be too numerous. For example, in our context:

Lemma 2.11. *No label sequence $0, 1, \dots, m$ or $1, \dots, m, m$ appears as a source sequence for any set of parallel edges in Υ . Moreover, if no label 0 appears in the target sequence, no label sequence $1, \dots, m$ can appear as a source sequence.*

Proof. We prove the last statement first. Consider the image of the label sequence $1, \dots, m$ in the target sequence $1, \dots, m, m, \dots, 1, 1, \dots, m$. If $\phi(1)$ is the ρ^{th} term in the target sequence, call $0 \leq \rho < 2m$, the *offset* of the function. Consider the possible offsets.

If the offset is trivial ($\rho = 0$) or $\rho = 2m$ then for all $1 \leq i \leq m$, $\phi(i) = i$. This means that for each label $1 \leq i \leq m$ there is an edge in Υ with that label at both ends. Looking at the other graph, this means that every vertex in Σ is the base vertex for a loop. An innermost loop then can have no vertices in its interior, i.e. it would be a trivial loop, contradicting Lemma 2.8.

Suppose $1 \leq \rho \leq 2m-1$. If ρ is even, then the label $m - \frac{\rho}{2}$ contradicts Lemma 2.9. If ρ is odd then the label $m - \frac{\rho-1}{2}$ contradicts Lemma 2.10.




FIGURE 7

The point is informally captured in the graph in Figure 7.

The case when the possible target sequence is $1, \dots, m, m, \dots, 1, 0, 1, \dots, m$ is essentially the same argument, but requires using the entire source sequence $0, 1, \dots, m$ or $1, \dots, m, m$.

□

It is easy to complete the proof of Proposition 2.7 in the case that no labels 0 appear (e. g. when \hat{P} is a sphere or is a special ∂ -reducing disk whose boundary is disjoint from β). Since $|\partial\alpha \cap \partial\beta| \geq 6$, every vertex in Υ has on its boundaries at least six disjoint label sequences $1, \dots, m$ or $m, \dots, 1$. According to Lemma 2.11 no such sequence can be entirely at the end of a parallel set of edges. It follows that there are at least six breaks between sets of parallel edges at each vertex, hence at least six separate sets of parallel edges incident to each vertex. This leads to a simple Euler characteristic contradiction, see [GL, Lemma 4.1] for the case in which \hat{Q} is a sphere.

The remaining case, when \hat{P} is a special ∂ -reducing disk and p_0 intersects $\partial\beta$, is only slightly more delicate. Note first that since β is separating, $|p_0 \cap \partial\beta|$ is even and hence ≥ 2 . This implies that each vertex of Υ has at least two labels 0 in its boundary. Each label 0 stands at the center of a sequence of edge labels

$$1, \dots, m, m, \dots, 1, 0, 1, \dots, m, m, \dots, 1$$

and each of these can be written as the end-point union of four label sequences

$$\begin{aligned} &1, \dots, m, m \\ &m, \dots, 1, 0 \\ &0, 1, \dots, m \\ &m, m, \dots, 1. \end{aligned}$$

Each of these four must have a break by Lemma 2.11. Of the two sets of 4 breaks thereby identified, at most two from each group can coincide since

$|\partial\alpha \cap \partial\beta| \geq 6$. Hence there are again at least $8m - 2 = 6$ breaks around each vertex, leading to the same contradiction. \square

3. WHEN MERIDIANS ARE NON-SEPARATING – AN INTRODUCTION

Much complication is added if one or both of the meridians α or β is non-separating. The most striking is that the graphs Σ and Υ might have trivial loops, as we now describe.

Just as in the proof of Proposition 2.7, no innermost trivial loop in Υ can have ends labelled $i, i + 1$, but there are three other possibilities: as before they could have ends labelled $1, 1$ or m, m . Or they could have ends labeled $1, m$ since, as α is non-separating, some arc of $\partial\beta - \partial\alpha$ could have its ends on opposite sides of $\partial\alpha$. Only the last possibility (ends labeled $1, m$) withstands closer scrutiny:

Lemma 3.1. *No innermost trivial loop in Υ has both ends labelled 1 (or, symmetrically, both ends labelled m).*

Proof. Consider the twice punctured torus $T^* = \partial W - \alpha$. The two punctures (that is, ∂ -components) of T^* can be identified with α_1 and α_m .

Claim 1: Any pair of arcs of $\partial\beta - \partial\alpha$ that have both ends at α_1 (or both ends at α_m) are parallel in T^* .

Proof of Claim 1: Two non-parallel such arcs with ends at α_1 would have complement in T^* a punctured disk, with puncture α_m . All arcs of $\partial\beta - \partial\alpha$ that have one end on α_m must then have their other end on α_1 in order to be essential in T^* . But then $\partial\beta$ would be incident to α_1 at least four more times than it is incident to α_m . This is absurd, since α_1 and α_m are parallel and $|\partial P \cap \partial Q|$ has been minimized up to isotopy. The contradiction proves Claim 1.

Claim 2: Any arc of $\partial\beta - \partial\alpha$ that has both ends at α_1 or both ends at α_m is meridional in T^* . That is, the arc, together with a subarc of α_1 (or α_m) joining the ends, form a meridian circle on the boundary of the solid torus $W - \alpha$.

Proof of Claim 2: Any outermost disk of β cut off by α intersects $T^* \subset \partial W$ in an arc γ which either has both ends at α_1 or both ends at α_m , say the former. The outermost disk is a meridian of the solid torus $W - \alpha$ so γ , together with a subarc of $\partial\alpha_1$ forms a meridian circle in the boundary of the solid torus $W - \alpha$. This and Claim 1 establish Claim 2 for arcs with both ends at α_1 .

Any arc of $\partial\beta - \partial\alpha$ with both ends at α_m (and a counting argument shows that there must be as many such arcs as there are with both ends at α_1) must be meridional since if it had any other slope, it would necessarily

intersect the meridional arc with both ends at α_1 . This establishes Claim 2 also for arcs with both ends at α_m .

Following the claims, consider the disk D cut off from Q by an innermost trivial loop λ in Υ with both ends labeled 1. The two claims guarantee that the curve ∂D intersects $\partial W - \alpha$ in a meridional arc. In particular, the union of D with the disk in \hat{P} cut off by λ and a subarc of ∂P is a disk D_+ in $M[\alpha]$ whose boundary in $W - \alpha$ is a meridian circle on the boundary of the solid torus $W - \alpha$. The union of D_+ with a meridian disk for $W - \alpha$ contradicts the assumption that M contains no non-separating spheres. \square

Following the Lemma, we have that any innermost trivial loop in Υ has ends labeled $1, m$. The disk $D \subset Q$ it cuts off can be used to ∂ -compress P to ∂W . This alters P , replacing the boundary curves α_1, α_m with a separating curve bounding a separating meridian α^* of W . Replace P then with a planar surface (still called P) whose boundary components consist of curves $\alpha_1, \dots, \alpha_k$ parallel to $\partial\alpha$ and curves $\alpha_1^*, \dots, \alpha_s^*$ parallel to $\partial\alpha^*$, each family labeled in order, and in such an order that one complementary component in ∂W is a pair of pants with boundary components α_s^*, α_1 and α_k . The construction guarantees that ∂Q intersects the pair of pants in at least one arc with ends on both α_1 and α_k , hence no arc with both ends on α_s^* . See Figure 8.




FIGURE 8

After this replacement, the vertices of Σ now are of two different types and the labeling of edges around any vertex of Υ is therefore more complicated. But by allowing two vertex types, choosing P and Q to minimize $|\partial P \cap \partial Q|$ now guarantees that the corresponding graph Υ in \hat{Q} has no trivial loops. Indeed, such a trivial loop would have its ends either:

- on α_1^* , which would allow s to be reduced by 1, via the argument of Lemma 2.8;
- on parallel copies of a meridian α or α^* , which would allow s to be reduced by 2, via the argument of Lemma 2.8;
- or on exactly two of the three meridians $\alpha_k, \alpha_1, \alpha_s^*$. In this case, a ∂ -compression removes the two meridians incident to the loop and adds a copy of the third meridian, reducing the total number of meridians and with it $|\partial P \cap \partial Q|$.

The whole construction can be done symmetrically if β is non-separating: Trivial loops could arise in Σ but be dealt with by identifying a separating meridian β^* and altering Q so that it intersects W in some curves parallel to $\partial\beta$ and some other curves parallel to $\partial\beta^*$.

Although, after this alteration, the graphs Σ and Υ have no trivial loops, the combinatorial argument now requires tracking four sets of circles. This means that there are two sets of labels in each graph and two possible labeling schemes around vertices in each graph. Also, the possible target sequences that arise are much more complicated than those that arise in the proof of Proposition 2.7. For example, a target sequence might contain

$$1, \dots, k, 1, \dots, k, 1, \dots, k, \dots$$

intermixed with sequences of the form

$$s^*, \dots, 1^*, 1^*, \dots, s^*.$$

As of this writing, such a combinatorial argument can be constructed for the case of reducing spheres. The argument appears likely to extend to the case of ∂ -reducing disks, but the addition of an extra boundary component (namely, the boundary component of the reducing disk) makes the final result (Conjecture 2 hence Conjecture 1) still uncertain.

Scott Taylor [Ta] has found a sutured manifold argument that avoids much of this combinatorial complication, in the same way that Gabai's [Ga] circumvented the combinatorial difficulties of [Sc3]. From slightly different assumptions he is able to prove an analogue of Conjecture 2 that is even stronger (namely, the first conclusion holds unless α and β are in fact disjoint) except in the most difficult but arguably the most interesting case: when $M = S^3$, α and β are both non-separating and at least one of $M[\alpha]$ or $M[\beta]$ is an unknotted torus.

REFERENCES

- [BS1] S. Bleiler and M. Scharlemann, A projective plane in R^4 with three critical points is standard. Strongly invertible knots have property P , *Topology* **27** (1988) 519–540.
 [BS2] S. Bleiler and M. Scharlemann, Tangles, property P , and a problem of J. Martin, *Math. Ann.* **273** (1986) 215–225.

- [CG] A. Casson and C. McA. Gordon, Reducing Heegaard splittings, *Topology and its applications* **27** (1987) 275-283.
- [CGLS] M. Culler and C. McA. Gordon, J. Luecke, P. Shalen, Dehn Surgery on Knots, *Annals of Mathematics* **125** (1987) 237-300.
- [EM1] M. Eudave Muñoz, Band sums of links which yield composite links. The cabling conjecture for strongly invertible knots, *Trans. Amer. Math. Soc.* **330** (1992) 463–501.
- [EM2] M. Eudave Muñoz, Prime knots obtained by band sums, *Pacific J. Math.* **139** (1989) 53–57.
- [EM3] M. Eudave Muñoz, Primeness and sums of tangles, *Trans. Amer. Math. Soc.* **306** (1988) 773–790.
- [Ga] D. Gabai, Genus is superadditive under band connected sum, *Topology* **26** (1987) 209–210.
- [GL] C. McA. Gordon, R. Litherland, Incompressible planar surfaces in 3-manifolds, *Topology and its Applications* **18** (1984) 121-144.
- [Go] C. McA. Gordon, Dehn filling: a survey, *Knot theory (Warsaw, 1995)* 129–144, Banach Center Publications **42**, Polish Acad. Sci., Warsaw, 1998.
- [Ja] W. Jaco, Adding a 2-handle to a 3-manifold: an application to property R , *Proc. Amer. Math. Soc.* **92** (1984) 288–292.
- [Ki] S. Kinoshita, On elementary ideals of polyhedra in the 3-sphere, *Pac. J. Math* **42** (1972) 89-98.
- [Sc1] M Scharlemann, Tunnel number one knots satisfy the Poenaru conjecture, *Topology and its Applications* **18** (1984) 235-258
- [Sc2] M. Scharlemann, Unknotting number one knots are prime, *Invent. Math.* **82** (1985) 37–55.
- [Sc3] M. Scharlemann, Smooth spheres in R^4 with four critical points are standard, *Invent. Math.* **79** (1985) 125–141.
- [Sc4] M. Scharlemann, Producing reducible 3-manifolds by surgery on a knot, *Topology* **29** (1990) 481-500.
- [ST1] M. Scharlemann and A. Thompson, Unknotting number, genus, and companion tori. *Math. Ann.* **280** (1988) 191–205.
- [ST2] M. Scharlemann and A. Thompson, Link genus and the Conway moves, *Comment. Math. Helv.* **64** (1989) 527–535.
- [Ta] S. Taylor, *personal communication; preprint to appear.*
- [Wa] F Waldhausen, Heegaard-Zerlegungen der 3-Sphäre, *Topology* **7** (1968) 195–203.

MARTIN SCHARLEMANN, MATHEMATICS DEPARTMENT, UNIVERSITY OF CALIFORNIA, SANTA BARBARA, CA USA

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