

One-sided Heegaard splittings of $\mathbb{R}P^3$

LORETTA BARTOLINI

J HYAM RUBINSTEIN

Using basic properties of one-sided Heegaard splittings, a direct proof that geometrically compressible one-sided splittings of $\mathbb{R}P^3$ are stabilised is given. The argument is modelled on that used by Waldhausen to show that two-sided splittings of S^3 are standard.

57M27; 57N10

1 Introduction

Since their formal introduction in 1978 [5], one-sided Heegaard splittings of 3-manifolds have been the subject of little study. This paucity of literature can largely be attributed to the lack of generality of such splittings, as compared with classical Heegaard splittings, and the invalidity of an analogue to Dehn's lemma and the loop theorem [6]. Various works, both prior and subsequent to Rubinstein [5], have addressed non-orientable surfaces in 3-manifolds such as Bredon and Wood [1], Hempel [3], Frohman [2] and Rannard [4], and classifications are made in the latter works when restricted to geometrically incompressible surfaces. However, in order to study one-sided splittings effectively, the existence and behaviour of geometrically compressible splittings must be considered.

Well known in two-sided Heegaard splitting theory, the stabilisation problem is also present for one-sided splittings. By its very nature, this issue demands an understanding of geometrically compressible splitting surfaces. To date, no connection has been drawn between geometric compressibility and stabilisation. Here, a direct correspondence is drawn for the simplest case: $\mathbb{R}P^3$.

The result is analogous to that of Waldhausen's for two-sided splittings of S^3 [7] and it is upon these original arguments that the proof is based. While there have been many subsequent proofs of the S^3 case using simpler arguments, in the absence of an analogue to Casson and Gordon's result on weak reducibility, such approaches are not currently viable for one-sided splittings.

We would like to thank Marc Lackenby for helpful discussions and feedback that assisted in the preparation of this paper.

2 One-sided Heegaard splittings

Throughout, let M be a closed, orientable 3-manifold and consider all manifolds and maps as PL.

Definition 2.1 A pair (M, K) is called a one-sided Heegaard splitting if K is a closed non-orientable surface embedded in M such that $H = M \setminus K$ is an open handlebody.

As with two-sided splittings, it is useful to consider *meridian discs* for (M, K) , which are taken to be the closure of meridian discs for the handlebody complement H in the usual sense. Due to the non-orientability of K , the boundaries of such discs can intersect themselves or one another in two distinct ways (see Figure 1).




Figure 1: Different intersection types for meridian discs of (M, K)

Definition 2.2 If $x = \partial d_i \cap \partial d_j$, where d_i, d_j are meridian discs for a one-sided splitting, and $B_\varepsilon(x)$ is a small ball centred at x , call x isolated if $d_i \cap d_j \cap B_\varepsilon(x) = x$. Call x non-isolated if $d_i \cap d_j \cap B_\varepsilon(x) = \alpha$, where α is an arc containing x .

2.1 Existence

Theorem 2.3 (Rubinstein [5]) *For any element $\alpha \neq 0$ in $H_2(M, \mathbb{Z}_2)$, there is a one-sided Heegaard splitting (M, K) with $[K] = \alpha$.*

The one-sided splitting technique is hence applicable to a large class of 3-manifolds, which can be easily identified using algebraic methods. Associated with any one-sided splitting is a double cover $p: \tilde{M} \rightarrow M$, where $\tilde{K} = p^{-1}(K)$ is the orientable double cover of K . The surface \tilde{K} gives a natural two-sided splitting of $\tilde{M} = p^{-1}(M)$, with handlebody components interchanged by the covering translation $g: \tilde{M} \rightarrow \tilde{M}$.

In order to consider the simplest surface representing a \mathbb{Z}_2 -homology class, a notion of incompressibility for non-orientable surfaces is required.

Definition 2.4 A surface $K \neq S^2$ embedded in M is *geometrically incompressible* if any simple, closed, noncontractible loop on K does not bound an embedded disc in M . Call K geometrically compressible if it is not geometrically incompressible.

The existence of such a one-sided splitting surface is not implied by existence of one-sided splittings in general. However, by restricting to the class of irreducible, non-Haken 3-manifolds, such a connection can be drawn.

Theorem 2.5 (Rubinstein [5]) *If M is irreducible and non-Haken, then there is a geometrically incompressible one-sided splitting associated with any nonzero class in $H_2(M, \mathbb{Z}_2)$.*

While little is known about general geometrically incompressible one-sided surfaces in 3-manifolds, a classification is available for Seifert fibered spaces. The Lens space case is discussed by the second author [5] and general Seifert fibered spaces in Frohman [2] and Rannard [4]. Considering \mathbb{RP}^3 as $L(2, 1)$, the former result is sufficient here.

Combining Theorem 2.3 and Theorem 2.5, any Lens space of the form $L(2k, q)$, where $(2k, q) = 1$, has geometrically incompressible one-sided Heegaard splittings. In [5], it is shown that any such space has a unique, geometrically incompressible splitting that realises the minimal genus of all one-sided splittings of the manifold. An algorithm is given by Bredon and Wood [1] for calculating this genus. Since $H_2(L(2k, q); \mathbb{Z}) = 0$ and all one-sided splitting surfaces of a Lens space are represented by the same \mathbb{Z}_2 -homology class, any splitting surface that is geometrically compressible must geometrically compress to the minimal genus surface.

2.2 Stabilisation

Definition 2.6 A one-sided splitting (M, K) is stabilised if and only if there exists a pair of embedded meridian discs d, d' for H such that $d \cap d'$ is a single isolated point.

Definition 2.7 A one-sided splitting of an irreducible manifold is called irreducible if it is not stabilised.

As stabilised one-sided splitting surfaces are inherently geometrically compressible, irreducibility is implied by geometric incompressibility. In future work, we hope to give evidence that geometric incompressibility of one-sided splitting surfaces is actually analogous to strong irreducibility in the two-sided case.

2.3 Stable equivalence

Definition 2.8 One-sided Heegaard splittings (M_1, K_1) and (M_2, K_2) are equivalent if there exists a homeomorphism from M_1 to M_2 that maps K_1 to K_2 .

As for two-sided splittings, there is a notion of stabilising distinct one-sided splittings until they are equivalent. Let (S^3, L) denote the standard genus 1 two-sided splitting of the 3-sphere and $(M, K) \# n(S^3, L)$ be the connected sum of (M, K) with n copies of (S^3, L) .

Definition 2.9 One-sided splittings (M_1, K_1) and (M_2, K_2) are stably equivalent if $(M_1, K_1) \# n(S^3, L)$ is equivalent to $(M_2, K_2) \# m(S^3, L)$ for some m, n .

Unlike two-sided splittings, stable equivalence does not hold for one-sided Heegaard splittings in general. However, a version applies to splitting surfaces represented by the same \mathbb{Z}_2 -homology class:

Theorem 2.10 (Rubinstein [5]) *If (M, K_1) and (M, K_2) are one-sided Heegaard splittings with $[K_1] = [K_2]$, then they are stably equivalent.*

Motivated by the fact that the little that is known about one-sided Heegaard splittings is largely restricted to geometrically incompressible splitting surfaces, we use these basic properties of one-sided splittings to broach geometric compressibility. Given any stabilised one-sided splitting is inherently geometrically compressible, it is natural to ask when geometric compressibility corresponds to stabilisation.

3 One-sided Heegaard splittings of \mathbb{RP}^3

Investigating any existence of a correlation between geometric compressibility and stabilisation, the simplest case to consider is \mathbb{RP}^3 , which corresponds to S^3 in the two-sided case. Here, the original arguments given by Waldhausen are adapted to show that all geometrically compressible splittings of \mathbb{RP}^3 are stabilised.

In brief, the approach is to take an unknown splitting and the known minimal genus splitting by \mathbb{RP}^2 and stabilise the two until they are equivalent. Keeping track of the disc systems introduced by this process, it is possible to arrange them such that the reverse process of destabilising to get the unknown splitting preserves dual pairs from the minimal genus splitting. Thus, dual discs exist for the original unknown splitting, hence it is stabilised.

Theorem 3.1 *Every geometrically compressible one-sided Heegaard splitting of \mathbb{RP}^3 is stabilised.*

Proof Take a geometrically compressible one-sided Heegaard splitting (M, K) of $M \cong \mathbb{RP}^3$ and let (M, P) be the splitting along $P \cong \mathbb{RP}^2$. Since $H_2(M; \mathbb{Z}_2) \cong \mathbb{Z}_2$, there is only one nontrivial \mathbb{Z}_2 -homology class so $[K] = [P]$. As P is the unique geometrically incompressible splitting surface of M , the unknown splitting surface K geometrically compresses to P .

By stable equivalence, each splitting surface can be stabilised a finite number of times until the two are equivalent. Represent this splitting by (M, K') and let $H = M \setminus K'$ be the handlebody complement. Let Δ^K be a set of meridian discs introduced by stabilisations of (M, K) , chosen such that $\Delta^K = \Delta_K \cup \Delta'_K$, where $\Delta_K = d_1, d_2, \dots, d_k$ and $\Delta'_K = d'_1, d'_2, \dots, d'_k$ are sets of disjoint discs with $|d_i \cap d'_i| = 1$ and $d_i \cap d'_j = \emptyset$ for $i \neq j$. Then $|\Delta^K| = 2k = (\text{genus}(K') - \text{genus}(K))$. Note that this number is always even, as each stabilisation increases the genus of the handlebody by 2.

Similarly, let $\Delta^P = \Delta_P \cup \Delta'_P$ be the set of discs introduced by stabilising (M, P) . Notice that since $M \setminus P$ is an open 3-cell, Δ^P is a complete disc system for H .

Consider the non-isolated intersections between discs in Δ_K, Δ'_K and Δ_P, Δ'_P . Let

$$\Lambda_0 = \{d \cap D\}, \quad \Lambda'_0 = \{d' \cap D'\}, \quad \Lambda_1 = \{d \cap D'\} \quad \text{and} \quad \Lambda'_1 = \{d' \cap D\}$$

be the collections of arcs of intersection between the given pairs for all $d \in \Delta_K, d' \in \Delta'_K, D \in \Delta_P, D' \in \Delta'_P$.

Stabilise (M, K') along $\Lambda_0, \Lambda'_0, \Lambda_1, \Lambda'_1$. Call the resulting splitting (M, K'') , with handlebody complement $H' = M \setminus K''$. Let

$$\begin{array}{cccccc} \overline{\Delta}_K & \Delta_K & \Lambda_0, \Lambda_1 & & \Delta'_K & \Lambda'_0, \Lambda'_1 \\ \overline{\Delta}_P & \text{be } \Delta_P & \text{cut } \Lambda_0, \Lambda'_1 & \text{plus the discs } \Delta'_P & \text{along } \Lambda'_0, \Lambda_1 & \\ \overline{\Delta}'_K & \Delta'_K & \text{along } \Lambda'_0, \Lambda'_1 & \text{dual to cuts of } \Delta_K & & \Lambda_0, \Lambda_1 \\ \overline{\Delta}'_P & \Delta'_P & \Lambda'_0, \Lambda_1 & & \Delta_P & \Lambda_0, \Lambda'_1 \end{array}$$

where a disc dual to a cut along an arc λ is a transverse cross-section of a closed regular neighbourhood of λ (see Figure 2). For such discs, use parallel copies for the K and P systems in order to retain dual pairs in each. Let $\overline{\Delta}^K = \overline{\Delta}_K \cup \overline{\Delta}'_K$ and $\overline{\Delta}^P = \overline{\Delta}_P \cup \overline{\Delta}'_P$. Notice that $\overline{\Delta}^P$ is again a complete disc system for H' .

The aim of this second stabilisation process is to remove all existing non-isolated intersections between Δ^P and Δ^K . Therefore, it is imperative that the disc systems are not moved once this second set of stabilisations is complete, as any moves may




Figure 2: Stabilising along an arc λ , where $d \in \Delta_K$ and $D \in \Delta_P$ or Δ'_P

introduce new intersections. Hence, the standard procedure of manipulating stabilising discs to get sets of disjoint dual pairs is not performed.

Order the $\bar{\Delta}_K, \bar{\Delta}'_K$ and $\bar{\Delta}_P, \bar{\Delta}'_P$ disc systems with respect to the nesting of arcs of stabilisation. For example, consider $\bar{d}_i, \bar{d}_j \in \bar{\Delta}_K$ that were split off $d \in \Delta_K$ by arcs λ_i, λ_j respectively. If λ_i is outermost with respect to the point $d \cap d'$, then $j < i$ (see Figure 3). Note that there is a rooted tree dual to the subdisc system for d , where the point of $d \cap d'$ is the root, which induces the ordering. Label the dual discs such that \bar{d}'_k is a transverse cross-section of λ_k , hence $\bar{d}'_k \in \bar{\Delta}'_K$ is dual to \bar{d}_k . Apply similar labelling to the $\bar{\Delta}_P, \bar{\Delta}'_P$ systems.




Figure 3: Discs \bar{d}_i, \bar{d}_j obtained by splitting d along λ_i, λ_j , where $j < i$

Consider the intersections between discs $\bar{d}_i \in \bar{\Delta}_K$ and $\bar{d}'_j \in \bar{\Delta}'_K$. By construction, $\partial\bar{d}_i \cap \partial\bar{d}'_i$ is a single isolated point and $\partial\bar{d}_i \cap \{\partial\bar{d}'_j \mid j = 1, 2, \dots, (i - 1)\} = \emptyset$. For $i \leq j$, points of $\partial\bar{d}_i \cap \partial\bar{d}'_j$ are isolated.

If $m = |\bar{\Delta}_K| = |\bar{\Delta}'_K|$, then $2m$ is the total change in genus from K to K'' . Construct the $2m \times 2m$ intersection matrix $\mathbf{M} = [m_{ij}]$ for discs in $\bar{\Delta}^K$. Define m_{ij} as follows,

where $|\partial\bar{d}_i \cap \partial\bar{d}_j|$ is given to be the number of isolated singularities of \bar{d}_i :

$$m_{ij} = \begin{cases} |\partial\bar{d}_i \cap \partial\bar{d}'_j|, & 1 \leq i, j \leq m \\ |\partial\bar{d}_i \cap \partial\bar{d}'_{j-m}|, & 1 \leq i \leq m, (m+1) \leq j \leq 2m \\ |\partial\bar{d}'_{i-m} \cap \partial\bar{d}'_j|, & (m+1) \leq i \leq 2m, 1 \leq j \leq m \\ |\partial\bar{d}'_{i-m} \cap \partial\bar{d}'_{j-m}|, & (m+1) \leq i, j \leq 2m \end{cases}$$

Since $\bar{\Delta}_K, \bar{\Delta}'_K$ are systems of embedded, disjoint discs, the off-diagonal blocks are zero. By symmetry, the diagonal blocks are mutually transpose. While initially this symmetry makes the full matrix unnecessary, the asymmetry of later moves requires the consideration of all entries as described. Given the discs are not to be manipulated after the second set of stabilisations, the matrix is the identity if and only if $\{\Lambda_0, \Lambda'_0, \Lambda_1, \Lambda'_1\} = \emptyset$. Thus:

$$\mathbf{M} = \left(\begin{array}{ccccc|ccccc} 1 & \star & \star & \dots & \star & 0 & \dots & \dots & \dots & 0 \\ 0 & 1 & \star & \dots & \star & & & & & \\ \vdots & \ddots & \ddots & \ddots & \vdots & \vdots & & \ddots & & \vdots \\ 0 & \dots & 0 & 1 & \star & & & & & \\ 0 & 0 & \dots & 0 & 1 & 0 & \dots & \dots & \dots & 0 \\ \hline 0 & \dots & \dots & \dots & 0 & 1 & 0 & \dots & 0 & 0 \\ & & & & & \star & 1 & 0 & \dots & 0 \\ \vdots & & \ddots & & \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ & & & & & \star & \dots & \star & 1 & 0 \\ 0 & \dots & \dots & \dots & 0 & \star & \dots & \star & \star & 1 \end{array} \right)$$

If $n = |\bar{\Delta}_P| = |\bar{\Delta}'_P|$, the $2n \times 2n$ intersection matrix \mathbf{N} for the discs in $\bar{\Delta}^P$ can be constructed similarly. This \mathbf{N} has a similar block structure to \mathbf{M} .

Let $D = \bar{D}_n \in \bar{\Delta}_P$, the disc corresponding to the last row of the upper half of \mathbf{N} , and let $D' \in \bar{\Delta}'_P$ be its dual. Thus D, D' are a dual pair disjoint from all other discs in $\bar{\Delta}^P$. However, several possibilities exist for how D, D' may intersect $\bar{\Delta}^K$:

- (a) Both D, D' are disjoint from $\bar{\Delta}^K$ or the pair intersect only one of $\bar{\Delta}_K, \bar{\Delta}'_K$;
- (b) One of D, D' is disjoint from $\bar{\Delta}^K$, while the other intersects both $\bar{\Delta}_K, \bar{\Delta}'_K$;
- (c) Both D and D' intersect $\bar{\Delta}^K$ and Case (a) does not apply.

In Case (a), compress along whichever of $\bar{\Delta}_K, \bar{\Delta}'_K$ is disjoint from D and D' . This results in (M, K) , without having affected D, D' , which remain a dual pair of embedded discs. Therefore, (M, K) is stabilised.

In Case (b), suppose D intersects both $\bar{\Delta}_K$ and $\bar{\Delta}'_K$, while $D' \cap \bar{\Delta}^K = \emptyset$. Since D' is disjoint from $\bar{\Delta}_K$, it can be used to remove intersections between D and $\bar{\Delta}_K$ by a process of band-summing:

Take $d \in \bar{\Delta}_K$, with $d \cap D \neq \emptyset$, such that there exists an arc $\alpha \subset \partial D$ with one endpoint at $D \cap D'$, the other at $d \cap D$ and $\alpha \cap \bar{\Delta}_K = \emptyset$. Join a parallel copy of D' to d by the boundary of a closed half-neighbourhood of α . This removes one point from $d \cap D$ and since $\alpha \cap \bar{\Delta}_K = \emptyset$, no additional intersections are created within $\bar{\Delta}_K$. Repeat this procedure for all discs in $\bar{\Delta}_K$ that intersect D , taking care to work in an order that does not introduce intersections. Thus, all intersections between D and $\bar{\Delta}_K$ can be removed without changing the intersection properties of $\bar{\Delta}_K$, resulting in Case (a) above.

In Case (c), both D and D' intersect $\bar{\Delta}^K$, with no immediate means by which to remove intersections. Any attempts at band-summing, as used for Case (b), would introduce intersections between $\bar{\Delta}_K$ and $\bar{\Delta}'_K$. Therefore, it is this case that requires significant attention.

Claim After modifying $\bar{\Delta}^K$, there exists a dual pair of discs $\bar{d}, \bar{d}' \in \bar{\Delta}^K$ such that $|\partial \bar{d} \cap \partial D| = 1$ and $|\partial \bar{d}' \cap \partial D| \leq 1$ (or vice versa), and $D \cap (\bar{\Delta}^K \setminus \{\bar{d}, \bar{d}'\}) = \emptyset$.

The proof of this claim requires two steps, each of which is technical in nature. In particular, in the first step the most vital, yet most subtle, part of the argument appears.

Step 1 Describe surgery on $\bar{\Delta}^K$ in order to make $|\partial d \cap \partial D| \leq 1$ for all $d \in \bar{\Delta}^K$.

Consider arcs contained in ∂D with endpoints on ∂d . Take a shortest arc $\alpha \subset \partial D$, with endpoints $\{a_0, a_1\}$ such that $a_i \in \partial d$ and $\alpha^\circ \cap d = \emptyset$. Such an arc can be chosen such that $\alpha \cap D' = \emptyset$. If $\beta_1, \beta_2 \subset \partial d$ are the arcs with $\partial \beta_i = \{a_0, a_1\}$, let $\beta = \beta_i$ such that $\beta \cap d'$ is a single point.

Subclaim 1 The loop γ formed by the arcs α and β bounds a disc in H' that is dual to d' .

In order to prove Subclaim 1, it is necessary to consider both isolated and non-isolated intersections between the $\bar{\Delta}^P$ and $\bar{\Delta}^K$ disc systems. For clarity, the subtleties are best captured by passing to the orientable double cover.

Take the orientable double cover (\tilde{M}, \tilde{K}'') corresponding to (M, K'') , with covering projection $p: \tilde{M} \rightarrow M$, covering translation $g: \tilde{M} \rightarrow \tilde{M}$ and handlebody components H_1, H_2 . Let $\tilde{d} = p^{-1}(d) \cap H_1$ and $\tilde{D} = p^{-1}(D) \cap H_2$, so an isolated intersection

between d and D will correspond to discs in opposite handlebodies meeting in a point on the splitting surface. Let $\tilde{\beta} = p^{-1}(\beta) \cap \tilde{d}$ and $\tilde{\alpha} = p^{-1}(\alpha) \cap \tilde{D}$. Hence, the loop $\tilde{\gamma}$, bounded by $\tilde{\alpha}$ and $\tilde{\beta}$, is on \tilde{K}'' and constitutes part of the boundary of a disc in each handlebody.

Since all non-isolated intersections between $\bar{\Delta}^P$ and $\bar{\Delta}^K$ have been removed, the intersection $(p^{-1}(\bar{\Delta}^P) \cap H_i) \cap (p^{-1}(\bar{\Delta}^K) \cap H_i) = \emptyset$ for $i = 1$ or 2 . Specifically, $\tilde{d} \cap g(\tilde{D}) = g(\tilde{d}) \cap \tilde{D} = \emptyset$, so the loop γ formed by α, β on K'' lifts to a pair of disjoint loops $\tilde{\gamma}, g(\tilde{\gamma})$ on \tilde{K}'' formed by $\tilde{\alpha}, \tilde{\beta}$ and $g(\tilde{\alpha}), g(\tilde{\beta})$ respectively.

As $\bar{\Delta}^P, \bar{\Delta}^K$ have no non-isolated intersections, \tilde{d} is disjoint from $p^{-1}(\bar{\Delta}^P) \cap H_1$, which is a complete disc system for H_1 . Thus the loop $\tilde{\gamma}$ bounds a disc \tilde{d}^1 in H_1 . Applying similar arguments to $g(\tilde{d})$ and $p^{-1}(\bar{\Delta}^P) \cap H_2$, the translated loop $g(\tilde{\gamma})$ bounds $g(\tilde{d}^1)$ in H_2 . Since $\tilde{\gamma}, g(\tilde{\gamma})$ are disjoint, $\tilde{d}^1, g(\tilde{d}^1)$ are discs in opposite handlebodies with disjoint boundaries, hence $\tilde{d}^1 \cap g(\tilde{d}^1) = \emptyset$. Projecting to (M, K'') , the disc $d^1 = p(\tilde{d}^1 \cup g(\tilde{d}^1))$ is embedded and dual to d' , by choice of β .

If α is disjoint from $\bar{\Delta}^K \setminus d$, replace d with d^1 , which has two fewer points of intersection with D than d . Repeat the process to remove all pairs of adjacent points in $d^1 \cap D$. Let d_α be the resulting disc and replace d with d_α in $\bar{\Delta}^K$.

Any remaining arc $\alpha \subset \partial D$ between points of intersection with d_α is interrupted by intersections with $\bar{\Delta}^K \setminus d_\alpha$. These points of intersection are necessarily isolated.

Subclaim 2 There is a disc d_0 with $\partial\alpha_0 \subset \partial d_0$ for some $\alpha_0 \subset \alpha$ such that the intersection $\alpha_0 \cap (\bar{\Delta}^K \setminus d_0) = \emptyset$.

Take $d_K \in \bar{\Delta}^K \setminus d_\alpha$ with $x \in (d_K \cap \alpha)$ and again lift to the orientable double cover. Let $\tilde{d}_K = p^{-1}(d_K) \cap H_1$, so $p^{-1}(x) \in (\tilde{d}_K \cap \tilde{\alpha}) \cup (g(\tilde{d}_K) \cap g(\tilde{\alpha}))$ since $\tilde{d}_K \cap g(\tilde{D}) = \emptyset$. Now both \tilde{d}_α and \tilde{d}_K intersect \tilde{D} . By the previous argument, α and part of $\partial\tilde{d}_\alpha$ bound a disc \tilde{d}_α^1 in H_1 . Since \tilde{d}_α^1 and \tilde{d}_K are discs in the same handlebody, their boundaries intersect in pairs of points. However, \tilde{d}_K does not intersect \tilde{d}_α , so both points of intersection lie on $\tilde{\alpha}$. Therefore, \tilde{d}_K intersects $\tilde{\alpha}$ in pairs of points. Similar arguments apply to $g(\tilde{d}_\alpha^1), g(\tilde{d}_K)$, thus d_K intersects α in pairs of points.

Applying the above argument to any discs intersecting the subarc $\alpha_K \subset \alpha$, where $\partial\alpha_K \subset \partial d_K$, yields that any arcs of intersection between $\bar{\Delta}^K \setminus d_\alpha$ and d_α^1 are nested. Therefore, there exists an innermost pair corresponding to intersections with the desired disc $d_0 \in \bar{\Delta}^K \setminus d_\alpha$ (see Figure 4).

Apply the previous surgery to split d_0 along α_0 and reduce the number of points of intersection with D . Continue this process, from edgemost arcs inwards, to remove all




Figure 4: Nested discs intersecting α

pairs of intersection points between D and $\bar{\Delta}^K \setminus d_\alpha$. Applying the previous surgery to d_α , the pair of intersection points that constitute the boundary of α can then be removed. Hence, the number of intersection points can be reduced to at most one.

Therefore, D intersects any disc in $\bar{\Delta}^K$ in at most one point. If D is disjoint from all such discs, then Case (a) above applies and the result holds.

Step 2 Reduce the number of discs in $\bar{\Delta}^K$ that have nonempty intersection with D to at most a dual pair.

Consider two discs, each intersecting D in a single point and let $\lambda \subset \partial D$ be the arc with an endpoint on each disc. Choose $d_a, d_b \in \bar{\Delta}^K$ to be such that $\lambda \cap (\bar{\Delta}^K \setminus \{d_a, d_b\}) = \emptyset$. Say that such discs are adjacent, since they are directly next to each other with respect to ∂D .

Do not perform surgery if the discs are a dual pair. Otherwise, take a parallel copy of whichever of d_a, d_b corresponds to a later stabilisation—say d_b . Join the copy of d_b to d_a by the boundary of a closed half-neighbourhood of λ . This forms a new disc \bar{d}_a with $\bar{d}_a \cap D = \emptyset$. Replace d_a in $\bar{\Delta}^K$ with \bar{d}_a . Note that any intersections of d_b with $\bar{\Delta}^K$ will be present in \bar{d}_a .

The effect of the surgery on the intersection matrix is to add the row of \mathbf{M} corresponding to d_b to that corresponding to d_a . If both d_a, d_b belong to one of $\bar{\Delta}_K, \bar{\Delta}'_K$, the surgery does not affect the off-diagonal blocks of \mathbf{M} . However, if $d_a = \bar{d}_k \in \bar{\Delta}_K, d_b = \bar{d}'_l \in \bar{\Delta}'_K$, where $k < l$, the k -th row of \mathbf{M} becomes:

$$\left(\underbrace{0 \dots 0}_k \mid 1 \star \dots \star \mid \star \dots \star \mid 1 \overbrace{0 \dots 0 \dots 0}^l \right)$$

Specifically, the $(m+k)$ -th entry of the k -th row remains 0. Therefore, throughout all surgery, discs in $\bar{\Delta}^K$ remain embedded. This allows the procedure to be iterated if necessary.

Perform surgery on all adjacent discs (except dual pairs) until there is, at most, a single pair of dual discs \bar{d}, \bar{d}' , each intersecting D in a single point. This pair corresponds to the latest stabilisation of any discs that had intersected D after Step 1. Note that this may not be the pair corresponding to the centremost rows of \mathbf{M} , as these discs may not have initially intersected D .

Having thus found discs \bar{d} and \bar{d}' that prove the claim, it is now possible to destabilise K'' in a useful manner:

Step 3 Replace \bar{d}' in $\bar{\Delta}'_K$ with D . Compress along D , thus destabilising K'' . Discard \bar{d} .

Since D is disjoint from $\bar{\Delta}^K \setminus \{\bar{d}, \bar{d}'\}$, all other discs in this system remain intact after the compression. Therefore, the remaining discs again form systems of embedded dual pairs that correspond to stabilisations of K and P , the latter of which is complete with respect to the newly destabilised splitting surface. As the original properties required for surgery on the discs systems are retained, Steps 1, 2 and 3 can be repeated for remaining discs in $\bar{\Delta}^P$. If the process is not terminated by the occurrence of Cases (a) or (b) as described previously, this process of destabilisation continues until it results in the original splitting (M, K) .

Since (M, P) has minimal genus, $|\bar{\Delta}^K| < |\bar{\Delta}^P|$ as $K \not\cong \mathbb{RP}^2$. Therefore, after destabilising (M, K'') to get (M, K) by the above process, there are dual pairs of discs remaining in $\bar{\Delta}^P$. Therefore, (M, K) is stabilised. \square

References

- [1] **G E Bredon, J W Wood**, *Non-orientable surfaces in orientable 3-manifolds*, Invent. Math. 7 (1969) 83–110 MR0246312
- [2] **C Frohman**, *One-sided incompressible surfaces in Seifert fibered spaces*, Topology Appl. 23 (1986) 103–116 MR855450
- [3] **J Hempel**, *One sided incompressible surfaces in 3-manifolds*, from: “Geometric topology (Proc. Conf., Park City, Utah, 1974)”, Springer, Berlin (1975) 251–258. Lecture Notes in Math., Vol. 438 MR0402748
- [4] **R Rannard**, *Incompressible surfaces in Seifert fibered spaces*, Topology Appl. 72 (1996) 19–30 MR1402233
- [5] **J H Rubinstein**, *One-sided Heegaard splittings of 3-manifolds*, Pacific J. Math. 76 (1978) 185–200 MR0488064
- [6] **J Stallings**, *On the loop theorem*, Ann. of Math. (2) 72 (1960) 12–19 MR0121796

- [7] **F Waldhausen**, *Heegaard-Zerlegungen der 3-Sphäre*, *Topology* 7 (1968) 195–203
MR0227992

*Department of Mathematics and Statistics, University of Melbourne
Parkville VIC 3010, Australia*

L.Bartolini@ms.unimelb.edu.au, H.Rubinstein@ms.unimelb.edu.au

Received: 31 August 2005 Revised: 6 July 2006