Mathematical Research Letters **1**, 613–630 (1994)

THIN POSITION AND THE RECOGNITION PROBLEM FOR S^3

Abigail Thompson

Abstract. We describe a modified version of Rubinstein's algorithm to detect the 3-sphere and use thin position combined with standard 3-manifold techniques to prove that the algorithm works.

1. Introduction

In the spring of 1992, Hyam Rubinstein gave a series of lectures at the Technion, Haifa, describing an algorithm to determine whether or not a triangulated 3-manifold is the 3-sphere. Rubinstein's proof that his algorithm works uses the language of PL minimal surface theory. The substance of this paper is a different approach to the proof that the algorithm works, using techniques from knot theory. In addition, we simplify the original algorithm by using a modified definition of *almost normal 2-sphere*.

There are two interesting points. One is simply that an algorithm exists, requiring for its proof only fairly standard normal surface theory and knot theory. The second is that PL-minimal surface theory and knot theory are so closely connected; a given tool or concept in one area seems to have its counterpart in the other. The correspondence between the two is likely to be a fruitful area for further study.

Theorem 1. [8] *There exists an algorithm to determine whether or not a compact 3-manifold is* S^3.

We give some definitions, followed by an outline of the proof.

Definitions. Let M be a closed orientable 3-manifold with fixed triangulation T. An arc in a 2-cell of T is *normal* if it connects distinct edges of the 2-cell. Let c be a simple closed curve on the boundary of a tetrahedron H in the triangulation, such that c intersects each face of H in normal arcs. c is a *normal curve*. The *length of* c is the number of times c crosses the edges of H. A closed orientable surface F imbedded in M is a *normal surface* if F intersects each tetrahedron in T in a collection

Received July 14, 1994.

The author was partially supported by the National Science Foundation and the Alfred P. Sloan Foundation.

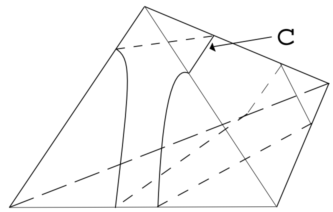

FIGURE 1. A normal curve of length eight.

of disks all of whose boundary curves are normal curves of length three or four, i.e., in normal triangles and quadrilaterals. F is *almost normal* if F intersects each tetrahedron in T in a collection of normal triangles and quadrilaterals and, in one of the tetrahedra, precisely one disk whose boundary is a normal curve c of length eight, and possibly some normal triangles.

Note. These definitions differ slightly from Rubinstein's; in particular, he considers a more general type of surface to be *almost normal*.

Outline of the proof of Theorem 1. Let Σ be a maximal collection of disjoint non-parallel (in (M, T)) normal 2-spheres in M. Notice that Σ is not empty, since the boundary of a small neighborhood of each vertex is a normal 2-sphere in M. Σ cuts M into three types of components M_0:

(1) M_0 is a 3-ball containing a single vertex.
(2) M_0 has more than one boundary component.
(3) M_0 has exactly one boundary component and is not of type 1.

Using simplicial homology, we can algorithmically check that $H_1(M; Z_2)$ is trivial to ensure that M contains no closed non-orientable or non-separating surfaces. In particular we assume from now on that all 2-spheres are separating.

Notice that components of type 2 and 3 contain no vertices of T. We prove the following lemmas:

Lemma 2. *A component of type 2 is a punctured 3-ball.*

Lemma 4. *A component of type 3 is a 3-ball if and only if it contains an almost normal 2-sphere.*

Assuming lemmas 2 and 4, the algorithm is as follows:

(1) Search for Σ, a maximal collection of disjoint, non-parallel normal 2-spheres.
(2) Consider each component of $M - \Sigma$. In each component with a single 2-sphere boundary, search for an almost normal 2-sphere.

(3) The manifold is S^3 if and only if each component considered in step 2 contains an almost normal 2-sphere.

The search for Σ is accomplished via an algorithm due to Haken [3], [4], [6], [7]. Haken's work does not describe how to find almost normal 2-spheres, but a modification of the algorithm suffices.

Outline of the paper: Section 2 will be devoted to the proof of lemma 2. Section 3 consists of an unknotting lemma, which is used in section 4 to prove lemma 4. Section 5 will discuss the necessary modification of Haken's algorithm.

We first give some additional definitions, state some facts about normal curves on the boundary of a tetrahedron and prove two useful claims.

Definitions. A separating surface F properly imbedded in a 3-manifold is *strongly compressible* if there exist disjoint compressing disks for F lying on opposite sides of F. If F is not strongly compressible it is *weakly incompressible.* Note that a separating surface incompressible to one or both sides is weakly incompressible.

Let F be a closed orientable surface imbedded in a 3-manifold M^3. Let X be a 1-complex properly imbedded in M, possibly puncturing F. We define *incompressible in the complement of X*, *strongly compressible in the complement of X* and *weakly incompressible in the complement of X* in the obvious manner.

A *parallel 2-sphere* is a 2-sphere punctured twice by a single edge of T, boundary parallel to that edge.

Facts about normal curves on the boundary of a tetrahedron H:

(1) *A normal curve of odd length has length 3.*
(2) *There are no normal curves of length 6.*
(3) *Any normal curve of length greater than 8 crosses some edge of H at least 3 times.*

Proof. First flatten the 1-skeleton of the tetrahedron out in the plane. The proofs then proceed by first drawing in the vertices of the 1-skeleton, then drawing in the normal curve c, which, w.l.o.g., can be drawn as a round circle, and then adding in the edges of the 1-skeleton, which now may look quite complicated (the more painful details are left to the reader):

(1) Any normal curve c of odd length must separate one vertex of H from the other three; such a curve can only have length three.

(2) Any normal curve c of even length separates the vertices of H in pairs. Suppose it separates v and v' from w and w'. Examine the disk $D \subset \partial$H bounded by c, containing v and v'. If an edge of H lies entirely inside D, then c has length 4. Suppose no edge of H lies entirely inside

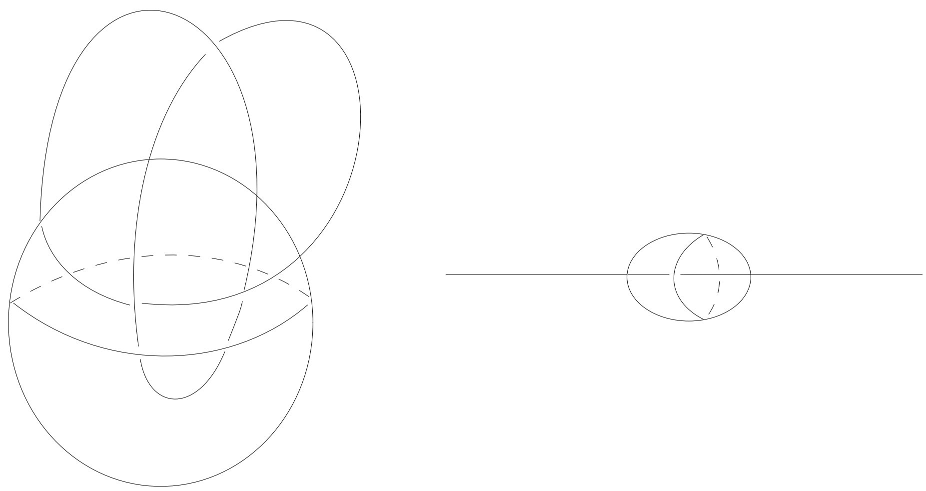

FIGURE 2. A 2-sphere in S^3, weakly incompressible in the complement of the trefoil, and a parallel 2-sphere.

D; by considering the complementary disk $D' = \partial H - D$, one can see that the length of c is at least 8.

(3) Suppose c is a normal curve of length greater than 8. Note that every edge of H connecting a vertex in D to a vertex in D' intersects c an odd number of times, and the edges connecting v to v' and w to w' each intersect c an even number of times. Counting suffices to show that some edge intersects c at least three times. □

The following claim says that normal surfaces in (M, T) are roughly the same as surfaces which are incompressible in the complement of the 1-skeleton. More precisely:

Claim 1.1. *A closed orientable normal surface F in (M, T) is incompressible in the complement of the 1-skeleton of T. If F is a closed orientable surface which is incompressible in the complement of the 1-skeleton of T, there exists an ambient isotopy rel the 1-skeleton of T from F to a normal surface $\tilde{F}$ or to a parallel 2-sphere.*

Proof. An innermost disk/outermost arc argument suffices for the first part of the claim. Now suppose F is a closed orientable surface in (M, T) which is incompressible in the complement of the 1-skeleton. Pick an innermost simple closed curve c of intersection between F and a face of one of the tetrahedra. The curve c bounds a disk D in the face with interior disjoint from F, and hence it must also bound a disk D' in F disjoint from the 1-skeleton. $D \cup D'$ together bound a 3-ball in M disjoint from the 1-skeleton.

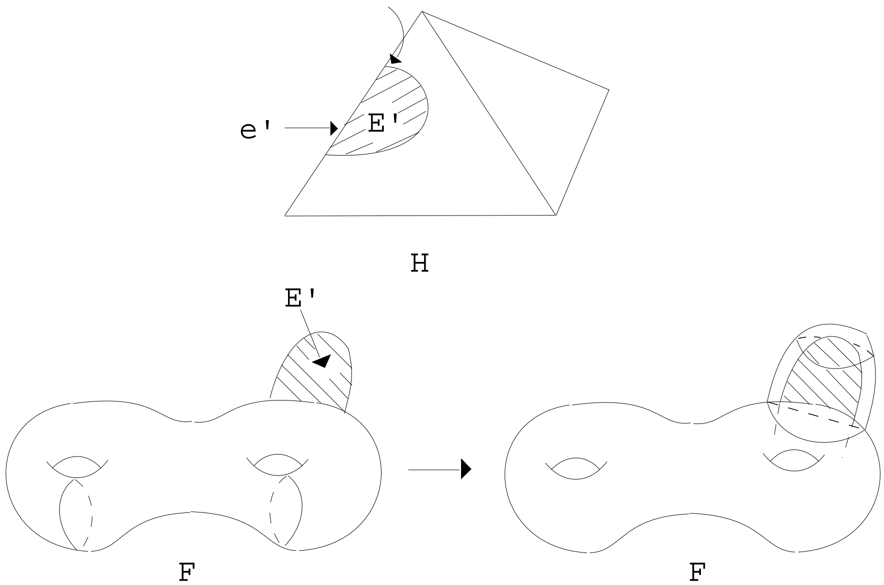

FIGURE 3. Doubling a boundary compression to yield a compressing disk.

Isotop F across the 3-ball to remove the curve of intersection c. Hence we can assume that F intersects any face of a given tetrahedron in arcs.

Let E be some face of a tetrahedron. Suppose some arc β of $F \cap E \subset E$ connects an edge of E to itself. Assume β is an outermost such arc in E. Then β cuts off a subdisk E' of E, which one can think of as a *boundary compressing disk for F with respect to the 1-skeleton*. One can construct a disk with boundary imbedded on F in the complement of the 1-skeleton by banding together two copies of E' across $\partial E' - \beta$.

We will call this operation *doubling a boundary compression* (see figure 3). Unless F is a 2-sphere parallel to the edge of E incident to α, this operation produces a compressing disk for F in the complement of T_0. Hence we can assume that every arc β of $F \cap E \subset E$ connects distinct edges of E.

Let H be any tetrahedron in T. Examine $F \cap H$. We need to show that if $F \cap H$ consists of anything but normal triangles and quadrilaterals, then F is compressible in the complement of the 1-skeleton.

Suppose $F \cap H$ contains a component $\tilde{F}$ which is not a disk. Then $F \cap H$ is compressible in H. A compressing disk for $F \cap H$ in H is a compressing disk for F in the complement of the 1-skeleton (one needs to check such a disk is not trivial when regarded as a compressing disk for F), hence F

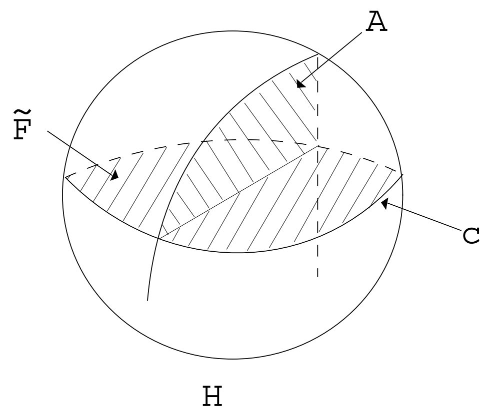

Figure 4

is compressible in the complement of the 1-skeleton. So we can assume all components of $F \cap H$ are disks.

Using facts 1, 2 and 3 about normal curves on the boundary of H, as well as the observation that a normal curve of length eight on the boundary of H must intersect some edge e of H at least twice, we see that if $F \cap H$ consists of anything but normal triangles and quadrilaterals, then there is some component $\tilde{F}$ of $F \cap H$ whose boundary is a normal curve c intersecting some edge E of H at least twice. Then there exists an imbedded disk A, with $\partial A = \alpha \cup \beta$, where $\alpha \subset E$ and $\beta \subset \tilde{F}$ (see figure 4).

If the interior of A is disjoint from F, we can either double A across α to produce a compressing disk for F in the complement of the 1-skeleton, or conclude that F is a 2-sphere parallel to α. If the interior of A intersects F, we first do an outermost arc/innermost disk argument on $A \cap F \subset A$, and then reach the same conclusion. $\square$

Let $(S, \partial S)$ be a connected orientable separating surface properly imbedded in a 3-manifold $(M, \partial M)$. S splits M into two pieces; call these M_1 and M_2. Suppose S is compressible into both M_1 and M_2. Let $\mathbf{D}_i$, $i = 1, 2$, be a minimal complete collection of compressing disks for S in M_i; each $\mathbf{D}_i$ is non-empty by hypothesis. Let W' be the 3-manifold obtained by attaching 2-handles to a small neighborhood $S \times I$ of S along the collection $\mathbf{D} = \mathbf{D}_1 \cup \mathbf{D}_2$. If any component of $\partial W'$ bounds a 3-ball in M disjoint from W', fill it in. Call the resulting object W. Then:

Claim 1.2. *At least one of the following must hold:*

(1) *$M - S$ is reducible (this includes the possibility that M has a 2-sphere boundary component disjoint from S).*
(2) *M_1 and M_2 are compression bodies, with $\partial_+ M_i = S$.*
(3) *S is strongly compressible in M.*
(4) *∂W is incompressible in M. Note that this implies that every component of ∂W is incompressible in M. Hence either*
    (a) *M contains an incompressible non-boundary parallel surface F. F is disjoint from S, and F is obtained by compressing S entirely to one side. or*
    (b) *∂W consists of a collection of components, each of which is parallel to some component of ∂M.*

Proof. This is an application of Casson-Gordon's work on Heegaard splittings (see Lemma 1.1 in [1]).

If $\partial W = \emptyset$, then (2) holds. Assume $\partial W \neq \emptyset$.

Suppose some component of ∂W is a 2-sphere. Then (1) holds.

Assume no component of ∂W is a 2-sphere. Then W has a natural Heegaard splitting with splitting surface S.

Assume ∂W is compressible in M. Since $\mathbf{D}_1$ and $\mathbf{D}_2$ are complete, ∂W is incompressible in $M - W$. Hence ∂W must be compressible in W. By [1], this implies that the Heegaard splitting of w by S is weakly reducible. Hence S is strongly compressible in M, and (3) holds.

If ∂W is incompressible in M, then each component of ∂W is incompressible in M, and either (4a) or (4b) holds. □

Notice that if X is a 1-complex properly imbedded in M, Claim 1.2 can be generalized to surfaces S which are compressible in the complement of X; if S is punctured by X, we can simply think of S as being properly imbedded in the manifold $M - n(X)$.

2. Proof of Lemma 2

Notation: Since all of the surfaces under consideration are punctured 2-spheres, we will sometimes use *incompressible* and *weakly incompressible* in place of *incompressible in the complement of the 1-skeleton* and *weakly incompressible in the complement of the 1-skeleton.*

Lemma 2. *Let M_0 be a component of $M - \Sigma$ with more than one boundary component. Then M_0 is a punctured 3-ball.*

Proof. Let T_0 be the remnants of the 1-skeleton of T in M_0. T_0 consists of imbedded arcs.

We need:

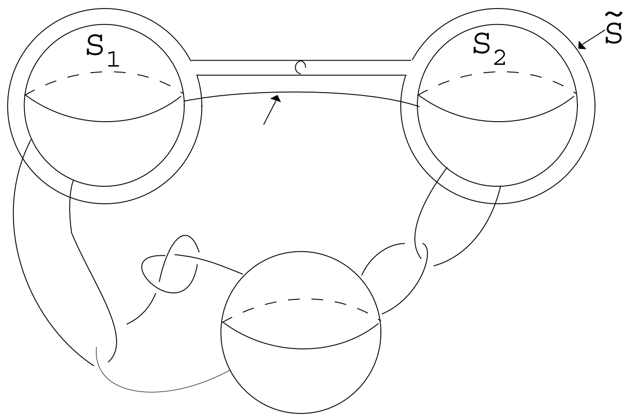

FIGURE 5

Subclaim 2.0.1. *There exists an arc α of T_0 connecting distinct boundary components S_1 and S_2 of boundary(M_0).*

Proof. Suppose not. Then $M_0 - n(T_0)$ is a connected manifold with more than one boundary component. Notice that the faces of the tetrahedra cut $M_0 - n(T_0)$ into 3-balls. But cutting a connected manifold with more than one boundary component open along disks yields a collection of components, at least one of which has more than one boundary component; in particular, it is not a collection of 3-balls. Hence such an α exists. □

Construct a new 2-sphere $\tilde{S}$ by tubing push-offs of S_1 and S_2 together via a tube parallel to α (see figure 5).

$\tilde{S}$ is compressible to both sides in the complement of T_0; a meridian disk $\mathbf{E}$ of the tube is the unique compressing disk to one side (the *inside*), while a disk enclosing the arc α is a compressing disk to the outside. Notice that $\tilde{S}$ is weakly incompressible in the complement of T_0, since the boundary of every compressing disk on the side containing α must run over the tube. Let $\mathbf{D}$ be a maximal collection of disjoint compressing disks for $\tilde{S}$ in M_0 in the complement of T_0 on the outside. Let $\tilde{M}$ be the manifold obtained by attaching 2-handles to a small neighborhood of $\tilde{S}$ along $\partial\mathbf{D} \cup \partial\mathbf{E}$. Notice that $\tilde{M}$ is homeomorphic to a punctured 3-ball. We aim to show that $\tilde{M}$ is (more-or-less) all of M_0.

Claim 2.1. *$\tilde{M}$ is homeomorphic to $M_0 - $ (3-balls), hence M_0 is a punctured 3-ball.*

Proof. The boundary of $\tilde{M}$ consists of a collection of punctured 2-spheres. The two inside $\tilde{S}$, S_1' and S_2', are parallel to S_1 and S_2. Let $\mathbf{S} = \partial\tilde{S} - (S_1' \cup$

S_2'). Apply Claim 1.2 to $\tilde{S}$. Options 1, 2 and 3 of Lemma 1.2 cannot occur, so $\partial\tilde{M}$ consists of a collection of punctured 2-spheres incompressible in the complement of T_0. By Claim 1.1, we can isotop $\tilde{M}$ in M_0 such that after the isotopy, $\partial\tilde{M}$ consists of a collection of normal 2-spheres and parallel 2-spheres. We will continue to call the components of $\partial\tilde{M}$ S_1', S_2' and $\mathbf{S} = \partial\tilde{S} - (S_1' \cup S_2')$. Each component of $\partial\tilde{M}$ is a separating sphere—otherwise we've just found a non-separating normal 2-sphere, which could not have been in our maximal collection Σ, a contradiction. S_1' and S_2' will remain parallel to S_1 and S_2; we can imagine *pushing them out* to S_1 and S_2. Let $S' \subset \mathbf{S}$. If S' is a parallel 2-sphere, it must be parallel to the *outside* (the side disjoint from $\tilde{S}$), and we can fill in the (3-ball+arc) that it bounds. If S' is a normal 2-sphere, it must be parallel to some component of ∂M_0, since Σ is maximal. It cannot be parallel to a boundary component of M_0 on the inside, since it contains at least two components (S_1 and S_2) of ∂M_0 on the inside. Hence it must be parallel to a component of ∂M_0 on the outside, i.e., we can push it out to ∂M_0 disjoint from the rest of $\tilde{M}$. Hence each component $\partial\tilde{M}$ in S can either be filled in with a 3-ball or pushed out to the boundary of M_0. Conversely, each component of ∂M_0 will have one component of $\partial\tilde{M}$ pushed onto it by this process. Hence $\tilde{M}$ is homeomorphic to $M_0 -$ (3-balls). □

The claim concludes the proof of Lemma 2.

3. An Unknotting Lemma

This variant on the light bulb trick is a key ingredient in the proof of Lemma 4. It implies that if you attach a circular (and very flexible!) light bulb to the ceiling via a knotted cord you can unknot the cord without detaching it from the ceiling or the bulb, if you're willing to distort the bulb. Generalizations and further details of this lemma can be found in [5].

3.1. The Fluorescent Light Bulb Trick.

Lemma 3. *Let K be a closed connected subset of S^3, and let v be a point in $S^3 - K$. Let α, β be arcs connecting v to a point p on ∂K, coincident near p, with interiors disjoint from K. Then there is a homeomorphism h of S^3, isotopic to the identity rel v, such that $h(K) = K$, $h(\alpha) = \beta$.*

Proof. There is a homeomorphism h_α (h_β) of S^3 rel v which shrinks α (β) to a very short arc, and straightens it out. It is helpful to imagine here that α (β) is made of very stiff wire threaded through a bead at v: h_α (h_β) first pulls α (β) almost completely through v, then straightens out the remaining short piece. Call this resulting straight arc $\delta_{\alpha(\beta)}$. A rotation r will take δ_α to δ_β. Then $h = (h_\beta^{-1}(r(h_\alpha)))$ is the desired homeomorphism. □

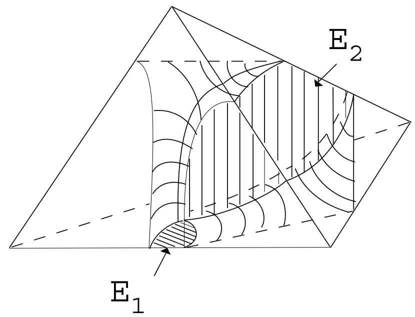

FIGURE 6

Corollary 3.1. *Let K be a closed connected subset of B^3 disjoint from ∂B^3, and let v be a point on ∂B^3. Let α, β be arcs connecting v to a point p on ∂K, coincident near p, with interiors disjoint from K. Then there is a homeomorphism of B^3, isotopic to the identity rel v, such that $h(K) = K$, $h(\alpha) = \beta$.*

Proof. Map B^3 into S^3 so that ∂B^3 maps to a point and use Lemma 3. □

4. Proof of Lemma 4

Lemma 4. *A component M_0 of type 3 is a 3-ball if and only if M_0 contains an almost normal 2-sphere.*

Proof.

Step 1: We first show that if M_0 contains an almost normal 2-sphere then it is a 3-ball.

Proof of step 1: Let T_0 be the remnants of the 1-skeleton of T in M_0. Let S be an almost normal 2-sphere in M_0. Examination of the single octagonal component of S yields the information that S is compressible to both sides in the complement of T_0, since one can double the obvious boundary compressions (see figure 6).

We need to show that S is weakly incompressible in the complement of T_0. S divides M_0 into two components, M_1 and M_2, with $\partial M_0 = M_1$. Let $E_1 \subset M_1$ and $E_2 \subset M_2$ be the boundary compressing disks for S arising from its octagonal component. Let $D_1 \subset M_1$ and $D_2 \subset M_2$ be compressing disks for S in the complement of T_0. Suppose $\partial D_1 \cap \partial D_2 = \emptyset$. After an isotopy which maintains $\partial D_1 \cap \partial D_2 = \emptyset$, we can assume that $D_1 \cup D_2$ intersects the 2-skeleton only in arcs. Minimize, up to isotopy maintaining $\partial D_1 \cap \partial D_2 = \emptyset$, the number of points in $(D_1 \cup D_2) \cap (S \cap \text{2-skeleton})$. Consider an outermost arc α in D_1 of $D_1 \cap$ (2-skeleton). α cuts off a subdisk D_1' of D_1. D_1' lies entirely in some tetrahedron H. If D' lies in a tetrahedron not containing the octagonal component, we could reduce the number of components in $(D_1 \cup D_2) \cap (S \cup 2\text{-skeleton})$, maintaining $\partial D_1 \cap \partial D_2 = \emptyset$, contradicting our hypothesis. Hence D_1' lies in the tetrahedron containing the octagonal component of S, and in fact D_1' must look like a push-off of E_1, disjoint from the edge incident to E_1. Applying the same argument to D_2 yields a contradiction, since D_1' and D_2' must intersect. Hence S is weakly incompressible in the complement of T_0.

Let $\mathbf{D} \subset M_1$ and $\mathbf{E} \subset M_2$ be minimal complete collections of compressing disks for S in the complement of T_0. The proof now is similar to the proof of Lemma 2. Let $\tilde{M}$ be the manifold obtained by attaching 2-handles to a small neighborhood of S along $\partial\mathbf{D} \cup \partial\mathbf{E}$. Notice that $\tilde{M}$ is homeomorphic to a punctured 3-ball. We aim to show that $\tilde{M}$ is (more-or-less) all of M_0.

Apply Claim 1.2 and Claim 1.1 to conclude, after an isotopy, that every component of $\partial\tilde{M}$ is either a parallel 2-sphere or a (separating) normal 2-sphere. Since Σ is maximal, any normal 2-sphere in M_0 is parallel to ∂M_0. There is exactly one component, S', of $\partial\tilde{M}$, that separates S from ∂M_0. Let S'' be any other component of $\partial\tilde{M}$. If S'' is normal, then it is boundary parallel, hence we have the almost normal 2-sphere S lying between two parallel normal 2-spheres ∂M_0 and S'', a contradiction. So S'' must be a parallel sphere, and we can fill it in with a (3-ball + arc). Filling in all such S'''s, we are left with the single component S' of $\partial\tilde{M}$, which cannot be a parallel 2-sphere, hence must be parallel to ∂M_0. Pushing this component out to ∂M_0 completes the argument.

Step 2: We now show that if M_0 is a 3-ball then M_0 contains an almost normal 2-sphere.

Proof of step 2: We adjust the definition of thin position for knots in the 3-sphere [2]: assume that T_0 is a collection of arcs properly imbedded in the 3-ball M_0. Let F be a foliation of M_0-(point) with 2-spheres, such that all but a finite number of the 2-spheres intersect T_0 transversely, and every 2-sphere in F has at most one point of tangency with T_0. Call the 2-spheres having a point of tangency with T_0 singular spheres, and all other 2-spheres in F transverse spheres. Between each adjacent pair of singular spheres, choose a transverse sphere S_0. Define the *width of* T_0 *with respect to* F to be the sum over i of [the number of times T_0 intersects S_i]. Define the *width of* T_0 to be the minimum width of T_0 with respect to F over all possible foliations F. If the foliation F realizes the width of T_0 then T_0 (rel F) is in *thin position.*

Let F be the foliation realizing the width of T_0. Beginning at ∂M_0, the foliation passes through critical levels with respect to T_0. We see a sequence of maxima with respect to F, then a sequence of minima, and so on. Call a transverse 2-sphere in the region where the sequence shifts from maxima to minima (minima to maxima) a *thick* (*thin*) 2-sphere.

Call a simple closed curve lying in the interior of a face of a tetrahedron a *simple curve*.

Let S be any transverse 2-sphere in F. Suppose S intersects the boundary of each tetrahedron in normal curves and simple curves.

Claim 4.1. *Let H be any tetrahedron in the triangulation of M. Then $S \cap \partial(H)$ contains no parallel curves of length greater than or equal to eight.*

Claim 4.2. *Let H be any tetrahedron in the triangulation of M. Then $S \cap \partial(H)$ contains no curve of length greater than eight.*

Claim 4.3. *Let H_1 and H_2 be distinct tetrahedra in the triangulation of M. Then $S \cap \partial(H_1)$ and $S \cap \partial(H_2)$ do not both contain curves of length eight.*

The proofs of Claims 4.1–4.3 proceed by exhibiting various violations of thin position.

Recall from Section 1 the following facts about normal curves on the boundary of a tetrahedron H:

(1) A normal curve of odd length has length 3.
(2) There are no normal curves of length 6.
(3) Any normal curve of length greater than 8 crosses some edge of H at least 3 times.

For each tetrahedron H, $S \cap H$ is almost determined by $S \cap \partial H$; we can picture $S \cap H$ as a collection of disks $\mathbf{D}$ with $\partial\mathbf{D} =$ the normal curves of $S \cap \partial H$, connected together and to the faces by tubes, which may, of course, be quite complicated, and may run through each other.

S divides M_0 into two pieces, the *exterior*, which contains ∂M_0, and the *interior*. Anything in the exterior of S lies *above* S, anything in the interior lies *below*.

An *upper* (*lower*) *disk* for S is a disk K such that $\partial K = \alpha \cup \beta$, where α is an arc imbedded in S, $\partial\alpha = \partial\beta$, β is a subarc of T_0, $K - \alpha$ intersects S transversely, and a small product neighborhood of α lies above (below) S. Call K a *strict upper* (*lower*) *disk* if β lies entirely above (below) S. By an innermost disk argument, one can see that such a disk K can be chosen to lie completely in the interior of M_0. K describes an isotopy of T_0 in which the arc β can be replaced by the arc α, and then pushed slightly below (above) S. See [9] for further details.

Proof of claim 4.1. Suppose $S \cap \partial H$ contains parallel normal curves c_1 and c_2 of length greater than or equal to eight. Choose c_1 and c_2 to be adjacent in ∂H, and such that c_2 bounds a disk J in ∂H containing only normal

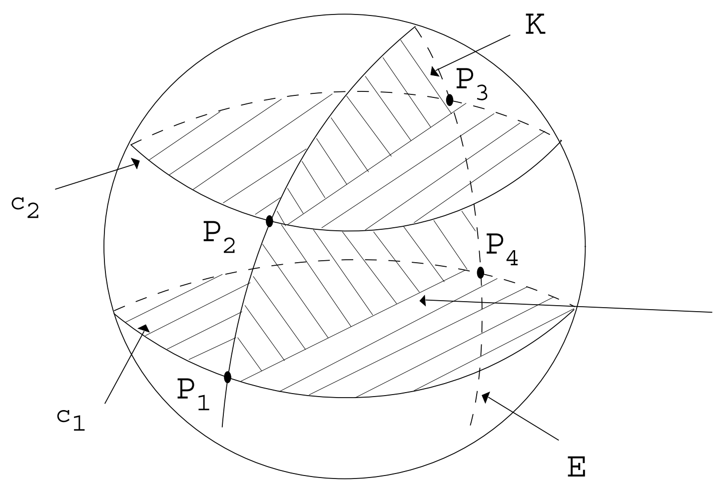

FIGURE 7

curves of length three. Let E be one of the edges crossed at least twice by c_1 and c_2. Let D_1 and D_2 be disks in $\mathbf{D}$ with $\partial D_1 = c_1$, $\partial D_2 = c_2$. Let p_1, p_2, p_3, p_4 be consecutive (on E) points of intersection between E and $c_1 \cup c_2$, where the segment of E between p_2 and p_3 lies in J. So p_1 and p_4 are points on c_1, while p_2 and p_3 are points on c_2. Let K be an imbedded disk in H, with ∂K consisting of the piece e of E connecting p_1 to p_4, containing p_2 and p_3, together with an imbedded arc α in D_1 connecting p_1 to p_4. Further, choose K so that K intersects D_2 in a single imbedded arc connecting p_2 to p_3 (see figure 7).

K acts as a simultaneous upper and lower disk for T_0; we now replace e with α via an isotopy across K, which takes place in M_0, reducing the width of T_0, a contradiction. Hence $S \cap \partial H$ contains no parallel normal curves c_1 and c_2 of length greater than or equal to eight. □

Proof of claim 4.2. Suppose $S \cap \partial H$ contains a normal curve c with length greater than eight. Let E be an edge of H intersecting c at least 3 times. Let $D \subset \mathbf{D}$ be the element of $\mathbf{D}$ with boundary c. Let p_1, p_2 and p_3 be consecutive points (on E) of intersection between E and c (see figure 8).

Then there exist upper and lower disks K_1 and K_2 for S with respect to T_0 which intersect only at p_2. The boundary of K_1 consists of the piece of E between p_1 and p_2, together with an arc α_1 in D connecting p_1 to p_2. The boundary of K_2 consists of the piece of E between p_2 and p_3, together with an arc α_2 in D connecting p_1 to p_2, with $\alpha_1 \cap \alpha_2 = p_2$. The idea is to replace the piece of E connecting p_1 to p_2 with α_1 and the piece connecting p_2 to p_3 with α_2. This can be done by an isotopy of E;

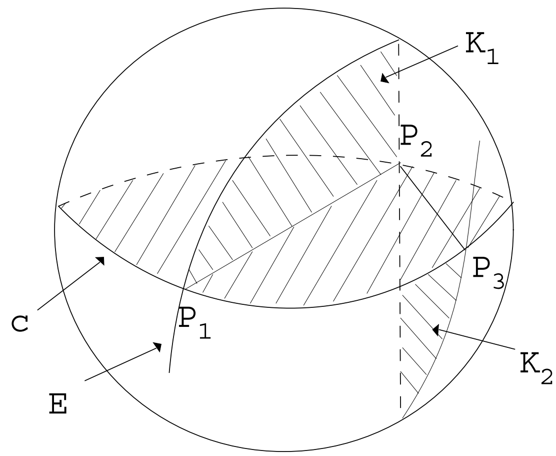

FIGURE 8

note that the isotopy in fact occurs entirely inside M_0, since no normal 2-sphere disjoint from S can intersect K_1 or K_2 non-trivially. This operation reduces the width of T_0, a contradiction. Hence $S \cap \partial H$ contains no normal curve c with length greater than eight. □

Proof of claim 4.3. Suppose $S \cap \partial(H_1)$ and $S \cap \partial(H_2)$ contain curves c_1 and c_2 of length eight. Let D_1 and D_2 be disks in $\mathbf{D}_1$ and $\mathbf{D}_2$ respectively with $c_i = \partial D_i$. c_1 crosses two edges of H_1 exactly twice. Let e_{1u} and e_{1l} be the segments of these edges disjoint from the vertices of H. Note that e_{1u} begins above S, e_{1l} begins below. Let K_{1u} be a disk imbedded in H with boundary consisting of e_{1u} together with an arc α_{1u} in D_1 connecting the ends of e_{1u}. Define K_{1l} similarly. We choose α_{1u} and α_{1l} to intersect in a single point. Define e_{2u}, e_{2l}, K_{2u}, K_{2l}, α_{2u} and α_{2l} similarly in H_2. We use K_{1u} and K_{2l} to reduce the width of T_0. If e_{1u} and e_{2l} are disjoint or intersect in a single (end)point, replace e_{1u} with α_{1u} and e_{2l} with α_{2l} via an isotopy. This reduces the width of T_0. The other possibility is that e_{1u}, say, is a proper subarc of e_{2l}. Then replacing e_{2l} with α_{2l} reduces the width of T_0, a contradiction. Hence $S \cap \partial(H_1)$ and $S \cap \partial(H_2)$ cannot both contain curves of length eight. □

Claims 4.1–4.3 imply that if a transverse S intersects the 2-skeleton only in normal and simple curves, then $S \cap$ (2-skeleton) consists of normal curves of lengths 3 and 4, simple curves and possibly one normal curve of length 8.

Claim 4.4. *There exists a transverse 2-sphere in the first thick region of F which intersects the 2-skeleton entirely in normal arcs and simple curves.*

Proof. This is a straightforward application of [2], Lemma 4.4. The idea is as follows: we can assume a transverse 2-sphere S' at the top of the thick region, i.e. immediately below a maximum of T_0, cuts off a boundary compressing disk, lying above S', in a face of one of the tetrahedra. This is an upper disk for S'. Similarly a transverse 2-sphere S'' immediately above the minimum occurring at the bottom of the thick region cuts off a boundary compressing disk in a face which lies below S'', forming a lower disk for S''. Because T_0 is in thin position, we never see upper and lower disks for a transverse S which are disjoint or which intersect at a single point in $S \cap T_0$; hence somewhere between S' and S'' we can find an S with no upper or lower disks contained in the faces of the tetrahedra. This S intersects the 2-skeleton in normal arcs and simple curves. □

By Claims 4.1–4.3, this sphere S intersects the boundary of each tetrahedron in simple curves, normal curves of length three and four, and possibly, on one tetrahedron, a single curve of length eight.

If $S \cap \partial H$ indeed contains a curve of length eight for some H, we are done; simply compress S as much as possible (to both sides, in several steps) along simple curves and disks interior to the tetrahedra. The resulting collection of 2-spheres will all be normal, with the exception of the component containing the curve of length eight. This will be the desired almost normal 2-sphere.

Claim 4.5. *$S \cap \partial H$ contains a normal curve of length eight for some H.*

Proof. Suppose not. Then compressing S as much as possible, as described above, yields a collection of normal 2-spheres. Since the only normal 2-sphere in M_0 is ∂M_0, S is n copies of ∂M_0 connected by tubes. The tubes may run through each other. We aim to show that this is impossible.

Case 1: Suppose S consists of two copies of ∂M_0 tubed together. Recall, however, that S is boundary compressible above. Let D be a meridian disk of the single tube. D lies above S. By an outermost arc argument, there exists a boundary compressing disk for S disjoint from D. But then there exists a boundary compressing disk for one of the two copies of ∂M_0, a contradiction.

Case 2: S consists of $n > 2$ (nested) copies of ∂M_0 connected by tubes. Beginning with the outermost copy of ∂M_0, we can number the copies of ∂M_0. Let R be the third copy of ∂M_0. Let P be the piece of R which is coincident with S. P is a connected planar surface (see figure 9). We will construct a new and thinner foliation on M_0.

Since R is parallel to ∂M_0, all the arcs exterior to R are parallel, say to a single arc α, and connect R to ∂M_0. Imagine that all the arcs exterior to R are contained in a thick copy of the arc α. Since P is a subsurface

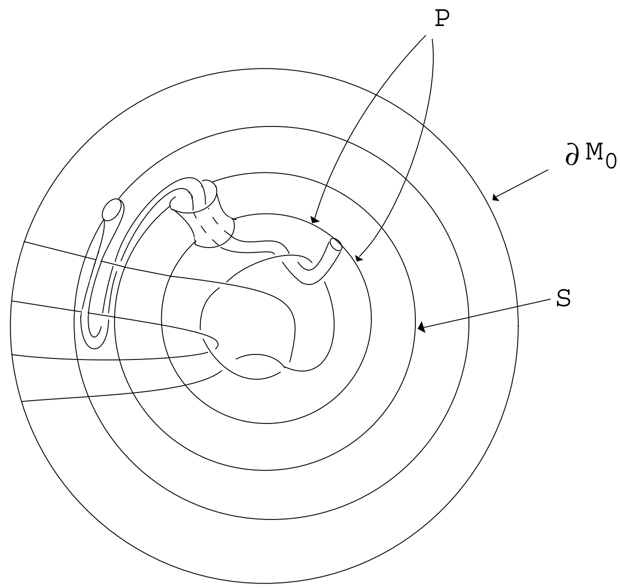

FIGURE 9

of the level sphere S, we can connect ∂M_0 to $P \cap \alpha$ via an arc β which intersects each 2-sphere in the foliation at most once. Now let $K = n(P \cup$ (arcs interior to R)), and apply Corollary 3.1 to replace α with β via a homeomorphism h of M_0. The original foliation F now describes a new foliation $\tilde{F}$ with respect to $h(T_0)$ in $M_0 -$ (point), with the properties:

(1) width(arcs interior to R with respect to F)= width(arcs interior to R with respect to $\tilde{F}$.)
(2) The arcs exterior to R contribute nothing to the width of $h(T_0)$.

Since any arc exterior to R contributed at least two to the width of T_0 with respect to F, $h(T_0)$ is thinner with respect to the foliation $\tilde{F}$. □

This completes the proof of Lemma 4.

5. Finding almost normal 2-spheres

Fix a tetrahedron H in M, and fix a normal curve c of length eight on the boundary of H (there are three choices for c). We modify standard normal surface theory algorithms to search for an almost normal 2-sphere S with octagonal component bounded by c.

Let $\mathcal{Z}$ be a finite system of linear equations such that a positive integral solution to $\mathcal{Z}$ corresponds to a (possibly immersed) surface F, where, except in H, F is composed of normal triangles and quadrilaterals. In H, F is composed of normal triangles and octagonal components with boundaries parallel to c. Just as for normal surfaces, there exists a finite collection $F_1, F_2, \ldots, F_n$ of *fundamental solutions* to $\mathcal{Z}$. Any solution to $\mathcal{Z}$ can be written as a finite sum of fundamental solutions, where the addition operation is a *regular alteration* (a double-curve sum, where the the direction of the sum is chosen to ensure the resulting surface will be normal) along curves of intersection. We would like to show that if M contains an almost normal 2-sphere, then one of the F_i's is an almost normal 2-sphere. Since we can find the F_i's algorithmically (see [4],[6]), this would finish the problem, provided we repeat the search procedure for each normal curve of length eight in each tetrahedron.

Unfortunately this doesn't quite work; we will use in addition the observation that we need only search for almost normal 2-spheres lying in M_0.

Lemma 5. *If there exists an almost normal 2-sphere S in M_0 then there exists one which is a fundamental solution to $\mathcal{Z}$.*

Proof. We note :

(1) Recall that $H_1(M;\mathbb{Z}_2)$ is trivial, hence M contains no closed non-orientable surfaces.
(2) Euler characteristic is additive under the addition operation.

Let S be an almost normal 2-sphere in M_0. If S is not a fundamental solution to $\mathcal{Z}$, then it can be written as the sum of fundamental solutions, so $S = G_1 + G_2 + \cdots + G_k$, where each G_i is a fundamental solution to $\mathcal{Z}$. One of the G_i's is an almost normal surface and the rest are normal surfaces. Since S lies in M_0, each of the G_i's also lies in M_0. By notes 1 and 2, at least one of the G_i's, say G_1, is a 2-sphere. If it is an almost normal 2-sphere, we are done. If not, it is a normal 2-sphere, hence parallel to ∂M_0.

There exists an isotopy i of M_0, preserving the normal structure or almost normal structure of the G_i's, such that $i(G_1)$ is disjoint from $i(G_2 \cup \cdots \cup G_k)$. To see the isotopy, in each tetrahedron, simply push the pieces of $G_2 \cup \cdots \cup G_k$ into M_0 disjoint from the (2-sphere) $\times I$ bounded by G_1 and ∂M_0.

Since i doesn't change the normal or almost normal structure of G_1 or $G_2 \cup \cdots \cup G_k$,

$$G_1 + G_2 + \cdots + G_k = i(G_1) + i(G_2) + \cdots + i(G_k).$$

But $i(G_1) + i(G_2) + \cdots + i(G_k)$ is not a connected surface. Hence S cannot be written as a non-trivial sum of solutions to $\mathcal{Z}$, hence S must be a fundamental solution to $\mathcal{Z}$. $\square$

Acknowledgements

I would like to thank Marty Scharlemann for helpful conversations, particularly on the proof of (the critical) Claim 4.5.

References

[1] A. Casson and C. McA. Gordon, *Reducing Heegaard splittings*, Topology and its Applications **27** (1987).
[2] D. Gabai, *Foliations and the topology of 3-manifolds III*, J. Diff. Geometry **26** (1987).
[3] W. Haken *Theorie der Normalflachen*, Acta. Math. **105** (1961), 245–375.
[4] G. Hemion *The classification of knots and 3-dimensional spaces*, Oxford University Press, 1992.
[5] J. Hass and A. Thompson *Fluorescent light bulbs and the unknotting of arcs in manifolds*, preprint.
[6] W. Jaco and U. Oertel, *An algorithm to decide if a 3-manifold is a Haken manifold*, Topology **23** (1984).
[7] W. Jaco and J. Tollefson, to appear.
[8] H. Rubinstein *The solution to the recognition problem for S^3*, lectures, Haifa, Israel, May 1992.
[9] M. Scharlemann and A. Thompson, *Thin position and Heegaard splittings of the 3-sphere*, preprint.

Department of Mathematics, University of California, Davis, CA 95616
E-mail address: thompson@math.ucdavis.edu