COMPARING HEEGAARD SPLITTINGS OF NON-HAKEN 3-MANIFOLDS

HYAM RUBINSTEIN[1] AND MARTIN SCHARLEMANN[2]

Abstract. Cerf theory can be used to compare two strongly irreducible Heegaard splittings of the same closed orientable 3-manifold. Any two splitting surfaces can be isotoped so that they intersect in a non-empty collection of curves, each of which is essential in both splitting surfaces. More generally, there are interesting isotopies of the splitting surfaces during which this intersection property is preserved. As sample applications we give new proofs of Waldhausen's theorem that Heegaard splittings of S^3 are standard, and of Bonahon and Otal's theorem that Heegaard splittings of lens spaces are standard. We also present a solution to the stabilization problem for irreducible non-Haken 3-manifolds: If $p \leq q$ are the genera of two splittings of such a manifold, then there is a common stabilization of genus $5p + 8q - 9$.

1. Background

In this paper, all 3-manifolds are assumed to be orientable and, except for handlebodies, to be closed as well. Much of the machinery developed works also for compact orientable manifolds split into compression bodies, but the arguments are more delicate and will appear elsewhere. A *handlebody* H is the boundary sum of a finite number of copies of $S^1 \times D^2$. Alternatively H is a homeomorph of the regular neighborhood of some finite graph in R^3. The image Ξ of the graph , to which H retracts, is called a *spine* of H. The retraction restricts to a map $\partial H \to \Xi$ whose mapping cylinder is itself homeomorphic to H. A properly imbedded essential disk in H is called a *meridian* of H. A collection of meridians is *complete* if its complement is a collection of 3-balls.

A *Heegaard splitting* $M = A \cup_P B$ of a 3-manifold consists of an orientable surface P in M, together with two handlebodies A and B into which P divides M. P itself is called the splitting surface. The genus of $A \cup_P B$ is defined to be the genus of P. A *stabilization* of $A \cup_P B$ is the Heegaard splitting obtained by adding to A a regular neighborhood of a proper arc in B which is parallel in B to an arc in P. A

[0]Mathematics subject classification (1991): *Primary 57N10, secondary 57M50.*
[1]Partially supported by a grant from the Australian Research Council
[2]Partially supported by NSF grant DMS 9203522.

stabilization has genus one larger and, up to isotopy, is independent of the choice of arc in B and is the same if the construction is done symmetrically to an arc in A instead.

If there are meridian disks D_A and D_B in A and B respectively so that ∂D_A and ∂D_B intersect transversally in a single point in P, then $A \cup_P B$ can be obtained by stabilizing a lower genus Heegaard splitting. We then say that $A \cup_P B$ is *stabilized* or *can be destabilized*. If there are meridian disks D_A and D_B in A and B respectively so that ∂D_A and ∂D_B are disjoint in P, then $A \cup_P B$ is *weakly reducible*. If there are meridian disks so that $\partial D_A = \partial D_B$, then $A \cup_P B$ is *reducible*. It is easy to see that reducible splittings are weakly reducible and that (except for the genus one splitting of S^3) any stabilized splitting is reducible. It is a theorem of Casson and Gordon [CG] that if $A \cup_P B$ is a weakly reducible splitting then either M contains an incompressible surface, or $A \cup_P B$ is reducible. It is a theorem of Haken [Ha] that any Heegaard splitting of a reducible 3-manifold is reducible and it follows from a theorem of Waldhausen [W] that a reducible splitting of an irreducible manifold can be destabilized.

This last theorem, that any positive genus Heegaard splitting of S^3 is standard, is the deepest. (For an updated proof, see [ST2].) The viewpoint we adopt here easily gives a new proof of this theorem (see 5.11). The other ingredient in our proof of 5.11 is the main theorem of [CG] which implies that any weakly reducible splitting of S^3 is reducible. A few early lemmas here are easier to state if we know 5.11, so we will put in [brackets] conditions which are not needed once 5.11 is known.

Any two Heegaard splittings of the same 3-manifold can be stabilized until they agree but it is uncertain how many stabilizations suffice. For lens spaces, no stabilization is needed [Bo], [BoO]. Our methods here give an easy alternative proof 6.3, 6.4. Examples exist [BO] for which one stabilization is necessary, and Johannson has shown [Jo, 40.5] that if M is Haken, then the number of stabilizations needed grows no more than polynomially with the genus of the two splitting surfaces. It is a consequence of what we show here that for irreducible non-Haken 3-manifolds the growth is linear. We suspect that this will generalize to Haken 3-manifolds as well, and that it can be derived from the machinery used here, together with that of [ST3].

I would like to thank the referee for several clarifying suggestions.

2. Main results

In sections 3 through 5 we develop the underlying machinery. If P and Q are two Heegaard splitting surfaces of the same closed orientable 3-manifold M then the splittings determine "sweep-outs" of M by $P \times I$ and $Q \times I$. Generically, copies of P and Q are transverse during the sweep-outs, but there are codimension one and two sets on which they are not transverse. When P and Q are strongly irreducible splittings there is a structure on these strata, viewed as a graphic in $I \times I$. For M irreducible and non-Haken, a Heegaard splitting is strongly irreducible if and only if it is irreducible.

At the end of section 5 we begin to develop the topological consequences. First we recover the main theorems already known for Heegaard splittings of non-Haken 3-manifolds:

Theorem 5.11 *Any positive genus Heegaard splitting of S^3 can be destabilized.*

Theorems 6.3 and 6.4 *Any lens space has a unique irreducible Heegaard splitting.*

This last result is an easy consequence of the following:

Theorem 6.2 *Suppose $X \cup_Q Y$ and $A \cup_P B$ are strongly irreducible Heegaard splittings of the same 3-manifold $M \neq S^3$. Then P and Q can be isotoped so that $P \cap Q$ is a non-empty collection of curves which are essential in both P and Q.*

This shows that P and Q can be put into a useful position in M. One can also find isotopies of P in M whose track across Q contains useful information. This is the content of the end of Section 6 through Section 7. The remainder of the paper shows how to use the isotopy to produce a bound on the number of stabilizations needed to make two splittings equivalent. The main technical result is this:

First define a *spine* of a closed orientable surface Q to be a 1-complex in Q whose complement consists entirely of disks.

Theorems 6.5 and 8.1 *Suppose $X \cup_Q Y$ and $A \cup_P B$ are strongly irreducible Heegaard splittings of the same 3-manifold M. Then P and Q may be put in general position so that*

(1) *all but one curve in $P \cap Q$ is essential in both P and Q. The other curve, if it exists, is inessential in both P and Q.*

(2) *for one of the splittings (say $A \cup_P B$) there is a complete collection of meridian disks Δ for A and B so that the 1-complex $Q \cap (P \cup \Delta)$ contains a spine of Q.*

This leads to the stabilization bound:

Theorem 11.5 *Suppose $X \cup_Q Y$ and $A \cup_P B$ are strongly irreducible Heegaard splittings of the same 3-manifold M and are of genus $p \leq q$ repectively . Then there is a genus $8q+5p-9$ Heegaard splitting of M which stabilizes both $A \cup_P B$ and $X \cup_Q Y$.*

Corollary 11.6 *Suppose $X \cup_Q Y$ and $A \cup_P B$ are Heegaard splittings of the same irreducible non-Haken 3-manifold M and are of genus $p \leq q$ repectively . Then there is a genus $8q + 5p - 9$ Heegaard splitting of M which stabilizes both $A \cup_P B$ and $X \cup_Q Y$.*

This bound is almost surely not the best possible. Two recent announcements of better bounds are Lu [Lu], who gives $4q - 3$ and Taimanov [Ta], who gives $p + q$. However, the status of these proofs remains uncertain.

3. Sweep-outs and their graphics

Suppose $A \cup_P B$ is a Heegaard splitting of M, and Ξ_A, Ξ_B are spines of A and B respectively. We may as well take spines in which each vertex has valence three.

Definition 3.1. *A* sweep-out *associated to the Heegaard splitting $A \cup_P B$ is a relative homeomorphism $H : P \times (I, \partial I) \to (M, \Xi_A \cup \Xi_B)$ which, near $P \times \partial I$, gives a mapping cylinder structure to a neighborhood of $\Xi_A \cup \Xi_B$.*

Given a sweep-out H and $0 < s < 1$, let P_s denote the splitting surface $H(P \times s)$, $P_{<s}$ denote the handlebody $H(P \times [0, s])$ and $P_{>s}$ denote the handlebody $H(P \times [s, 1])$.

If $M = X \cup_Q Y$ is another Heegaard splitting of M and Q is in general position with respect to $\Xi_A \cup \Xi_B$ and the sweep-out H is generic with respect to Q, then, for small values of ϵ, $P_{<\epsilon} \cap Q$ is a (possibly empty) collection of meridian disks of A and $P_{>1-\epsilon} \cap Q$ is a (possibly empty) collection of meridian disks of B. Generically, $P_t \cap Q$ is a disjoint collection of simple closed curves in Q.

We are interested in analyzing intersection patterns which arise in simultaneous sweepouts P_s, Q_t of M corresponding to different Heegaard splittings. Cerf theory (see [C]) says that for generic sweep-outs, the interior of the square $I \times I = \{(s, t) | 0 \leq s, t \leq 1\}$ decomposes into four strata:

Regions:: The set of values (s,t) for which P_s and Q_t intersect transversally comprise an open subset of $int(I \times I)$. A component of this two-dimensional stratum is called a *region*.

Edges:: The set of points (s,t) for which P_s and Q_t intersect transversally except for a single non-degenerate tangent point comprise a 1-dimensional subset of $int(I \times I)$. A component of this 1-dimensional stratum is called an *edge*.

Crossing vertices:: These are the points (s,t) for which P_s and Q_t have exactly two non-degenerate points of tangency but are otherwise transverse. Such points are isolated in $I \times I$.

Birth-death vertices:: These are the isolated set of points at which P_s and Q_t intersect transversally except for a single degenerate tangent point locally modelled on $P_s = \{(x,y,z)|z=0\}$ and $Q_t = \{(x,y,z)|z = x^2 + y^3\}$.

The set of edges and vertices form a 1-complex Γ called the *graphic* in the interior of $I \times I$. An edge is *adjacent* to a region if its contained in the closure of the region. Two regions are adjacent if there is an edge which is adjacent to both of them. We similarly define edges to be adjacent if they terminate in the same vertex. A crossing vertex has valence 4 in Γ, for it represents a point where an edge in the graphic associated to one tangent point crosses an edge corresponding to another. A birth-death vertex has valence two, with one adjacent edge corresponding to a saddle and the other corresponding to a cancelling center. Locally there is a parameterization (λ, μ) of (s,t)-space so that, if P_s is $\{(x,y,z)|z=0\}$, then $Q_t = \{(x,y,z)|z = x^2 + \lambda + \mu y + y^3\}$ (see [C, II.2])

The graphic Γ naturally extends to a properly imbedded 1-complex in all of $I \times I$: A point $(0,t)$, say, on $\{0\} \times I \subset \partial(I \times I)$ represents simultaneously the spine Ξ_A of handlebody A (since $s = 0$) and the surface Q_t. Generically these are transverse, implying that P_ϵ and Q_t are transverse for ϵ small. There are two types of exceptions: For finitely many values of t, Ξ_A is tangent to Q_t at a single point in the interior of one of its edges. At finitely many other values of t, Q_t crosses a vertex of Ξ_A. Since each vertex of Ξ_A is of valence three, this changes the number of intersection points with Ξ_A by ± 1. Call these non-generic points (and similar points on the other three sides of $I \times I$) *boundary vertices* of Γ. For $(0,t_0)$ such a boundary vertex, consider nearby points in the interior of $I \times I$. As Q_t sweeps across the point a where Ξ_A and Q_{t_0} are tangent, consider how Q_t sweeps across P_ϵ for small ϵ. There are two nearby values $t_\pm$ of t so that P_ϵ is tangent to each of $Q_{t_\pm}$ at a single point. Between the values $t_\pm$, Q_t sweeps across the meridian of P_ϵ at a and these two values of t are the first and last values for which Q_t intersects this meridian. At one of $t_\pm$ the tangency to P_ϵ is a center and at the other a saddle. In the graphic, this means that the boundary vertex $(0, t_0)$ abuts two edges in the graphic Γ. Similarly, a boundary vertex corresponding to a sweep of Q_t across a vertex of Ξ_A abuts an edge of Γ corresponding to a saddle tangency of Q_t with P_ϵ near the vertex. The same argument applies at each boundary vertex, so Γ can be completed to a 1-complex in $I \times I$ by adjoining all boundary vertices. We continue to call this 1-complex the graphic Γ.

4. Essential and inessential curves of intersection

Consider a region of $I \times I - \Gamma$ as defined above. The collection of curves $P_s \cap Q_t$ is, up to isotopy, independent of the choice of (s, t) in a given region and we'll often suppress the subscripts when they are clear from the context. Our first goal is to find a region in which this collection contains curves which are essential in both P and Q. To that end we define certain subcollections of curves.

Definition 4.1. *For (s, t) in a region of $I \times I - \Gamma$, P and Q intersect transversally in a collection $\mathcal{C}$ of simple closed curves. Let $\mathcal{C}_P$ (resp. $\mathcal{C}_Q$) denote the set of these curves which are essential in P (resp. Q). A curve c in $\mathcal{C}_P$ is further defined to be in $\mathcal{C}_A$ if it bounds a disk in $Q - \mathcal{C}_P$ which, near c, lies in A. We similarly define $\mathcal{C}_B \subset \mathcal{C}_P$ and $\mathcal{C}_X, \mathcal{C}_Y \subset \mathcal{C}_Q$.*

Definition 4.2. *A curve $c \in \mathcal{C}$ is* remote *from $c' \in \mathcal{C}$ in P (resp Q) if no component of $P - Q$ (resp $Q - P$) has both c and c' on its boundary.*

Lemma 4.3. *If $c \in \mathcal{C}_A$ then c bounds a disk in A. Moreover, if no curve in $\mathcal{C}_X$ or $\mathcal{C}_Y$ is remote from c in P, then c bounds a disk in A which intersects Q only in inessential circles. Symmetric statements hold for $c \in \mathcal{C}_B, \mathcal{C}_X, \mathcal{C}_Y$.*

Proof. If the disk D which c bounds in $Q - \mathcal{C}_P$ has interior disjoint from $\mathcal{C}$ then a slight push-off would be a disk in $A - Q$ as required. Since $int(D)$ is disjoint from $\mathcal{C}_P$ at worst it intersects P in circles which are inessential in P.

Let $c_1, \ldots, c_n$ be the components of $int(D) \cap P$ that are outermost in $int(D)$. Let D_i, E_i be the disks bounded by c_i in $int(D)$ and P respectively. Then the desired disk is $E = (D - \bigcup_{i=1}^n \eta(D_i)) \cup \bigcup_{i=1}^n E_i'$, where E_i' is a copy of E_i pushed into A in such a way that if $E_i \subset E_j$ then E_i' is pushed slightly further into A than E_j'. The interior of each E_i' intersects Q only in curves which are remote in P from c, since c_i separates them. Then E is a disk in A bounded by c which intersects Q only in curves parallel to curves of $P \cap Q$ which are remote from c in P and which bound disks in P. If one of these curves is essential in Q then an innermost such in E lies in $\mathcal{C}_X \cup \mathcal{C}_Y$. □

Corollary 4.4. *If in any region both $\mathcal{C}_A$ and $\mathcal{C}_B$ are non-empty then $A \cup_P B$ is weakly reducible.*

Lemma 4.5. *Suppose $\mathcal{C}_P$ and $\mathcal{C}_Q$ are empty and there is a meridian disk in A which intersects Q only in inessential circles. If A also contains an essential curve of Q then $A \cup_P B$ is weakly reducible [or M is S^3]. Symmetric statements hold for A replaced with B, or for (A, P, Q) replaced with (X, Q, P) or (Y, Q, P).*

Proof. Since $\mathcal{C}_P \cup \mathcal{C}_Q = \phi$ any curve in $P \cap Q$ bounds a disk in both P and Q. A standard innermost disk argument provides an isotopy of Q which makes P and Q disjoint. This isotopy affects neither the essential curve of Q lying in A nor the existence of a meridian disk in A intersecting Q only in inessential circles. After the isotopy we conclude that Q must lie entirely in A and that A has a meridian disk disjoint from Q. Attach to B a maximal collection of 2-handles which ∂-reduce A in the complement of Q. The resulting 3-manifold B' has boundary a surface P' lying entirely in either X or Y, say X.

If P' consists of 2-spheres, then M can be obtained from B' by attaching some 3-handles, one of which must contain Q and hence all of Y. Since X is irreducible it follows that the boundary of the 3-handle bounds also a ball containing B', so M has a Heegaard splitting of genus 0 and so is S^3.

If P' contains a non-spherical component then that component, since it lies in X, must be compressible in X. The compressing disk can't lie outside B' by definition of B', so it must compress in B'. This implies that $A \cup_P B$ is weakly reducible [CG]. □

5. Labelling regions of the graphic

Motivated by the above discussion, we label a region of $I \times I - \Gamma$ according to the following scheme. If $\mathcal{C}_A$ (resp. $\mathcal{C}_B$, $\mathcal{C}_X$, $\mathcal{C}_Y$) is non-empty we label it A (resp. B, X, Y). If $\mathcal{C}_P$ and $\mathcal{C}_Q$ are both empty and A (resp. B) contains an essential curve of Q label the region b (resp. a) (*sic*) and if X (resp. Y) contains an essential curve of P label the region y (resp. x). Notice that no region can have both labels a and b (or both labels x and y) since, if some essential curve of Q lies in B and another lies in A then these must be separated in Q by some essential curve from $P \cap Q$ so $\mathcal{C}_Q$ would be non-empty. So the label a actually implies that some spine of Q lies in B, and similarly for labels b, x, y. By Corollary 4.4 , if $A \cup_P B$ (resp. $X \cup_Q Y$) is strongly irreducible, no region can have both labels A and B (resp. X and Y). Finally, no region can have both an upper case label and a lower case label, for the former implies that one of $\mathcal{C}_P$ or $\mathcal{C}_Q$ is non-empty, while the latter assumes that both are empty.

Consider how labels can change as we cross an edge in Γ. Each such edge corresponds to a non-degenerate tangent point between P and Q, and crossing the edge is equivalent to pushing P across Q at that point. In particular, if the tangent point is a "center", a single circle of intersection, inessential in both P and Q is either created or destroyed, and there is no effect on the labeling. If the tangent point is a "saddle" then there can be an effect on the labelling, for passing through the saddle has the effect of banding together two curves of $P \cap Q$ into one, or vice versa.

To understand the effect of this move, suppose curves c_0, c_1 of $P \cap Q$ are banded together to make the curve c. The "figure 8" component of $P \cap Q$ containing the saddle tangency has a regular neighborhood in P (resp. Q) which is a pair of pants. Each of the three boundary components of the neighborhood is parallel in P (resp. Q) to one of c_0, c_1 or c. So if c and c_0, say, are both essential in P and one bounds a disk in A and the other in B, then $A \cup_P B$ is weakly reducible. Hence

Corollary 5.1. *If, in two adjacent regions of $I \times I - \Gamma$, both labels A and B (resp. X and Y) appear, then $A \cup_P B$ (resp. $X \cup_Q Y$) is weakly reducible.*

Similarly, suppose labels a and b occur on opposite sides of the edge. This requires first of all that $\mathcal{C}_P$ and $\mathcal{C}_Q$ be empty throughout, so in particular c_0, c_1 and c are inessential in both P and Q. Secondly it requires that some essential curve of Q lies in A before passing through the saddle and a perhaps different essential curve of Q lies in B after passing through the saddle. But if c_0, c_1 and c are inessential, passing through the saddle has no effect on whether or not such essential curves exist, so there must simultaneously be essential curves of Q in both A and B, and so an essential curve of $P \cap Q$ in Q which separates them, contradicting $\mathcal{C}_Q = \phi$. We have then:

Corollary 5.2. *In two adjacent regions of $I \times I - \Gamma$, labels a and b (resp. x and y) cannot both appear.*

We have earlier noted that no region can have both an upper case letter and a lower case letter. But, under certain circumstances, adjacent regions may have labels of different cases:

Lemma 5.3. *Suppose, in $I \times I - \Gamma$, a region labelled A (or B) is adjacent to a region labelled with a lower case letter. Then the edge represents a saddle tangency in which a band which is essential in P and inessential in Q is attached to an intersection curve which is inessential in both P and Q. (And symmetrically, when (A, P, Q) is replaced with (X, Q, P) or (Y, Q, P).)*

Proof: In the region R_l labelled with a lower case letter, all intersection curves are inessential in both surfaces, whereas in the adjacent region R_A labelled A there is at least one intersection curve which is essential in P (and it is inessential in Q). So the edge must represent a saddle tangency. As described above, the saddle tangency corresponds in each surface to a band move which divides a single component c of into two components, c_0 and c_1. If c_0 and c_1 were curves of intersection in the region R_l, then they would be inessential in P and so c would be also. Since in fact an essential curve is created passing from R_l to R_A, it must be that c is a curve of intersection in the region R_l and c_0, c_1 curves of intersection in the region R_A. Furthermore, at least one of c_0, c_1 is essential in P and inessential in Q. But since c_0 and c_1 are made by a band move on an inessential curve, they must be parallel in both P and Q. □

Corollary 5.4. *Suppose, in $I \times I - \Gamma$, a region labelled A is adjacent to a region labelled b. Then either $A \cup_P B$ is weakly reducible [or M is S^3]. (And symmetrically, if (A, b) is replaced with (B, a) or $(A, b, A \cup_P B)$ is replaced with $(X, y, X \cup_Q Y)$ or $(Y, x, X \cup_Q Y)$.)*

Proof: Using the notation of the previous proof: since $\mathcal{C}_Q$ is empty in R_l and all three curves c, c_0, c_1 are inessential in Q, it follows that $\mathcal{C}_Q$ is empty in R_A. Then 4.3 applied in R_A shows that c_0 bounds a disk in A which intersects Q only in inessential circles. The result then follows from Lemma 4.5 applied in R_l. □

We have used upper and lower cases of the same letter because of the similarities of 5.1, 5.2, and 5.4. To further exploit this similarity we will let A mean a label which could be either a or A, and similarly for B, X, and Y. For example, 5.1, 5.2, 5.4 can all be summarized by:

Corollary 5.5. *If labels* A *and* B *appear in adjacent regions of $I \times I - \Gamma$ then either $A \cup_P B$ is weakly reducible [or M is S^3.]*

Now make a similar analysis around vertices in Γ. Consider first a birth-death vertex. One of the two edges incident to any birth-death vertex corresponds to a center tangency between P and Q and we know that on opposite sides of such an edge labels don't change. So an edge incident to a birth-death vertex has the same labels on both sides.

Now consider a crossing vertex v in Γ at which four edges meet. The four edges divide a neighborhood of v in $I \times I$ into four quadrants, each lying in some region. If an incident edge corresponds to a center tangency, so will the edge opposite to it across v. Such edges will have the same labels on both sides, so they are really invisible in our labelling scheme. Suppose both pairs of opposite edges at v correspond to saddle tangencies, and the two saddle points lie on different singular components of $P \cap Q$ at v. Then the arguments above apply separately across each edge. In particular, if both labels A and B (resp. both labels X and Y) appear in quadrants of such a vertex then either $A \cup_P B$ (resp. $X \cup_Q Y$) is weakly reducible or M is S^3.

The remaining case is that the two saddle points lie on the same singular component of $P \cap Q$ at v. The behavior of $P \cap Q$ in the four quadrants near the vertex can then be described as follows: Among the curves of $P \cap Q$ determined by one quadrant (called the north) is a component c_n to which bands corresponding to the two saddles are attached. In each of the two adjacent quadrants (the east and west) is a pair of curves in $P \cap Q$ obtained by attaching one of the two bands. We denote the pairs respectively as $c_{e_\pm}$ and $c_{w_\pm}$. In the remaining quadrant (the south) each of the pair of curves $c_{e_\pm}$ and $c_{w_\pm}$ are banded together by one of the saddles to produce either three curves or one curve of $P \cap Q$, depending on how the bands are situated. We call this curve (these curves) c_s.

Lemma 5.6. *In each of P and Q, either c_n can be isotoped off c_s or $c_{e_\pm}$ can be isotoped off of $c_{w_\pm}$.*

Proof: If the two bands are attached to the same end of a collar of c_n then the opposite side of the collar is an isotope of c_n which persists after both bands are attached to make c_s. If the bands are attached on opposite ends, then c_n separates $c_{e\pm}$ from $c_{w\pm}$. $\square$

Lemma 5.7. *If all four letters* A, B, X, *and* Y *appear in quadrants of a crossing vertex of* Γ *then either two opposite quadrants are unlabelled, or one of* $A \cup_P B$ *or* $X \cup_Q Y$ *is weakly reducible [or* M *is* S^3*].*

Proof: It follows from 5.5 that if the conclusion does not hold (that is, if all four letters appear, at least two adjacent quadrants are labelled, M is not S^3 and both splittings are strongly irreducible) then each of the four letters appears in a different quadrant, with A opposite B and X opposite Y. If all four letters are upper case then, following 5.6, there is either a curve in $\mathcal{C}_A$ disjoint from a curve in $\mathcal{C}_B$ or a curve in $\mathcal{C}_X$ disjoint from a curve in $\mathcal{C}_Y$. But the former would imply that $A \cup_P B$ is weakly reducible and the latter would imply that $X \cup_Q Y$ is weakly reducible. Hence we conclude that in at least one quadrant there is only a lower case letter, say x. But then in that quadrant $\mathcal{C}_P$ and $\mathcal{C}_Q$ are empty, so either Q is a sphere (making $M = S^3$) or an essential curve in Q lies in one of A or B. This would force the label a or b on that quadrant and thereby ensure that a letter A lies across an edge from a letter B which, via 5.5 completes the proof. $\square$

The following variant of 5.7 is only used in the proof of 5.11 below:

Lemma 5.8. *If labels* A *and* B *and some lower case letter all appear as labels of quadrants of a crossing vertex, then* $A \cup_P B$ *is either weakly reducible or it can be destabilized.*

Proof: From 5.1 we may as well assume that A and B are in opposite quadrants and that both saddle tangencies lie on the same singular component of $P \cap Q$ at the crossing vertex. The lower case label is in one of the other quadrants R_l. According to 5.3 the move from the lower case quadrant to the quadrants labelled A or B is accomplished by attaching a band to an inessential component c of $P \cap Q$. The bands are disjoint, since they correspond to simultaneous saddles at the vertex, so a curve produced by one band intersects a curve produced by the other in either one point (if the bands are attached along the same side of c and the ends of the two bands are linked in c) or none (otherwise). But when one band is attached the curves bound a meridian of A and when the other is attached the curves bound a meridian of B. $\square$

Suppose that $A \cup_P B$ and $X \cup_Q Y$ are strongly irreducible and $M \neq S^3$. Consider the labelling of the regions adjacent to $\partial(I \times I)$. Suppose, for example, that (s,t) is a generic point with s near 0. Then P_s is the boundary of a small regular neighborhood of a spine Ξ of A. If Ξ intersects Q, then Q intersects $P_{<s}$ in meridian disks, so the region should be labelled A. If Ξ is disjoint from Q, then so is $P_{<s}$, and P_s lies in either X or Y. It follows that the region is labelled y or x and, since Q is not a sphere, also labelled a. Similarly, regions adjacent to $\{1\} \times I$ are either labelled B or labelled b and one of x or y, regions adjacent to $I \times \{0\}$ are either labelled X or labelled x and one of a or b, and regions adjacent to $I \times \{1\}$ are either labelled Y or labelled y and one of a or b. Any region whose closure contains a vertex on $\partial(I \times I)$ also has one of these four types of labellings. If the boundary vertex abuts only one edge in Γ this is obvious. If it abuts two, then one of the edges only corresponds to a center tangency, so the regions on either side of that edge will have the same label, and one is fully adjacent to $\partial(I \times I)$. Regions adjacent to the four corners of $I \times I$ must then be labelled, respectively, $(a,x), (a,y), (b,x)$ and (b,y). Under these conditions we have:

Proposition 5.9. *There is an unlabelled region in $I \times I - \Gamma$.*

Proof: Amalgamate edges of Γ which are incident to the same birth-death vertex, so that all vertices of Γ have valence 4. Let Λ be the dual complex to Γ in $I \times I$. Then the labelling of the regions of $I \times I - \Gamma$ gives a labelling of vertices of Λ and, since each vertex of Γ is of valence 4, each face of Λ is 4-sided. Let Λ_l denote the subcomplex of Λ consisting of vertices which are labelled, and edges and faces of Λ which are incident only to labelled vertices. It follows from 5.5 that the labelling defines a simplicial map ϕ from the 1-skeleton of Λ_l to the 1-skeleton of the complex K shown in Fig. 1. Explicitly, ϕ assigns to a vertex of Λ_l the identically labelled vertex of K. ϕ extends to any (4-sided) face of Λ_l, since the only essential 4-cycle in K (namely A-X-B-Y) can't appear around such a face, by 5.7. On the other hand, we've just seen that the circuit in Λ coming from regions and edges whose closures intersect $\partial(I \times I)$ lies in Λ_l. In fact, the description above of the labels on this circuit shows that ϕ maps the circuit to K with winding number 1. But Λ is contractible, so ϕ can't extend to all of Λ. This means that $\Lambda_l \neq \Lambda$, so some regions are unlabelled. □

A path in $I \times I$ is *generic* if it is in general position with respect to Γ. That is, it never goes through a vertex of Γ and is transverse to each edge of Γ. Let U

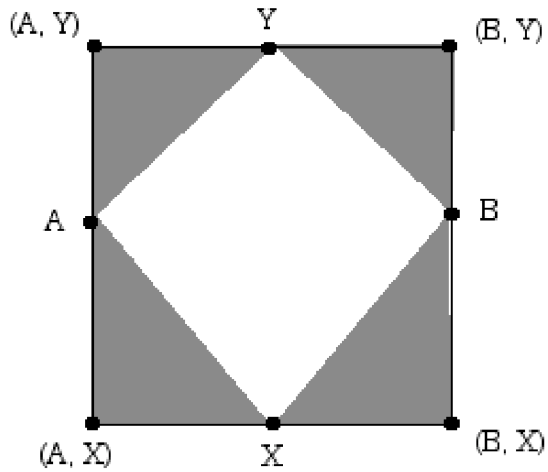

FIGURE 1

denote the union of the unlabelled regions, together with all their adjacent edges and vertices.

Proposition 5.10. *For one of the pairs of letters* A, B *or* X, Y *(say the latter) there is a generic path in U which begins at an edge adjacent to a region labelled* X *and ends at an edge adjacent to a region labelled* Y.

Proof: We continue with the same notation. Suppose R is an unlabelled region. There is a generic path in U from the interior of R to an edge adjacent to a labelled region S. Give R the labels of S and any other such labelled region which can be reached by a generic path in U from R. If R ends up with both labels A and B or both labels X and Y then the ends of the paths which give these labels can be connected in R to give the path we are looking for. This gives a labelling scheme for R which we can apply to every other previously unlabelled region so that either we can find the required path, or

$*$ each region is labelled and no region has both labels A and B or both labels X and Y.

Then each region in $(I \times I) - \Gamma$ has labels which correspond to some vertex in K and, by the definition of the labelling rule and assumption $*$, adjacent regions still satisfy 5.5.

Now suppose that all four labels A, B, X, and Y appear in the four quadrants at a vertex. In order to satisfy 5.5, which we've shown remains true for the new labelling, each quadrant must carry precisely one of the four labels. But according to our new labelling rule, any region which was previously unlabelled will have all the labels of adjacent regions. This would mean that all four regions were among those regions which were labelled to begin with, and this is forbidden by 5.7. We conclude that 5.7 remains true in the new labelling.

Since 5.5 and 5.7 still hold in the new labelling, the map ϕ of 5.9 can be extended to all of Λ, which, as observed in the proof of 5.9 is absurd. We conclude that $*$ is false, and the required path exists. □

Suppose now that M is S^3, P is not the 2-sphere, but Q is. Then the labels a, b, X, and Y never appear, since Q is simply connected. Also, every region must have some label, either label A or B if a curve of intersection is essential in P, or label x or y if no curve of intersection is essential in P.

Theorem 5.11. *Any positive genus Heegaard splitting of S^3 can be destabilized.*

Proof: If the theorem is false, then there is a least genus counterexample. Let $A \cup_P B$ be such a counterexample. First we show that $A \cup_P B$ is weakly reducible. If it isn't, then from 5.1 and 5.8 we conclude that no two adjacent regions can be labelled A, B or x, y, and around no crossing vertex can all labels A, B, x, y occur. Since every region is labelled, this leads to the same sort of contradiction as in the proof of 5.9. Here we use a map $\phi : \Lambda \to K'$, where the 1-complex K' is just a square with its four corner vertices labelled $A - x - B - y$ in order around K'.

Since $A \cup_P B$ is weakly reducible and S^3 contains no incompressible surfaces, it follows from [CG] that $A \cup_P B$ is reducible. That is, some 2-sphere intersects P in a single essential circle. Then $A \cup_P B$ can be viewed as the connected sum along this 2-sphere of two Heegaard splittings of S^3, each of positive genus, but of lower genus than P. By choice of $A \cup_P B$ each of these lower genus Heegaard splittings can be destabilized. This implies that $A \cup_P B$ also can be destabilized, a contradiction. □

6. Interpreting the graphic

We continue with the hypotheses of 5.9 and 5.10: $A \cup_P B$ and $X \cup_Q Y$ are strongly irreducible and $M \neq S^3$. These propositions mean, first, that P and Q can be positioned in a particularly interesting way in M and secondly that there is an isotopy of P with useful properties. To be precise, begin with

Definition 6.1. *A pair of surfaces (P, Q) in M is* compression-free *if P and Q are in general position and each curve of $P \cap Q$ is either essential in both P and Q or is* *inessential in both P and Q. A curve of the former type is called an* essential *curve of intersection and one of the latter type is called an* inessential *curve of intersection.*

An isotopy $F : P \times I \to M$ is compression-free *with respect to Q if it is in general position with respect to Q and, at every regular value t, $F(P \times \{t\})$ and Q are compression-free.*

With these definitions we have:

Corollary 6.2. *P may be isotoped in M so that P and Q are in general position and intersect in a non-empty family of curves, each of which is essential in both P and Q.*

Proof: Consider the positions of P and Q corresponding to an unlabelled region of the graphic (5.9). First note that P and Q are compression-free, for if, say, some intersection curve were essential in P but inessential in Q then an innermost such curve would either lie in $\mathcal{C}_A$ or $\mathcal{C}_B$ forcing the label A or B onto the region. Moreover, if $P \cap Q$ consisted only of inessential curves, then $\mathcal{C}_P$ and $\mathcal{C}_Q$ would be empty, and any essential curve in P could be made disjoint from $P \cap Q$ in P, and so would lie in either X or Y. This would force the label x or y on the region. So $P \cap Q$ must contain some essential curves. A standard innermost disk argument then gives an isotopy of P which eliminates all inessential curves of intersection without eliminating the essential curves of intersection. □

From 6.2 one can immediately deduce the central theorems of [Bo] and [BoO] which together classify Heegaard splittings of the lens spaces.

Corollary 6.3. *Any two genus one Heegaard surfaces in a lens space are isotopic.*

Proof: Let P and Q be two genus one Heegaard surfaces in a lens space, separating the lens space, as usual, into solid tori A and B and solid tori P and Q respectively. According to 6.2, P and Q may be isotoped so that they intersect in a non-empty family of essential circles. Further assume that they've been isotoped to minimize the number $n > 0$ of such circles. Since the surfaces are separating, n is even. An easy outermost arc argument on the intersection of Q with a meridian disk of A or B shows that n can always be reduced, so in fact $n = 2$. Then $P_X = P \cap X, P_Y = P \cap Y, Q_A = Q \cap A$ and $Q_B = Q \cap B$ are all annuli which are boundary parallel in their respective solid tori. That is, P_X is parallel in X to one of Q_A or Q_B and symmetrically for the other three annuli. Together, these four statements imply that P_X is parallel to one of Q_A or Q_B and P_Y is parallel to the other. This means P is parallel to Q. □

Corollary 6.4. *Any irreducible Heegaard splitting of a lens space has genus one.*

Proof: Let $A \cup_P B$ be a genus one Heegaard splitting of a lens space L and $X \cup_Q Y$ be a splitting of higher genus. Since L contains no incompressible surfaces, it suffices to show that Q is weakly reducible. According to 6.2, P and Q may be isotoped so that they intersect in a non-empty family of essential circles. As in 6.3 assume that they've been isotoped to minimize the number $n > 0$ of such circles and let $P_X = P \cap X, P_Y = P \cap Y, Q_A = Q \cap A$ and $Q_B = Q \cap B$. Very explicit information is known about the structure of Q_A and Q_B ([MR]), but this is a deeper result than we will need here, so we proceed with a direct argument.

Case 1: Q_A and Q_B both contain components which aren't annuli.

Consider the families of annuli P_X and P_Y in their respective handlebodies. None can be boundary parallel, since n has been minimized. With no loss of generality, suppose P_X is incompressible in X and ∂-compresses to Q_A. This implies that there is a meridian disk of X which lies in A. If there were also a meridian disk of Y lying in B then Q would be weakly reducible. We conclude that also P_Y ∂-compresses to Q_A (and not to Q_B) or compresses in A. Symmetrically, P_X then can't ∂-compress to Q_B. But an outermost arc argument on the intersection of Q with a meridian of B shows that one of P_X and P_Y must ∂-compress to Q_B, a contradiction. We are reduced to

Case 2: Q_A or Q_B (say the former) consists entirely of annuli.

In this case, as in 6.3, we can assume that $n = 2$ and Q_A is a single annulus, so P_X and P_Y are each a single annulus as well. With no loss of generality, we can assume that P_X is incompressible in X and Q_A and P_X are parallel via $A \cap X$. Consider the annulus $P_Y \subset Y$. It must be ∂-compressible in the handlebody Y. If it ∂-compresses to an arc in Q_B, then the combination of this ∂-compression and an isotopy of P_X to Q_A defines an isotopy of P which makes $P - Q$ a disk D. Then $X \subset B$ and A is a solid torus summand of Y. That is, Q gives a Heegaard splitting of the solid torus B and such splittings are easily shown to be either genus one or reducible (see [ST2]).

If P_Y ∂-compresses to an arc in Q_A then these two are parallel as well, via $A \cap Y$, so we can switch the roles of X and Y in the above argument. An outermost arc argument on the intersection of a meridian of B with Q shows that Q_B ∂-compresses to either P_Y (and we are done as above) or to P_X. In the latter case, switch the roles of X and Y. □

For $F : P \times I \to M$, an isotopy of P which is in general position with respect to $Q \subset M$, let $f_t : P \to M$ denote $F|P \times \{t\}$ and P_t denote $f_t(P)$.

Proposition 6.5. *For one of the pairs of letters A, B or X, Y (say the latter), there is an isotopy $F : P \times I \to M$ so that*

1. *F is compression-free with respect to Q*
2. *every component of $P_0 \cap Q$ and $P_1 \cap Q$ is essential*
3. *there is a meridian disk of X which is disjoint from P_0*
4. *there is a meridian disk of Y which is disjoint from P_1.*

Proof: Consider the path in the graphic given by 5.10. The path begins at an edge separating an unlabelled region R_u from a region R_{X} labelled X. If R_{X} is in fact labelled x, then extend the path slightly into R_{X}. If it is labelled X, then truncate the end slightly so that the path begins in R_u. Similarly extend or truncate the other end of the path. The resulting path lies entirely in regions which are either unlabelled or have lower case labels. The path defines simultaneous isotopies of P and Q in M. Extend the isotopy which this gives of Q in M to an ambient isotopy of all of M and then compose the simultaneous isotopies of P and Q with the inverse of this ambient isotopy of M. This maneuver makes Q stationary throughout, and we can focus on the resulting isotopy of P. Since the path never enters a region with an upper case label, this isotopy of P is compression-free with respect to Q.

Consider how P_0 intersects Q. Suppose first that R_{X} is labelled X. Then in that region there is a curve c of intersection which is essential in Q and bounds a disk in P_0 containing no other essential curve in Q. After a single saddle tangency we enter region R_u, where the intersection is compression-free. This implies that c is altered by a band move at the saddle. After the saddle move, $\mathcal{C}_A \cup \mathcal{C}_B$ is empty, so before the saddle move (i. e. in R_{X}) no component of $\mathcal{C}_A$ or $\mathcal{C}_B$ can be remote from c in Q. Then c bounds a disk in X which intersects P_0 only in inessential curves (4.3). Let c' be a curve in Q which is parallel to c and on the side of c opposite to that on which the saddle is attached. Then c' will be unaffected by the saddle move. So after the saddle move (i. e. in R_u), c' bounds a disk in X which intersects P_0 only in inessential circles.

Now suppose instead that R_{X} is labelled x, so the path begins in this region. In R_{X} all components of $P_0 \cap Q$ are inessential in both P_0 and Q. Moreover, the label x

means that a spine of P_0 lies in Y. This implies that a meridian of X chosen so that its boundary is disjoint in Q from $P_0 \cap Q$ intersects P_0 only in inessential circles.

We have shown that, regardless of whether R_X is labelled x or X, there is a meridian disk of X which intersects P_0 only in inessential circles. In order to guarantee that the meridian is in fact disjoint from P_0, we first describe how to alter the isotopy so that at the beginning of the isotopy all curves of intersection are essential.

A standard innermost disk argument shows that we can eliminate all inessential circles of $P_0 \cap Q$ by an isotopy of P. Recall the argument: Let D be the disk in Q bounded by an inessential component c of $P \cap Q$ that is innermost on Q. Let E be the disk in P bounded by c. The sphere $E \cup D$ can be pushed off P and hence bounds a ball B whose interior is disjoint from P. Hence P is isotopic to $P' = (P - E) \cup D$. This isotopy eliminates c and maybe other curves from $P \cap Q$; continue until all inessential components of $P \cap Q$ are eliminated.

Notice that this isotopy moves only a neighborhood of disks in P, and the only part of Q through which parts of P are moved is the part lying inside the ball B. But $B \cap Q$ consists of disks since, by assumption, the curves $\partial B \cap Q = E \cap Q$ are inessential in Q and no spine of Q can lie inside B (since $M \neq S^3$). This means that the isotopy is compression-free. Furthermore, the isotopy only deletes, and never creates, curves of intersection of P with the meridian of X found earlier. But once this isotopy has eliminated all inessential components of $P \cap Q$, it follows immediately that the meridian in fact bounds a disk in X which is disjoint from P. So we precede our original isotopy with the reverse of the isotopy just defined. Then the isotopy remains compression-free and begins from a position in which a meridian of X is disjoint from P. Then at the beginning of the isotopy all curves of intersection are essential and a meridian of X is disjoint from P.

The same argument and construction can be applied at the end of the isotopy. $\square$

It may be worth noting that there is nothing which prevents F from being constant. That is, there is no reason why it cannot simultaneously be true that $P \cap Q$ contains only essential curves, and that there are meridians of X and Y which are disjoint from P. Of course the boundaries of these meridians must intersect, since Q is strongly irreducible, and, for the same reason, $P \cap Q$ can't be empty (cf 4.5).

7. Finding simple isotopies

In analogy with 6.2, it would be good if we could somehow limit the number of inessential curves of intersection which appear during the isotopy constructed in 6.5.

Of course a center tangency creates or destroys an inessential curve of intersection, so it would be restrictive indeed not to allow any inessential curves of intersection at all during the isotopy. The aim of this section is to do the next best thing.

Definition 7.1. *An isotopy of P which is compression-free with respect to Q is called* simple *if for each generic value t there is no more than one inessential curve in $P_t \cap Q$.*

We will show that the isotopy of 6.5 can be used as a model to construct a similar isotopy which is simple. The first lemma shows that this is true in a special case.

For an isotopy of P which is in general position with respect to Q, let C_t denote $P_t \cap Q$ and let q_t denote the pre-image of C_t in P. For all but a finite number of critical values of t, C_t is a collection of simple closed curves. At the critical values, C_t may contain a single point of tangency, either a saddle point (lying in a component of C_t homeomorphic to the figure 8) or a center, which is an isolated point.

Lemma 7.2. *Suppose there is an isotopy $F : P \times I \to M$ so that*

1. *F is in general position with respect to Q*
2. *each component of C_0 and C_1 is essential in both P and Q.*
3. *for every regular value of t, C_t is the union of C_0 and a collection of curves which are inessential in both P and Q.*

Then there is an isotopy $F' : P \times I \to M$ from f_0 to f_1, so that for all t, ${f'}_t(P) \cap Q = C_0$.

Proof: Extend q_0 to a spine ζ of P. We first define the isotopy F' on ζ so that during the isotopy, ${f'}_t^{-1}(Q)$ just remains $q_0 \subset \zeta$. That is, during the isotopy, the part of the spine ζ away from q_0 never intersects Q.

Away from saddle tangencies of P_t and Q, $F'|\zeta$ will be the composition of F with an isotopy $i_t : \zeta \to P$. Construct the isotopy i_t as follows. Let $t_1, t_1, ..., t_n$ be the levels at which there are saddle tangencies of P_t with Q and let $t_0 = 0$ and $t_{n+1} = 1$. For t near $t_0 = 0$, just let i_t be the inclusion. Suppose, for $i = 0, ..., n$, the isotopy i_t has been defined on $[0, t_i + \epsilon]$. On the following interval, from $t_i + \epsilon$ to $t_{i+1} - \epsilon$, let i_t be an isotopy of ζ in P chosen to avoid the family of curves q_t. This is possible, since q_t changes only by isotopy and the addition or deletion of inessential curves (corresponding to center tangencies).

Near a saddle tangency, i. e. in the interval $[t_i - \epsilon, t_i + \epsilon], i = 1, ..., n$, it may be impossible to define i_t so as to avoid q_t, since the core of the band associated to the saddle tangency at t_i may essentially intersect $i_{t_i-\epsilon}(\zeta)$. Since the saddle tangency only involves curves which are inessential in P, it is possible to isotope all of q_{t_i} lying in a component of $P - q_0$ into a subdisk of the component. Thus there is an isotopy of $i_{t_i-\epsilon}(\zeta)$ in $P - q_0$ to a new position which is well away from q_{t_i}. Define $i_{t_i+\epsilon}(\zeta)$ to be this imbedding of ζ. Unfortunately, the isotopy in $P - q_0$ from $i_{t_i-\epsilon}(\zeta)$ to $i_{t_i+\epsilon}(\zeta)$ may involve pushing arcs of ζ across inessential curves in $q_{t_i-\epsilon}$, and so push ζ across Q. Perform the isotopy anyway, with the following modification: When an arc of ζ is supposed to be pushed across a disk in P bounded by an inessential curve c of $q_{t_i-\epsilon}$, push the arc of ζ instead across a disk parallel to the disk which c bounds in Q. Then ζ never is pushed across Q and ends up in a position on $f_{t_i-\epsilon}(P)$ which is distant from the saddle tangency. After this push, we can pass through the saddle tangency without forcing any of ζ across Q. This completes the definition of F' on ζ.

Extend $F' : \zeta \times [0, 1-\epsilon] \to M$ to all of P using a neighborhood of ζ in P. Since ζ never crosses Q during the isotopy, ${f'}_t^{-1}(Q)$ remains just q_0. Furthermore at the end of the isotopy, as we have constructed it, $f'_{1-\epsilon}$ carries ζ to a spine of $P_{1-\epsilon}$ which is isotopic in $P_{1-\epsilon}$ to $f_{1-\epsilon}(\zeta)$ rel q_0. Follow F' with this isotopy, ambiently extended across $P_{1-\epsilon}$ rel q_0. The isotopy will not push any of $P_{1-\epsilon}$ across Q, since $P_{1-\epsilon}$ intersects Q only in q_0. Afterwards, $f'_1 = f_1$ on ζ; ambiently extend this equality to a neighborhood η of ζ. Now $f'_1(P - \eta)$ is disjoint from Q, so the collection of disks $f'_1(P - \eta)$ lies in $M - Q$. Since $M - Q$ is aspherical, a standard innermost disk argument can be used to isotope $f'_1(P - \eta)$ to $f_1(P - \eta)$ in $M - Q$ by an isotopy fixing ζ. □

Theorem 7.3. *Suppose the isotopy $F : P \times I \to M$ is compression-free with respect to Q and each component of C_0 and C_1 is essential.*

Then there is a simple isotopy $F' : P \times I \to M$ from f_0 to f_1 so that, for any regular value of t, the collection of essential curves in ${f'}_t(P) \cap Q$ consists precisely of the essential curves of C_t.

Proof: **Case 1:** No critical point of the isotopy involves essential curves of intersection.

This case is essentially Lemma 7.2

Case 2: There is just one critical point which involves essential curves of intersection.

Let t_0 be the critical level. With no loss of generality we can assume that two curves c', c'' in $C_{t_0 - \epsilon}$ are fused to create c in $C_{t_0+\epsilon}$, and that at most one of the three curves is inessential (since otherwise all three would be). In particular, we can assume that c' is essential, so that the singular component of C_{t_0} is not contained in the interior of any disk in P_{t_0} or Q bounded by an inessential curve of intersection. Now a standard innermost disk argument (as in 6.5) provides a compression-free isotopy which eliminates all inessential circles of intersection in C_{t_0}. We can incorporate this isotopy just before t_0 and its inverse just after and thereby assume that C_{t_0} has no inessential circles. A bicollar of P_{t_0} in M then defines an isotopy $G : P \times [t_0 - \epsilon, t_0 + \epsilon] \to M$. During the isotopy G at most one curve (either c'' or c) is inessential in P and Q. If there is such an inessential curve, extend the isotopy by isotoping the disk the curve bounds in P to the disk it bounds in Q and then incorporate a center tangency which removes the component. The construction shows that the pair of imbeddings f_0 and $g_{t_0-\epsilon}$ are connected by an isotopy satisfying the hypotheses of Lemma 7.2, as are the pair $g_{t_0+\epsilon}$ and f_1. Then the conclusion of Lemma 7.2 provides isotopies from f_0 to $g_{t_0-\epsilon}$ and from $g_{t_0+\epsilon}$ to f_1 during which no inessential curve of intersection is introduced. Combining the three isotopies, we get an isotopy from f_0 to f_1 during which at most one inessential curve of intersection (either c'' or c) is introduced and it is then immediately eliminated.

Case 3: The general case.

The proof of the general case is by induction on the number of critical values whose critical point involves essential curves. If there is only one, we are done by the previous case. Otherwise, let t_0 be a regular value between the first two such critical values. A standard innermost disk argument (as in 6.5 gives an isotopy G from f_{t_0} to a map $f : P \to M$ so that during the isotopy no essential curves of intersection with Q are affected, but all inessential curves of intersection are removed. In particular, all curves in $f(P) \cap Q$ are essential in both surfaces.

Now alter the isotopy F by inserting the isotopy G followed by its reverse $\overline{G}$ near the level t_0. After this alteration F is the product of an isotopy from f_0 to f having one critical value involving essential curves and an isotopy from f to f_1 having one fewer such critical point than F did. Apply the inductive assumption to each of these isotopies independently, and then adjoin the result. □

8. Spinal intersections with meridian systems

Theorem 8.1. *Suppose $X \cup_Q Y$ is strongly irreducible and there is a generic ambient isotopy $F : P \times I \to M$ so that*

(1) *there is a meridian disk of X which is disjoint from P_0 and*
(2) *there is a meridian disk of Y which is disjoint from P_1.*

Then there are complete collections of meridian disks Δ_A, Δ_B for A and B respectively and a generic extension of F to the 2-complex $K = P \cup \Delta_A \cup \Delta_B$ so that for each t in some sub-interval of I, the 1-complex $\kappa_t = f_t(K) \cap Q$ in Q contains an entire spine of Q.

Proof: Let $\mu_i, i = 0, 1$ be the given meridian disks for X and Y which are disjoint respectively from $P_i, i = 0, 1$. Choose complete collections of meridian disks Δ_A, Δ_B and the extension of F so that

a) no component of $A - \Delta_A$ (resp. $B - \Delta_B$) is adjacent to both sides of the same disk in Δ_A (resp. Δ_B)
b) for $i = 0, 1, f_i(K)$ is also disjoint from μ_i .

This can be done by first choosing any complete collections of meridian disks Δ_A and Δ_B which satisfy a) and any extension of F to K, then modifying them near $t = i$ via the reverse of an isotopy rel P_i which shrinks μ_i very small.

Suppose, for a generic t, some component R of the complement of κ_t in Q is essential in Q. Then R lies in a component W of the 3-manifold $M - P_t(K)$. By definition of K, W is a ball. No spine of Q could lie in W, for otherwise M would be a 3-sphere. Hence some component of ∂R is essential in Q. It also bounds a disk in $\partial W \simeq S^2$. A component of $\partial W \cap Q$ which is innermost in ∂W among all components of $\partial W \cap Q$ which are essential in Q then bounds a disk lying entirely in X or Y. We've thus shown that if some component of the complement of κ_t in Q is essential in Q then some boundary of some such component bounds a disk in X or Y. Call such a disk a κ_t *compressing disk* in X or Y. There can't simultaneously be a κ_t compressing disk in both X and Y, since Q is strongly irreducible.

At the beginning of the isotopy there's no κ_0 compressing disk in Y, since it and μ_0 would give a weak reduction of $X \cup_Q Y$. Similarly, at the end of the isotopy there's no κ_1 compressing disk in X. Hence either there is a generic value t_0 for which no κ_{t_0} compressing disk exists (the desired conclusion), or there is a critical value t_0 at which there is a switch, say, from a $\kappa_{t_0-\epsilon}$-compressing disk in X to a $\kappa_{t_0+\epsilon}$-compressing disk in Y.

There are two possible types of non-generic behavior at t_0. There may be a point of tangency between Q and a point of $f_t(K)$ away from $\partial\Delta_A \cup \partial\Delta_B$ or there may be a point of tangency between Q and the attaching circles $\partial\Delta_A \cup \partial\Delta_B$. In the latter case, we can assume by general position that P and Q are not also tangent at that point, so the effect of the tangency of the attaching circle on κ_{t_0} is merely to add or remove a small inessential arc near the point of tangency. In particular, this sort of singularity can't create or destroy a $\kappa_{t_0-\epsilon}$ compressing disk. This is also true when Q has a center tangency with a point in K away from $\partial\Delta_A \cup \partial\Delta_B$.

The remaining possibility is that Q has a saddle tangency at a point in K away from $\partial\Delta_A \cup \partial\Delta_B$. Such a saddle tangency can indeed simultaneously destroy a $\kappa_{t_0-\epsilon}$-compressing disk in X and create a $\kappa_{t_0+\epsilon}$-compressing disk in Y. But the curves created and destroyed by a single saddle tangency in Q can be isotoped in Q to be disjoint, for our requirement a) of Δ_A and Δ_B guarantees that the curves lie in distinct components of $M - K$. So there would persist a meridian disk of X whose boundary is disjoint from the meridian disk of Y created at t_0, and this would contradict the strong-irreducibility of $X \cup_Q Y$. □

9. Destabilizing annular 1-handles

Much of this section was inspired by more delicate arguments used in [K] to understand families of annuli and tori in Heegaard splittings.

Suppose $\mathcal{A}$ is a finite set of ∂-compressible annuli embedded in a handlebody H of genus p, γ is a set of spanning arcs for $\mathcal{A}$ and τ is a regular neighborhood of γ in H. We view τ as a collection of 1-handles added to $P = \partial H$, each corresponding to an annulus in $\mathcal{A}$. Let H' denote the closure of $H - \tau$ and P' denote $\partial H'$. Since a spanning arc of a ∂-compressible annulus in H is parallel to an arc on ∂H, it's apparent that H' is a handlebody of genus $p + |\mathcal{A}|$.

Proposition 9.1. *Suppose $\partial\mathcal{A}$ is essential in ∂H and Δ is a complete collection of meridian disks for H which intersect $\mathcal{A}$ only in spanning arcs. Then there is an ordering $A_1, A_2, \cdots A_n$ of $\mathcal{A}$ and, for each of all but at most $2p-2$ of the A_i there is a properly imbedded disk E_i in H' so that the E_i are all disjoint and have the following properties:*

1. *∂E_i is disjoint from the 1-handles corresponding to the annuli $A_k, k > i$*

(2) *∂E_i runs exactly once across the 1-handle corresponding to A_i.*

(3) *∂E_i is disjoint from any component of $P - \partial\mathcal{A}$ which is not an annulus.*

Moreover, we may find E_i among the components of $\Delta - \mathcal{A}$.

Proof: Some component A_1 is ∂-compressible in H via a ∂-compressing disk whose interior is disjoint from $\mathcal{A}$. A useful way to see this is to use as the ∂-compressing disk the component cut off from Δ by an outermost arc of $\Delta \cap \mathcal{A}$. Inductively define A_i as a component of $\mathcal{A} - \cup\{A_j, j < i\}$ which ∂-compresses in the complement of $\mathcal{A} - \cup\{A_j, j < i\}$. Choose a ∂-compressing disk D_i at each stage so that ∂D_i is disjoint from the disks in H which are the remains of the $A_j, j < i$ after ∂-compression. This ordering of $\mathcal{A}$ and choice of disks D_i is called a *compressing system* for $\mathcal{A}$.

A compressing system for $\mathcal{A}$ defines a directed graph Γ as follows: Each vertex of Γ corresponds to a component of $H - \mathcal{A}$ and each edge of Γ corresponds to an annulus of $\mathcal{A}$. The ends of an edge in Γ corresponding to an annulus A_i are adjacent to the vertex or vertices in Γ which correspond to the component or components of $H - \mathcal{A}$ adjacent to A_i. In other words, if the components V and V' of $H - \mathcal{A}$ lie on either side of A_i, and vertices v and v' are the corresponding vertices of Γ, then the edge of Γ corresponding to A_i runs between v and v'. Direct the edge toward the vertex which represents the component of $H - \mathcal{A}$ on which D_i abuts A_i.

Notice the effect of the ∂-compressions on the topology of the components of $H - \mathcal{A}$: ∂-compressing A_1 changes the topology only of the component V in which D_1 lies. In V, D_1 is non-separating, so V is changed by a single non-separating ∂-reduction. This increases its Euler characteristic by one. More generally, let $\mathcal{A}_i$ denote the remains of $\mathcal{A}$ after the ∂-compressions to $A_1, \cdots, A_i$, which converts each of these annuli into a disk. Then the ∂-compressing disk D_{i+1} for A_{i+1} lies in a single component of $H - \mathcal{A}_i$ and the effect of the ∂-compression is then to alter precisely this component by a single non-separating ∂-reduction, raising its Euler characteristic by one.

For any vertex $v \in \Gamma$, let $n_-(v)$ denote the number of edges in Γ which are oriented into v. Following the previous discussion, we see that if V is the corresponding component in $H - \mathcal{A}$ then $\chi(V) \leq 1 - n_-(v)$, since after $n_-(v)$ non-separating ∂-reductions the Euler characteristic is at most $1 = \chi(B^3)$. In particular, if $\chi(V) = 0$ then $n_-(v)$ is at most one.

If a vertex v is the base of a loop in Γ, then the corresponding component V of $H - \mathcal{A}$ has $\chi(V) \leq -1$. Indeed, the annulus corresponding to a loop must be non-separating, so after a ∂-compression of the annulus (i. e. a ∂-reduction of V), V would still contain a non-separating disk. This means V admits two independent ∂-reductions, so $\chi(V) \leq -1$.

Consider a component V with corresponding vertex v for which $\chi(V) = 0$ and $n_-(v) = 1$, so V can be ∂-reduced. Then V must be a solid torus. If A_i is the annulus corresponding to the edge pointing into v, then the disk D_i of the compressing system becomes a meridian disk for V after V is expanded by the earlier boundary compressions of $A_j, j < i$. The disk D_i intersects A_i in a single spanning arc and intersects no $A_k, k > i$. In particular, A_i is a longitudinal annulus in V and, since $\partial\mathcal{A}$ is essential in ∂H, any $A_j, j < i$ incident to the solid torus V must also have been a longitudinal annulus. It follows that the subdisk of $D_i - \mathcal{A}$ which abuts A_i has all the properties we seek for E_i. It remains to do a count of how many annuli in $\mathcal{A}$ don't satisfy these conditions, i. e. the edge corresponding to the annulus points into a vertex w in Γ with $n_-(w) > 1$ or for which the corresponding component W in $H - \mathcal{A}$ has $\chi(W) < 0$.

An easy way to do this count is to collapse any edge in Γ with the property that it is the unique edge oriented into some vertex, and for which the component of $H - \mathcal{A}$ corresponding to that vertex is a solid torus. (We've already seen that such an edge can't be a loop.) The collapse removes the edge and one vertex, which we take to be the vertex corresponding to the solid torus, i. e. the head of the arrow. The number of edges pointing into the remaining vertex (i. e. the tail of the arrow) is unchanged. (The graph is equivalent to the graph we would get if we removed from $\mathcal{A}$ the annulus corresponding to the edge.) After collapsing all such edges we are left with a graph Γ' so that each vertex is either a source or corresponds in $H - \mathcal{A}$ to a component with negative Euler characteristic. In particular, if Υ is the set of vertices in Γ' which are not sources, then $|\Upsilon| \leq -\chi(H) = p - 1$. Furthermore, for a vertex $v \in \Upsilon$ corresponding via Γ to a component V of $H - \mathcal{A}$ we still have $1 - n_-(v) \geq \chi(V)$. Summing over all vertices in Υ:

$$\Sigma_{v\in\Upsilon}(1 - n_-(v)) \geq \chi(H) = 1 - p.$$

Now clearly $\Sigma_{v\in\Upsilon}(n_-(v))$ is just the number e of edges in Γ', so we have $|\Upsilon| - e \geq 1 - p$ or $e \leq |\Upsilon| + p - 1 \leq 2p - 2$. $\square$

10. Spinal intersections of splittings surfaces

Definition 10.1. *Suppose Δ_A and Δ_B are (not necessarily complete) collections of meridian disks for A and B respectively and K is the 2-complex $P \cup \Delta_A \cup \Delta_B$. Then K has* pre-spinal intersection *with Q if K and Q are in general position and the 1-complex $\kappa = K \cap Q$ contains an entire spine of Q. K has* spinal intersection *with Q if, in addition, for each disk $D \in (\Delta_A \cup \Delta_B)$, $D \cap Q$ is a single arc. We say that P is* pre-spinal *(resp.* spinal*) with respect to Q if there is some collection of meridian disks whose union with P has pre-spinal (resp. spinal) intersection with Q.*

Theorem 10.2. *Let P and Q be Heegaard splitting surfaces in M of genus $p \geq 2$ and $q \geq 2$ respectively. Suppose the pair (P, Q) is compression-free and only a single component of $P \cap Q$ is inessential. If P is pre-spinal with respect to Q then after at most $7q + 4p - 9$ stabilizations of P, P is spinal with respect to Q.*

Proof: Since κ contains an entire spine of Q it follows that no circle of intersection of $\Delta_A \cup \Delta_B$ with Q can be essential in Q. An innermost disk argument in Q then allows us to remove such circles of intersection by an isotopy of $\Delta_A \cup \Delta_B$. Afterwards, Δ_A and Δ_B intersect Q only in arcs and κ becomes the union of $P \cap Q$ and these arcs; let γ be a minimal collection of arcs in $Q \cap (\Delta_A \cup \Delta_B)$ so that $(P \cap Q) \cup \gamma$ is a spine of Q. Denote this spine ζ.

Alter P by a 1-surgery along each arc of γ. That is, remove from A a neighborhood of each arc of γ that lies in Δ_A and attach the neighborhood to B. This creates a new 1-handle in B whose 2-disk cocore intersects Q in a single arc. Similarly, remove from B a neighborhood of each arc of γ lying in Δ_B and attach the neighborhood to A. Every arc of γ is parallel in Δ_A or Δ_B to a subarc of P, so this operation stabilizes P. Denote the resulting stabilized Heegaard splitting surface P', and the set of co-cores of the new 1-handles Δ'. Let K' denote the 2-complex $P' \cup \Delta'$ and κ' denote $K' \cap Q$. It is easy to see how κ' is obtained from κ: When P is stabilized, each arc of γ in Q is replaced by a band. That is, a figure I neighborhood of the arc in Q becomes a figure][. When the cocore of the new 1-handles are added to P' to make K', the figure][becomes a figure H. The combination, which replaces I with H, has no effect on the topology of the complement of ζ; in particular, the 1-complex remains a spine, which we call ζ'. So after this stabilization, P' is spinal with respect to Q.

One simplification of this picture is immediate: If there are three curves in $P \cap Q$ which are parallel in Q then each of the two annular components of $Q - P$ which they cut out contains a spanning arc from γ. When P is stabilized along these arcs, as above, the effect is to replace these three components of $P \cap Q$ by a single isotopic curve of intersection. So ζ' is a spine even if we don't include the cocores dual to these 1-handles in Δ'. Generalizing from this observation, consider the arcs of γ lying in a single collection of adjacent annuli in $Q - P$. From the new stabilizing 1-handles which correspond to these arcs, we need to include at most one cocore 2-disk in Δ' to ensure that ζ' remains a spine of Q. In fact, if the number of annuli in the collection is odd, we don't have to add any cocore 2-disks. (These observations remain true even if the inessential component of $P \cap Q$ lies in one of the annuli. That is, even if we allow into the collection of parallel curves those which are parallel ignoring the inessential component of intersection.) To summarize, P' is obtained from P by stabilizing along $|\gamma|$ arcs. For some subcollection $\hbar$ of the stabilizing 1-handles, one from each arc lying in a non-annular component of $Q - P$, and at most one from each collection of annuli which are adjacent in Q, include the cocore in $\Delta' \subset K'$. Then $\kappa' = Q \cap K'$ is a spine of Q, so P' is spinal with respect to Q.

There is no apparent bound to the genus of P', because there is no bound on the number of annular 1-handles, that is 1-handles added to P along arcs of γ spanning annuli of $Q - P$. The number of 1-handles in $\hbar$ is bounded, however, since at most one needs to be chosen from any contiguous set of annuli. We will show that $|\hbar| \leq 7q - 7$.

First notice that since there is at most one disk component of $Q - P$, there are at most $2q - 1$ components of $Q - P$ which have negative Euler characteristic. In each component of $Q - P$ there will be at most one disk of $Q - \zeta$, for otherwise we could reduce ζ by removing an edge between two such disks. Hence the complex consisting of the circles $P \cap Q$ together with all arcs of γ lying in components of $Q - P$ with negative Euler characteristic has itself Euler characteristic no less than $\chi(Q) - (2q-1) - 1 = 2 - 4q$. Hence it includes at most $4q - 2$ arcs of γ. Similarly, in Q the total number of families of parallel curves of $P \cap Q$ is at most $3q - 3$. (Since only one curve of $P \cap Q$ is inessential, only essential curves of intersection appear in parallel families, so for this last calculation we can ignore the inessential intersection curve.)

In the absence of a bound on the number of annuli in $Q - P$, and hence to the number of annular 1-handles which stabilize P, we instead will use a collection of destabilizing disks, found via 9.1, to cancel all but at most $4p-4$ annular 1-handles not in $\hbar$. This bounds the total number of stabilizations needed to make P spinal by $|\hbar| + 4p - 4 \leq 7q + 4p - 9$.

As a preliminary move, separately isotope $\partial\Delta_A$ and $\partial\Delta_B$ near annular components of $P-Q$ so that $\partial\Delta_A$ and $\partial\Delta_B$ don't intersect in any such annulus component. Use a collar of $P \cap Q$ in Q to taper this isotopy so it's only visible effect on ζ in Q is to alter the ends of γ near some components of $P \cap Q$ by a fractional Dehn twist. In particular ζ remains a spine and so nothing is lost. What is gained is that now, because of property 9.1.3, the destabilizing disks defined in Proposition 9.1 for the annular components $\mathcal{Q}_A$ of $Q \cap A$ will be disjoint from the destabilizing disks defined for the annular components $\mathcal{Q}_B$ of $Q \cap B$, since we may take the destabilizing disks to lie in Δ_A and Δ_B respectively. So we may use the disks to destabilize all but at most $4p-4$ of the annular 1-handles. Of course we don't destabilize an annulus whose 1-handle is in $\hbar$ nor is it immediately apparent we can destabilize across those solid torus components of $A - \mathcal{Q}_A$ or $B - \mathcal{Q}_B$ which may contain other components of $Q - P$.

In fact, the number of solid torus components of $A - \mathcal{Q}_A$ or $B - \mathcal{Q}_B$ which can contain a non-annular component of $Q - P$ is shown in [MR] to be at most one, so this last problem is minor. But more directly, it's easy to argue as in the proof of 9.1, that even in such torus components there is still a destabilization disk for the annular 1-handle. The only difference is that its boundary may run over the 1-handles in $\hbar$ which come from arcs of γ in the non-annular components of $Q - P$. □

11. Spinal intersections and stabilization

Definition 11.1. *An oriented splitting surface P in M has* spinal intersection *with an oriented splitting surface Q in M if*

1. *P and Q are in general position except at a finite number of saddle tangencies*
2. *at the points where P and Q are tangent the orientations of P and Q in M coincide*
3. *the resulting 1-complex $\kappa = P \cap Q$ contains a spine of Q.*

Lemma 11.2. *If P is spinal with respect to Q then P may be isotoped so that it has spinal intersection with Q.*

Proof: Let Δ be a set of meridian disks in A and B so that the complex $K = P \cup \Delta$ has spinal intersection with Q. Consider a disk $D \in \Delta$ and the single arc $D \cap Q$ lying in D. The arc cuts D into two disks, either of which can be used to isotope Q near D so that $D \cap Q$ becomes a single point in ∂D at which P and Q are tangent. The effect on $\kappa = K \cap Q$ is to replace a neighborhood in Q of the arc $D \cap Q$, which looks like a figure I, with a saddle tangency of P and Q, which looks like a figure X. Choose one of the two disks in $D - Q$ so that P and Q have the same orientation at the tangency point.

After this isotopy is done at every disk in Δ, then $K \cap Q = P \cap Q$ is still a spine of Q, and at every tangency point the two surfaces have the same orientation. $\square$

Lemma 11.3. *If P has spinal intersection with Q then there are neighborhoods $\eta_P(\kappa)$ and $\eta_Q(\kappa)$ of κ in P and Q respectively so that, after a small ambient isotopy of M rel κ, $\eta_P(\kappa) = \eta_Q(\kappa)$.*

Proof: Since P and Q have a saddle tangency near each vertex in κ, a small isotopy carries a disk neighborhood in P of each vertex to a disk neighborhood in Q. Nowhere on κ will there be a point at which the normal vector to P is directly opposed to the normal vector to Q, for this would give rise to a saddle tangency and by hypothesis the orientations coincide at all such tangencies. Hence along an edge in κ, the winding number of the normal vector to P with respect to the normal vector to Q must be trivial. Then an isotopy near the edge will rotate a neighborhood of the edge in P tangent to a neighborhood of the edge in Q. After this is done on all of κ, a neighborhood $\eta_P(\kappa)$ can be isotoped down to $\eta_Q(\kappa)$ via orthogonal projection to $\eta_Q(\kappa)$. $\square$

Proposition 11.4. *Suppose P and Q are oriented splitting surfaces of genus p and q respectively. If P has spinal intersection with Q then P and Q have a common stabilization of genus $p + q$.*

Proof: Following the previous lemma, isotope a neighborhood $\eta_P(\kappa)$ of $\kappa = P \cap Q$ in P so that it coincides with a neighborhood $\eta_Q(\kappa)$ of κ in Q. Let Ξ be a spine of the handlebody X with a single vertex and q edges. Since Ξ can be isotoped into Q, it can be isotoped into $\eta_Q(\kappa) = \eta_P(\kappa) \subset P$. After Ξ has been moved into P, push a small interior arc of each of the q edges of Ξ into B and off of P. The union H of A and a relative regular neighborhood of these arcs in B is a handlebody obtained by adding q trivial handles to A, so ∂H is a q-fold stabilization of P.

Now imagine pulling more of each arc of Ξ into B until all of Ξ except the vertex has been pulled into B. This defines an isotopy of H after which ∂H is apparently also a Heegaard splitting of the handlebody Y obtained by removing a neighborhood of Ξ from M. Any Heegaard splitting of a handlebody is just a stabilization of the boundary [ST1, 2.7], so ∂H is then also a stabilization of $Q = \partial Y$. $\square$

Theorem 11.5. *Suppose $A \cup_P B$ and $X \cup_Q Y$ are strongly irreducible Heegaard splittings of M and are of genus $p \leq q$ repectively . Then there is a genus $8q+5p-9$ Heegaard splitting of M which stabilizes both $A \cup_P B$ and $X \cup_Q Y$.*

Proof: By 5.11, 6.3, and 6.4 we may as well assume that $p \geq 2$. According to 6.5 and 7.3 there is an isotopy of P in M which is simple with respect to Q and satisfies the hypotheses of 8.1 (but possibly symmetrically, with (A, B, P) instead of (X, Y, Q)). Then the conclusion of 8.1 gives a location for P so that, according to 10.2, after at most $7q+4p-9$ stabilizations of P (so that P has genus $7q+5p-9$), P is spinal with respect to Q. Then according to 11.2 there is a small isotopy of P which gives it spinal intersection with Q. Finally, 11.4 then says that the new P and the old Q have a common stabilization of genus $8q + 5p - 9$. $\square$

Corollary 11.6. *Suppose $X \cup_Q Y$ and $A \cup_P B$ are Heegaard splittings of the same irreducible non-Haken 3-manifold M and are of genus $p \leq q$ repectively . Then there is a genus $8q + 5p - 9$ Heegaard splitting of M which stabilizes both $A \cup_P B$ and $X \cup_Q Y$.*

References

[Bo] F. Bonahon, Difféotopies des espaces lenticulaires, *Topology*, **22** (1983), 305-314.

[BoO] F. Bonahon and J. P. Otal, Scindements de Heegaard des espaces lenticulaires, *Ann. scient. Éc. Norm. Sup.*, **16** (1983), 451-466.

[BO] M. Boileau and J. P. Otal, Groupes des difféotopies de certaines varietés de Seifert, *C. R. Acad. Sci. Paris*, **303-I** (1986), 19-22.

[CG] A. Casson and C. McA. Gordon, Reducing Heegaard splittings, *Topology and its applications*, **27** (1987), 275-283.

[C] J. Cerf, *Sur les difféomorphismes de la sphère de dimension trois ($\Gamma_4 = 0$)*, Lecture Notes in Math, **53** Springer-Verlag, Berlin and New York, 1968.

[Ha] W. Haken, Some results on surfaces in 3-manifolds, *Studies in Modern Topology*, Math. Assoc. Am., Prentice Hall, 1968, 34-98.

[H] J. Hempel, *3-manifolds*, Annals of Math Studies **86** (1976), Princeton University press.

[J] W. Jaco, *Lectures on Three-manifold Topology*, Regional Conference Series in Mathematics **43** (1981), Amer. Math. Soc.

[Jo] K. Johannson, Computations in 3-manifolds, preprint, U. of TN., 1991.

[Lu] N. Lu, The stable equivalence of Heegaard decompositions of 3-manifolds, *Ph. D. Thesis*, Rice University, 1990.

[K] H.-Z. Kowng, Straightening Tori in Heegaard splittings, *Ph. D. Thesis*, U. C. Santa Barbara, 1994.
[MR] Y. Moriah and H. Rubinstein, Heegaard structures of negatively curved 3-manifolds, to appear.
[ST1] M. Scharlemann and A. Thompson, Heegaard splittings of (surface) x I are standard, *Math. Ann.*, **295** (1993), 549-564.
[ST2] M. Scharlemann and A. Thompson, Thin position and Heegaard splittings of the 3-sphere, to appear in *Jour. Diff. Geom.*.
[ST3] M. Scharlemann and A. Thompson, Thin position for 3-manifolds, to appear in Proceedings of the 1992 Haifa Topology Conference.
[Ta] I. A. Taimanov, On the stable equivalence of Heegaard splittings, preprint, Novosibirsk.
[W] F. Waldhausen, Heegaard-Zerlegungen der 3-Sphäre, *Topology*, **7** (1968), 195-203.

Department of Mathematics
University of Melbourne
Parkville, Vic 3052
Australia

Department of Mathematics
University of California
Santa Barbara, CA 93106
USA