Meridional almost normal surfaces in knot complements

ROBIN TODD WILSON

Suppose K is a knot in a closed 3–manifold M such that $M - N(K)$ is irreducible. We show that for any integer n there exists a triangulation of $M - N(K)$ such that any weakly incompressible bridge surface for K of n bridges or fewer is isotopic to an almost normal bridge surface.

57M99

1 Introduction

It was shown independently by M Stocking [14] and J H Rubinstein [11] that any strongly irreducible Heegaard splitting for an irreducible 3–manifold is isotopic to an almost normal surface. Also see S King [10] for another proof of this result. The concept of a bridge surface for a knot complement is analogous to the idea of a Heegaard surface for a 3–manifold in that the bridge surface is a splitting surface that separates the knot complement into two equivalent and fairly elementary submanifolds. In addition, the fact that a bridge surface lifts to a Heegaard surface in the 2–fold branched cover of a knot complement gives another important connection between bridge surfaces for knot complements and Heegaard surfaces for 3–manifolds.

In the study of bridge surfaces for knots and links the idea of a weakly incompressible splitting surface is immediately analogous to the idea of a strongly irreducible Heegaard surface for a 3–manifold. In this paper we prove an analog to the main theorem of Stocking [14] and Rubinstein [11, Theorem 3]. We show that any weakly incompressible bridge surface in a 3–manifold is isotopic to an almost normal bridge surface.

Main Theorem 1 *Let K be a knot in a closed, orientable, irreducible 3–manifold M. Let $N(K)$ be an open regular neighborhood of K and suppose that the complement $M_K = M - N(K)$ is irreducible. Then for any integer n there is a triangulation $\mathcal{T}$ of M_K such that if S is a bridge surface for K of n bridges or fewer that gives an irreducible Heegaard splitting of M and $S_K = S - N(K)$ is weakly incompressible, then S_K is properly isotopic in M_K to an almost normal surface with respect to $\mathcal{T}$.*

Published: 8 October 2008 DOI: 10.2140/agt.2008.8.1717

The proof will be similar in spirit to that of [14] but the proof here fills in a few missing cases and simplifies the argument by making greater use of edge slides. Closely related results have been proven by David Bachman [1] and Alexander Coward [4]. In Section 2 we briefly introduce some definitions and notation. The proof of the main theorem is contained in Section 3.

This research was done while under the support of the UC President's Postdoctoral Fellowship Program and the Department of Mathematics at UC Santa Barbara. I would like to thank Martin Scharlemann for all of the helpful conversations and many valuable comments, as well as Scott Taylor for his insightful comments about the proof of Lemma 24.

1.1 Acknowledgments

Research supported by UC President's Postdoctoral Fellowship Program and Department of Mathematics at the University of California, Santa Barbara.

2 Preliminaries

Notation If K is a properly embedded 1–manifold in a 3–manifold M then let $M_K = M - N(K)$. If X is any surface in M transverse to K such that $K \cap X \neq \varnothing$, then let $X_K = M_K \cap X$. For $\mathcal{T}$ a triangulation of a 3–manifold M, let $\mathcal{T}_{\partial M}$ denote the restriction of $\mathcal{T}$ to ∂M.

The following definition is from Tomova [16] and is based on the definition of a K–compression body given in Bachman [2].

Definition 2 ([16]) A properly embedded arc K in a 3–manifold M is *boundary parallel* if there is a disk D in the 3–manifold so that ∂D is the end point union of K and an arc in ∂M. The disk D is called a *cancelling disk* for K. A *K–handlebody* (A, K) is a handlebody A containing a finite collection of boundary parallel arcs K. When there is little risk of confusion we will also refer to $A_K = A - N(K)$ as a K–handlebody. For our purposes, a *K–compression body* (W, K) is a compression body W containing a finite collection of arcs K properly embedded in W such that each arc is either boundary parallel or each arc has one end on each of $\partial_+ W$ and $\partial_- W$ and is vertical in the product region $\partial_- W \times I \subset W$.

Remark 3 Two sets K and K' of boundary parallel arcs in a handlebody A or vertical arcs in a compression body are properly isotopic in A if they have the same cardinality, ie, $|K| = |K'|$.

Definition 4 A *spine of a handlebody* A is a graph Σ_A properly embedded in A such that $A-\Sigma_A$ is a product $\partial A \times I$. Let (A, K) be a K–handlebody and suppose that Σ_A is a spine for handlebody A and K is a collection of boundary parallel arcs in A. Let α be a collection of $|K|$ arcs, each connecting Σ_A to a single arc of K. Then a regular neighborhood $A' = N(\Sigma_A \cup \alpha)$ is again a K–handlebody, and K intersects it in a boundary parallel set of arcs $K' \subset A'$. If the closure of the region $(A-A')_{(K-K')}$ between them is a product $\partial A_K \times I$ then $\Sigma_{(A,K)} = \Sigma_A \cup \alpha$ is called a *spine of the* K*–handlebody* (A, K).

Definition 5 A *spine of a compression body* W is a graph Σ_W with endpoints in $\partial_- W$ such that $W-(\Sigma_W \cup \partial_- W)$ is a product $\partial_+ W \times I$. Let (W, K) be a K–compression body and suppose that Σ_W is a spine for W and K is a collection of boundary parallel and vertical arcs in W. Let k denote the collection of boundary parallel arcs in K. Let β be a collection of $|k|$ arcs, each connecting Σ_W to a single arc of k. Then a regular neighborhood $W' = N(\Sigma_W \cup \beta \cup \partial_- W_K)$ is again a K–compression body and K intersects it in a set of boundary parallel and vertical arcs $K' \subset W'$. If the closure of the region $(W-W')_{(K-K')}$ between them is a product $\partial_+ W_K \times I$ then $\Sigma_{(W,K)} = \Sigma_W \cup \beta$ is called a *spine of the* K*–compression body* (W, K).

Remark 6 It is relatively easy to find such a spine for a K–handlebody or K–compression body (A, K). Choose a spine for Σ of handlebody (compression body) A and isotope K so that in the product structure $A-N(\Sigma) = \partial A \times I$, each (boundary parallel) arc of K has a single maximum. Let α (β) be a collection of vertical arcs in this product structure, connecting each maximum of K to Σ.

Definition 7 (Scharlemann–Tomova [13]) Let K be a knot in a closed, orientable 3–manifold M and let S be a Heegaard surface for M. That is, $M = W \cup_S W'$, where W and W' are handlebodies in M. If in addition, W_K and W'_K are K–handlebodies then we call S a *bridge surface* for M_K. (We will often abuse notation and call the punctured surface S_K a bridge surface as well.) We call the decomposition $M_K = W_K \cup_{S_K} W'_K$ a *bridge splitting* of the 3–manifold M_K and we say that K is in *bridge position* with respect to bridge surface S.

Definition 8 ([13]) Let K be a 1–manifold embedded in M and suppose that F is a properly embedded surface in M so that F is transverse to K. A simple closed curve on F_K is *essential* if it doesn't bound a disk or a once punctured disk in F_K. An embedded disk $D \subset M_K$ is a *compressing disk* for a surface F_K if $D \cap F_K = \partial D$ and ∂D is an essential curve in F. A surface F in M is a *splitting surface* for M if we can express M as the union of two 3–manifolds along F. If F is a splitting surface for M then we say that the surface F_K is *weakly incompressible* if any pair of compressing disks on opposite sides of the surface intersect. If F_K compresses on both sides but is not weakly incompressible then it is called *strongly compressible*.

The study of normal surfaces was first developed by Haken [7]. The concept of an almost normal surface that is used in this paper first appeared in [11].

Definition 9 Let S be a triangulated surface and let c be a curve on S. Assume that c is transverse to the 1–skeleton of the triangulation. A curve c in S is called *normal* if the intersection of c with any triangle of the triangulation contains no closed curves and no arcs with both endpoints on the same edge.

Definition 10 ([7]) Let M be a triangulated 3–manifold. A *normal triangle* in a tetrahedron of the triangulation is an embedded disk that meets three edges and three faces of the tetrahedron. A *normal quadrilateral* is an embedded disk in a tetrahedron that meets four edges and four faces of the tetrahedron. Normal triangles and quadrilaterals are called *normal disks*. Normal disks meet the faces of the boundary of a tetrahedron in normal curves.

Definition 11 ([11]) Let M be a triangulated 3–manifold. An embedded surface $S \subset M$ is a *normal surface* if it meets each tetrahedron in a disjoint collection of normal disks. A surface S is *almost normal* if S meets each tetrahedron of the triangulation in a collection of normal disks, but in one tetrahedron there is exactly one exceptional piece. This exceptional piece is either a normal octagon, or it is an annulus consisting of two normal disks with a tube between them that is parallel to an edge of the 1–skeleton.

The proof of the main theorem relies heavily on the idea of thin position, first introduced by Gabai [6].

Definition 12 ([14]) Let $M_K = W_K \cup_{S_K} W'_K$ denote a bridge splitting of M_K. Given spines $\Sigma_{(W,K)}$ and $\Sigma_{(W',K)}$ for the K–handlebodies (W, K) and (W', K) respectively, there is a diffeomorphism $S_K \times (0,1) \simeq M_K - N(\Sigma_{(W,K)} \cup \Sigma_{(W',K)})$. For $t \in (0,1)$ denote the surface corresponding to $S_K \times \{t\}$ by $S_t \subset M_K$. A *standard singular foliation* F of $M_K = W_K \cup_{S_K} W'_K$ extends this structure to all of M_K by adding two singular leaves S_0 and S_1, called the top and bottom leaves. All leaves meet the torus ∂M_K in the standard foliation in meridian circles. The top and bottom singular leaves consist of the union of the spines of the K–handlebodies W_K and W'_K respectively and the meridian circles of ∂M_K corresponding to each of the n endpoints of $\Sigma_{(W,K)}$ and $\Sigma_{(W',K)}$. There is a height function $h\colon M \to [0,1]$ associated with the standard singular foliation given by the map that sends all points on a leaf S_t together with the incident meridian disks of $N(K)$ to the point t in $[0,1]$.

Definition 13 (Thompson [15] and Stocking [14]) Assume that T is a collection of arcs properly embedded in M_K and is in general position with respect to F, a standard singular foliation of M_K. That is, all but a finite number of leaves of F intersect T transversally, every leaf in F has at most one point of tangency with T and T is disjoint from the singular subarcs of the singular leaf. If a leaf has a point of tangency with T call it a *tangent leaf* and all other leaves *transverse leaves*. Between each two adjacent tangent leaves choose a transverse leaf L_i. Define the *width of a fixed embedding of* T *with respect to* F to be the sum over i of the number of times T intersects L_i. If T has been properly isotoped to minimize its width with respect to F then we say that T is in *thin position* with respect to F.

Definition 14 ([15; 14]) Let T be in thin position with respect to a standard singular foliation F. Then as we move down the foliation from the top the arcs will form a sequence of maxima with respect to F, then a sequence of minima and so on. We will call a leaf in a region where the sequence shifts from maxima to minima a *thick leaf* and we will call such a region a *thick region*. An *upper (lower) disk* D for a transverse leaf L of F is a disk in $\mathrm{int}(M)-T$ such that $\partial D=\alpha\cup\beta$ where α is an arc embedded in L, β is a subarc of T, $\partial\alpha=\partial\beta$, $D-\alpha$ intersects L transversely and a small neighborhood of α lies above (below) L.

For the proof of the main theorem we will need the following Lemmas. The first Lemma is proved by Stocking in [14].

Lemma 15 ([14, Lemma 1]) *Let S be an almost normal surface in an irreducible 3–manifold. Suppose that S is incompressible to one side. Then S is isotopic to a normal surface that does not intersect S and that does not contain S to the incompressible side.*

A version of the following theorem was proved for strongly irreducible Heegaard surfaces by Casson and Gordon in [3] and has been adapted to the situation of weakly incompressible bridge surfaces by Tomova in [16].

Lemma 16 ([16, Corollary 6.3]) *Let K be a knot in a closed, orientable, irreducible 3–manifold M. Let S_K be a weakly incompressible splitting surface for M_K and let S'_K be a surface that is obtained from S_K by compressing S_K to one side. Then S'_K is incompressible to the other side.*

3 Almost normal bridge surfaces

The proof of the main theorem follows from an application of ideas from [14] where the Heegaard surface is replaced with a bridge surface. An important difference between the two arguments is that the leaves of the foliations in this case are surfaces with boundary as opposed to closed surfaces. The argument relies heavily on Lemma 17, Lemma 24 and Lemma 25 whose statements are close to those of [14, Lemmas 4 and 5]. The proofs of these lemmas are similar in spirit to the originals, but differ in detail. Several cases missed in the original proof in [14] are included here, more extensive use is made of edgeslides and the arguments have been adapted to our situation.

Proof of Main Theorem Let K be a knot in a closed, orientable 3–manifold M and assume that M and M_K are both irreducible. Suppose that the knot K is in n–bridge position with bridge surface S so that $M = W \cup_S W'$ is an irreducible Heegaard splitting of M and the punctured surface S_K is weakly incompressible. Also assume that S_K separates M_K into the two K–handlebodies W_K and W'_K. Let $M_K = W_K \cup_{S_K} W'_K$ denote the bridge splitting of M by S. We can foliate $M_K = W_K \cup_{S_K} W'_K$ with a standard singular foliation that intersects the torus ∂M_K in meridian circles. The top singular leaf of the foliation, L_{top}, is a 1–complex given by the union of a spine $\Sigma_{(W,K)}$ of W_K and one meridian circle of ∂M_K for each of the n endpoints of $\Sigma_{(W,K)}$ on K. Similarly, the bottom singular leaf of the foliation, L_{bot}, is a 1–complex given by the union of a spine $\Sigma_{(W',K)}$ of W'_K and one meridian circle of ∂M_K for each of the n endpoints of $\Sigma_{(W',K)}$ on K. Thus there is a symmetric picture near the top and bottom leaves of the foliation.

Consider a nearby leaf of the frontier of a regular neighborhood of L_{top} (resp. L_{bot}) in M_K. It can be viewed as consisting of two parts. See Figure 1. The first is a collection Γ^{top} (resp. Γ^{bot}) of n boundary parallel annuli. Secondly, these annuli are tubed together via the boundary t^{top} (resp. t^{bot}) of a regular neighborhood of $\Sigma_{(W,K)}$ (resp. $\Sigma_{(W',K)}$). Topologically Γ^{top} (resp. Γ^{bot}) consists of n once-punctured annuli and t^{top} (resp. t^{bot}) consists of an n–punctured copy of S_K. Since t^{top} and t^{bot} arise as the boundary of a regular neighborhood of a 1–complex it is natural to refer to them as collections of "tubes". We will refer to the collection of annuli Γ^{top} and Γ^{bot} as Γ. Throughout the paper the foliation we refer to will always be the standard product foliation of $S^1 \times I$ on each component between the top and bottom annuli on ∂M_K.

Next we will describe how to triangulate M_K so that the collection Γ of annuli is normal with respect to the triangulation. See Figure 2. Consider the collection of $2n$ meridional annuli on ∂M_K parallel to Γ. Start by choosing a longitude of ∂M_K. Next, in each annulus choose a meridian circle. View the intersection of each meridian and

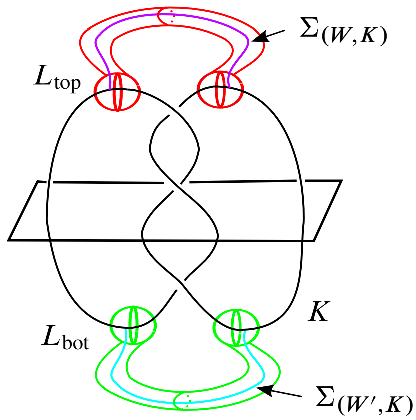

Figure 1: An example with $M = S^3$ and K a trefoil

the longitude as a vertex, view each of these meridians as the union of a vertex and an edge and view the longitude as a union of $2n$ edges and $2n$ vertices. This divides ∂M_K into rectangles. Now subdivide each rectangle by adding a diagonal edge connecting two adjacent vertices. This gives a triangulation of ∂M_K.

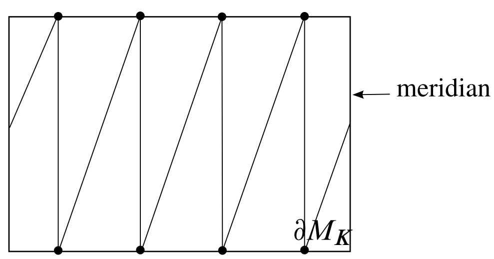

Figure 2: Triangulation of ∂M_K

By Jaco–Letscher–Rubinstein [8] (also see Jaco–Rubinstein [9, page 56]) we can extend the triangulation of ∂M_K to a triangulation of all of M_K without adding any vertices. Denote this triangulation of M_K by $\mathcal{T}$. All of the vertices of this triangulation are contained in a neighborhood of the top and bottom leaves of the foliation. The collection Γ of $2n$ annuli is normal with respect to this triangulation and Γ contains all of the $2n$ vertices of the triangulation to one side, separating them from the rest of M_K. Also note that $\mathcal{T}$ has all of its vertices on ∂M_K, so it has no vertex-linking 2–spheres. However, $\mathcal{T}$ may contain normal 2–spheres disjoint from Γ that are not vertex-linking.

Let Λ denote a maximal collection of non-parallel disjoint normal 2–spheres in M_K disjoint from Γ. Cutting M_K open along Λ results in several components, but only one component will contain the torus boundary of M_K. Call this component ${M_0}^+$. Note that ${M_0}^+$ may have multiple 2–sphere boundary components along with ∂M_K. Since M_K is irreducible each 2–sphere in Λ must bound a 3–ball to the opposite side of ${M_0}^+$.

Since the annuli Γ are normal and they contain the vertices of the triangulation to one side, each normal 2–sphere in $\partial {M_0}^+$ is connected to Γ via an edge of $\mathcal{T}^1$. If an edge of $\mathcal{T}^1$ connects two 2–sphere boundary components of ${M_0}^+$ then by Thompson [15, Lemma 2] it follows that ${M_0}^+$ must be a punctured 3–ball with a torus boundary component, which it clearly is not. Thus we can conclude each 2–sphere component of $\partial {M_0}^+$ is connected to Γ^{top} or Γ^{bot} by an edge of $\mathcal{T}^1 \cap {M_0}^+$. Assume without loss of generality that an edge connects the 2–sphere to Γ^{top}. Taking a tube that lies in the interior of a tetrahedron parallel to this edge connecting Γ^{top} to the normal 2–sphere gives an almost normal annulus isotopic to an annulus in Γ^{top}. By Lemma 15 this surface is isotopic to a normal surface giving a new collection of normal annuli that we will call $\Gamma^{\text{top}'}$. We can isotope the tubes t so that they lie in ${M_0}^+$. We can do this for each normal 2–sphere in ${M_0}^+$ and then replace the original singular leaf L_{top} with the singular leaf $L'_{\text{top}} = \Gamma^{\text{top}'} \cup \Sigma_{(W,K)}$. Let M_0 be the side of $\Gamma^{\text{top}'}$ that lies in ${M_0}^+$.

Let $K_0 = K \cap M_0$. Isotoping the bridge surface S_K to be disjoint from Γ' induces a splitting of M_0 into two K_0–compression bodies W_0 and W'_0. Continue to call this splitting surface S_K. Since M_0 is a deformation retract of M_K the surface S_K is a weakly incompressible splitting surface for M_0. We can foliate M_0 with a standard singular foliation F_0 with leaves isotopic to S_K. The top leaf of the foliation is L_{top} and the bottom leaf is L_{bot}. Let $\mathcal{T}_0^1$ denote the part of the 1–skeleton of $\mathcal{T}$ that lies in the interior of M_0. Put $\mathcal{T}_0^1$ into thin position with respect to F_0. Let Σ_0 denote the pair of spines $\Sigma_{(W_0,K_0)}$ and $\Sigma_{(W'_0,K_0)}$ of the K_0–compression bodies W_0 and W'_0 respectively. Note that $\Sigma_{(W_0,K_0)} \subset \Sigma_{(W,K)}$ and $\Sigma_{(W'_0,K_0)} \subset \Sigma_{(W',K)}$. If $\mathcal{T}_0^1$ intersects Σ_0 then isotope $\mathcal{T}_0^1$ slightly off of Σ_0. Let Γ_0 denote the pair $\Gamma^{\text{top}'}$ and $\Gamma^{\text{bot}'}$.

The triple $(M_0, \Sigma_0, \Gamma_0)$ is the first step in an iterative process. Each later step will consist of a triple $(M_i, \Sigma_i, \Gamma_i)$ with the following properties: $M_i \subset M_{i-1}$ will be a submanifold of M_K for each $i \geq 0$. The surface $\Gamma_i = \partial M_i - \partial M$ will be a pair Γ_i^{top} and Γ_i^{bot} of properly embedded normal surfaces with respect to the triangulation $\mathcal{T}$ given above. Let $K_i = K \cap M_i$. The submanifold M_i has a weakly incompressible splitting surface S_K that separates M_i into two K_i–compression bodies W_i and W'_i. Let Σ_i^{top} and Σ_i^{bot} denote the spines $\Sigma_{(W_i,K_i)}$ and $\Sigma_{(W'_i,K_i)}$ of K_i–compression bodies W_i and W_i' respectively. The spine Σ_i^{top} (resp. Σ_i^{bot}) can be extended to give a spine of $M - \Sigma_{W'} \simeq W$ (resp. $M - \Sigma_W \simeq W'$). Let Σ_i denote the pair of spines Σ_i^{top} and Σ_i^{bot} of W_i and W_i' respectively. As usual we define Σ_i up to isotopy and slides of edges over other edges and over $\partial_- W_i = \Gamma_i^{\text{top}}$. The complement of a regular neighborhood $N(\Sigma_i^{\text{top}})$ (resp. $N(\Sigma_i^{\text{bot}})$) in M_i is foliated by copies of S_K with singular leaf $\Gamma_i^{\text{bot}} \cup \Sigma_i^{\text{bot}}$ (resp. $\Gamma_i^{\text{top}} \cup \Sigma_i^{\text{top}}$). Thus M_i can be foliated with a singular foliation F_i by copies of S_K with its top singular leaf consisting of $\Gamma_i^{\text{top}} \cup \Sigma_i^{\text{top}}$ and its bottom singular leaf consisting of $\Gamma_i^{\text{bot}} \cup \Sigma_i^{\text{bot}}$. Now consider the edges of $\mathcal{T}^1$ that do not lie on ∂M_K and let $\mathcal{T}_i^1$ denote their intersection with M_i. Put $\mathcal{T}_i^1$ into thin position with respect to F_i.

In general we will obtain the submanifold M_i from M_{i-1} for each $i \geq 1$ by showing that M_{i-1} contains an almost normal surface obtained by compressing the splitting surface S_K to one side. This almost normal surface is then isotopic to a normal surface by Lemma 15 and Lemma 16 which we will call Γ_i^{top} (or Γ_i^{bot} depending on whether we compressed S_K above or below). The surface Γ_i^{top} determines a submanifold M_i of M_0 with $\Gamma_i^{\text{top}} \subset \partial M_i$. The surface S_K separates M_i into two K_i–compression bodies, W_i and W_i' with $\partial_- W_i = \Gamma_i^{\text{top}}$ and $\partial_- W_i' = \Gamma_i^{\text{bot}}$. The spines of these K_i–compression bodies are now what we will call Σ_i^{top} and Σ_i^{bot} respectively.

Denote by t_i^{top} and t_i^{bot} the boundary of a regular neighborhood $N(\Sigma_i)$ in M_i which we continue to call "tubes". A regular leaf near the top (resp. bottom) singular leaf is then obtained by attaching t_i^{top} (resp. t_i^{bot}) to a punctured copy of Γ_i^{top} (resp. Γ_i^{bot}).

The surface S_K is a weakly incompressible splitting surface for M_i that separates it into two K–compression bodies W_i and W_i' for each $i \geq 0$. To see this observe that if S_K is a weakly incompressible splitting surface for M_{i-1} then since M_i was constructed by compressing along disks in W_{i-1} and W_{i-1}' to obtain W_i and W_i' it follows that the collection of compressing disks in W_i and W_i' are a subset of the collection of compressing disks in W_{i-1} and W_{i-1}'. Since every pair of compressing disks on opposite sides of S_K intersect in M_{i-1} it follows that every pair of compressing disks on opposite sides of S_K in M_i must intersect as well.

Here is a sketch of the iterative process that we will describe in detail later. If it happens that at the first step of this process the surface S_K is isotopic to the collection of annuli Γ_0 then we can conclude that K is the unknot. Otherwise, start with $(M_i, \Sigma_i, \Gamma_i)$. If, without loss of generality, $\chi(\Gamma_i^{\text{top}}) = \chi(S_K)$ and Γ_i^{top} is isotopic to an almost normal surface then since Γ_i^{top} is obtained by compressing a leaf of the foliation it follows that Γ_i^{top} is isotopic to S_K and we are done. Otherwise apply either Lemma 17 or Lemma 24 to obtain a new collection of normal and almost normal surfaces in M_i. If there is an almost normal surface G in the collection with $\chi(G)=\chi(S_K)$ then again because the almost normal surface comes from compressing a leaf of the foliation we know G must be isotopic to S_K and we are done. Otherwise since by Lemma 16 the surface G is incompressible to one side we can use Lemma 15 to isotope the almost normal surfaces (if any) in the collection to be normal. This collection of normal surfaces now becomes either $\Gamma_{i+1}^{\mathrm{top}}$ or $\Gamma_{i+1}^{\mathrm{bot}}$, depending on whether we compressed S_K above or below. Then, using this new collection of normal surfaces we can cut (M_i,Σ_i,Γ_i) along this collection to obtain $(M_{i+1},\Sigma_{i+1},\Gamma_{i+1})$, which will also satisfy the above properties. It turns out that we only need to repeat the recursive step a finite number of times before obtaining an almost normal surface isotopic to S_K. This completes the sketch.

Now, consider the arcs $\mathcal{T}_i^1$ in M_i in thin position with respect to F_i. Recall that all ends of $\mathcal{T}_i^1$ lie on Γ_i^{top} or Γ_i^{bot}, part of the top or bottom singular leaves of F_i. One possibility is that there is a maximum of $\mathcal{T}_i^1$ that is above a minimum of $\mathcal{T}_i^1$ which implies that there is a thick region of $\mathcal{T}_i^1$ in M_i. The other possibility is that all of the minima of $\mathcal{T}_i^1$ are above all of the maxima of $\mathcal{T}_i^1$ and so there is no thick region. In this situation we will consider separately the following two possibilities. The first is that there is no thick region and there is some arc of $\mathcal{T}_i^1$ with both ends on Γ_i^{top} or both ends on Γ_i^{bot}. The second possibility is that there is no thick region and each arc each arc of $\mathcal{T}_i^1$ has one endpoint on Γ_i^{top} and the other endpoint on Γ_i^{bot}. We will consider each of the three possibilities in turn.

The first possibility is that there a thick region of $\mathcal{T}_i^1$ with respect to F_i.

Lemma 17 (cf [14, Lemma 5]) *If there is a thick region for $\mathcal{T}_i^1$ in M_i, then there is a collection of normal and almost normal surfaces in M_i obtained from a leaf of the foliation by compressing the leaf to one side. At most one surface in the collection can be almost normal. Not all of the surfaces are boundary parallel.*

Proof of Lemma 17 The proofs of Claim 18, Claim 19, Claim 20 and Claim 21 can be found in [15]. Let (M_i,Σ_i,Γ_i) be as described above, where $\mathcal{T}_i^1$ is in thin position with respect to the foliation F_i of M_i. Since there is a thick region of $\mathcal{T}_i^1$ in F_i we can apply [15, Claim 4.5].

Claim 18 ([15, Claim 4.5]) *There exists a transverse leaf L in the first thick region of F_i which intersects the 2–skeleton entirely in normal arcs and simple closed curves disjoint from the 1–skeleton.*

Let L be a leaf of F_i in a thick region intersecting the 2–skeleton in normal arcs and simple closed curves disjoint from $\mathcal{T}^1$ as is guaranteed by Claim 18. Then we can apply Claim 19, Claim 20 and Claim 21 to the leaf L.

Claim 19 ([15, Claim 4.1]) *Let H be any tetrahedron in the triangulation $\mathcal{T}$ of M_K. Then $L \cap \partial H$ contains no parallel curves of length greater than or equal to eight.*

Claim 20 ([15, Claim 4.2]) *Let H be any tetrahedron in the triangulation $\mathcal{T}$ of M_K. Then $L \cap \partial H$ contains no curve of length greater than eight.*

Claim 21 ([15, Claim 4.3]) *Let H_1 and H_2 be distinct tetrahedra in the triangulation of M_K. Then $L \cap \partial H_1$ and $L \cap \partial H_2$ do not both contain curves of length eight.*

The above claims together imply that this leaf L of the foliation F_i intersects the 2–skeleton only in simple closed curves disjoint from the 1–skeleton and normal curves of lengths three, four and at most one of length eight. Compressing the simple closed curves in $L \cap \mathcal{T}^2$ as well as any compressions in the interior of the tetrahedra gives a collection of normal surfaces with at most one almost normal surface. The almost normal surface, if it exists, must be a normal octagon since we have compressed any almost normal annuli in the interior of the tetrahedra. We can think of the leaf L as this collection of normal and almost normal surfaces with tubes attached. Since our triangulation has no normal 2–spheres we can conclude that this collection will contain no almost normal 2–spheres as well since any almost normal 2–sphere can be isotoped to give a normal one by Lemma 15. Also notice that in this collection we will not have a normal surface with a tube attached to the side opposite the side that we compressed to or we get a contradiction to the leaf being weakly incompressible. As long as there are no normal 2–spheres in this collection of normal surfaces then the collection of disks of $L \cap \partial H$ that arise in each tetrahedron H of $\mathcal{T}$ correspond to actual compressing disks of L. Also, since this collection of disks is disjoint these compression disks must all be to the same side of L otherwise we again contradict the fact that the surface L is weakly incompressible. In general, by Lemma 16 compressing L to one side gives a surface that is incompressible to the opposite side. Therefore all compressions must be to one side of L. This completes the proof of Lemma 17. □

Remark 22 There is no choice in the direction in which the tubes of L compress, however Lemma 16 implies that after compressing L, the remaining collection of normal and almost normal surfaces is incompressible in the direction opposite to which we have compressed.

For the proof of the Lemma 24 we will need the following claim.

Claim 23 *A properly embedded, orientable, normal surface is incompressible and boundary incompressible in the complement of the 1–skeleton $\mathcal{T}^1$.*

Proof This claim follows from a standard innermost disk and outermost arc argument. □

The second possibility is that there is no thick region and some arc of $\mathcal{T}_i^1$ has both endpoints on Γ_i^{top} or both endpoints on Γ_i^{bot}.

Lemma 24 (cf [14, Lemma 4]) *If there is no thick region for $\mathcal{T}_i^1$ in M_i and some arc of $\mathcal{T}_i^1$ has both endpoints on Γ_i^{top} (resp. has both endpoints on Γ_i^{bot}), then there is an almost normal surface in M_i that is isotopic to a surface obtained from a leaf of F_i by compressing the leaf above (resp. below).*

Proof We will prove the Lemma for arcs of $\mathcal{T}_i^1$ with both endpoints on Γ_i^{top}. The argument for arcs of $\mathcal{T}_i^1$ with both endpoints on Γ_i^{bot} is symmetric. Let L be a leaf of the foliation near the top singular leaf $\Gamma_i^{\text{top}} \cup \Sigma_i^{\text{top}}$, so that L consists of the normal surface Γ_i^{top} punctured and attached to the tubes $t_i^{\text{top}} = \partial N(\Sigma_i^{\text{top}})$. See Figure 1 for the case $i = 0$ and $M = S^3$. Let β be an arc of $\mathcal{T}_i^1$ with both endpoints on Γ_i^{top}. Since there is no thick region for $\mathcal{T}_i^1$, β has only a single minimum and it is parallel to an arc on L, so there is a lower disk E whose boundary is the union of β and an arc α in L. We will show that after some edge slides and isotopies of Σ_i^{top}, α runs once over exactly one tube of t_i^{top} and that this tube connects two normal disks in a tetrahedron, therefore is part of an almost normal surface.

Define the complexity of E to be (a, b), lexicographically ordered, where a is the number of points of $\Sigma_i^{\text{top}} \cap \mathcal{T}^2$ to which α is also incident and b is the number of components in which E meets the 2–skeleton of $\mathcal{T}$. We will assume that the complexity of E has been minimized over all choices of E.

Observe that the arc α of ∂E can't lie entirely in Γ_i^{top}. Otherwise E would be a boundary compressing disk in the complement of the 1–skeleton of the normal surface Γ_i^{top} which contradicts Claim 23. Our strategy will be to show that there is a sequence of proper isotopies of Σ_i^{top} and slides of ends of arcs of Σ_i^{top} over each other and over Γ_i^{top} (neither of which affect the isotopy class of $\Gamma_i^{\text{top}} \cup \Sigma_i^{\text{top}}$) so that afterwards α is incident to a single edge z in Σ_i^{top}, α runs along this edge once and E lies entirely inside a single tetrahedron. Then $\Gamma_i^{\text{top}} \cup \partial(N(z))$ is the required almost normal surface obtained from L by compressing all other tubes of Σ_i^{top}. Now that we have established some notation, for the rest of the proof we will consider the intersections of the disk E with the 2–skeleton of $\mathcal{T}$ and we will show that any intersection violates the minimality of (a, b).

When we consider the arcs of intersection between E and the 2–skeleton, we can get four types of components of $E \cap \mathcal{T}^2$ in E. See Figure 3. Components of Type I are simple closed curves in E. Components of Type II are arcs with both endpoints on α. Components of Type III are arcs with both endpoints on β and components of Type IV are arcs with one endpoint on α and the other endpoint on β. Next we will describe how each type of component of intersection of $E \cap \mathcal{T}^2$ can be removed, violating minimality of (a, b).

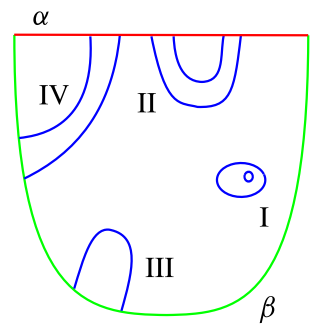

Figure 3: The disk E

Removing components of Type I Components of Type I are simple closed curves in E. A component of Type I that is innermost in a 2–simplex of $\mathcal{T}^2$ can be removed by substituting the disk it bounds in $\mathcal{T}^2$ for the disk it bounds in E. This reduces the number of times that E meets the 2–skeleton, thus reducing b and contradicting the minimality of E.

Removing components of Type II A component of Type II corresponds to an arc in a face σ of some tetrahedron of $\mathcal{T}$ that either has both endpoints on distinct components of $(\Sigma_i^{\text{top}} \cup \Gamma_i^{\text{top}}) \cap \sigma$, or has both endpoints on the same component of $(\Sigma_i^{\text{top}} \cup \Gamma_i^{\text{top}}) \cap \sigma$. Suppose γ is an arc of intersection between $\mathcal{T}^2$ and E that is outermost in E and is of Type II. See Figure 4(a). Let δ be a subarc of α such that $\gamma \cup \delta$ is the boundary of a disk D in $E - \mathcal{T}^2$. In what follows we will use edge slides of Σ_i^{top} to remove components of $\mathcal{T}^2 \cap E$ and $\mathcal{T}^2 \cap \Sigma_i^{\text{top}}$. Recall that t_i^{top} is the boundary of a neighborhood of Σ_i^{top} and we will abuse notation and consider $\delta \subset \alpha$ as an arc on Σ_i^{top} when we really mean that δ is an arc on t_i^{top}.

Any two ends of edges of Σ_i^{top} that meet the same normal disk in Γ_i^{top} can be isotoped together so that there is at most one edge incident to each normal disk. Since Σ_i^{top} can be extended to give a spine of W, any cycle in Σ_i^{top} gives a cycle in Σ. Since $W \cup_S W'$ is an irreducible Heegaard splitting of M it follows from Frohman [5] (also see [12, Proposition 2.5]) that no cycle of $\Sigma_{(W,K)}$ lies in a 3–ball. Hence for any tetrahedron H of $\mathcal{T}$, $\Sigma_i^{\text{top}} \cap H$ cannot contain a cycle. Thus $\Sigma_i^{\text{top}} \cap H$ is a union

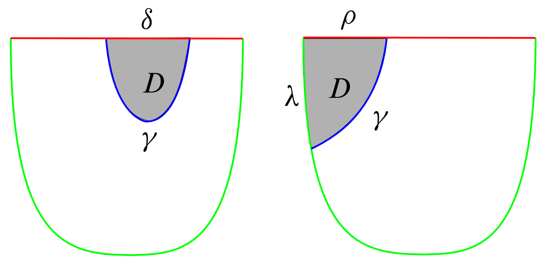

Figure 4: Arcs of Type II and IV in E

of trees for each tetrahedron H in $\mathcal{T}$. Each component of $\Sigma_i^{\text{top}} \cap H$ is a tree and each component of $(\Sigma_i^{\text{top}} \cup \Gamma_i^{\text{top}}) \cap H$ is a tree with disks attached and so is simply connected.

Case 1 The arc γ has both endpoints on distinct components of $(\Sigma_i^{\text{top}} \cup \Gamma_i^{\text{top}}) \cap \sigma$.

Case 1.1 Let H be the tetrahedron in $\mathcal{T}$ that contains δ and let q denote the component of $\Sigma_i^{\text{top}} \cap H$ that contains δ. If q is a single arc with both endpoints on the same face of H then D describes an isotopy that removes two points of intersection of q with $\mathcal{T}^2$. This reduces the number of points of $\Sigma_i^{\text{top}} \cap \mathcal{T}^2$ to which α is incident, thus reducing a, which is a contradiction.

Case 1.2 Now suppose that q is not an arc and δ is a path in q that begins at point x in q that is not in a normal disk but is in some face σ of H. See Figure 5. Let z be the edge of Σ_i^{top} containing x. Then δ describes a series of edge slides of z which culminate by introducing an extra point of intersection between Σ_i^{top} and σ. However, after the edge slides the disk D runs only over the edge z. Hence we can reduce the number of intersections of edges of Σ_i^{top} incident to α that meet $\mathcal{T}^2$ by two as in Case 1.1 reducing a by at least one and contradicting the minimality assumptions.

Case 1.3 If δ is an arc on L that has both endpoints of $\mathcal{T}^2 \cap E$ on normal disks of L, then either δ must run over some edges of Σ_i^{top} or it lies in a normal disk. If δ lies on a normal disk in H then it bounds a subdisk D' of the normal disk. Together D and D' bound a 3–ball in H that can be used to isotope E into the next tetrahedron removing γ and reducing complexity.

So we can assume that δ runs over some edges of Σ_i^{top}. Say that δ runs from normal disk D_1 to normal disk D_2. Since Σ_i^{top} is incident to D_1 in only a single point, δ is incident to ∂D_1 in a single point. It follows that $D_1 - N(\Sigma_i^{\text{top}})$ is an annulus and δ intersects the annulus in a single spanning arc. Thus δ runs precisely once along the

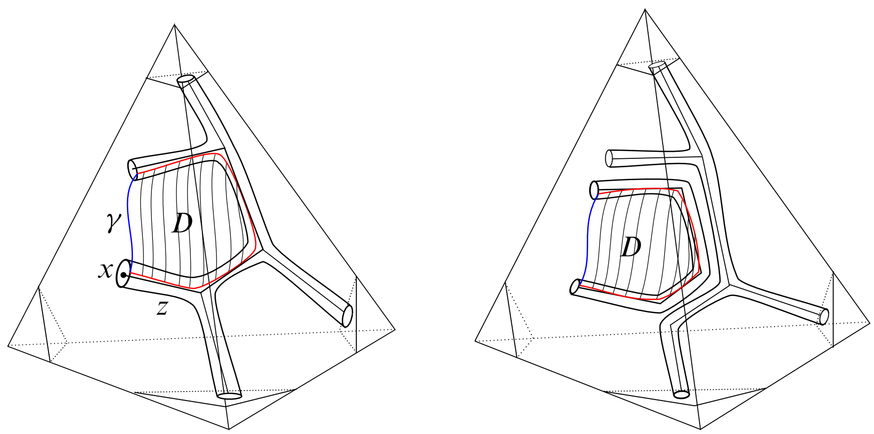

Figure 5

edge z that is incident to D_1. Then, as above, D describes a slide and isotopy of z that carries it to the arc γ in a simplex of $\mathcal{T}^2$. But then a subdisk of that face describes a parallelism between z and a subarc of $\mathcal{T}^1$. In particular, attaching a tube to Γ_i^{top} along z gives an almost normal surface.

Case 2 The arc γ has both endpoints on the same component of $(\Sigma_i^{\text{top}} \cup \Gamma_i^{\text{top}}) \cap \sigma$.

In this case the arc γ has both endpoints on the same component of $q \cap \partial H$. Let x denote the endpoint of q in the face σ of tetrahedron H. We assume that $x \in \Sigma_i^{\text{top}}$, so γ forms a loop based at x in σ bounding a disk A in σ; the case where both ends of γ lie on a normal disk is similar.

Case 2.1 If $\text{int}(A) \cap \Sigma_i^{\text{top}} = \varnothing$, then construct new disks E' and E'' by cutting the disk D along γ and attaching a copy of A to each piece. One of the disks E' or E'' will still be a lower disk and it will meet $\mathcal{T}^2$ in fewer components than E, contradicting the minimality of E.

Case 2.2 Now suppose that $A \cap \Sigma_i^{\text{top}} \neq \varnothing$. Since Σ_i^{top} is a union of trees in H, we know that a neighborhood of each component q of $\Sigma_i^{\text{top}} \cap H$ is a 3–ball. So there is a disk D' in $N(q)$ whose boundary is the union of δ and a diameter δ' of a small disk ϵ with which $N(q)$ meets ∂H at x.

Isotope the leaf L by compressing δ to δ' via the disk D' in $N(q)$, splitting the disk ϵ in two. See Figure 6. The effect on the spine is a possibly complicated series of edge slides. The overall effect is that the number of components of $\Sigma_i^{\text{top}} \cap \sigma$ increases by one when ϵ splits and $D \cup D'$ becomes a disk disjoint from Σ_i^{top} and parallel to A.

The disk A now contains at least two points of $\Sigma_i^{\mathrm{top}} \cap \sigma$. Now push $D \cup D'$ across A to remove γ, thus reducing b which is a contradiction.

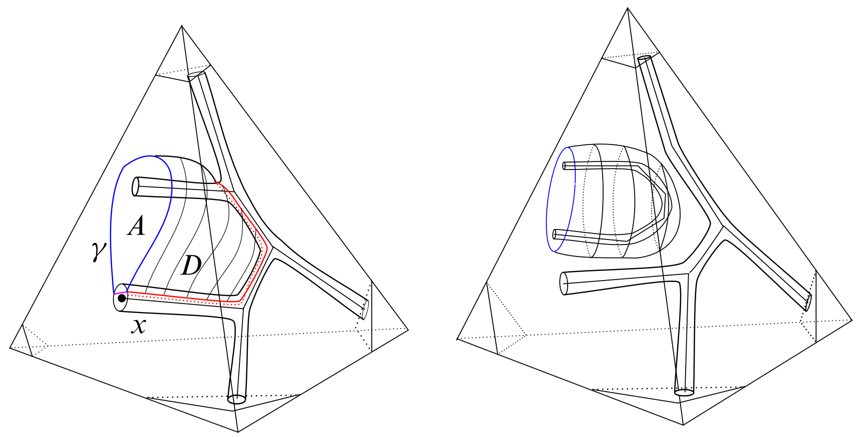

Figure 6

Removing components of Type III and IV To see how to remove components of Types III and IV it will be helpful to view the arc β that runs along the edge e of $\mathcal{T}^1$ as an arc that lies on $\partial N(e)$. As an arc on $\partial N(e)$, β may wind around the edge e. If the winding is not monotone a priori then we can reduce the number of components in which the disk E meets the faces of $\mathcal{T}^2$, contradicting minimality. Thus we may assume that the curve β winds monotonically around the edge e, creating a "barber pole" effect shown in Figure 7 and Figure 8. This implies that there are no curves of Type III since the existence a curve of Type III means that the arc β must 'double back' as it winds around e, contradicting monotonicity.

Let γ be an outermost arc component of Type IV, and let D be the corresponding outermost sub-disk of E. Let σ be the face of $\mathcal{T}^2$ that contains γ. See Figure 4(b). The disk D is co-bounded by a sub-arc ρ of α, a sub-arc λ of β, and γ. Let H denote the tetrahedron containing D in its interior and with σ as a face.

There are two cases that we will consider separately. The first case is when the arc γ of $E \cap \sigma$ that runs from the edge e to L ends on a normal disk η of L. The second case is when γ ends on a tube (neighborhood of Σ_i^{top}) of L.

Case 3.1 Suppose first that γ ends on a normal disk η. See Figure 7. In this situation there are two subcases. Either $\Sigma_i^{\mathrm{top}} \cap \rho = \varnothing$ or $\Sigma_i^{\mathrm{top}} \cap \rho \neq \varnothing$.

Case 3.1.1 Suppose that $\Sigma_i^{\mathrm{top}} \cap \rho = \varnothing$. In this case the arc ρ runs over only normal disks and does not meet any tubes of Σ_i. See Figure 7(a). Observe that there is a disk

D' in $\partial N(e)$ whose boundary consists of 3 arcs. The first arc is the subarc λ of β. The second arc is an arc on $\partial N(e)$ that corresponds to where the face σ meets $\partial N(e)$. The third arc is a subarc of a meridian of $\partial N(e)$ that corresponds to where η meets $\partial N(e)$. In this case $D' \cup \eta \cup D \cup \sigma$ bounds a 3–ball in H that we can use to isotope D across σ and into the next tetrahedron. If no components of Σ_i^{top} are contained in the 3–ball then isotoping D across σ and into the next tetrahedron will remove γ and reduce b, thus reducing the complexity of E, giving a contradiction. If components of Σ_i^{top} are contained in the 3–ball then the overall effect of isotoping D into the next tetrahedron will remove γ thus reducing b and will not increase the number of components of $\Sigma_i^{\text{top}} \cap \mathcal{T}^2$, and again the complexity of E has been reduced giving a contradiction.

Case 3.1.2 Suppose now that $\Sigma_i^{\text{top}} \cap \rho \neq \varnothing$. See Figure 7(b). Since Σ_i^{top} is a union of trees, each component of $N(\Sigma_i^{\text{top}})$ is a 3–ball. In particular, there is a disk Δ in $N(\Sigma_i^{\text{top}})$ whose boundary is the union of a sub-arc of ρ and a diameter d of the disk ϵ with which $N(\Sigma_i^{\text{top}})$ intersects the normal disk η. Isotope $N(\Sigma_i^{\text{top}})$ by compressing d to ρ in $N(\Sigma_i^{\text{top}})$, splitting the disk ϵ in two. The effect on Σ_i^{top} is a series of edge slides that results in a new component of $\Sigma_i^{\text{top}} \cap H$ that is on the same side of ρ as e. We can repeat this process until we have removed all components of $\Sigma_i^{\text{top}} \cap \rho$. Now proceed as in Case 3.1.1.

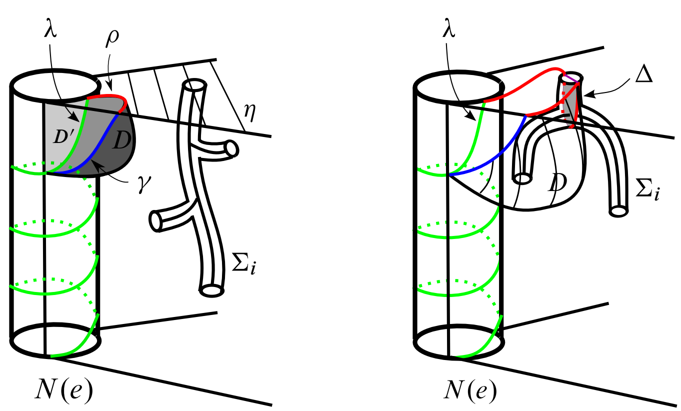

(a) (b)

Figure 7: Arcs of type IV

Case 3.2 Now suppose that γ ends on a tube of L. See Figure 8. The core of this tube is an edge τ that may connect to other edges of Σ_i^{top} in $\Sigma_i^{\text{top}} \cap H$, and Σ_i^{top} connects to a normal disk η. We will describe in two steps a slide of τ and an isotopy of D that will remove a component of $\Sigma_i^{\mathrm{top}} \cap E$, reducing a, and thereby reducing the complexity of E.

First, since τ connects to other edges of Σ_i in $\Sigma_i \cap H$, ρ describes an edge slide of τ that keeps $\tau \cap \sigma$ fixed but slides the opposite end of the edge τ off of Σ_i and onto the normal disk η. See Figure 8(b). We continue to slide τ along η following ρ until it almost meets $\partial N(e)$. Now we can use the disk D to isotope all of τ until it lies close to $\lambda \cup \gamma$. See Figure 8(c). At this point the entire disk D and tube τ lie very close to $\beta \cup \gamma$ in H.

For the second step recall the disk D' in $\partial N(e)$ that is bounded by λ, a copy of part of the edge e that bounds the face σ and a copy of a meridian of $\partial N(e)$. Together the disks D and D' describe an isotopy of τ across the face σ and into the next tetrahedron, removing the component γ from $\sigma \cap E$ and, in particular, removing the component of intersection between τ and σ in $\mathcal{T}^2$, reducing a, which is a contradiction. See Figure 8(d). Thus there can be no arcs of Type IV. Therefore the arc β of $\mathcal{T}_i^1$, the edges of Σ_i^{top} that α runs along and the disk E are all contained in one tetrahedron. An argument similar to the one given in the last paragraph of Case 1 above shows that α is incident to a single edge in Σ_i^{top} and α runs along this edge exactly once. Attaching a tube to Γ_i^{top} along this edge and compressing all other tubes of Σ_i^{top} in L gives an almost normal surface. By Lemma 16 compressing L to one side gives a surface that is incompressible to the opposite side. Therefore all compressions must be to one side of L, namely all compressions are either above or below. □

The third possibility is that there is no thick region and each arc of $\mathcal{T}_i^1$ has one endpoint on Γ_i^{top} and the other endpoint on Γ_i^{bot}.

Lemma 25 *If there is no thick region for $\mathcal{T}_i^1$ in M_i and each arc of $\mathcal{T}_i^1$ has one endpoint on Γ_i^{top} and has the other endpoint on Γ_i^{bot} then M_i is a product region.*

Proof Recall that $\partial M_i = \Gamma_i^{\mathrm{top}} \cup \Gamma_i^{\mathrm{bot}}$. Since Γ_i^{top} and Γ_i^{bot} are normal with respect to $\mathcal{T}$ it follows that there are two possibilities for how the region M_i between Γ_i^{top} and Γ_i^{bot} can intersect a face of the 2–skeleton. Either the region bounded by Γ_i^{top} and Γ_i^{bot} is a trapezoid region or a hexagon region. See Figure 9.

Suppose that there is a hexagon region of $M_i \cap \mathcal{T}^2$. Then three edges of the hexagon are arcs of $(\Gamma_i^{\mathrm{top}} \cup \Gamma_i^{\mathrm{bot}}) \cap \mathcal{T}^2$ and the other three edges are arcs of $\mathcal{T}_i^1$ connecting the three components of $\Gamma_i \cap \mathcal{T}^2$. But this implies that some arc of $\mathcal{T}_i^1$ connects either Γ_i^{top} to Γ_i^{top} or Γ_i^{bot} to Γ_i^{bot} which is a contradiction. Therefore there cannot be any hexagonal components and all regions of intersection between M_i and $\mathcal{T}^2$ are trapezoids.

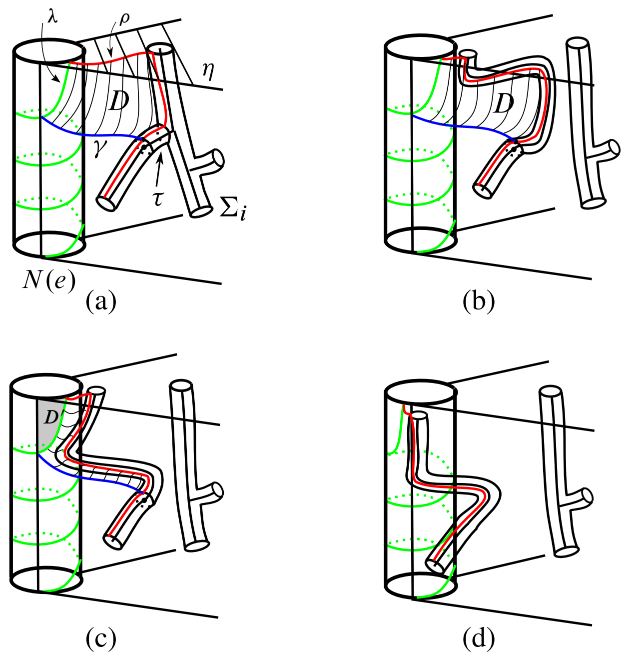

Figure 8: Arcs of type IV

Because we know that there are no hexagonal components of intersection between M_i and $\mathcal{T}^2$ this implies that the only possibilities for components of intersection between M_i and the tetrahedra of the 3–skeleton are triangular product regions and quadrilateral product regions. Each triangular and quadrilateral product region is bounded on one side by a normal disk of Γ_i^{top} and on the other by a normal disk of Γ_i^{bot}. Since each component of $M_i \cap H$ is a product region with one end on each of Γ_i^{top} and Γ_i^{bot} for each tetrahedron H in $\mathcal{T}$ we can conclude that M_i is itself such a product region. □

We can now complete the proof of Main Theorem 1. We will prove the theorem by describing a recursive process that will end when it produces an almost normal surface isotopic to the bridge surface S_K. Recall that we began with a knot K in a closed 3–manifold M with the assumptions that M and M_K are irreducible. We foliated M_K by copies of the bridge surface S_K with two singular leaves and triangulated M_K so that the annuli Γ are normal and the vertices of $\mathcal{T}$ are to one side of Γ. Cutting along a maximal family of non-parallel normal 2–spheres tubed to the normal annuli Γ we obtained the submanifold M_0 of M_K. The surface S_K induces a splitting of M_0 into K_0–compression bodies W_0 and W_0' and M_0 is foliated by copies of the bridge surface S_K and where the top (resp. bottom) leaf of the foliation is given by

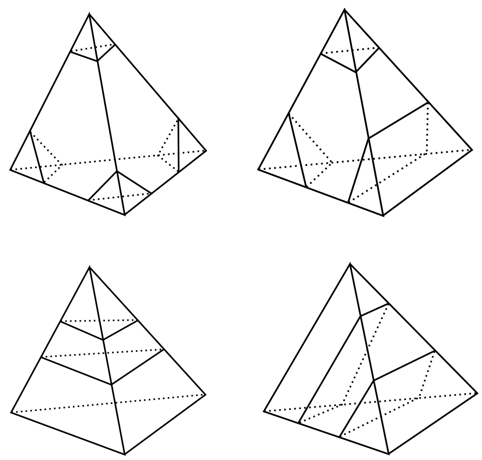

Figure 9

the union of the spine Σ_0^{top} (resp. Σ_0^{bot}) of W_0 (resp. W_0') and Γ_0^{top} (resp. Γ_0^{bot}). Here $\Gamma_0^{\text{top}} = \Gamma^{\text{top}'}$ (resp. $\Gamma_0^{\text{bot}} = \Gamma^{\text{bot}'}$) are the normal annuli in ∂M_0. The triple $(M_0, \Sigma_0, \Gamma_0)$ is the beginning of the recursive process. Each later step will produce a triple $(M_i, \Sigma_i, \Gamma_i)$ such that $M_i \subset M_{i-1}$ and for each i the surface S_K is a weakly incompressible splitting surface for M_i separating it into two K_i–compression bodies W_i and W_i', where $K_i = K \cap M_i$. The spine Σ_i^{top} (resp. Σ_i^{bot}) of W_i (resp. W_i') is contained in some spine for W (resp. W') and $\Gamma_i = \partial M_i - \partial M$ is a pair of collections of normal surfaces Γ_i^{top} and Γ_i^{bot} in ∂M_i.

The top (bottom) leaf of a singular foliation F_i is given by the union of Γ_i^{top} (resp. Γ_i^{bot}) and the intersection $\Sigma_i^{\text{top}} \subset \Sigma^{\text{top}}$ with M_i (resp. $\Sigma_i^{\text{bot}} \subset \Sigma^{\text{bot}}$ with M_i). Put $\mathcal{T}_i^1$, the part of $\mathcal{T}^1$ lying in $M_i - \partial M$, in thin position with respect to F_i. As mentioned earlier, either the arcs of $\mathcal{T}_i^1$ all have one endpoint on Γ_i^{top} and one endpoint on Γ_i^{bot}; or there is some arc that either has both endpoints on Γ_i^{top} or both endpoints on Γ_i^{bot}. If there is a thick region of $\mathcal{T}_i^1$ in M_i then we are in a position to apply Lemma 17. Otherwise we are in a position to apply either Lemma 24 or Lemma 25.

We will describe the step that takes us from $(M_i, \Sigma_i, \Gamma_i)$ to $(M_{i+1}, \Sigma_{i+1}, \Gamma_{i+1})$. First consider the initial step. If at the first step we encounter a thick region in M_0, then start with a leaf L_0 in a thick region of F_0 intersecting $\mathcal{T}^2$ in normal arcs and simple closed curves as is guaranteed by Claim 18. Applying Lemma 17 we obtain a collection G_0 of normal surfaces and at most one almost normal surface obtained by compressing L_0 to one side. If G_0 contains an almost normal surface and L_0 is incompressible above and below then $G_0 = L_0$ is an almost normal surface isotopic to a leaf and we are done. If G_0 does not contain an almost normal surface isotopic to L_0 then without loss of generality let $G_0 = \Gamma_1^{\text{top}}$ and proceed as below.

Henceforth assume without loss of generality that L_0 compresses above to give G_0. Otherwise we can invert the picture and declare Γ_i^{bot} to be the "top" leaf. Since G_0 has been obtained by compressing above, Lemma 16 implies that G_0 is incompressible below. By Lemma 15 we can isotope the almost normal surface G_0 to be normal. This gives a new collection Γ_1^{top} of normal surfaces isotopic to G_0. Cut M_0 along the collection Γ_1^{top} and keep the component to the incompressible side below Γ_1^{top} that contains part of ∂M_K. Call this submanifold M_1. Observe that $\Gamma_1^{\text{top}} \subset \partial M_1$. The cores of the tubes of the thick leaf that were compressed to give the almost normal surface $G_0 \simeq \Gamma_1^{\text{top}}$ form the required 1–complex Σ_1^{top}. Let Σ_1 denote the pair Σ_1^{top} and $\Sigma_1^{\text{bot}} = \Sigma_0^{\text{bot}}$ and let Γ_1 denote the pair Γ_1^{top} and $\Gamma_1^{\text{bot}} = \Gamma_0^{\text{bot}}$. This completes the first step.

The remainder of the proof falls into the following three cases.

Case 1 M_i contains a thick region of $\mathcal{T}_i^1$ with respect to F_i.

In this case using Claim 18 start with a leaf L_i in a thick region of the foliation F_i intersecting $\mathcal{T}^2$ in normal arcs and simple closed curves disjoint from the 1–skeleton. Applying Lemma 17 we obtain a collection G_i of normal surfaces and at most one almost normal surface obtained by compressing L_i to one side. Lemma 16 implies that G_i is incompressible to the opposite side. If L_i is incompressible then $G_i = L_i$ and since L_i is isotopic to a leaf we are done. So suppose without loss of generality that L_i is compressible above to give G_i. The cores of the tubes of L_i that are compressed above to give G_i will make up the spine $\Sigma_{i+1}^{\text{top}}$. Let Σ_{i+1} denote the pair $\Sigma_{i+1}^{\text{top}}$ and $\Sigma_{i+1}^{\text{bot}} = \Sigma_i^{\text{bot}}$. By Lemma 15 we can isotope the almost normal surface G_i to give a new collection $\Gamma_{i+1}^{\text{top}}$ of normal surfaces. Cut M_i along the collection $\Gamma_{i+1}^{\text{top}}$ and keep the component to the incompressible side below $\Gamma_{i+1}^{\text{top}}$ that contains part of ∂M_K. Call this new submanifold M_{i+1}. Let Γ_{i+1} denote the pair $\Gamma_{i+1}^{\text{top}}$, $\Gamma_{i+1}^{\text{bot}} = \Gamma_i^{\text{bot}}$.

If on the other hand G_i is compressible below then the cores of the tubes of L_i that are compressed below to give G_i will make up the spine $\Sigma_{i+1}^{\text{bot}}$. By Lemma 15 we can isotope G_i to be normal and call the new collection of normal surfaces $\Gamma_{i+1}^{\text{bot}}$. Let Γ_{i+1} denote the pair $\Gamma_{i+1}^{\text{top}} = \Gamma_i^{\text{top}}$ and $\Gamma_{i+1}^{\text{bot}}$. Cut M_i along the collection of normal surfaces $\Gamma_{i+1}^{\text{bot}}$ and keep the component to the incompressible side above $\Gamma_{i+1}^{\text{bot}}$. Call this new submanifolds M_{i+1}. This completes the recursive step in this case.

Case 2 M_i contains no thick region of $\mathcal{T}_i^1$ and some arc of $\mathcal{T}_i^1$ either has both endpoints on Γ_i^{top} or has both endpoints on Γ_i^{bot}.

Without loss of generality suppose that there is an arc of $\mathcal{T}_i^1$ with both endpoints on Γ_i^{top}. Applying Lemma 24, starting with a leaf L_i of F_i near the top singular leaf above all of the minima we obtain an almost normal surface G_i in M_i by compressing the leaf L_i above. It follows from Lemma 16 that G_i is incompressible below. Moreover, $\chi(G_i) = \chi(\Gamma_i^{\text{top}}) - 2$. Using Lemma 15 isotope the almost normal surface G_i to give a normal surface $\Gamma_{i+1}^{\text{top}}$. Cut M_i along $\Gamma_{i+1}^{\text{top}}$ and keep the component to the incompressible side below $\Gamma_{i+1}^{\text{top}}$ to obtain the submanifold M_{i+1}. Denote by Γ_{i+1} the pair $\Gamma_{i+1}^{\text{top}}$ and $\Gamma_{i+1}^{\text{bot}} = \Gamma_i^{\text{bot}}$. The spine $\Sigma_{i+1}^{\text{top}}$ of M_{i+1} consists of the cores of the tubes of L_i that are compressed above to give G_i. Denote by Σ_{i+1} the pair $\Sigma_{i+1}^{\text{top}}$ and $\Sigma_{i+1}^{\text{bot}} = \Sigma_i^{\text{bot}}$.

In both Cases 1 and 2 the new surface G_i isotopic to $\Gamma_{i+1}^{\text{top}}$ (resp. $\Gamma_{i+1}^{\text{bot}}$) in M_i is not parallel as a normal surface to the normal surfaces Γ_i^{top} (resp. Γ_i^{bot}). The reason depends on whether Lemma 17 or Lemma 24 was applied. If the surface, without loss of generality say $\Gamma_{i+1}^{\text{top}}$, comes from compressing a thick leaf via Lemma 17 then there is a subarc of $\mathcal{T}^1$ lying between Γ_i^{top} and $\Gamma_{i+1}^{\text{top}}$ with both ends on $\Gamma_{i+1}^{\text{top}}$. Hence Γ_i^{top} and $\Gamma_{i+1}^{\text{top}}$ are not parallel. If $\Gamma_{i+1}^{\text{top}}$ comes via Lemma 24 then $\chi(\Gamma_{i+1}^{\text{top}}) = \chi(\Gamma_i^{\text{top}}) - 2$ so the surfaces are not parallel. If Γ_i^{top} and Γ_j^{top} are parallel then all leaves Γ_k^{top} where $i \le k \le j$ are parallel as well. In particular, then $\Gamma_{i+1}^{\text{top}}$ is parallel to Γ_i^{top} which cannot happen as we have just seen above. Therefore it follows that Γ_i^{top} is non-parallel to Γ_j^{top} for all $i < j$.

Case 3 M_i contains no thick region of $\mathcal{T}_i^1$ and each arc of $\mathcal{T}_i^1$ has one endpoint on Γ_i^{top} and one endpoint on Γ_i^{bot}.

In this case by Lemma 25 M_i is a product. Suppose $i \neq 0$. The submanifold $\overline{M_i} = M_i \cup N(K_i)$ is a product as well and has the surface S as a Heegaard surface that gives an irreducible Heegaard splitting of $\overline{M_i}$. By [12] it follows that the splitting surface is isotopic to Γ_i^{top} and Γ_i^{bot}, one of which is in turn isotopic to the almost normal surface G_{i-1} and so we are done.

If $i = 0$ then the argument above shows that the surface S_K consists of a collection of annuli. However the surface S_K is a bridge surface for K so it is connected. Therefore S_K consists of one annulus and K must be the unknot.

It follows from a well known result of Haken that there are only a finite number of non-parallel, disjoint, normal surfaces in M_K. See Haken [7]. Therefore we will only have to apply Lemma 24 and Lemma 17 a finite number of times before we either reach a situation where we apply Lemma 25 and obtain an almost normal surface isotopic to the bridge surface S_K or we exhaust all of the non-parallel, disjoint, normal surfaces in M_K and we obtain an almost normal surface isotopic to S_K. □

References

[1] **D Bachman**, *Heegaard splittings with boundary and almost normal surfaces*, Topology Appl. 116 (2001) 153–184 MR1855961

[2] **D Bachman**, *Thin position with respect to a Heegaard surface*, preprint

[3] **A J Casson**, **C M Gordon**, *Reducing Heegaard splittings*, Topology Appl. 27 (1987) 275–283 MR918537

[4] **A Coward**, *Algorithmically detecting the bridge number of hyperbolic knots* (2006) arXiv:0710.1262

[5] **C Frohman**, *The topological uniqueness of triply periodic minimal surfaces in* $\mathbf{R}^3$, J. Differential Geom. 31 (1990) 277–283 MR1030674

[6] **D Gabai**, *Foliations and the topology of* 3*-manifolds. III*, J. Differential Geom. 26 (1987) 479–536 MR910018

[7] **W Haken**, *Theorie der Normalflächen*, Acta Math. 105 (1961) 245–375 MR0141106

[8] **W Jaco**, **D Letscher**, **H Rubinstein**, *One vertex, ideal, and efficient triangulations of 3–manifolds*, in preparation

[9] **W Jaco**, **J H Rubinstein**, *0–efficient triangulations of 3–manifolds*, J. Differential Geom. 65 (2003) 61–168 MR2057531

[10] **S King**, *Almost normal Heegaard surfaces* (2003) arXiv:math.GT/0303377

[11] **J H Rubinstein**, *Polyhedral minimal surfaces, Heegaard splittings and decision problems for 3–dimensional manifolds*, from: "Geometric topology (Athens, GA, 1993)", AMS/IP Stud. Adv. Math. 2, Amer. Math. Soc., Providence, RI (1997) 1–20 MR1470718

[12] **M Scharlemann**, **A Thompson**, *Heegaard splittings of* (surface) $\times I$ *are standard*, Math. Ann. 295 (1993) 549–564 MR1204837

[13] **M Scharlemann**, **M Tomova**, *Uniqueness of bridge surfaces for 2–bridge knots*, Math. Proc. Cambridge Philos. Soc. 144 (2008) 639–650 MR2418708

[14] **M Stocking**, *Almost normal surfaces in 3–manifolds*, Trans. Amer. Math. Soc. 352 (2000) 171–207 MR1491877

[15] **A Thompson**, *Thin position and the recognition problem for* S^3, Math. Res. Lett. 1 (1994) 613–630 MR1295555

[16] **M Tomova**, *Multiple bridge surfaces restrict knot distance*, Algebr. Geom. Topol. 7 (2007) 957–1006 MR2336246

Department of Mathematics and Statistics, California State Polytechnic University Pomona, 3801 West Temple Ave, Pomona, CA 91768, USA

robinwilson@csupomona.edu

Received: 6 October 2007