

THERE ARE NO UNEXPECTED TUNNEL NUMBER ONE KNOTS OF GENUS ONE

MARTIN SCHARLEMANN

ABSTRACT. We show that the only knots that are tunnel number one and genus one are those that are already known: 2-bridge knots obtained by plumbing together two unknotted annuli and the satellite examples classified by Eudave-Muñoz and by Morimoto and Sakuma. This confirms a conjecture first made by Goda and Teragaito.

1. INTRODUCTION AND OVERVIEW

There are many useful ways of indexing the complexity of knot types: crossing number, bridge number, tunnel number, genus, etc. Often the relationship between these indices is unclear, and sometimes it is clear that there is no relationship. Thus, for example, tunnel number one knots may be of arbitrarily high genus (e.g. torus knots) and genus number one knots may be of arbitrarily high tunnel number (e.g. doubles of complicated knots.) Given two indices of complexity, it's natural to ask a sort of complementary question: how unusual is it for a knot (other than the unknot) to be of minimal complexity with respect to both indices? For example, how rare is it that a knot have both genus one and tunnel number one?

It's easy to construct examples of knots of this type. Plumb together two twisted unknotted annuli. The boundary is typically knotted and the union of the annuli is visibly a genus one Seifert surface. If we imagine hanging the union of annuli from a single peg, we see that its boundary is naturally a 2-bridge knot and therefore has tunnel number one. See Figure 1. It is known that these are the only 2-bridge knots of genus one (cf. [BZ, Proposition 12.25]).

Are there other examples of genus one tunnel number one knots? Morimoto and Sakuma [MS] and independently Eudave-Muñoz [EM] classified satellite knots which have tunnel number one. They have a concrete description and can be naturally indexed by a 4-tuple of integers. In [GT], Goda and Teragaito determined which of these satellite knots have genus one, and made the conjecture that these knots complete the list of knots that have both genus one and tunnel number one. The conjecture was confirmed by Matsuda [Ma] for any knot that admits a $(1, 1)$ decomposition; that is, for any knot which is 1-bridge on an unknotted torus.

The central objective of this paper is to prove the Goda-Teragaito conjecture in complete generality. (An overview of the proof, with much technical detail

Received by the editors July 24, 2001 and, in revised form, July 25, 2002.

2000 *Mathematics Subject Classification*. Primary 57M25.

This research was supported in part by an NSF grant, the Miller Institute, and RIMS Kyoto.

©2003 American Mathematical Society



FIGURE 1.

suppressed, can be found in [Sc2].) The strategy will be to use thin position to show that any tunnel for a genus one knot can either be isotoped to lie on a genus one Seifert surface, or isotoped to form an unknotted loop. In the latter case, it is shown that the knot admits a $(1, 1)$ decomposition, and Matsuda's argument applies. In the former case, it follows from work of Eudave-Muñoz and Uchida [EU] that K is a 2-bridge knot.

In retrospect, [GST] and [ST1] can be viewed as the first two steps of the program. In [GST] we show how thin position can be used to understand unknotting tunnels for tunnel number one knots. In particular, we show that if K and then λ are put in the thinnest possible position, then λ is level. If λ is a loop, then Matsuda's theorem applies. Otherwise, we show in [ST1], either K is 2-bridge or there is a well-defined invariant $\rho \in \mathbb{Q}/2\mathbb{Z}$ which, unless it is 1, ensures that the tunnel can be moved onto a minimal genus Seifert surface. So all that remains is to consider the case in which $\rho = 1$, which we do here.

In the case that $\rho = 1$ and the tunnel γ is not an unknotted loop, it will be shown that there is another useful way of describing K on the boundary of the genus 2 handlebody $H = \eta(K \cup \gamma)$. That is, there is a different spine for H , namely a Θ -curve θ , with these properties:

- The graph θ can be put in general position in S^3 in such a way that $K \subset \partial\eta(\theta)$ remains in thin position.
- K intersects each meridian of each edge of θ always with the same orientation.
- $\theta \subset S^3$ is thinner than the graph $K \cup \gamma$.
- A minimal genus Seifert surface F for $K \subset H$ intersects H only in $K = \partial F$.

A combinatorial argument will show that if K intersects each meridian more than once then $\text{genus}(F) \geq 2$. Assuming that $\text{genus}(F) = 1$ and γ cannot be isotoped to F , the program then will be to find the thinnest spine satisfying the conditions above (plus a more technical condition called the “wave condition”). For such a graph, we know that K intersects one of the meridians in only one point. We will argue, via thin position, that the cycle obtained by deleting this meridian

is unknotted. It will follow that K has a $(1, 1)$ decomposition, so Matsuda's result applies.

2. INTERSECTING (p, q) QUASI-CABLES WITH SPHERES

Consider a graph θ as just described. Notice that the condition on meridians of edges of θ can be interpreted as follows: K and each edge of θ can be oriented so that K always runs along a given edge in the direction consistent with that edge's orientation. In particular, K intersects these meridians algebraically as well as geometrically in some $p, q, p+q$ points. With this in mind, we establish the more general definition (and notation):

Definition 2.1. Suppose θ is a Θ -curve in S^3 with edges e^+, e^-, e^\perp . In $H = \eta(\theta)$, denote the corresponding meridians by μ^+, μ^-, μ^\perp . Suppose $K \subset \partial H$ is a primitive curve in ∂H (i.e., it intersects some essential disk in H in a single point) and K intersects each of the meridians μ^+, μ^-, μ^\perp always with the same orientation and so that some minimal genus Seifert surface F for K intersects H only in $K = \partial F$. Arrange the labelling and orientations of the edges and meridians so that, geometrically as well as algebraically,

- $K \cap \mu^- = q \geq 1$,
- $K \cap \mu^+ = p \geq q$,
- $K \cap \mu^\perp = p + q$.

Then we say that K (or (K, F)) is presented on θ as a (p, q) quasi-cable.

The fact that K is primitive ensures that p and q are relatively prime, so $p > q$ unless $p = q = 1$. Given p, q , there is a straightforward algorithm to describe the order in which K intersects the three meridians (see for example [OZ]): Consider a line in \mathbb{R}^2 of slope p/q that is disjoint from the lattice \mathbb{Z}^2 . Choose a segment σ that projects to a simple closed curve in the torus $\mathbb{R}^2/\mathbb{Z}^2$. Then the order in which σ intersects respectively lines of the form $y \in \mathbb{Z}, x \in \mathbb{Z}, x+y \in \mathbb{Z}$ is the order in which K intersects respectively the meridians μ^+, μ^-, μ^\perp . This has the useful corollary:

Corollary 2.2. Suppose K is presented on θ as a (p, q) quasi-cable, with $p > q \geq 2$. Then there are at least two arcs of $K - \{\mu^+, \mu^-\}$ that are oriented from μ^+ to μ^- (and of course two then oriented from μ^- to μ^+).

Proof. Since $q \geq 2$, the corresponding arc $\sigma \subset \mathbb{R}^2$ crosses at least two vertical lines $x \in \mathbb{Z}$. Since $p > q$, in between such crossings σ must cross at least one horizontal line $y \in \mathbb{Z}$. \square

In order to appreciate the point of Definition 2.1, it's useful to observe that any pair (K, F) can be presented as a (p, q) quasi-cable for any relatively prime non-negative pair (p, q) . Consider the following construction. There is a natural embedding of a punctured torus T_0 in S^3 with two properties:

- $K \subset T_0$,
- F is transverse to T_0 with $F \cap T_0 = \partial F$.

For example, $F - \eta(K)$ is a copy of F intersecting $\partial\eta(K)$ in a longitudinal copy of K ; just let T_0 be the complement of a disk in $\partial\eta(K)$.

In the punctured torus T_0 , choose two non-parallel normally oriented essential arcs σ^+ and σ^- . Once $K \cap (\sigma^+ \cup \sigma^-)$ is minimized by isotopy, K will intersect each arc σ^\pm , always with the same orientation. Indeed, given (p, q) non-negative and

relatively prime, it's easy to find such arcs and to choose their normal orientation so that $\sigma^+ \cdot K = p$ and $\sigma^- \cdot K = q$. One of the two choices σ^\perp for a third essential arc in T_0 that is not parallel to σ^+ or σ^- will have the property that $\sigma^\perp \cdot K = p + q$. Let H be the genus two handlebody obtained by thickening $T_0 \subset S^3$ slightly, so $H \cong (T_0 \times I)$. We can then regard H as the neighborhood of a Θ -graph $\theta \subset S^3$, with two vertices (one for each component of $T_0 - (\sigma^+ \cup \sigma^- \cup \sigma^\perp)$) and three edges e^-, e^+, e^\perp , each dual to its cognate arc. The natural meridians for H , namely $\mu^+ \cong \sigma^+ \times I$, $\mu^- \cong \sigma^- \times I$ and $\mu^\perp \cong \sigma^\perp \times I$, are the meridians required to give θ the structure that presents (K, F) as a (p, q) quasi-cable.

So if any pair (K, F) can be presented as a (p, q) quasi-cable, what is the point of the construction? The point will be to use general position between θ and a surface in S^3 as a shorthand way of describing in which meridians we will allow H and the surface to intersect. If we view H in this way, then any time the 1-complex θ is put in general position with respect to a surface $S \subset S^3$, it will automatically be true that $S \cap H$ is a collection of meridian disks, each parallel to one of μ^+, μ^-, μ^\perp .

The first lemma may clarify the point:

Lemma 2.3. *Suppose the pair (K, F) is presented as a (p, q) quasi-cable on θ with both $p > q \geq 2$. Suppose there is a sphere $P \subset S^3$ so that*

- θ is in general position with respect to P ,
- P intersects both of the cycles $e^\perp \cup e^+$ and $e^\perp \cup e^- \subset \theta$, and
- each component of the 1-manifold $P \cap F$ is essential in F .

Then $\text{genus}(F) \geq 2$.

Remark. One might conjecture that this is the simplest case of a more general result, perhaps describing how the continued fraction expansion of q/p determines a lower bound for the genus of F .

Proof. We have seen in Corollary 2.2 that $K \subset \eta(\theta)$ at least twice switches from crossing μ^+ (perhaps repeatedly) to crossing μ^- (and vice versa).

To simplify the number of cases we need to consider, note first that we can slit open the part of H corresponding to $e^\perp \subset \theta$, lengthening e^\pm while shortening e^\perp until e^\perp is so short that it is disjoint from P . So, with no loss of generality, we may as well assume that P intersects both of the segments e^\pm , but not e^\perp .

Once this is done, the components of $\theta - P$ consist of three types: subarcs of e^+ , subarcs of e^- , and a single component θ^\perp that contains e^\perp together with all four ends of the two edges e^\pm . Then P intersects each of the segments e^\pm in an even number of points. Let κ^+ denote the component of $e^+ - P$ that is exactly half-way along e^+ as measured by intersections with P . That is, an arc in e^+ starting from a point in κ^+ and ending in θ^\perp will intersect P in the same number of points no matter which way along e^+ it runs. Denote by κ^- the analogous point in e^- .

The knot $K \subset H$ is similarly split up into segments by P , some parallel to segments of $e^+ - P$, some to segments of $e^- - P$ and some lying on $\eta(\theta^\perp)$. Any subsegment of K that is a union of components of the first (resp. second) type will be said to be *colored +* (resp. *-*). Remembering that K is oriented, each segment of $K \cap \eta(\theta^\perp)$ can be described as one of three types:

- (1) components of $K - P$ that run from an end of e^+ to an end of e^- ,
- (2) components of $K - P$ that run from an end of e^- to an end of e^+ ,
- (3) components of $K - P$ that run from an end of e^+ to the other end of e^+ .




FIGURE 2.

(There are no components of $K \cap \eta(\theta^\perp)$ that run from e^- to e^- , since $p > q$.) We will say that these three types of segments of $K - P$ are colored $0_+^+, 0_-^+, 0_+^+$ respectively. The notation is meant to suggest that in clockwise rotation, the color changes from $+$ to $-$, etc.

Since all components of $F \cap P$ are essential in F , F cannot be a disk. We will suppose $genus(F) = 1$ and arrive at a contradiction.

There are various ways that the 1-manifold $F \cap P$ can lie in F . If any component is a closed curve, then the complement of the closed curve is a simple pair of pants (i.e., a 3-punctured sphere), so all arc components of intersection must be parallel to each other. If there are no closed curves, the arcs of $F \cap P$ fall into (at most) three classes of parallel arcs in F . We will assume for the purposes of this argument that $F \cap P$ consists of three such classes of parallel arcs in F ; if there are fewer classes of parallel arcs, the same method works, but more easily.

Abstractly, the three families of parallel arcs of $F \cap P$ in F give $K = \partial F$ the structure of a hexagon R , in which opposite sides are connected via arcs of intersection that are parallel in F . See Figure 2. Each end of such an arc of intersection lies in a meridian disk of H , corresponding to a point of $\theta \cap P$; if two ends of arcs of intersection lie in the same meridian, we say that the ends *have the same label*.

Claim 1. *Opposite ends of the same arc of $F \cap P$ cannot have the same label.*

Proof of Claim 1. Since $K = \partial F$ always crosses each meridian with the same orientation, a normal orientation induced on the intersection arc by a normal orientation of $P \subset S^3$ would have to have opposite direction at each end of the intersection arc. □

Claim 2. *Suppose α_1 and α_2 are intersection arcs parallel in F , connecting opposite sides s_1 and s_2 of the hexagon R . Suppose further that the labels of α_1 at s_1 and α_2 at s_2 are the same. Then both are $+$ -labels, and all labels lying between the ends of the α_i on one of the s_i are $+$ -labels. On the other side, between the ends of the α_i , there is exactly one subsegment of $-$ -labels.*

Proof of Claim 2. If there were a counterexample, choose α_1 and α_2 to be as close as possible (among parallel arcs of intersection in F) among all such counterexamples.




FIGURE 3.

They cannot be the same intersection arc, by Claim 1. We now show they cannot be adjacent intersection arcs in F . For if they were, then the segments of s_1 and s_2 that lie between them would correspond to parallel segments of $K - P$ on H . This is obvious unless the component of $\theta - P$ on which the segments of s_i lie is θ^\perp , i.e., the segments are colored 0. But even if the segments are colored 0, then the fact that this is a counterexample forces both segments to be colored 0^- , or both 0^+ or both 0_+^+ , so they are in fact parallel on $\partial H - P$. Now follow a standard argument that traces its origins to [GL] or [Sc]: Consider the union of the sphere P , the single 1-handle subsection of H corresponding to the segments in the s_i between the intersection arcs, and a 2-handle whose core is the rectangle in F lying between the two intersection arcs. (See Figure 3.) This defines a Heegaard splitting of a punctured lens space $L(2, 1) = RP^3$, lying in S^3 , clearly an impossibility.

Since α_1 and α_2 are not adjacent, we can consider the intersection arcs α'_1 and α'_2 adjacent to α_1 and α_2 but closer together. Since the intersection arcs α'_1 and α'_2 are not a counterexample, either the α'_i have different labels at s_i , $i = 1, 2$, or they have the same labels but the count of segments colored 0^- or 0^+ between them changes. Either outcome is only possible if the segments k_i between α_i and α'_i on s_i are colored 0, for both $i = 1, 2$ and at least one, say k_1 , is colored 0^- or 0^+ , say 0^- . Then k_2 is colored either 0_+^+ (and the labels of the α'_i at s_i are different for $i = 1, 2$) or k_2 is also colored 0^- . Consider first the former case, k_2 is colored 0_+^+ . This immediately implies that both of the initial labels were $+$ -labels. If another segment colored 0^- lies between α_1 and α_2 on s_2 , then an intersection arc adjacent to it would be a counterexample closer to α_1 than α_2 is. So we deduce that all labels on s_2 between the α_i are $+$ -labels. Because α_1 and α_2 are a counterexample, another segment colored 0^- must lie between α_1 and α_2 on s_1 . Then there is also one colored 0_+^+ between α'_1 and α'_2 . In order for an intersection arc adjacent to it and the intersection arc α'_2 not to be a counterexample, there must be a segment on s_1 even closer to α'_2 that is colored 0^- . Between it and the segment colored 0_+^+ lies every label colored $+$. Opposite to this segment on s_2 every label is colored $+$, since no 0^- or 0_+^+ color appears on s_2 between α_1 and α_2 . So for any of these intersection arcs, whatever the label is on s_2 , there's a parallel intersection arc with that label on s_1 , and between them lies no label 0^- or 0_+^+ . This creates a closer together pair of intersection arcs that are a counterexample, a contradiction. See Figure 4.

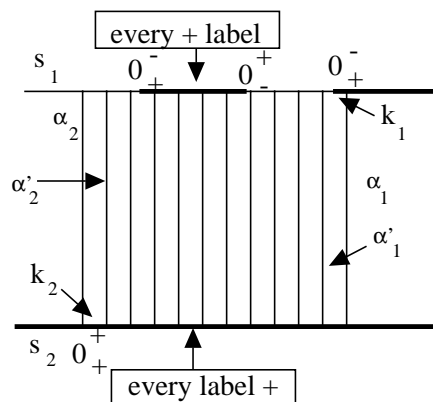


FIGURE 4.

If instead k_2 , like k_1 , is colored 0_+^- , then the labels of the α'_i at s_i are the same for $i = 1, 2$, and, since these arcs are not a counterexample, they must be $+$ -labels, all the labels on one side between the α'_i must be $+$ -labels, but on the other side there must be a switch to $-$ -labels. But this produces exactly the same contradiction as before: all the labels on one side are $+$ -labels and an entire sequence of $+$ -labels appears on the other side. Hence there can be no counterexamples, completing the proof of Claim 2. \square

Claim 3. *Segments corresponding to κ^+ (or κ^-) cannot appear on opposite sides of R (including corners).*

Proof of Claim 3. If κ^- appeared on opposite sides, it would immediately contradict Claim 2. Suppose κ^+ appeared on opposite sides of R . Choose those occurrences that are closest together (as measured by arcs between them) and let α_1 and α_2 denote the arcs adjacent to those occurrences of κ^+ that are closer together. Let $m = |e^+ \cap P|$. If the number of arcs between the α_i is less than $m/2$ it would contradict Claim 2, since neither side could then have any $-$ -labels. If the number of arcs between them is no less than $m - 2$, then the fact, from Claim 2, that one side consists entirely of $+$ -labels would ensure that another label κ^+ occurs even more closely to the opposite κ^+ , contradicting the choice of κ^+ segments that are closest together. Suppose finally that there are between $m/2$ and $m - 2$ arcs between them. Then, by Claim 2, on one side between them will be a segment colored 0_+^- or 0_-^+ and on the other a segment colored 0_+^+ . Moreover, the arcs α'_i adjacent to these segments and closer together would have the same $+$ -label but would have no $-$ -labels between them on either side, again contradicting Claim 2. \square

Claim 4. *Two segments, one corresponding to each of κ^\pm cannot occur on the same side of R (including corners).*

Proof of Claim 4. Let $n = |e^- \cap P|$. If both types of κ occur on one side, say s_1 , then that side is incident to at least $(m + n)/2$ arcs and contains a segment of type 0_+^- or 0_-^+ . Also, from Claim 3, the opposite side s_2 can be incident to neither type of κ interval. Since the s_i are incident to the same number of arcs, this means that s_2 contains some segment of the form 0_+^- , 0_-^+ or 0_+^+ , say 0_+^- . In fact, it must be of

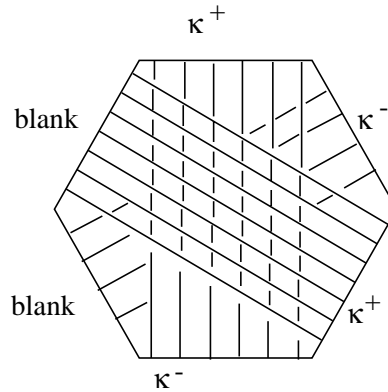


FIGURE 5.

the form 0_+^+ , since otherwise its having length $> (m+n)/2$ would force a κ segment to appear. For the same reason, the 0_+^+ segment in s_2 must be opposite a segment in s_1 that lies between a κ_+ label and the 0_-^+ label. The arcs α_i adjacent to the 0_-^+ and 0_+^+ segments and closer together would have the same $+$ -label but would have no $-$ -labels between them on either side, again contradicting Claim 2. \square

Following Claim 4, we can think of each side of R as either a $+$ -side, a $-$ -side or blank, depending on whether a copy of κ_+, κ_- or no κ at all appears. Moreover, by Claim 3 opposite sides can't both be $+$ -sides or both be $-$ -sides, and (a crucial point) following Corollary 2.2, at least 4-sides have signs, alternating around R as $+, -, +, -$. Combining these facts, the only possible signing of the sides of R is, in order, *blank, blank, +, -, +, -* (with some orientation of ∂R). Now the fact that the two adjacent blank sides have no sign whereas their adjacent sides have different signs means that the total number of arcs intersecting those two blank sides must be less than $(m+n)/2$. On the other hand, the adjacent sides opposite these blank sides have signs $+$ and $-$, which guarantees that their combined length is greater than $(m+n)/2$. The contradiction proves the lemma. See Figure 5. \square

The requirement that $p, q \geq 2$ in the above lemma is central to the proof, of course, since it guarantees the repetitions in patterns around ∂F that lead to the combinatorial contradiction. Nonetheless, there are important situations in which the results of Lemma 2.3 hold true even when $q = 1$.

The easiest example to see requires a preliminary construction. Suppose (K, F) is presented as a (p, q) quasi-cable on θ with regular neighborhood H , and consider the 4-punctured sphere Σ obtained from ∂H by removing copies of the meridians μ^\perp and μ^- . Then the boundary of Σ consists of two copies of $\partial\mu^\perp$ and two copies of $\partial\mu^-$, one on each side of the circle $\partial\mu^+ \subset \Sigma$. The arcs of $K \cap \Sigma$ are oriented to flow from one side of $\partial\mu^+$ to the other. The slopes of these arcs naturally define another circle $\mu^{\perp'} \subset \Sigma$ with the property that each arc of $K \cap P$ intersects $\mu^{\perp'}$ exactly once and $|\mu^{\perp'} \cap \mu^+| = 2$. See Figure 6. (Actually, there are two candidates for $\mu^{\perp'}$; the other is obtained by vertical reflection.) If we discard the meridian μ^+ of H and replace it with $\mu^{\perp'}$, then the associated Θ -graph θ' is one obtained from θ by a Whitney move. (K, F) is still presented as a quasi-cable on θ' , but now of




FIGURE 6.

type $(p + q, q)$. With the new structure, the old \mp is discarded, the old μ^- becomes also the new $\mu^{-'}$, and the old μ^\perp becomes the new $\mu^{+'}$.

There is a similar move in which H is cut up along μ^+ and μ^\perp , the meridian μ^- is discarded and the new graph θ' presents (K, F) as a $(p + q, p)$ quasi-cable.

Definition 2.4. The moves on θ just described, which change the presentation of (K, F) from that of a (p, q) quasi-cable to, respectively, a $(p + q, q)$ quasi-cable or a $(p + q, p)$ quasi-cable are called *standard Whitney moves* on e^+ and e^- respectively.

Lemma 2.5. *Suppose the pair (K, F) is presented as a (p, q) quasi-cable on θ with $p > q \geq 1$. Suppose there is a sphere $P \subset S^3$ so that*

- θ is in general position with respect to P ,
- P intersects both of the edges e^\perp and e^+ but is disjoint from e^- , and
- each component of the 1-manifold $P \cap F$ is essential in F .

Then $\text{genus}(F) \geq 2$.

Proof. Perform a standard Whitney move on e^- so that afterwards the new Θ -graph θ' presents (K, F) as a $(p + q, p)$ quasi-cable. Since P was disjoint from e^- , this has no effect on $P \cap F$, and, since $p + q \geq p \geq 2$, Lemma 2.3 applies. \square

In a related but more complicated case, the combinatorics is so close to that of an actual cable knot that the arguments are considerably easier than those above.

Lemma 2.6. *Suppose the pair (K, F) is presented as a (p, q) quasi-cable on θ , $p \geq q \geq 1$. Suppose there is a sphere $P \subset S^3$ so that*

- θ is in general position with respect to P ,
- P intersects e^\perp but is disjoint from e^\pm , and
- each component of the 1-manifold $P \cap F$ is essential in F .

Then $\text{genus}(F) \geq 2$.

Proof. Suppose, as in the proof of Lemma 2.3, $\text{genus}(F) = 1$. The case $q \geq 2$ is already settled, so we assume $q = 1$. Set $n = |e^\perp \cap P|$, necessarily even; then $|K \cap P| = n(p + 1) \geq 4$ and the number of edges of $F \cap P$ is $n(p + 1)/2$. Of the

intervals in $K - P$, $(p + 1)(n - 1)$ lie on segments of $e^\perp - P$, p contain e^+ in their interior, and exactly one contains e^- . Call the last the “special” component. We now consider how these intervals are distributed around the hexagon R described in the proof of Lemma 2.6 above. First note that the two vertices corresponding to the corners of the hexagon (when reidentified to give F) lie on either the same or opposite sides of P , depending on whether $n(p + 1)/2$ is even or odd, and this in turn determines whether the number of intersections of any edge of R with P is even or odd. The upshot is that every edge of R intersects P with the same parity, and that parity is determined by the parity of $n(p + 1)/2$.

The combinatorial argument gets easier as p gets larger, since intervals corresponding to the same components of $\theta - P$ appear more often. So for brevity we just do the case $p = 1, 2$ (and of course $q = 1$) and merely outline the argument. (In fact the argument for these cases also instantly gives as well all cases in which $p + 1$ is a multiple of 2 or 3.)

When $p = 1$ there are $2n$ segments in ∂F and it follows that, except perhaps for the special component, centers of opposite edges in R represent the same segment of e^\perp , for there are exactly as many intersection points with P (namely n) going one way around K between them as the other. At most one of these three opposite pairs contains the special component, so the other two display lens spaces $L(2, 1)$ in S^3 , a contradiction. The only way to avoid this contradiction is if only two (opposite) sides of R intersect K , so all arcs in $F \cap P$ are parallel in F , and the special component appears in the center of this band of parallel arcs. This only transfers the contradiction: if we let A be the annulus which is the complement in F of the single band containing all the parallel arcs of $P \cap F$, then it is easy to see that ∂A lies on the boundary T of the punctured solid torus obtained by attaching a single component of $\eta(e^\perp) - P$ to P . Moreover, the components of ∂A are oriented so that they form parallel circles in T . This implies that A , together with an annulus in T , forms a Klein bottle in S^3 , another (related) contradiction.

When $p = 2$ then $|K \cap P| = 3n$, so the number of edges in $K \cap P$ is $3n/2 \geq 3$. There are two cases:

If each edge of R intersects P the same number of times (as our parity discussion guarantees will happen when $n = 2$), then each triple of corners of the hexagon (corresponding to a vertex of the punctured torus when it’s reassembled) represents the same interval of e^\perp , or perhaps the special segment. At least one of the triples doesn’t contain the special segment, and the hexagon in F cut off from F by the arcs of $P \cap F$ adjacent to these corners, together with P and the segment of e^\perp the hexagon is incident to, describes the spine of a lens space $L(3, 1) \subset S^3$, a contradiction.

If some edges of R intersect P more often than others (so $n \geq 4$ and so $K \cap P \geq 12$), then consider a longest pair $\rho_i, i = 1, 2$, of opposite edges of R (length here is shorthand for the number of points of intersection with P). Picturing these opposite sides as the top and bottom of the hexagon, consider the distance in R between the left ends of the ρ_i . Our choice of edge guarantees that it is less than $1/3$ the circumference of R , that is, $K \cap P$. Similarly for the right-hand ends. It follows readily that there are at least two rectangles in F , cut off by a pair of adjacent arcs of $P \cap F$ running between the ρ_i that either make up part of the spine of an $L(2, 1) \subset S^3$ or contain the special component. At most one can contain the special component, leading to the same contradiction. \square

The same argument, with the roles of e^+ and e^\perp switched and p replacing $p+1$, shows

Lemma 2.7. *Suppose the pair (K, F) is presented as a (p, q) quasi-cable on θ , $p > q \geq 1$. Suppose there is a sphere $P \subset S^3$ so that*

- θ is in general position with respect to P ,
- P intersects e^+ but is disjoint from e^\perp and e^- , and
- each component of the 1-manifold $P \cap F$ is essential in F .

Then $\text{genus}(F) \geq 2$.

Essentially the same argument applies in a slightly different setting:

Lemma 2.8. *Suppose the pair (K, F) is presented as a (p, q) quasi-cable on θ . Suppose there is a sphere $P \subset S^3$ so that*

- θ is in general position with respect to P ,
- P intersects each of e^\pm in a single point, and
- each component of the 1-manifold $P \cap F$ is essential in F .

Then either also $|e^\perp \cap P| = 1$ or $\text{genus}(F) \geq 2$.

Proof. Here the intervals of $K - P$ that are incident to $P \cap e^-$ constitute two adjacent “special” components. Since we can take $|e^\perp \cap P| \geq 3$ (it’s necessarily odd, since the hypothesis guarantees that its ends lie on opposite sides of P), then $K \cap P \geq 8$. Then the combinatorial argument of Lemma 2.6 applies with little change. \square

We can switch the roles of e^+ and e^\perp in the above proofs, at the cost of raising p to 2:

Lemma 2.9. *Suppose the pair (K, F) is presented as a (p, q) quasi-cable on θ with $p > q \geq 1$. Suppose there is a sphere $P \subset S^3$ so that*

- θ is in general position with respect to P ,
- P intersects e^\perp and e^- in a single point, and
- each component of the 1-manifold $P \cap F$ is essential in F .

Then either also $|e^+ \cap P| = 1$ or $\text{genus}(F) \geq 2$.

Proof. We need only consider the case $q = 1$. Let $n = |e^+ \cap P|$, necessarily odd. Then $|K \cap P| = p(n+1) + 2$. In fact, $p(n-1)$ segments lie parallel to segments of $e^+ - P$, $2p+2$ segments (in sequential pairs) are incident to the point $e^\perp \cap P$, and one of these sequential pairs (called the “special pair”) is incident to the point e^- . Much as before, we set $p = 2, 3$ as the most difficult but also roughly representative cases.

When $p = 2$ there are an odd number of arcs in $P \cap F$, since $|K \cap P|/2$ is odd and so there are an odd number of arcs incident to each edge. The extra “special” segments in K mean that the intersection arc of $K \cap F$ that lies in the center of each pair of opposite sides of R does not have its ends at the same point of intersection of θ with P (an immediate contradiction via its normal orientation), but it does mean that adjacent to each such central arc, at opposite ends, are segments of $K - P$ that are parallel to the same segment of $\theta - P$. This would exhibit, as usual, the absurd $L(2, 1) \subset S^3$. This contradiction is only avoided if each side of R intersects P in exactly one point. But this means that $|K \cap P| = 6$, so $n = 1$, as required.

When $p = 3$, then, since the number of arcs in $P \cap F$ is $1 \pmod 3$, not all edges of R intersect P the same number of times. We may as well restrict to the case $n \geq 3$, so there are at least 7 arcs in $F \cap P$. If there is a single pair of longest edges, each must then intersect P at least 5 times. It's easy to see, in this case, that wherever the adjacent special edges lie, they cannot disrupt the existence of at least one pair of opposite intervals of $K \cap P$ (among those lying in these longest sides) that correspond to the same component of $e^+ - P$ and so constitute part of a lens space $L(2, 1) \subset S^3$. Similarly, if there are two pairs of opposite longest sides, then each must be of length at least 3 and, wherever the adjacent special components lie, they cannot disrupt the existence of a similar lens space contradiction from at least one pair of opposite longest sides. \square

3. KNOTS THINLY PRESENTED ON HANDLEBODIES

In [GST, Section 2], we extended Gabai's notion of thin position for knots to include also certain types of graphs in 3-space. We briefly review (and incidentally somewhat extend) that development here, since it will be an important ingredient of our argument.

Choose a height function $h : S^3 - \{x, y\} = S^2 \times \mathbb{R} \rightarrow \mathbb{R}$ and let $P(t) = h^{-1}(t)$.

Definition 3.1. A finite trivalent graph $\Gamma \cup S^3 - \{x, y\}$ is in *normal form* with respect to h if

- (a) for each edge $e \subset \Gamma$ the critical points of $h|_e$ are nondegenerate and lie in the interior of e ,
- (b) the critical points of $h|_{edges}$, and the vertices of Γ , all occur at different heights, and
- (c) at each (trivalent) vertex v of Γ either two ends of incident edges lie above v (we say v is a *Y-vertex*) or two ends of incident edges lie below v (we say v is a *λ -vertex*)

Standard Morse theory shows that any finite trivalent graph in S^3 can be infinitesimally isotoped so that it is in normal form.

Definition 3.2. The *maxima* of Γ consist of all local maxima of $h|_{edges}$ and all λ -vertices. Similarly, the *minima* of Γ consist of all local minima of $h|_{edges}$ and all Y -vertices. A maximum (resp. minimum) that is not a λ -vertex (resp. Y -vertex) will be called a *regular* maximum (resp. minimum). The union of the maxima and minima (hence including the vertices) are called the *critical points* of Γ and their heights the *critical values* or *critical heights*.

Definition 3.3. Let $t_0 < \dots < t_n$ be the successive critical heights of Γ and suppose t_{v_1}, \dots, t_{v_j} are that subset of levels at which vertices occur. Let $s_i, 1 \leq i \leq n$, be generic levels chosen so that $t_{i-1} < s_i < t_i$. Define the *width* of Γ to be

$$W(\Gamma) = 2 \left(\sum_{i \notin v_1, \dots, v_j} |P(s_i) \cap (\Gamma)| \right) + \left(\sum_{i \in v_1, \dots, v_j} |P(s_i) \cap (\Gamma)| \right).$$

Definition 3.4. A *thin position* of a graph $\Gamma \subset S^3$ is a normal form (with respect to h) which minimizes the width of Γ .

Remark. In practice, the chief property of a thin positioning of a graph Γ that we will need is this: The positioning becomes thinner if a maximum is pushed below a minimum, but the width is unaffected by pushing one maximum above or below

another maximum, or one minimum above or below another minimum. See [GST, Section 3] for details.

A graph $\Gamma \subset S^3$ in normal form with respect to h can be thickened slightly to give a solid handlebody $\eta(\Gamma) \subset S^3$. Standard techniques allow us to take a neighborhood so thin that the height function $h|_{\partial(\eta(\Gamma))}$ has the obvious Morse structure: very near any regular maximum (resp. minimum) of Γ there are two non-degenerate critical points of $h|_{\partial(\eta(\Gamma))}$, one a saddle just below (resp. above) and one a maximum (resp. minimum) just above (resp. below). Similarly, just above (resp. below) a Y -vertex (resp. λ -vertex) there is a single saddle singularity. When we refer to a regular neighborhood $\eta(\Gamma)$ of Γ we will always mean a thickening with this property. Slightly abusing notation, $S^3 - \eta(\Gamma)$ will denote the closed complement of $\eta(\Gamma)$. We will be concerned with simple closed curves on $\partial\eta(\Gamma)$ and with properly imbedded surfaces in $S^3 - \eta(\Gamma)$.

Definition 3.5. Suppose Γ is a graph, in normal form with respect to h , and $K \subset \partial\eta(\Gamma)$ is a simple closed curve. Then K is in normal form on $\partial\eta(\Gamma)$ if each critical point of h on K is non-degenerate, and occurs near an associated critical point of Γ in $\partial\eta(\Gamma)$. Furthermore, the number of critical points of K has been minimized via isotopy of K in $\partial\eta(\Gamma)$.

Definition 3.6. A properly imbedded surface

$$(F, \partial F) \subset (S^3 - \eta(\Gamma), \partial\eta(\Gamma))$$

is in *normal form* if

- (1) each critical point of h on F is nondegenerate,
- (2) ∂F is in normal form with respect to h ,
- (3) no critical point of h on $\text{int}(F)$ occurs near a critical height of h on Γ ,
- (4) no two critical points of h on $\text{int}(F)$ occur at the same height,
- (5) the minima (resp. maxima) of $h|_{\partial F}$ at the minima (resp. maxima) of Γ are also local extrema of h on F , i.e., ‘half-center’ singularities, and
- (6) the maxima of $h|_{\partial F}$ at Y -vertices and the minima of $h|_{\partial F}$ at λ -vertices are, on the contrary, ‘half-saddle’ singularities of h on F .

Remark. The meaning of “near” in (3) is probably best thought of informally, but the technical requirement (for, say, the critical height of a maximum v of Γ) is this: No critical point of h on the interior of F occurs at a height between the levels of the maxima of ∂F (if any) near v and the level of the saddle point of $\partial\eta(\Gamma)$ near v . Standard Morse theory ensures that, for Γ in normal form, any properly imbedded surface $(F, \partial F)$ can be put in normal form.

Definition 3.7. Γ is in *bridge position* if there is a level sphere, called a *dividing sphere* for the bridge position, that lies above all minima of Γ and below all maxima.

Definition 3.8. Given Γ in normal form and P a level sphere for h at a generic height, let B_u and B_l denote the balls which are the closures of the region above P and below P respectively. An *upper disk* (resp. *lower disk*) for P is a disk $D \subset S^3 - \eta(\Gamma)$ transverse to P such that $\partial D = \alpha \cup \beta$, where α is an arc imbedded on $\partial\eta(\Gamma)$, $\beta = \partial D \cap P$ is an arc properly imbedded in $P - \eta(\Gamma)$, $\partial\alpha = \partial\beta$, and a small product neighborhood of ∂D in D lies in B_u (resp. B_l), i.e., it lies *above* (resp. *below*) P .

Note that $\text{int}(D)$ may intersect P in simple closed curves. An innermost such simple closed curve cuts off a disk that lies either above or below P . Such a disk is called an *upper cap* or *lower cap*. For the moment, these caps will be unimportant.

A natural occurrence of upper (or, symmetrically, lower) disks is this: According to Definition 3.6, a maximum of ∂F near a maximum of Γ is a half-center singularity on ∂F . In particular, a sphere P just below this maximum will cut off an upper disk from F .

Definition 3.9. Suppose Γ is a graph, in normal form with respect to the height function h , and K is a normal form simple closed curve on $\partial\eta(\Gamma)$. If K is also in thin position (as a knot in S^3) with respect to h , then we say that K is *thinly presented on $\eta(\Gamma)$* .

Lemma 3.10. *Suppose Γ is a graph, in normal form with respect to the height function h , $F \subset S^3 - \eta(\Gamma)$ is an incompressible surface in normal form and $K = \partial F$ is thinly presented on $\eta(\Gamma)$. Suppose a maximum and a minimum of K occur respectively at heights u and l with $l < u$. Then there is a generic level sphere $P = P(t)$, $l < t < u$, so that every arc component of $P \cap F$ is essential in F .*

Proof. We have seen that a sphere just below $P(u)$ cuts off an upper disk from F and that a sphere just above $P(l)$ cuts off a lower disk from F . The seminal point of thin position (see [G]) is that there cannot simultaneously (even at a critical point of h on the interior of F) be both an upper and a lower disk, for these disks could be used to push a maximum of K below a minimum, thinning K . Hence there is a generic height t between l and u for which the level sphere $P = P(t)$ cuts off neither an upper nor a lower disk from F . But this means there can be no arcs of $P \cap F$ which are inessential, for an outermost such inessential arc would cut off either an upper or a lower disk from F . \square

Combining Lemma 3.10 with the central lemma of the previous section, we have this corollary:

Corollary 3.11. *Suppose the pair (K, F) is thinly presented as a (p, q) quasi-cable on θ with both $p, q \geq 2$. Then $\text{genus}(F) \geq 2$.*

Proof. Let M^+ and M^- be the highest maxima of, respectively, the cycles $e^\perp \cup e^+$ and $e^\perp \cup e^-$. Similarly, let m^+ and m^- be the respective lowest minima of these cycles. Let $u = \min\{M^+, M^-\}$ and $l = \max\{m^+, m^-\}$. Since e^\perp is in both cycles, we know that any point on e^\perp lies below u and above l , so $l < u$. Since $p, q > 0$, u and l are (near) the heights of, respectively, maxima and minima of K . Choose a level sphere as in Lemma 3.10 between u and l . By construction, such a level sphere lies at a height between the maximum and minimum of each of the cycles $e^\perp \cup e^+$ and $e^\perp \cup e^-$, and so intersects both of them. Now apply Lemma 2.3. \square

Lemma 3.12. *Suppose K is a knot which is thinly presented as a (p, q) quasi-cable on $\theta \subset S^3$. Suppose θ has been made as thin as possible subject to this condition. Suppose furthermore that the complement $S^3 - \eta(\theta)$ is a genus two handlebody. Then θ is in bridge position.*

Proof. The proof is analogous to that of [GST, Proposition 4.4]. Suppose that θ is not in bridge position. Then there is a level sphere P that lies between a sequential pair of critical levels for θ , a maximum just below P and a minimum just above P . Maximally compress $P - \theta$ in the complement of θ . The resulting meridional

planar surface \tilde{P} is incompressible in the handlebody $S^3 - \eta(\theta)$, so each component is parallel to a subsurface of $\partial\eta(\theta)$ (see [Mo]). Since the boundary components of \tilde{P} are meridians of $\eta(\theta)$, each component of \tilde{P} can be completed to a sphere in S^3 , and the piece of θ lying in one of the balls bounded by that sphere is an unknotted tree (possibly just an arc) in the ball.

For concreteness, choose an innermost such ball B and suppose it lies above \tilde{P} . Then every arc of K lying in B has at least one maximum in B . If $\theta \cap B$ is a single arc α , then isotope that arc to lie in \tilde{P} . The arc in \tilde{P} can be chosen to be disjoint from those disks which are the results of the compressions that created \tilde{P} from P , so in fact then α lies in P . After this isotopy, the width of θ is reduced (since a maximum has been pushed below whatever minimum lay just above P) and that of K is not increased (since each arc of $K \cap B$ still has at most one maximum). This argument shows more: any arc of $K \cap B$ must have *exactly* one maximum in B , for otherwise a disk of parallelism between that arc and \tilde{P} (guaranteed by [Mo]) could be used to reduce the number of critical points on K , contradicting the assumption that K is in thin position, hence in minimal bridge position.

If the tree $\theta \cap B$ contains only maxima (including perhaps λ -vertices), then pushing one to \tilde{P} would push it below the minimum that we know lies (elsewhere) just above \tilde{P} , again thinning θ . So we know that \tilde{P} contains at least one minimum and, since it can't be a minimum of K , it must be a Y -vertex, with ends of e^\pm descending into it. Similarly, if both vertices are in B , then both must be Y -vertices. Since $\theta \cap B$ is a tree, at most one of e^\pm lies in B , so we can assume that, say, e^- intersects ∂B . Also e^+ intersects ∂B , since otherwise K would contain a minimum in B . Let α be the arc in $\theta \cap B$ consisting of the end of e_- and the end of e_+ at the Y -vertex. Then the parallelism between α and an arc on \tilde{P} , guaranteed by Morimoto's theorem [Mo]), describes how to pull the end of e_- down to change the vertex into a λ -vertex. This thins θ . \square

4. KNOTS PRESENTED AS p -EYEGASSES

We will need a second way in which K can be viewed as lying on a neighborhood of a normal form graph in S^3 . Let \bowtie be the "eyeglass" graph, obtained from two circles e_l and e_r by attaching an edge e_b running between them. If \bowtie is imbedded in S^3 , then a regular neighborhood $\eta(\bowtie)$ of \bowtie is a genus two handlebody that can be described as follows: Take two solid tori T_l and T_r with cores the loops e_l and e_r and meridian disks μ_l and μ_r respectively, and join them together by a 1-handle with core e_b and meridian μ_b .

We will sometimes refer to e_b as the bridge between the cycles e_l and e_r .

Definition 4.1. Given a normal form eyeglass $\bowtie \subset S^3$, a normal form knot $K \subset \partial\eta(\bowtie)$ is *presented as a p -eyeglass on \bowtie* if

- $|K \cap \mu_r| = 1$,
- $|K \cap \mu_b| = 2$,
- $|K \cap \mu_l| = p \geq 1$, and
- K always intersects μ_l with the same orientation.

Less formally, K can be described as the band-sum, via a band running once along the bridge 1-handle, of a longitude of T_r and a (p, q) cable of T_l . Note that if K is presented as a p -eyeglass on \bowtie and $S^3 - \eta(\bowtie)$ is a handlebody, then K is tunnel number one: Since K goes just once through a meridian of e_r , we can




FIGURE 7.

isotope $H = \eta(\infty)$ in S^3 , “vacuuming” up K with e_r until K is simply a longitude of e_r .

It’s easy to see that any knot presented as a p -eyeglass is also presented as a $(p - 1, 1)$ quasi-cable on the same underlying handlebody H . The difference is in how meridians are chosen to define the graph that H is a neighborhood of. The correspondence is given by $\mu_r = \mu^-$, $\mu_l = \mu^\perp$ while the meridian disks μ_b and μ^+ intersect in a single arc. That is, the difference of the two graphs is a simple Whitney move. The θ graph obtained from ∞ this way is called the *associated* $(p - 1, 1)$ *quasi-cable* and the disk in ∞ just described that becomes the meridian μ^+ will be called the *pre-cable disk* in ∞ .

Recall that a graph is in bridge position if every maximum lies above every minimum. Following [ST2], we will extend this notion in the case of an eyeglass graph.

Definition 4.2. Suppose a height function is defined on S^3 . A cycle in S^3 is *vertical* if it has exactly one minimum and one maximum. An eyeglass graph is in *extended bridge position* if any minimum that does lie above a regular maximum (resp. maximum that lies below a regular minimum) is a Y -vertex at the minimum (resp. λ -vertex at the maximum) of a vertical cycle. A vertical cycle whose minimum is a Y -vertex is called an *extended maximum*. One whose maximum is a λ -vertex is called an *extended minimum*. Such a Y -vertex or λ -vertex is called a *base vertex* of the extended maximum (resp. minimum). A level sphere that lies above all minima and extended minima (except perhaps a base Y -vertex) and below all maxima and extended maxima (except perhaps a base λ -vertex) is called a *dividing sphere* for the extended bridge position. See Figure 7.

Proposition 4.3. *Suppose K is thinly presented as a p -eyeglass on ∞ , whose closed complement is also a genus two handlebody. Suppose that ∞ is made as thin as possible, subject to the condition that it thinly presents K . Then either*

- (1) e_l is vertical, or
- (2) $p = 1$ and e_r is vertical, or
- (3) ∞ is in extended bridge position.

Moreover, in the last case, e_b is disjoint from some dividing sphere for ∞ .

Proof. The proof follows the same line of argument as the proof of the main theorem of [ST2]. We only need to verify that the argument there does not interfere with the thin presentation of K . In fact, the argument here is simpler because some of the more complicated steps in [ST2] are required only after a step that, in our case, clearly thins K or shows that, e.g., e_l is vertical.

So suppose \bowtie thinly presents K but is not in extended bridge position. We've noted above that K is a tunnel number one knot, so we know that K is in bridge position (cf. [T]). So if a maximum of \bowtie lies below a minimum, either the maximum is a λ -vertex or the minimum is a Y -vertex, or both. So there are at most two level spheres with the property that each lies just below a minimum and just above a maximum. Let Q be the sphere or pair of spheres with this property. Compress Q as much as possible in the complement of \bowtie , and call the result Q' .

A path in $\partial H = \eta(\bowtie)$ between meridians is *regular* if the corresponding path in \bowtie is embedded. It is shown in [ST2] that there is a disk F in S^3 whose boundary is the union of a path α in \bowtie and an arc β in Q' . Moreover, the interior of F is disjoint from Q' , and either α is a regular path that is disjoint from some meridian of e_b , or α has both its ends at the same point p of $e_b \cap Q'$ and runs once around either e_r or e_l . Consider each possibility in turn.

Case 1: α is a regular path that is disjoint from some meridian of $e_b \subset \bowtie$.

Say F lies below $\beta \subset Q'$. By general position we can assume that β is disjoint from the disks in Q' which are the remains of the compressing disks of Q , so in fact β lies on Q . If α does not pass through a vertex of \bowtie , then just use F to isotope the arc of $\bowtie - Q$ that contains α to $\beta \subset Q$. During the isotopy, as α perhaps passes through Q (though not through Q'), \bowtie may get thicker, but once it reaches β it will have been thinned, since all that remains of its internal critical points is one minimum, which will have been brought up above the level of the maximum just below Q . The move similarly cannot thicken K .

Essentially the same argument applies even when α passes through a vertex. F is used to slide an end of one of the edges incident to the vertex down to β , perhaps thereby just extending e^b and not affecting the bridge structure of K . In any case, \bowtie is thinned and K is not thickened.

Case 2: α has both its ends at the same point q of $e_b \cap Q'$ and runs once around either e_r or e_l .

Much as in the previous case, F can be used to move the cycle e_r or e_l together with the end of e_b between q and the cycle to Q . Unless the cycle was already vertical, this move (once the cycle is tilted again to restore genericity) will thin \bowtie and will not thicken K . So we can assume the cycle is vertical. This means we are done, unless in fact the cycle is e_r and $p \geq 2$. This case only arises if Q intersects e_l but not e_r since if it is disjoint from both, we can appeal to [Mo] directly to get a disk as in Case 1. (Unless, of course, some component of Q intersects \bowtie in exactly one point of e_b . But then \bowtie would be planar.)

Now note that, unless the maximum just below Q is a λ -vertex maximum of e_r , the move on e_r just described pushes a minimum of K (on e_r) past a maximum of K , contradicting the thin position of K . We deduce that the one and only maximum just below Q is in fact a λ -vertex maximum of e_r . If Q consists of more than one sphere, we could repeat the same argument, just using the component to which we have not just pushed e_r . But that would lead to the contradiction that the λ -vertex maximum of e_r also lies just below the other plane. We deduce that Q is a single plane. We have shown then that, aside from the base λ -vertex of e_r , only minima lie below Q . Moreover, just above Q is a minimum, and at no other level does a maximum lie below a minimum. It follows that \bowtie is in extended bridge position.

It remains to show that e_b is disjoint from some dividing sphere. Much of the proof mimics [ST2]. We suppress most of the technical details, except to note that many of the technical problems do not arise in our context. Most importantly, if P is a dividing sphere and there are disjoint lower and upper caps, then pushing a vertical cycle which is, say, an extended maximum down past a minimum would immediately thin K , even if (as discussed in [ST2]) passing other maxima might thicken ∞ . We deduce that in our context disjoint lower and upper caps cannot arise for elementary reasons.

In any case, the upshot of the argument in [ST2] is that there is a dividing sphere P that cuts off from a meridian disk E of $S^3 - \eta(\infty)$ both an upper disk D_u and a lower disk D_l . Moreover, the interior of each is disjoint from P .

Consider the components C_u and C_l of $\eta(\infty) - P$ to which D_u and D_l are incident. If neither C_u nor C_l contain vertices, or if they have no ends in common, or if together they contain at most one vertex and they have a single end in common, then it is easy to use D_u and D_l to push a maximum down past a minimum, contradicting the thinness of ∞ or K . Note in particular that if the boundary of one of the disks, say D_l , goes once around e_r , although the move described may thicken ∞ (cf. [ST2]), it does push a minimum of K (namely the minimum of e_r) past a maximum of K and so would violate the assumption that ∞ thinly presents K .

We now proceed to dispose of the other cases. Suppose that, say, C_u contains a vertex and that C_u and C_l have two end meridians in common. We can assume C_u contains only one vertex, else e_b is disjoint from P and we are done. Then $C_u \cup C_l$ contains a vertical cycle. If that cycle is e_l or $p = 1$ we are done, so we'll assume it's e_r and that $p \geq 2$. Either D_u and D_l can be used to make ∞ (indeed $K!$) thinner, or they can be used to isotope e_r into P . This last move not only makes ∞ non-generic, but it may thicken K if K winds around e_r . Nonetheless, we persist, inspired by the proof of [GST, Theorem 5.14]. That argument shows that, once e_r is level, so the solid torus neighborhood T_r divides P into two disks, an innermost disk component of $E \cap P$ in E or a disk cut off by an outermost arc of $E \cap P$ in E can be used to push a maximum (resp. minimum) of ∞ (possibly the maximum near the end of e_b at e_r) down (resp. up) through the level of P . Afterwards, e_r can be tilted slightly to restore genericity and thereby to remove the extra bridges of K that may have been introduced when e_r was made perfectly level. Since a maximum has been pushed down (or a minimum up) past e_r , it follows that ∞ (indeed K , since $p \geq 2$) has been thinned, the usual contradiction.

The possibility remains that C_u and C_l each have a single vertex and they also have a single end in common (they can't have two ends in common since the result would be a cycle in ∞ containing both vertices.) Their common end must be a point of $e_b \cap P$, since in ∞ that is the only arc that connects the two vertices; in particular, e_b is monotonic. In this case, the disks D_u and D_l either could be used to thin ∞ (an immediate contradiction) or they can be used to make e_b level. If neither component C_u or C_l contains all of e_r , then the move simply levels e_b . Once again, this move makes ∞ no longer generic and may also thicken K , for K may wind many times around the edge e_b . But we continue anyway, inspired this time by the proof of [GST, Theorem 6.1, Subcase 3b]. The argument has a number of subcases, but all result in the following conclusion: a maximum (say) of e_l or e_r (possibly contiguous to an end of e_b) can be pushed down to the level of e_b

(or below, if it is not contiguous). Once this is achieved, tilt e_b slightly to restore genericity, but leave the pushed-down maximum at (or below) the lower end of e_b . The result is a thinning of \bowtie (indeed K).

The final possibility is that the move just described levels all of $e^b \cup e^r$ because, say, C_u contains e_r . In this case it seems that the move might pull C_u past other maxima lying below it, thickening \bowtie so that, after $e^b \cup e^r$ is tilted to restore genericity, \bowtie actually ends up thicker. Nevertheless, it is argued in [ST2] that in fact this does not happen, or at least, if it does, the extra thickness (and more) can immediately be removed by a move analogous to that described above when e^r was levelled. This is established by a somewhat complicated combinatorial argument on ∂E . We won't repeat the argument here. The upshot is that either \bowtie ends up thinner, a contradiction, or there were in fact no maxima between e_r and P . But in this last case, all the minima (including the base Y -vertex of e_r) lie below all the maxima, so \bowtie is in non-extended bridge position. (Only the base vertex of e_r lies between P and the level plane P' for this bridge presentation.) Moreover, all of e_b lies below P' , verifying the proposition in this case as well. \square

Corollary 4.4. *Suppose K is thinly presented as a p -eyeglass on \bowtie , $S^3 - \eta(\bowtie)$ is a handlebody, and \bowtie has been made as thin as possible. Then either*

- (1) e_l is vertical, or
- (2) $p = 1$ and e_r is vertical, or
- (3) $\text{genus}(K) \geq 2$, or
- (4) $p \geq 2$ and K is also thinly presented as a $(p - 1, 1)$ quasi-cable on the associated Θ -graph θ and θ is no thicker than \bowtie , or
- (5) $p = 1$ and the graph $K \cup \gamma$ is no thicker than \bowtie .

Remark. The last two possibilities are essentially the same and are only distinguished by the value of p . It's convenient to restrict the terminology (p, q) quasi-cable to the case $p, q \geq 1$. If we were to extend that definition to $(1, 0)$ quasi-cable, then the original graph $K \cup \gamma$ would be its natural meaning. See the beginning of Section 5.

Proof. Following Proposition 4.3, we only need to consider the case in which \bowtie is in possibly extended bridge position and e_b is disjoint from a dividing sphere. Suppose first that \bowtie is in fact in (non-extended) bridge position and, with no loss of generality, suppose e_b lies above the dividing sphere. Then e_b ascends from the lower λ -vertex and either is monotonic, or has one internal maximum and descends into the other λ -vertex as well. It's easy to move from one position to the other without affecting the width of either \bowtie or K , so we'll assume for concreteness that e_b is monotonic. By (perhaps) twisting around the other two ends at the lower λ -vertex we can ensure that the meridian disk for the associated Θ -graph, namely the pre-cable disk in \bowtie that runs the length of e_b , is disjoint from the descending disk incident to e_b given by the bridge structure. See Figure 8. Then the Whitney move has no effect on the bridge structure (hence the width) of K , nor the width of \bowtie : a pair of λ -vertices with an edge between them is replaced by exactly the same thing.

Now assume that \bowtie has an extended maximum, say. Since the extended maximum contains a vertical cycle, we are done immediately unless the vertical cycle is e_r and $p \geq 2$. Since e_b is disjoint from the dividing sphere (say it lies above), it runs monotonically from a λ -vertex to the Y -vertex base of e_r . Since $p \geq 2$, the




FIGURE 8.

knot K has a maximum at the λ -vertex. Now find, somewhere below the λ -vertex, a level sphere P as in Lemma 3.10, so every component of $P \cap F$ is essential in the Seifert surface F . Then Lemma 2.6 applied to the associated $(p-1, 1)$ quasi-cable shows that $genus(K) \geq 2$. \square

5. THINNING QUASI-CABLES

Suppose K is a tunnel number one knot with γ an unknotting tunnel. According to [ST1, Proposition 4.2] there is a minimal genus Seifert surface F for K that is disjoint from γ . We now pursue the line of argument used in [GST] and [ST1] to analyze the relation between F and the pair (K, γ) . The philosophy will be to view the pair (K, γ) as an incipient case of K being thinly presented as a (p, q) quasi-cable, though with $(p, q) = (1, 0)$. Here the graph θ would be $K \cup \gamma$, with γ playing the role of e^- and the two segments of K into which the ends of γ divide K playing (interchangeably) the roles of e^+, e^- . Roughly, the idea is this: inspired by [GST], we will consider the thinnest graph $\theta \subset S^3$ that thinly presents K as a (p, q) quasi-cable, and, inspired by [ST1], ask how F and the height function h interact with “splitting” spheres for the handlebody $\eta(\theta) \subset S^3$.

We begin by setting some terminology and notation. In analogy to the notation used for quasi-cables, denote the meridian of $H = \eta(K \cup \gamma)$ corresponding to a point of γ by μ^- , and meridians corresponding to points in the two edges of $K - \gamma$ by

μ^+ and μ^\perp . In general, for H a handlebody in S^3 whose closed complement is also a handlebody, a *splitting sphere* S for H is a sphere that intersects ∂H in a single essential circle. In other words, it is a reducing sphere for the Heegaard splitting $S^3 = H \cup_{\partial H} (S^3 - \text{int}(H))$. A splitting sphere S is best viewed as the union of two disks, $D = S \cap H$ and $E = S - \text{int}(H)$, that have a common boundary in ∂H . The exterior disk E will, much as above, suggest possible thinning moves for $\theta \supset K$. The interior disk D will give useful information about how ∂E can behave, since $\partial D = \partial E$.

Definition 5.1. Let $\{\mu_i\}$ be a (not necessarily complete) family of pairwise disjoint meridian disks for H , and let c be a simple closed curve on ∂H isotoped so as to minimize $|c \cap (\cup_i \mu_i)|$. Suppose an arc component c_0 of $c - \{\mu_i\}$ has both its ends on a single meridian μ_i in the family, and the union of c_0 and a subarc of $\partial \mu_i$ bound a disk in H . Then c_0 is a *wave* of c with respect to $\{\mu_i\}$. The wave is said to be *based* at the meridian μ_i .

In particular, if D is an essential disk in H , isotoped so as to minimize $|D \cap (\cup_i \mu_i)|$, then all components of this intersection are arcs, and an outermost arc of intersection in D cuts off a disk $D_0 \subset D$ so that $\partial D_0 \cap \partial H$ is a wave of ∂D with respect to $\{\mu_i\}$. The disk D_0 is called a *wave disk*.

Definition 5.2. Suppose θ thinly presents the pair (K, F) as a (p, q) quasi-cable with $q \leq p$. Consider the family of meridians $\{\mu^-, \mu^+, \mu^\perp\}$ for the handlebody $H = \eta(\theta)$. Then an essential disk D in H *satisfies the wave condition* if there is a wave of ∂D based at either μ^- or μ^+ (so in particular the wave is disjoint from μ^\perp).

Similarly, if S is a splitting sphere for H , then S satisfies the wave condition if $D = S \cap H$ does.

Note that D satisfies the wave condition if and only if some outermost disk D_0 of D (hence all outermost disks) cut off by $\{\mu^-, \mu^+, \mu^\perp\}$ is cut off by an arc lying in either μ^- or μ^+ .

With this terminology, [ST1, Corollary 5.3] can be reinterpreted as follows (noting that “ ρ is finite” translates to “ S has a wave at μ^- ”):

Lemma 5.3. *Suppose K is a tunnel number one knot with γ an unknotting tunnel, and F is a minimal genus Seifert surface F for K that is disjoint from γ . As described above, let $\theta = K \cup \gamma$ present (K, F) as a $(1, 0)$ quasi-cable. Then γ may be slid and isotoped so that either*

- γ lies on F , or
- there is a splitting sphere for $\eta(\theta)$ so that, with respect to the set of meridians $\{\mu^-, \mu^+, \mu^\perp\}$, a wave of D is based at μ^- .

In fact more of [ST1] can be reinterpreted in this setting. If K is not a 2-bridge knot, an invariant $\rho(K, \gamma) \in \mathbb{Q}/2\mathbb{Z}$ is defined and, if $\rho \neq 1$, it is shown that γ can be isotoped onto F . The case $\rho = 1$ translates to this statement: For the pair of meridians $\{\mu^+, \mu^\perp\}$ (which are interchangeable in this context) and any splitting sphere S there are waves of $D = S \cap H$ with the property that each wave disk intersects μ^- in a single arc. In particular, a wave disk at μ^+ , say, can be glued to a subdisk of μ^+ to get a non-separating meridian μ^{wave} of H that is disjoint from the wave disk and intersects μ^- in a single arc. Then the meridians $\{\mu^+, \mu^\perp, \mu^{\text{wave}}\}$ give H the structure of a Θ -graph θ' which presents (K, F) as a $(1, 1)$ quasi-cable

satisfying the wave condition. In θ' , the meridian μ^+ has become the meridian of the edge e'^- , μ^\perp the meridian of the edge e'^+ , and μ^{wave} the meridian of the edge e'^\perp . The two graphs θ and θ' differ by a standard Whitney move on γ . So we have:

Lemma 5.4. *Suppose K is a tunnel number one knot with γ an unknotting tunnel, and F is a minimal genus Seifert surface for K that is disjoint from γ . Then either there is a graph θ that presents (K, F) as a $(1, 1)$ quasi-cable and a splitting sphere that satisfies the wave condition, or γ may be isotoped to lie on F .*

The next ingredient to throw into the mix is thin position. Of course in Lemma 5.3 there is no obstacle to having θ thinly present K —just begin with K in thin position. It's not obvious that the process that leads from Lemma 5.3 to Lemma 5.4 preserves the property that K is thinly presented, but we now show that it does (or γ can be made into an unknotted loop). Or, more accurately, we show that this follows immediately from the results of [GST] and [ST1].

Lemma 5.5. *Suppose K is a tunnel number one knot with γ an unknotting tunnel, and F is a minimal genus Seifert surface for K that is disjoint from γ . Then either*

- (1) *there is a graph θ , no thicker than the graph $K \cup \gamma$, that thinly presents (K, F) as a $(1, 1)$ quasi-cable and a splitting sphere that satisfies the wave condition, or*
- (2) *γ may be isotoped to lie on F , or*
- (3) *γ can be slid and isotoped to form an unknotted loop with its ends at the same point of K .*

Proof. Following [GST, Theorem 3.5], γ can be slid to become either an unknotted loop, and we are done, or γ is a level edge, with its ends incident to the top two maxima (say) of the thinly presented K . (In the latter case, regain a generic positioning by slightly perturbing γ from its level position, changing it to a monotone edge connecting two λ -vertices.) If K is 2-bridge, it's easy to see that F contains an isotopic copy of γ (cf. [BZ, Remark 12.26]), and we are done. If K is not 2-bridge and the invariant $\rho(K, \gamma) \in \mathbb{Q}/2\mathbb{Z}$ then defined in [ST1] is not 1, it follows from [ST1, Theorem 5.2] that γ may be isotoped into F , and we are done.

Consider finally the case $\rho(K, \gamma, S) = 1$ for S a splitting sphere as in Lemma 5.3. The fact that $\rho = 1$ means that the pair of meridians μ^\perp, μ^+ cuts off a wave of $D = S \cap H$ based at the meridian μ^+ , say, and that wave intersects μ^- in a single arc. In particular (as above), the wave disk can be glued to a subdisk of μ^+ to get a non-separating meridian μ^{wave} of H that is disjoint from the wave disk and intersects μ^- in a single arc. Appropriately twist the two arcs of K that descend from the bottom λ -vertex of γ (equivalently, choose an appropriate set of descending disks) as discussed in Corollary 4.4 (see Figure 8) so that the descending disk is disjoint from μ^{wave} . Then the standard Whitney move on γ , using μ^{wave} , not only converts θ to a graph θ' that thinly presents K as a $(1, 1)$ quasi-cable, it does it without thickening θ , for one pair of λ -vertices is just replaced with another. (See again Figure 8). \square

We consolidate our results a bit more:

Definition 5.6. Suppose K is a knot, F is a minimal genus Seifert surface for K and θ is a Θ -graph such that θ thinly presents (K, F) as a (p, q) quasi-cable, $p \geq q \geq 1$, and such that there a splitting sphere for $\eta(\theta)$ that satisfies the wave

condition. We say that θ is an *appropriate* Θ -graph (for the pair (K, F) and splitting sphere S).

Proposition 5.7. *Suppose K is a tunnel number one knot and γ is an unknotting tunnel for K . Then either*

- (1) γ can be slid and isotoped to form an unknotted loop with its ends at the same point of K , or
- (2) there is a minimal genus Seifert surface F for K with this property: either
 - γ can be slid and isotoped to lie on F , or
 - there is an appropriate Θ -graph for (K, F) (and some splitting sphere).

In the last case, if the Θ -graph is the thinnest one with the required property, then it is in bridge position.

Proof. As noted above, according to [ST1, Proposition 4.2] there is a minimal genus Seifert surface F for K that is disjoint from γ . The rest follows from Lemmas 5.5 and 3.12. \square

We will expand on the last possibility, but it will be useful to have the following general lemma:

Lemma 5.8. *Let $\theta \subset S^3$ be a Θ -graph with edges e_1, e_2, e_3 . Suppose θ is in bridge position, P is a dividing sphere, and the cycle $e_1 \cup e_2$ is vertical. Then either*

- (1) all cycles in θ are vertical, or
- (2) there are bridges above and below P made up of interior subarcs of e_3 , or
- (3) one of e_1, e_2 , say e_1 , is disjoint from P and $e_1 \cup e_3$ is vertical.

Proof. Suppose first that the vertices lie on the same side of a dividing sphere P , so, with no loss of generality, both are Y -vertices, say. Then each edge intersects P in an even number of points. Since $e_1 \cup e_2$ is vertical, together the edges intersect P in 2 points. Hence one of these two edges, say e_1 , is disjoint from P . This means that e_1 is the edge that descends from the higher Y -vertex. e_3 can't also descend from this vertex, so e_3 intersects P . If e_3 intersects P in two points then $e_1 \cup e_3$ is vertical. If e_3 intersects P in four or more points, then there are bridges above and below P made up of interior subarcs of e_3 , as required.

Suppose next that the vertices lie on opposite sides of P , so one vertex is a λ -vertex and the other is a Y -vertex, and each edge intersects P in an odd number of points. Since $e_1 \cup e_2$ is vertical, each of these edges intersects P in a single point. If e_3 also intersects P in a single point, then every cycle is vertical. If it intersects P in three or more points, then there are bridges above and below P made up of interior subarcs of e_3 , as required. \square

Proposition 5.9. *Suppose K is a knot with Seifert surface F and there is an appropriate Θ -graph for (K, F) and some splitting sphere. Let θ be a thinnest such Θ -graph. Suppose that $\text{genus}(F) = 1$. Then θ is in bridge position and (thinly) presents K as a $(p, 1)$ quasi-cable, for some $p \geq 1$.*

If furthermore one of the cycles $e^- \cup e^+$ or $e^- \cup e^\perp$ is vertical and a dividing sphere intersects both edges of the vertical cycle, then either

- $e^+ \cup e^\perp$ is vertical, or
- $p = 1$ and the cycle $e^- \cup e^\perp$ is vertical.

Proof. Corollary 3.11 notes that $q = 1$. Proposition 5.7 shows that θ is in bridge position.

Suppose first that the cycle $e^- \cup e^+$ is vertical. We apply Lemma 5.8, using e^-, e^+ for e_1, e_2 . If all cycles are vertical, then of course we are done. By hypothesis, neither e^- nor e^+ is disjoint from a dividing sphere. So we may assume, following Lemma 5.8, that there are bridges above and below a dividing sphere made up entirely of interior subarcs of e^\perp . We can of course arrange that the lowest maximum and highest minimum are these bridges. Now choose a level sphere P as in Lemma 3.10. Since $e^- \cup e^+$ is vertical and P intersects both edges, P intersects each edge e^-, e^+ in a single point. The result then follows from Lemma 2.8.

Similarly, suppose the cycle $e^- \cup e^\perp$ is vertical. Again apply Lemma 5.8, this time using e^-, e^\perp for e_1, e_2 . If all cycles are vertical, we are done. By hypothesis, neither e^- nor e^\perp is disjoint from a dividing sphere, so we may arrange that the lowest maximum and highest minimum are from bridges that lie entirely in e^+ . Now choose a level sphere P as in Lemma 3.10. Since $e^- \cup e^\perp$ is vertical and P intersects both edges, P intersects each edge e^-, e^\perp in a single point. The result then follows from Lemma 2.9. \square

Definition 5.10. Suppose θ is an appropriate Θ -graph for the pair (K, F) and splitting sphere S . Suppose there is a level sphere P at a generic height for θ such that P cuts off both an upper and a lower disk from $E = S - \text{int}(H)$. (As usual, P may lie at a critical height of $\text{interior}(E)$). Then P is called a *critical sphere* for θ .

Lemma 5.11. *Suppose θ is a thinnest Θ -graph appropriate for (K, F) . Then either*

- (1) *Some dividing sphere for θ is also a critical sphere, or*
- (2) *all the cycles in θ are vertical, or*
- (3) *the edge e^w (one of e^\pm) on whose meridian a wave is based is disjoint from a dividing sphere, and both of the cycles containing e^w are vertical.*

Proof. The conclusions make sense, since we know from Proposition 5.7 that θ must be in bridge position. If a sphere just above the highest minimum cuts off a lower disk, and a sphere just below the lowest maximum cuts off an upper disk, then some level sphere between them is a critical sphere. This condition is guaranteed unless the highest minimum (or the lowest maximum) is a Y -vertex (resp. λ -vertex) with, via the wave condition, an end of e^w descending (resp. ascending) from the vertex.

So suppose the highest minimum, say, is a Y -vertex with an end of e^w descending. Let P be any dividing sphere for θ . If any component of $\theta - P$ below P is a simple arc, its minimum could be pushed higher than the Y -vertex, eliminating the problem, so we can assume that all components of $\theta - P$ below P contain vertices. If there is only one such component and it contains a single vertex, then each of the cycles has at most one minimum, and so each is vertical. If there is only one such component and it contains both vertices, then the edge between them must be incident to the higher vertex from below; hence that edge is e^w . Moreover, both cycles containing e^w have exactly one minimum, and so both are vertical.

The remaining case is when the vertices are in separate components, each lying below P , i.e., both of them Y -vertices. Let v denote the higher Y -vertex. Push the regular minimum on the component containing v up to a height just below v . Now slide the end of e^\perp ascending from v down to the regular minimum and back up the other side. This has no effect on the width of K (since, for example, it doesn't change the number of bridges), but it alters the arrangement of the edges around v .




FIGURE 9.

In particular, afterwards the end of e^w ascends from v , so, by the wave condition, we can assume that just above v , a level sphere cuts off a lower disk from E , and so somewhere between that level and that of the lowest maximum there is a critical sphere, as required. See Figure 9. \square

It will be useful to assume that we always perform the move described at the end of the proof of Lemma 5.11 when it is possible. That is, if

- θ is in bridge position,
- a dividing sphere P cuts off a component of θ containing a single vertex (say a Y -vertex), and
- the end of e^w descends from that Y -vertex,

then we slide the end of e^\perp at that vertex over the contiguous regular minimum of e^w . This ensures that e^w ascends from every such Y -vertex (and, symmetrically, descends from any such λ -vertex) and has no effect on the width of anything, since we are just rearranging the heights of minima (or the heights of maxima).

Lemma 5.12. *Suppose θ is a thinnest Θ -graph appropriate for (K, F) and θ has a dividing sphere that is a critical sphere. Then either*

- (1) *one of the edges of θ is disjoint from the critical sphere (hence from a dividing sphere), or*
- (2) *the edge e^w of θ on whose meridian a wave is based is monotonic, and one of the two circuits in θ containing e^w is vertical.*

Proof. Let P be the critical sphere. Let C_l and C_u be the components of $\theta - P$ to which the lower and upper disks D_l and D_u are incident. Suppose first that one of these components, say C_l , has no vertex (so C_l is a regular minimum). If C_u also has no vertex, then K could be thinned, a contradiction. If C_u has two vertices, then it contains an edge disjoint from P , and we are done. If C_u has one vertex, then, since every circuit in θ has two vertices, C_l and C_u have at most one common end point on P . Moreover, by the wave condition, the path $\beta_u = \partial D_u \cap P$ has at least one end incident to an end of e^w in C_u . Then D_u and D_l describe how to slide the end of e^w in C_u down to P while simultaneously isotoping all of C_l up to P . There are two possible routes which the end of e^w could take during the slide, but either leads to a contradiction, as we now illustrate. If the slide were down an end of e^\perp , that would thin K . If the slide were down the other edge, then the effect of the slide is to sew together the ends of e^\pm . This effectively extends the end of e^\perp




FIGURE 10.

down to or below the level of the minimum at D_l . (See Figure 10.) This thins θ (though not necessarily K), possibly by changing a λ -vertex into a Y -vertex. This contradicts the assumption that θ is a thinnest such graph.

The remaining case is that C_l and C_u both contain a single vertex. As above, D_l and D_u can be used to slide the ends of e^w in C_l and C_u up and, respectively, down, until they lie in P . This would thin θ unless e^w crossed P in exactly one point (an end of both C_u and C_l). In the latter case, e^w was monotonic and the slide would level e^w by isotoping it into P . If the levelled e^w becomes a loop, then originally it was part of a vertical cycle, and we are done. This is also true if one of the other edges of θ has a single interior critical point. If neither of these cases occurs, then the argument of [GST, Theorem 6.1, Subcase 3b] can be used (as it was in Proposition 4.3) to move the levelled e^w up (or down) to connect two maxima or two minima. Afterwards K is no wider but θ is thinned, a contradiction completing the proof. \square

Corollary 5.13. *Suppose that θ is a thinnest Θ -graph appropriate for (K, F) . Then either*

- (1) *some dividing sphere is a critical sphere and the critical sphere is disjoint from one of the edges of θ , or*
- (2) *$e^+ \cup e^\perp$ is vertical, or*
- (3) *$p = 1$ and the cycle $e^- \cup e^\perp$ is vertical, or*
- (4) *$\text{genus}(F) \geq 2$.*

Proof. Suppose first that some dividing sphere is a critical sphere. Then, following Lemma 5.12, if the critical sphere intersects all the edges, then some cycle involving one of e^\pm is vertical. If it's the cycle $e^+ \cup e^\perp$, we're done. If it's either of the other two cycles, apply Proposition 5.9, observing that either some edge is above a dividing sphere or the dividing sphere will intersect both edges of any vertical cycle.

If no dividing sphere is a critical sphere, then apply Lemma 5.11. If all the cycles in θ are vertical, then of course we are done. If the wave is based on e^+ so $e^w = e^+$, then Lemma 5.11 says $e^+ \cup e^\perp$ is vertical, as required. Finally, suppose $e^w = e^-$. If $p = 1$, then we are done, since $e^- \cup e^\perp$ is vertical. So suppose $p = 2$ and, following Lemma 5.11, e^- is disjoint from a dividing sphere. Find a dividing sphere P as in Lemma 3.10, and apply Lemma 2.5. \square

6. REGULAR ANNULI IN HANDLEBODY COMPLEMENTS

Definition 6.1. Let Γ be a trivalent graph in S^3 , in normal form with respect to a height function h . Then a normal form simple closed curve c on $\partial\eta(\Gamma)$ is *regular* if c never “back-tracks” along an edge, traversing the edge twice in opposite directions, after looping around a vertex at the end of the edge.

More precisely, if $\{\mu_i\}$ is a collection of meridian disks in $\eta(\Gamma)$, one for each 1-handle corresponding to an edge of Γ , then no subsegment of c with interior disjoint from these meridians has its ends incident to the same side of the same meridian. In particular, no minimum of c occurs near a λ -vertex of Γ and no maximum of c occurs near a Y -vertex of Γ .

Lemma 6.2. *Suppose Γ is a trivalent graph in S^3 in bridge position with respect to a height function h and A is a properly imbedded annulus in $S^3 - \eta(\Gamma)$ in normal form with respect to h . Denote the components of ∂A by $\partial_{\pm}A$. Suppose some dividing sphere $P_0 = P(t_0)$ intersects ∂_+A more often than it intersects ∂_-A . Suppose finally that the lowest maximum and the highest minimum of Γ are regular critical points, on edges that are incident to ∂_+A . Then some dividing sphere cuts off from A both an upper disk D_u and a lower disk D_l , both of which are incident to ∂_+A (as opposed to ∂_-A).*

Proof. Since the lowest maximum is a regular maximum and ∂_+A runs along the edge that contains it, a dividing sphere $P(y)$ just below the lowest maximum cuts off an upper disk that is incident to ∂_+A . Similarly, a dividing sphere $P(x)$ just above the highest minimum cuts off a lower disk that is incident to ∂_+A . Since some (hence every) dividing sphere intersects ∂_+A more often than it intersects ∂_-A , every dividing sphere cuts off some disk, either upper or lower, incident to ∂_+A . Since at x there is a lower one, and at $y > x$ there is an upper one, and at every height between there is one or the other, it follows that at some height there is both an upper and a lower disk incident to ∂_+A . (As usual, this height may be at a saddle tangency of P with an interior point of A .) \square

In the next section we will see that such a useful annulus A can often be found. In this section we examine how, by exploiting upper disks and lower disks lying in A , we can thin a Θ curve that presents K as a $(p, 1)$ quasi-cable. Until we begin to use Lemma 6.2 there is nothing special about using A ; any properly embedded surface in the graph complement would do, though it is important that the upper and lower disks themselves are incident to θ only along regular curves disjoint from K . For example, the external disk E used above cannot generally be used for the purposes of this section, because its boundary is typically not a regular curve.

Lemma 6.3. *Suppose there is a Θ -graph appropriate for (K, F) and, for a thinnest one θ , there is a dividing sphere P that cuts off from A disjoint upper and lower disks D_u and D_l so that the arcs $\alpha_u = D_u \cap \partial\eta(\theta)$ and $\alpha_l = D_l \cap \partial\eta(\theta)$ are regular curves disjoint from K .*

If the edge e^+ is disjoint from P , then either

- (1) α_u runs from a point of $e^- \cap P$ to a point of $e^{\perp} \cap P$, traversing e^+ once, or
- (2) $\text{genus}(K) \geq 2$, or
- (3) $e^+ \cup e^{\perp}$ is unknotted.




FIGURE 11.

Proof. We may as well assume e^+ lies above P , so the vertices are λ -vertices. By thinness we can assume that α_u lies on the 4-punctured sphere component Σ of $\partial\eta(\theta) - P$ lying above P . We can think of the components $K \cap \Sigma$ as edges of a graph on Σ , with vertices the four meridian boundary components. That is, if we label the meridian components of $\partial\Sigma$ as $\mu_l^-, \mu_r^-, \mu_l^\perp, \mu_r^\perp$ in the obvious way, then the components of $K \cap \Sigma$ consist of a single arc connecting μ_l^- to μ_l^\perp , a single arc connecting μ_r^- to μ_r^\perp , and p arcs connecting μ_l^\perp to μ_r^\perp . The p arcs either are all parallel or comprise two families of parallel arcs, separating in Σ the points μ_l^-, μ_r^- . See Figure 11.

Consider the possibilities for α_u : There is only one path in Σ that is disjoint from K and has ends at meridians μ_l^- and μ_l^\perp . And it's parallel in Σ to an arc of K . Hence if α_u were that path, K could be thinned, using D_u and D_l . Similarly for paths from μ_r^- to μ_r^\perp . There are (at most) two paths from μ_l^- to μ_r^\perp disjoint from K , each traversing e^+ once, and similarly two paths from μ_r^- to μ_l^\perp . If α_u is any of these paths, then the first conclusion of the lemma follows.

Now suppose α_u is the path in Σ (available only if all p arcs are parallel) that is disjoint from K and runs from μ_l^- to μ_r^- . Suppose, to begin with, that α_l does not also run from μ_l^- to μ_r^- but rather has at least one end at another point of $\theta \cap P$. Consider the $p + 1$ eyeglass graph \bowtie obtained from θ by a Whitney move along e^+ , using as the new meridian a neighborhood of $\mu_l^- \cup \alpha_u \cup \mu_r^-$ (see Figure 12). Then D_u describes a descent of the new bridge edge e_b for \bowtie down to P . Simultaneously, D_l describes how to move a minimum of (now) \bowtie above or at least to the level of P . In particular, \bowtie is in bridge position, and a dividing sphere necessarily intersects e_b . Moreover, \bowtie is thinner than θ since, in effect, a maximum has been pushed below a minimum. Now suppose \bowtie is as thin as possible, still requiring that it thinly present K as a $(p + 1)$ -eyeglass. According to Corollary 4.4 (exploiting the fact here that \bowtie is a $p + 1 > 1$ eyeglass), either e_l becomes vertical (which implies that $e^+ \cup e^\perp$ was unknotted), or $genus(K) \geq 2$, or the associated Θ -graph, namely




FIGURE 12.

θ , can be made as thin as ∞ . The last contradicts the hypothesis, and the first two possibilities are the conclusions we seek.

Now suppose α_l , like α_u , runs from μ_l^- to μ_r^- (so, in particular, $e^+ \cup e^-$ is vertical). We could construct ∞ as just described; the construction places $e_r \subset \infty$ level in P . Unfortunately, when e_r is tilted to restore genericity, P is again a level sphere for ∞ and no thinning will have occurred. So a different argument will be used, this one exploiting the fact that α_u and α_l are disjoint from K . Consider the situation once D_u and D_l have been used to put e_r into the plane P . Let T_r be a thickened regular neighborhood of e_r and consider the two longitudinal circles $\{\lambda_1, \lambda_2\} = P \cap \partial T_r$. It is easy to see what they are: Before the edge e_r is levelled it intersects P in two points (actually the points $e^- \cap P$). The λ_i are obtained by banding the corresponding pair of meridians to itself using both α_u and α_l . The important point for our purposes is that $|K \cap (\lambda_1 \cup \lambda_2)| = 2$. Consider where these points lie. If they both lie on the same longitude, then an arc of K they cut off is inessential in the annulus component of $T_r - (\lambda_1 \cup \lambda_2)$ in which it lies, and so it can be removed by an isotopy. On the other hand, if one point lies on each longitude, consider the algebraic intersection of K with the disk component P_1 of $P - T_r$ bounded by λ_1 , say. One point is the point $\lambda_1 \cap K$. All others come from intersections of e_l , i.e., intersections of the old e^\perp . Each point $e_l \cap P_1$ contributes $p + 1$ points to $K \cap P_1$, and they are all of the same sign. So the total algebraic intersection of K with P_1 is $\pm 1 \not\equiv 0 \pmod{p + 1}$. This contradicts the fact that K bounds F in $S^3 - H$. We are left with the conclusion that indeed K can be isotoped off of the two longitudes, so K only intersects the top of ∂T_r . But in that case, consider $F \cap (P_1 \cup P_2)$. It's easy to see that any component of intersection that is inessential in F can be removed (else K could be thinned). So every component of $P \cap F$ is essential. Now simply attach the bottom annulus of T_r to $P_1 \cup P_2$ to obtain a sphere intersecting the original θ only in e^\perp . Then Lemma 2.6 shows that $\text{genus}(K) \geq 2$.

The remaining case is if α_u has one end at each of the meridians μ_l^\perp and μ_r^\perp . Exclude any such arc that is parallel to a subarc of K , since if α_u were such an arc it would either violate the thinness of K or (if α_l has ends at the same meridians) exhibit that $e^+ \cup e^\perp$ is vertical. But the only way that α_u can be disjoint from K , connect μ_l^\perp to μ_r^\perp , and not be parallel to a subarc of K is if all p arcs of K with ends at these meridians are parallel and α_u is one of the other two paths connecting μ_l^\perp and μ_r^\perp . Although these paths are not parallel to a component of $K \cap \Sigma$ in




FIGURE 13.

Σ , they are sufficiently parallel in the component of $\eta(\theta) - P$ on whose boundary Σ lies to derive the same contradiction. Here is the argument: If the ends of α_l are also at μ^+_l and μ^+_r , then together α_r and α_l show that the cycle $e^+ \cup e^\perp$ is vertical. So suppose no end of α_l is at μ^+_l , say. One arc κ of K running from μ^+_l to μ^+_r is visibly isotopic to α_u in the 3-ball component of $\theta - P$ on whose boundary Σ lies. (See Figure 13.) The isotopy moves the end of κ across the meridian μ^+_l and so destroys the property that K lies on $\partial(\theta)$. Nonetheless, once κ is moved to α_u , then there is no obstruction to pushing κ below P via D_u while simultaneously pushing arcs of $K - P$ parallel to α_u above P using D_l . The result is a thinning of K , violating the hypothesis. \square

Lemma 6.4. *Suppose there is a Θ -graph appropriate for (K, F) and, for a thinnest one θ , there is a dividing sphere P that cuts off disjoint upper and lower disks D_u and D_l from A so that the arcs $\alpha_u = D_u \cap \partial\eta(\theta)$ and $\alpha_l = D_l \cap \partial\eta(\theta)$ are regular curves disjoint from K .*

If e^\perp is disjoint from P , then either

- (1) *genus(K) ≥ 2 , or*
- (2) *$e^+ \cup e^\perp$ is unknotted, or*
- (3) *$p = 1$, and $e^- \cup e^\perp$ is unknotted.*

Proof. The proof is mostly similar to that of Lemma 6.3. The relevant figure is modified as shown (Figure 14), with $p - 1$ arcs running between meridians μ^+_l and μ^+_r . There is only one regular path from μ^+_l to μ^-_r (or from μ^+_r to μ^-_l) that is disjoint from K . These paths do not cross μ^\perp , and so, if α_u is such a path, we could use D_u and D_l to thin θ (without altering the wave condition), extending e^\perp down to P . The only regular path between μ^+_l and μ^-_l (or μ^+_r and μ^-_r) is parallel to an arc of K , so α_u cannot be such a path. If $p \geq 2$ and α_u runs between μ^+_l and μ^+_r , we use the same argument as was used for paths from μ^+_l and μ^+_r previously.

Suppose finally that α_u runs between μ^-_l and μ^-_r (or, symmetrically, $p = 1$ and α_u runs between μ^+_l and μ^+_r .) We would like to use the same trick as was used previously, namely, let D_u describe a Whitney move that converts θ into an eyeglass graph. There is a subtle complication, however. Note that the eyeglass graph \bowtie that is created by this move is in fact a p eyeglass, not a $(p + 1)$ eyeglass as before. In particular, the associated Θ -graph θ' to \bowtie is not θ , which presented K as a $(p, 1)$ quasi-cable. Rather θ' presents K as a $(p - 1, 1)$ quasi-cable (or just as $K \cup \gamma$ if $p = 1$). Nonetheless, we are still in a position to get the same contradiction with Corollary 4.4 (for, after all, θ was chosen to be thinnest among all appropriate Θ -graphs, and it was shown that such a graph is no thicker than $K \cup \gamma$), as long as we verify that θ' is still appropriate. In other words, we need to verify that θ'




FIGURE 14.

still satisfies the wave condition. The pre-cable disk is easy to identify: whereas in the previous argument it was (essentially) a thickened vertical arc in D_u together with a meridian of e^+ , here it is obtained from a thickened vertical arc in D_u and a meridian of e^\perp by adding a half-twist. (See Figure 15.) Notice that K intersects this meridian in $(p - 1)$ points, so it is the meridian of $e^{+'}$ (or γ if $p = 1$) and, whether the wave was based at e^- or at e^+ , the extra half-twist guarantees that it is afterwards based at the new meridian, as required. (The language when $p = 1$ is: the slope of the wave is still finite.) \square

It will be useful to provide notation for arcs in the 4-punctured sphere Σ discussed in the proof of Lemmas 6.3, 6.4. As motivation for the notation we will use, suppose, as above, that the knot K is thinly presented as a $(p, 1)$ quasi-cable on a Θ -graph θ in S^3 . Suppose further that θ is in bridge position with respect to a height function h , and there is a dividing sphere P that is disjoint from one of the edges. In particular, one of the components of $\theta - P$ is a tree C with 4-ends (whose regular neighborhood intersects $\partial\eta(\theta)$ in the 4-punctured sphere Σ) and all the other components of $\theta - P$ are simple arcs.

Of course the fundamental group of θ is free on two generators. The natural generators of $\pi_1(\theta)$ are loops that traverse each of the two edges not disjoint from P exactly once. To be concrete, orient K and suppose e^\perp is disjoint from P ; then a loop in θ , based at a point in C , will give rise to a word in letters a and b (with inverses \bar{a} and \bar{b}), where a corresponds to traversing e^- once, and b to traversing e^+ once, each in the same direction as K . Similarly, if one of e^\pm is disjoint from P , then a will correspond to traversing e^\mp once and b to traversing e^\perp once. We will only be interested in such presentations for regular simple closed curves on $\partial\eta(\theta)$ that are disjoint from K . In this case, the cyclic permutation class of the word in a, b, \bar{a}, \bar{b} corresponding to the regular simple closed curve $\sigma \subset \partial\eta(\theta) - K$ determines




FIGURE 15.

the isotopy class of σ in $\partial\eta(\theta)$ almost precisely. The only ambiguity is in how σ intersects Σ .

So consider isotopy classes of regular arcs in Σ that are disjoint from K . Representative are those illustrated and labelled in Figure 16. Two of the figures show (via a heavy line from top to bottom) the meridian of the edge that's disjoint from P . This meridian is relevant for determining that an arc is regular: a regular arc can cross the meridian at most once. The meridians of the arcs intersecting P , i.e., the boundary components of Σ , are shown as circles. The arcs are labelled by where they would occur in a word in a, b, \bar{a}, \bar{b} ; thus if the sequence $\dots \bar{a}\bar{b}\dots$ (or, inversely, $\dots b\bar{a}\dots$) occurs in the word, the corresponding arc in Σ would be one labelled $\bar{a}\bar{b}$. Notice, as one example, that there are two arcs labelled $\bar{b}a$, indicating ambiguity in how such an arc may run through Σ . There is less ambiguity in the lowest figure (corresponding to the edge e^\perp disjoint from P). The heavy oriented horizontal curves correspond to arcs in K , and they are labelled in the same manner. In the first picture, to avoid crowding only one arc labelled $\bar{b}\bar{b}$ (or, equivalently, $b\bar{b}$) is shown; it's a subarc of K . One of the two others is shown in the lowest figure; the other arc labelled $\bar{b}\bar{b}$, which is only relevant to the first figure, is obtained from the second by reflection through the vertical arc μ^+ . The special case $p = 1$ and e^\perp disjoint from P is not shown.

With this labelling, another way of stating the first possibility in Lemma 6.3 would then be: α_u is of type $\bar{a}\bar{b}$ or $\bar{b}a$.

Under the hypotheses of the lemma, it is natural to define

Definition 6.5. The *geometric length* $l(w)$ of a word w in a, b, \bar{a}, \bar{b} representing a regular loop $\sigma \subset \partial\eta(\theta)$ is $|\sigma \cap P|$, for P a dividing sphere.




FIGURE 16.

Given $l(a)$ and $l(b)$, the geometric length of w is clearly $ml(a) + nl(b)$, where m and n are the total number of occurrences of, respectively, a and \bar{a} , b and \bar{b} . We will say that a word w is *positive in a, b* if neither \bar{a} nor \bar{b} occur in w (e.g., when $w = \emptyset$). We then have

Corollary 6.6. *Suppose there is a Θ -graph θ appropriate for (K, F) and θ is a thinnest one. Suppose further that*

- e^+ is disjoint from a dividing sphere,
- there is a properly imbedded normal form annulus A in $S^3 - \eta(\theta)$ whose boundary components $\partial_{\pm}A$ are disjoint from K and regular in $\partial\eta(\theta)$, and
- a dividing sphere P cuts off upper and lower disks from A , and the component(s) of ∂A to which these disks are incident represents a word that is positive in a, b .

Then either $e^+ \cup e^{\perp}$ is unknotted or $\text{genus}(K) \geq 2$

Proof. The upper disk cannot be of type $a\bar{b}$ or $\bar{b}a$, because a word that is positive in a, b contains neither \bar{b} nor \bar{a} . The result then follows from Lemma 6.3. \square

Similarly, from Lemma 6.4 we derive this corollary:

Corollary 6.7. *Suppose there is a Θ -graph θ appropriate for (K, F) , and θ is a thinnest one. Suppose further that*

- e^{\perp} is disjoint from a dividing sphere,

- there is a properly imbedded normal form annulus A in $S^3 - \eta(\theta)$ whose boundary components $\partial_{\pm}A$ are disjoint from K and regular in $\partial\eta(\theta)$, and
- a dividing sphere P cuts off upper and lower disks from A .

Then either $e^{\perp} \cup e^+$ is unknotted, or $\text{genus}(K) \geq 2$, or $p = 1$ and $e^{\perp} \cup e^-$ is unknotted.

7. WHEN K HAS GENUS ONE

Our interest in regular annuli comes from the following observation. Suppose F is genus one. Then an outermost disk cut off by $F \cap E$ in E provides a boundary-compression of F to an essential annulus $A \subset S^3 - \eta(\theta)$. When viewed on $\partial\eta(\theta)$ the relation between ∂A and K is this: K is banded to itself via a subarc $\omega \subset \partial E - K$ whose ends are incident to the same side of K . If $D \subset \eta(\theta)$ satisfies the wave condition, then any subarc of $\partial E = \partial D$ that is disjoint from K is regular. So we are assured that ∂A is regular in $\eta(\theta)$.

This leads to

Theorem 7.1. *Suppose in a thinnest Θ -graph θ appropriate for (K, F) , the edge e^+ is disjoint from a dividing sphere. Suppose further that $\text{genus}(F) = 1$ and that the wave for ∂D is based at μ^- . Then either the cycle $e^+ \cup e^{\perp}$ is unknotted, or $p = 1$ and the cycle $e^- \cup e^{\perp}$ is unknotted.*

Proof. That $q = 1$ follows from Proposition 5.9. We will show that the annulus A obtained from ∂ -compressing F to $\eta(\theta)$, using the disk E from the splitting sphere, gives rise to an annulus satisfying the conditions of Corollary 6.6. The result then follows from that corollary, possibly by way of Proposition 5.9.

Let $\omega \subset \partial\eta(\theta) - K$ be the arc described above, which we may as well slide to minimize intersections with the meridians of $\eta(\theta)$. In particular, for P a level sphere between the lowest maximum and the highest minimum of θ , the ends of ω will lie on the 4-punctured sphere component Σ of $\partial\eta(\theta) - P$. Let w be the word in a, b, \bar{a}, \bar{b} represented by ω . Because the wave of ∂D is based at μ^- ,

- any occurrence of the letter b (resp. \bar{b}) in w is preceded and followed by the letter a (resp. \bar{a}),
- any occurrence of the letter a in w is followed by a or b and preceded by a or b , and
- any occurrence of the letter \bar{a} in w is followed by \bar{a} or \bar{b} and preceded by \bar{a} or \bar{b} .

In particular, by a choice of orientation for w , we can assume that w is positive (say) in a and b (e.g., perhaps $w = \emptyset$). Exploiting these facts, together with the symmetries of the diagram, we have three essentially different ways in which the ends of w can lie in Σ . These are shown in Figure 17.

Now orient w from left to right, and read off the words corresponding to $\partial_{\pm}A$ (boundaries oriented to be parallel in A , not antiparallel). All three cases can be expressed in one of the following two forms, with the details depending on where the ends of ω are incident to the word given by K , namely ab^{p+1} . Note that, in all cases, the word w begins and ends with the letter a . The choice of labelling of the components $\partial_{\pm}A$ of ∂A is made so that the word corresponding to ∂_+A is positive in a, b .




FIGURE 17.

Form 1:

- $\partial_+A \leftrightarrow wb^i,$
- $\partial_-A \leftrightarrow w\bar{b}^j\bar{a}\bar{b}^k.$

Here $j, k > 0$ and $i + j + k = p + 1$.

Form 2:

- $\partial_+A \leftrightarrow wb^j ab^k,$
- $\partial_-A \leftrightarrow w\bar{b}^i.$

Here $j > 0, i, k \geq 0$ and $i + j + k = p + 1$.

Note that the geometric length of ∂_+A is greater than that of ∂_-A exactly when, for annuli of the first form, $(i - j - k)l(b) > l(a)$ and, for annuli of the second form, $(i - j - k)l(b) < l(a)$. In either case, ∂_+A contains an occurrence of the letter b , so ∂_+A runs along e^\perp . If e^\perp intersects a dividing sphere P just twice, then $e^+ \cup e^\perp$ is vertical, and we are done. If $|e^\perp \cap P| \geq 4$, then a regular maximum and minimum (which we can take, respectively, to be the lowest maximum and the highest minimum) lie in the interior of e^\perp . It follows, then, from Lemma 6.2 and Corollary 6.6 that we are done if $(i - j - k)l(b) > l(a)$ for annuli of the first form,

or $(i - j - k)l(b) < l(a)$ for annuli in the second form. So we now consider only the alternative possibilities.

Case 1: $(i - j - k)l(b) = l(a)$.

In this case, it follows roughly from the same argument used in Lemma 6.2 that there is a dividing sphere P so that no arc of $P \cap A$ cuts off an outermost disk incident to $\partial_+ A$. Indeed, as above, a high dividing sphere cuts off an upper disk incident to $\partial_+ A$ and a low dividing sphere cuts off a lower disk incident to $\partial_+ A$. There can't be both an upper and a lower such disk incident to $\partial_+ A$ by Corollary 6.6. So at some level there is a dividing sphere P that cuts off no outermost disk incident to $\partial_+ A$. But since the geometric lengths of $\partial_{\pm} A$ are equal, this implies that P intersects A only in spanning arcs.

We will argue that this is impossible. Suppose (with no loss) that e^+ and hence Σ lie above (not below) a dividing sphere. The spanning arcs determine a correspondence between subintervals of $\partial_- A$ and $\partial_+ A$. To be precise, say that a component of $\partial_- A - P$ and a component of $\partial_+ A - P$ are *opposite* if there is a (“square”) component of $A - P$ incident to both. More generally, a segment of $\partial_+ A$ and a segment of $\partial_- A$ are opposite each other if a spanning arc of $P \cap A$ runs between the beginning of each and another spanning arc runs between the end of each. For example, we first observe that no segment of $\partial_- A$ that is part of a \bar{b} interval (meaning that it comes from an occurrence of a letter \bar{b} , i.e., that it runs along e^\perp with an orientation opposite to that of K) can lie opposite a segment of $\partial_+ A$ that is part of a b interval. For if this occurred, then it is easy to see that somewhere on the entire length of the b and \bar{b} intervals there would be components of $\partial A - P$ that are opposite to each other on A but lie on the same component κ of $e^\perp - P$. Since they have opposite orientation in $\partial_- A$ and $\partial_+ A$, the square component of $A - P$ connecting them can be attached to the punctured solid torus $P \cup \eta(\kappa)$ to create a punctured lens space $L(2, 1) \subset S^3$, which is absurd. Similarly no segment of $\partial_- A$ that is part of an \bar{a} interval can lie opposite a segment of $\partial_+ A$ that is part of an a interval. This immediately rules out the first form above (again, only under the assumption that $(i - j - k)l(b) = l(a)$), since, following these observations, the only possible segment opposite the transition segment from \bar{b} to \bar{a} in $\partial_- A$ would be exactly a transition segment from a to b in $\partial_+ A$, and that would lead to the same contradiction.

Ruling out the second form (in which no letter \bar{a} appears) is only slightly more complicated. We will focus on the segments of $\partial A - P$ that lie on the middle component of $e^\perp - P$; that is, on the arc component that is equidistant (measuring distance by intersection with P) from both ends of e^\perp . Note that this segment of $e^\perp - P$ lies below P . Label corresponding segments of $\partial A - P$ by κ . Note that none of these can be opposite to a segment of ∂A lying on Σ (e.g., those segments in ∂A that correspond to the transition between different letters), since the segments on Σ lie above P . This remark allows us to be a bit casual about length arguments in the next few paragraphs, since it means that inequalities will usually imply strict inequalities.

The first observation is that no label κ occurs opposite to any part of a \bar{b}^p interval in $\partial_- A$, for this would allow us to display a lens space $L(2, 1)$ in S^3 , as noted above.

Let $\sigma \subset \partial_- A$ be the segment between the first and last labels κ in $\bar{b}^i \subset \partial_- A$. Then $l(\sigma) = (i - 1)l(b) \geq l(a)$. Since no label κ lies opposite to \bar{b}^p , it follows that opposite to σ is part of a segment in $\partial_+ A$ corresponding to a^q , $q \geq 1$. Notice that

if a label κ in ∂_-A , say, is opposite to any part of a b interval in ∂_+A , then the relation is reciprocal: the label κ in the b interval on ∂_+A is opposite to the b interval in ∂_-A containing the original label κ . (This is not deep, just a reflection that we have taken κ to lie half way along e^\perp and so it appears half way along each b or \bar{b} interval.) Because $j + k > 0$ there is at least one more label κ in ∂_+A than there are κ labels in ∂_-A , not counting the labels κ in \bar{b}^i . It follows that some label κ in ∂_+A is opposite a part of an a interval in ∂_-A , so an entire half of a b interval in ∂_+A is opposite to a subsegment τ of a single a interval in ∂_-A , for $l(a) \geq l(b)$. And, as we've seen, another copy of τ lies opposite to a subsegment of σ . This works just as well for constructing a lens space in S^3 as having the half of the b -segment itself opposite to σ . (Two rectangles are glued together along the boundary interval they share on a component of $e^- - P$ corresponding to part of the a intervals.) So we arrive at the same contradiction as previously.

Case 2: For an annulus of the first form, $(i - j - k)l(b) < l(a)$, or, for one of the second form, $(i - j - k)l(b) > l(a)$.

We will arrive at the same sort of contradiction, though the argument is a bit more complicated. Again, with no loss, we assume that e^+ lies above a dividing sphere and that both the lowest maximum and the highest minimum of θ are regular critical points on e^\perp . This implies that just below the lowest maximum (i.e., at a high dividing sphere) the dividing sphere cuts off an upper disk from A that is incident to a regular maximum (namely the lowest maximum). Then no outermost disk cut off from A by this high dividing sphere can be a lower disk, by thin position. On the other hand, a low dividing sphere does cut off a lower disk from A . So there is a height y such that a dividing sphere just below y cuts off a lower disk from A but just above y a dividing sphere cuts off no lower disk. But any dividing sphere must cut off some outermost disk, since the geometric lengths of the words represented by the two boundary components of A are different. It follows from Corollary 6.6 then that either $e^+ \cup e^\perp$ is unknotted, or just above y all outermost disks cut off of A by the level sphere P are incident to ∂_-A and, moreover, for each such disk D_u the arc $\alpha_u = \partial D_u \cap \partial H$ is of the form $a\bar{b}$. That is, there are at most two outermost disks in A , and they are incident to the subarcs α_1 and α_2 labelled $a\bar{b}$ and $\bar{b}a$ of ∂_-A .

In this position, the total number of non-spanning arcs in A , all of them incident to ∂_-A and each of them cutting off a disk containing either α_1 or α_2 , is $x = |(i - j - k)l(b) - l(a)|/2$, since each arc has two ends. Since x is less than the distance between the ends of $\bar{b}^j \bar{a}\bar{b}^k$ (first form) and less than the distance between the ends of \bar{b}^i (second form), each non-spanning arc cuts off a disk containing exactly one of α_1 or α_2 and so each arc is parallel to one of the α_i . Let x_i denote the number parallel to α_i , so $x_1 + x_2 = x$. Let σ denote the segment in \bar{b}^p that is still incident to spanning arcs, and let s be its length. See Figure 18. For obvious pictorial reasons, we'll refer to the part of A containing the collection of arcs parallel to α_i as the x_i *peninsula*.

Subcase 2a: ∂A is of the first form.

In this case note that $s + x = (j + k)l(b) + l(a)$, so $s = x + il(b)$.

Subcase 2a.i: The entire \bar{a} interval of ∂_-A is disjoint from σ .

Say the entire \bar{a} interval lies on the x_1 peninsula. Then σ is entirely made up of powers of \bar{b} and its length is at least $x_1 + il(b) \geq l(a) + (i + j)l(b) \geq l(a) + 2l(b)$.




FIGURE 18.

Since σ is made up of \bar{b} intervals, at most $l(b)$ of the length of the segment opposite to σ can lie in b intervals (half at each end), so in particular, the segment opposite to σ contains an a -segment longer than $l(a)$. In particular, if τ denotes the terminal segment of a at the end of the x_1 peninsula, there is also a copy of τ lying opposite σ . Since across the ends of the x_1 peninsula the orientations of a and b coincide (that is, the orientations of a and \bar{b} are reversed by the folding along x) whereas across from σ they disagree, we obtain the standard lens space contradiction. See Figure 19i.

Subcase 2a.ii: σ is completely contained in the \bar{a} interval.

We have $s = x + il(b)$ and the longest b -segment in $\partial_+ A$ is of length $il(b)$, since the wave assumption ensures there are no proper powers of b in w (i.e., no powers greater than one). Hence the total length of a -segment(s) opposite σ is at least x . In particular, the length of the a -segment(s) across from σ must be greater (by at least $(j+k)l(b)$) than the length of the ends of a not contained in σ . This implies that some subsegment of σ is opposite a copy of itself, leading to the standard lens space contradiction. See Figure 19ii.

Subcase 2a.iii: The \bar{a} interval is completely contained in σ .

The argument of the previous subcase applies as long as $s < l(a) + l(b)$, so assume that $s \geq l(a) + l(b)$. To avoid the standard lens space contradiction, across from the \bar{a} interval in σ is a segment comprised entirely of b intervals. The largest power of b is b^i and σ is even longer than that, so at least one end of \bar{a} is across from a copy of b that lies completely in the segment across from σ . If the ends of \bar{a} and that copy of b coincide, we get a lens space contradiction via the terminal segment of x . If the copy of b extends out beyond $\bar{a} \subset \sigma$, then we get a lens space contradiction with the end of \bar{b} adjacent to \bar{a} in σ . See Figure 19iii.

Subcase 2a.iv: One end of \bar{a} is contained in x_1 , say, and the other end is contained in σ .




FIGURE 19.

Suppose first that the segment opposite the end of σ at x_1 is part of a b interval. Since the longest b interval comes from b^i and $s = x + il(b)$, it follows that the b interval ends somewhere in σ and is followed by an a interval. We get the standard lens space contradiction with either \bar{a} or \bar{b} , depending on whether the b interval ends across from a point in \bar{a} or a point in \bar{b} . See Figure 20i.

Next suppose the segment opposite the end of σ at x_1 is part of an a interval and let α denote that part of the *single* a interval that lies across from σ . (So α is followed either by a b interval or another a interval.) Abusing notation somewhat, let $\bar{a} \cap x$ denote that part of the \bar{a} interval that lies on the x_1 peninsula. If α is longer than $\bar{a} \cap x$, we get a lens space contradiction between α and $\bar{a} \cap \sigma$. See Figure 20ii.

If α is shorter than $\bar{a} \cap x$, then, since $s = x + il(b)$ and b^i is the highest power of b , there is more a segment across from σ than just α . If there are some b intervals between α and the additional a segment, then the far (right-hand in the figure) end of the b -segment gives the same lens space contradiction. So we conclude that α is immediately followed by another copy of a , which we'll call a_1 .

If α is shorter than $\bar{a} \cap \sigma$, as must happen if most of \bar{a} lies in σ , we get a lens space contradiction between a_1 and the end of $\bar{a} \cap \sigma$. See 20iii. If α is longer than $\bar{a} \cap \sigma$ we get a lens space contradiction, comparing the end of α across from a \bar{b} segment with the end of the x_1 peninsula. See Figure 20iv.

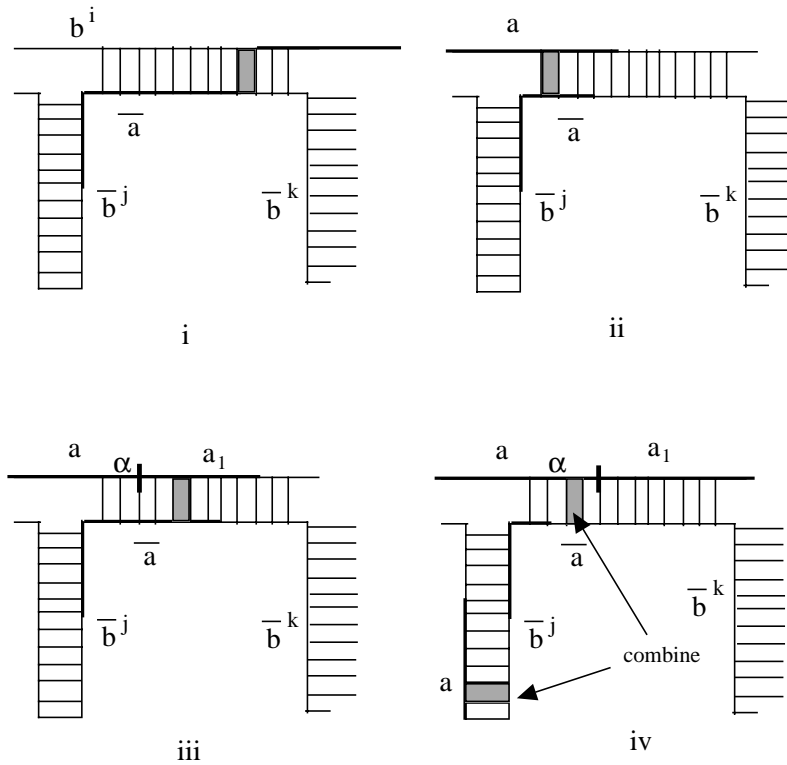


FIGURE 20.

Case 2b: ∂A is of the second form.

Observe then that $s + x = il(b)$, so $s = x + (j + k)l(b) + l(a) > l(b) + l(a)$. It follows that the distance between the outermost labels κ (see Case 1) is greater than $l(a)$. The rest now follows almost exactly as for the second form in Case 1. □

When the wave is based at μ^+ the result is less ambitious. We need the following lemma:

Lemma 7.2. *If $|\partial A \cap \mu^-| = 1$, then $e^\perp \cup e^+$ is unknotted; and symmetrically, if $|\partial A \cap \mu^+| = 1$, then $e^\perp \cup e^-$ is unknotted.*

Proof. Suppose $|\partial A \cap \mu^-| = 1$. Take two parallel copies of μ^- and band them together along the part of $\partial_- A$ that does not lie between them. The result is a disk $E \subset H$ that is disjoint from ∂A and separates H , leaving one of $\partial_\pm A$ in the boundary of each of the solid tori components of $H - E$. Label these solid tori (correspondingly) L_\pm and denote by L the link whose core circles are $L_- \cup L_+$. Note that $\partial_- A$ is a longitude of L_- and $\partial_+ A$ is a (p, q) cable of L_+ , for some q . L is visibly a non-hyperbolic (because of A) tunnel number one link (the tunnel is dual to E). These have been classified (cf. [EU]): In particular, L_+ is the unknot. But the core of L_+ is $e^\perp \cup e^+$, as required.

If $|\partial A \cap \mu^+| = 1$ (so $p = 1$), the argument is symmetric, interchanging μ^- and μ^+ . □

Proposition 7.3. *Suppose, in a thinnest Θ -graph θ appropriate for (K, F) , the edge e^+ is disjoint from a dividing sphere. Suppose also that F is of genus one and that the wave for ∂D is based at μ^+ . Then either $e^\perp \cup e^+$ is unknotted, or a Whitney move on θ changes it to an equally thin Θ -graph θ' that is appropriate for (K, F) .*

If θ presented K as a $(p, 1)$ quasi-cable, then θ' presents it as a $(p + 1, 1)$ quasi-cable.

Proof. With no loss of generality, assume that e^+ lies above the dividing sphere and that e^+ is monotonic.

The proof now has the same features as the proof of Theorem 7.1, and we use similar notation. Let $\omega \subset \partial\eta(\theta) - K$ be the arc as previously, again slid to minimize intersections with the meridians of $\eta(\theta)$. Let Σ again be the 4-punctured sphere lying in $\partial\eta(\theta) - P$, on the neighborhood of the component of $\theta - P$ that lies above P and contains e^+ .

For the purposes of the argument, we will assume that all p arcs of $K \cap \Sigma$ that run between the two copies of μ^\perp in $\partial\Sigma$ are parallel. If in fact there are two families of parallel arcs, the argument is essentially identical, except for one difference which is noted below.

Then $K \cap \Sigma$ consists of three families of arcs. One family of p arcs runs between the two copies of meridian μ^\perp in Σ ; two arcs each run from a copy of μ^\perp to a copy of μ^- . It is natural to parameterize slopes of proper arcs on Σ using these arcs of K . Indeed, the discussion will now, in some sense, be parallel to that of [ST1]. We declare the family of p arcs to have slope 0 and the second pair to have slope ∞ . An outermost disk of $D \subset \eta(\theta)$ cut off by the pair of meridians μ^-, μ^\perp defines a wave in Σ ; the wave assumption guarantees that such an outermost disk also intersects μ^+ , so we conclude that the wave has finite slope u/v in the coordinates just defined by the arcs $K \cap \Sigma$. Moreover, u is odd, since a wave in Σ will be based at each copy of a single meridian (either μ^- or μ^\perp) (see [ST1] for details). An argument will now show that either $e^\perp \cup e^+$ is unknotted, or $u/v = \pm 1$.

The arc ω is disjoint from the wave. Suppose to begin that ω intersects both meridians μ^- and μ^\perp . Then some arc component β of $\omega \cap \Sigma$ has one end on a copy of each of μ^- and μ^\perp in $\partial\Sigma$. Then the slope r/s of β is odd and can't differ from u/v , the slope of the waves, since if it did its ends would have to run between the base of both waves, i.e., different copies of the same meridian. On the other hand, since β is disjoint from K , which has one parallel family of arcs of slope 0 and two non-parallel arcs of slope ∞ , we have $|r| \leq 2$ (hence $r = \pm 1$) and $|s| \leq 1$. Since we are given that $u/v \neq \infty$, it follows that $u/v = r/s = \pm 1$, as claimed.

Next suppose that ω intersects μ^- but never μ^\perp . Then any component β of $\omega \cap \Sigma$ that has both ends on copies of μ^- in Σ will have slope 0 (since it's disjoint from K). The two terminal segments of ω in Σ will then each have one end on different copies of μ^- . But then they can't have their other end (i.e., the end points of ω) on the same side of K . For if they did, then either the arcs cross in the "square" component of $\Sigma - (K \cup \beta)$ in which they lie, or one must be part of a segment of $\partial D \cap \Sigma$ of slope ≥ 1 and the other of slope ≤ -1 . See Figure 21. (If not all p arcs of $K \cap \Sigma$ that run between the two copies of μ^\perp in $\partial\Sigma$ are parallel, the slopes of both these arcs could be ± 1 , still sufficient to deduce that this is the slope of the wave.)




FIGURE 21.

The only remaining possibility (to avoid the conclusion that $u/v = 1$) would be that ω is disjoint from μ^- . But that would imply that the boundary of the annulus A obtained by ∂ -compressing F to $\partial\eta(\theta)$ intersects μ^- in a single point. Then by Lemma 7.2 $e^\perp \cup e^+$ is unknotted.

So we continue, assuming that the slope of the waves in Σ is ± 1 . Now apply a Whitney move, replacing the meridian of e^+ (slope ∞) with the disk whose boundary has slope u/v . This redefines the Θ -curve as θ' , presenting K as a $(p+1, 1)$ quasi-cable. Moreover, it is no thicker than θ and it satisfies the wave condition, since the meridian of the new edge e'^\perp has been chosen to be disjoint from the wave. As usual, we can ensure that the Whitney disk has no effect on the bridge structure of K , so K remains in thin position. Recall Figure 8. \square

Proposition 7.4. *Suppose in a thinnest Θ -graph θ appropriate for (K, F) the edge e^- is disjoint from a dividing sphere, and suppose $\text{genus}(F) = 1$. Then $p = 1$ and either one of $e^\perp \cup e^\pm$ is unknotted, or a Whitney move on θ changes it to an equally thin Θ -graph θ' that is appropriate for (K, F) . Moreover, θ' presents K as a $(2, 1)$ quasi-cable.*

Proof. Without loss we assume e^- lies above a dividing sphere and e^- is monotonic. Suppose $p \geq 2$. Then K has a maximum at the lowest vertex. Find a level sphere as in Lemma 3.10. The result contradicts Lemma 2.5. Having established that $p = 1$, switch the labels of e^\pm and apply Proposition 7.3 or Theorem 7.1, depending on where the wave is. \square

Propositions 7.3 and 7.4 focus attention on the single remaining case to consider: when e^\perp is disjoint from a dividing sphere.

Theorem 7.5. *Suppose, in a thinnest Θ -graph θ appropriate for (K, F) , the edge e^\perp is disjoint from a dividing sphere. Suppose further that $\text{genus}(F) = 1$. Then either the cycle $e^+ \cup e^\perp$ is unknotted, or $p = 1$ and the cycle $e^- \cup e^\perp$ is unknotted.*

Proof. That $q = 1$ follows from Proposition 5.9. As previously, let A be the annulus obtained from ∂ -compressing F to $\eta(\theta)$ using the disk E from the splitting sphere. Without loss of generality, assume e^\perp lies above the dividing sphere. If there are no regular maxima of $\theta - e^\perp$, then $e^+ \cup e^\perp$ is vertical, and we are done. If there is a regular maximum of $\theta - e^\perp$, we can assume it's the lowest maximum. In that case, a level sphere just below the lowest maximum cuts off an upper disk from A and a level sphere just above the highest minimum cuts off a lower disk from A . So either some dividing sphere cuts off both an upper disk and a lower disk, or some




FIGURE 22.

dividing sphere P intersects A only in essential arcs. In the former case Corollary 6.7 finishes the proof.

The rest of the proof is an extended proof by contradiction. We will show that it is impossible for a dividing sphere to intersect A only in spanning arcs.

Let $\omega \subset \partial\eta(\theta) - K$ be the arc as previously, again slid to minimize intersections with the meridians of $\eta(\theta)$. Let P again be a level sphere between the highest minimum and the lowest maximum, and Σ again be the 4-punctured sphere lying in $\partial\eta(\theta) - P$, again supposing (with no loss) that both vertices are λ -vertices, so Σ lies above P . Let w be the word in a, b, \bar{a}, \bar{b} represented by ω . The wave condition now guarantees that (with the right choice of direction for w) w is positive in a and \bar{b} (including as usual the possibility that w is the empty word). If the wave is based at μ^- , then no proper power of \bar{b} occurs in w (i.e., no power greater than one); if it's based at μ^+ , then no proper power of a occurs in w . Exploiting these facts, together with the rotational symmetry of the diagram, we have several essentially different ways in which the ends of w can lie in Σ . Representative samples indicating that w can begin or end on any letters are shown in Figure 22. We have oriented w from left to right. Symmetric figures in which $w = \bar{b} \dots a$ are not shown. The waves themselves are also not shown, but they are described (except for details of how their ends lie near μ^\pm) by the requirement that they are disjoint from μ^\perp , which is shown. Note also that the number of arcs of K connecting the copies of μ^+ is now $p - 1$.

The case $p = 1$ is special, since in this case there are no arcs connecting the copies of μ^+ . Variants that arise in this case are shown in Figure 23.




FIGURE 23.

The resulting words corresponding to ∂_-A and ∂_+A can be put in two forms:

- (1)
 - $\partial_+A \leftrightarrow wb^i ab^j$,
 - $\partial_-A \leftrightarrow w\bar{b}^k$;
- (2)
 - $\partial_+A \leftrightarrow w\bar{b}^i \bar{a}b^j$,
 - $\partial_-A \leftrightarrow wb^k$.

Here $i, j, k \geq 0, i + j + k = p$.

Now apply thin position. It follows immediately that there is a level sphere P so that no arc of $P \cap A$ cuts off an outermost disk incident to ∂A , since there can't be both an upper and a lower such disk incident to ∂A simultaneously by Lemma 6.4. In particular, $P \cap A$ consists only of spanning arcs, so $l(\partial_-A) = l(\partial_+A)$. This immediately implies that $(k - i - j)l(b) = l(a)$. Thus it also forces $k > 0$. In the second form above, observe then that by regularity of ∂_-A , w begins and ends with the letter a , so by regularity of ∂_+A , $i, j > 0$.

In fact we will show that the second category above does not arise and the first is limited to the case $i = j = 0, k = p$.

Lemma 7.6. *The letters b and \bar{a} do not occur in the words determined by ∂A .*

Proof. Consider first the form

- $\partial_+A \leftrightarrow wb^i ab^j$,
- $\partial_-A \leftrightarrow w\bar{b}^k$.

We will show that $i = j = 0$, so $k = p$. Suppose, with no loss of generality, that $i > 0$, so the length of the segment $b^i ab^j$ (the same as the length of the segment \bar{b}^k) is at least $l(a) + l(b)$. The easy case is when the wave is based at μ^- , so there

are no repeating \bar{b} 's in w . Then opposite (in ∂_+A) to the segment \bar{b}^k is a complete collection of a -labels (perhaps separated by a single letter \bar{b}). On the other hand, opposite (in ∂_-A) to b^i must be part of an a segment. Combining the two easily gives a lens space contradiction.

It seems to be harder to establish a lens space contradiction in the case when the wave is based at μ^+ , so there are no repeating a 's in w . For the first time we need to use the graph G in the dividing sphere P whose vertices are the points of intersection of e^\pm with P (we will call these points the a - and b -vertices in G) and whose edges are the arcs $P \cap A$, viewed both as spanning arcs of A and as edges in G . Ends of edges at the same vertex in G will be said (usually in ∂A) to have the same label. The label will be an a -label, b -label, \bar{a} -label or \bar{b} -label depending on whether ∂A at that point (as oriented so that w is positive in a and \bar{b} as above) is passing through an a - or b -vertex in the direction that the θ -graph is oriented or in the opposite direction. Our use of the graph G in this lemma will be modest, mostly as a bookkeeping device. Once the lemma is established we will need to examine G much more seriously.

Orient all edges in G to point from ∂_+A to ∂_-A . We first claim there is an oriented path that begins with an edge incident to a b -label and ends with an edge incident to a \bar{b} -label. To see this, remove the b -vertices from G (but not their incident edges), and let G_b denote those edges, and the a -vertices they pass through, that are part of an oriented path beginning with an edge incident to a b -label. If there are m a -vertices in G_b and there are q occurrences of a in the word w , then G_b has $(q+1)m + (i+j)l(b)$ ends of edges from ∂_+A . Yet only qm ends of edges in ∂_-A could be in G_b but not at \bar{b} -labels. So some ends of the edges must be at \bar{b} -labels, as claimed. Consider then the shortest oriented paths beginning with a b -label and ending with a \bar{b} -label. Among all shortest such paths, pick a path ρ whose ends are closest in e^+ as measured by the number of components of $e^+ - P$ that lie between them. We claim that that number is one; i.e., ρ begins and ends at the opposite ends of the same interval of $e^+ - P$.

First note that the ends of ρ can't be at the same b -vertex, because ρ would then be a loop in P whose normal I -bundle, as pieced together from neighborhoods of the edges in A , would not be oriented. So call the initial b -vertex b_i and the terminal b -vertex b_t . To be concrete, suppose that, in a single b -letter of ∂_+A , b_t precedes b_i (we'll say that b_t lies to the left of b_i in the oriented ∂A). Unless the labels are precisely adjacent in ∂_+A (which is our claim), we can construct a better path ρ' as follows: Start at the b -label just to the left of the origin of ρ and construct a path by always using the edge that is one to the left (in A) of the edge in ρ .

Notice first of all that the collection of edges ρ' we have just described is indeed a path in G : Suppose α_1 and α_2 are successive edges of ρ' and the edges to their left in A are α'_1 and α'_2 . We need to show that the end of α'_1 in ∂_-A is at the same vertex as the end of α'_2 in ∂_+A . (See Figure 24.) This is obvious unless the end of α_1 at ∂_-A (and so the end of α_2 in ∂_+A) is the first label of an a segment. But if it were, then, since there are no repeating a 's in w , the end of α'_1 would in fact be a \bar{b} -label, and we would have found a shorter path. Having established that ρ' is in fact a path in G , notice that it ends just to the left of the label b_t in \bar{b} , hence to the right of the label b_t in b . Hence we have found a path of equal length but with ends closer together.




FIGURE 24.

Having established that the ends of ρ represent adjacent intersections of e^+ with P , carry through the above construction of ρ' . Consider the sequence of squares that lie between the two paths. When glued together along the components of $e^+ - P$ that their edges in ∂A represent, the result is a Möbius band whose boundary lies on the level sphere P . This is impossible (and represents a different way of viewing a lens space contradiction).

Next consider the form

- $\partial_+ A \leftrightarrow w\bar{b}^i\bar{a}\bar{b}^j$,
- $\partial_- A \leftrightarrow wb^k$.

Since $kl(b) = l(a) + (i + j)l(b) \geq l(a) + 2l(b)$ and, by the lens space argument, no part of the b^k interval can lie across from any part of a \bar{b} interval, it follows that no part of the b^k interval can lie across from any part of the $\bar{b}^i\bar{a}\bar{b}^j$ segment. In particular, across from the b^k interval must be an a segment longer than $l(a)$. This means that there are repeated a 's in w , hence no repeated \bar{b} 's. What then lies across from the \bar{a} interval? None of it can be part of an a segment, by the lens space argument, so it must be part of a single \bar{b} interval. Then $l(b) \geq l(a)$. On the other hand, $l(a) = (k - i - j)l(b) \geq l(b)$. We conclude that $l(a) = l(b)$ and immediately across from \bar{a} is precisely a single \bar{b} . But if this is the case, then each segment corresponding to a letter in one boundary component will lie exactly opposite a segment corresponding to a single letter in the other boundary component.

This is clearly a very special case, and resolving it will begin the process of understanding how to use the graph G effectively. We've already identified (across from b^k) a term in $\partial_+ A$ of the form $a^n, n \geq 3$. Across from that same term a^n in $\partial_- A$ must be three letters, at least one of which is also an a . There are then two letters a exactly aligned opposite each other, exhibiting that each a -vertex is the base of a loop in G . Ask then what lies in $\partial_- A$ exactly opposite \bar{b}^i . If any of it is an a interval, then this fact, together with the established fact that across from the \bar{a} interval is a \bar{b} interval, gives a lens space contradiction. If any of it is a \bar{b} interval, then we will have exhibited that every b -vertex in G is also the base of a loop. Then, since every vertex is the base of a loop, some such loop will contain no vertices in its interior. The following lemma shows that this is impossible. \square

Lemma 7.7. *Any loop in the graph G must have vertices in both disks into which the loop divides P .*

Proof. A loop without such vertices in its interior would give a problematic ∂ -compression of A . To see the problem, consider the base of an innermost such

loop, that is, the subarc σ of the meridian to which A is ∂ -compressed, an arc in $\partial(\eta(\theta)) - K$. The arc σ could not be a simple cocore of the band along ω ; otherwise F would have been compressible. On the other hand, if σ were incident to a copy of ω in A from the side in $\partial\eta(\theta)$ opposite to the band ω , then the ∂ -compression would turn A into an essential disk in $S^3 - \eta(\theta)$ whose boundary is disjoint from K , which is also absurd, for if we attach a neighborhood of the disk to $\eta(\theta)$, K would lie on the resulting unknotted torus and F would be a Seifert surface in the solid torus complement, forcing K to be trivial. The only remaining possibility is that σ has both ends incident to the parts of ∂A that come from K . But K crosses each meridian always in the same direction, so two such crossings can't be the ends of a ∂ -compression. \square

Lemma 7.7 completes the proof of Lemma 7.6. We conclude that the words corresponding to the boundary components of the annulus A are exactly

- $\partial_- A \leftrightarrow wa$,
- $\partial_+ A \leftrightarrow w\bar{b}^p$.

In fact, more can be said. The fact that there is no occurrence of the letter b in $\partial_- A$ or $\partial_+ A$ means that neither end of w can lie on the segments of $K \cap \Sigma$ connecting the two copies of μ^+ (when $p \geq 2$). Then

Corollary 7.8. *Suppose $p \geq 2$. If the $p - 1$ segments $\{\kappa_1, \dots, \kappa_{p-1}\}$ of $K \cap \Sigma$ that connect the two copies of μ^+ are not all parallel in Σ , then w begins and ends with the letter \bar{b} . If they are all parallel, then w begins or ends (perhaps both) with the letter \bar{b} .*

Proof. If some κ_i were incident to both ends of ω , then one component of ∂A would represent the word w and the other one $w\bar{b}^i a \bar{b}^j$. This is impossible, since these words have different lengths. If some κ_i contained a single end of ω , then one component of ∂A would contain an occurrence of the letter b , contradicting Lemma 7.6. So the ends of w lie, one each, on the two components of $K \cap \Sigma$ that are not among the κ_i . The result follows easily (see Figure 22). \square

The argument now proceeds by considering every possible type of word w . We begin by considering short words, then long words, then words of intermediate length.

Lemma 7.9. *If $w = \bar{b}^m$, $m \geq 0$ (e.g., $w = \emptyset$), then $e^\perp \cup e^+$ is unknotted.*

If $w = a^m$, $m \geq 1$, then $p = 1$ and $e^\perp \cup e^-$ is unknotted.

Proof. If $w = \bar{b}^m$ (or is empty), then ∂A intersects the meridian μ^- in exactly one point, a point in $\partial_- A$. If $w = a^m$, $m \geq 1$, then $p = 1$ by Corollary 7.8. Then ∂A intersects the meridian μ^+ in exactly one point, a point in $\partial_+ A$. In both cases the result follows from Lemma 7.2. \square

In view of Lemma 7.9 we can and will restrict our attention only to words that contain both letters a and \bar{b} .

To deal with longer words it will be useful to generalize Lemma 7.7. To appreciate how, we examine the local structure of G . The key to organizing the information is to orient each edge of G , as was done briefly above, so that the edge, when viewed in A , points from $\partial_+ A$ to $\partial_- A$. This has the obvious consequence that any a -vertex has at least one edge pointing into it, since the word wa contains the letter




FIGURE 25.

a , and each b -vertex has at least p edges pointing out, since the word $w\bar{b}^p$ contains at least p occurrences of the letter \bar{b} . Beyond these ends of edges, though, is one more pair at each a -vertex (resp. b -vertex) for each occurrence of a (resp. \bar{b}) in w . (One of the pair is identified with the occurrence of the letter in $w \subset \partial_+ A$ and the other with the occurrence in $w \subset \partial_- A$.) Call these ends of edges the w -ends. At any a -vertex there is a single non- w edge pointing in plus a sequence of w -ends alternating between pointing in and pointing out. At any b -vertex there are p non- w edges pointing out and a collection of w -ends, the latter coming in pairs of adjacent ends, one pointing in and one pointing out. See Figure 25.

Between each pair of w -ends in $\partial\mu^-$ or $\partial\mu^+$ is an arc ω_\perp that is a cocore of the band along ω . Put another way, banding together the boundary components of ∂A along ω_\perp would recover F from A . Call that side the w -side of the w -end. Since occurrences of a or \bar{b} in w occur always with the same sign, w crosses μ^- always in the same direction, and similarly for μ^+ . It follows that, at any given vertex of G , the w -side of any (oriented) w -end is always to the right (or always to the left) of the end, as the end is oriented by the orientation of its edge. Moreover, since the ends of ω_\perp are incident to the same side of A , that normal direction is well-defined, so if both ends of the same edge in G are w -ends, then the w -side is the same side at both ends. Combining these facts, we discover that, throughout any one component G_0 of G , the w -side of any w -end at any vertex lies always to the same side (say always to the right as the edge is oriented) of the w -end. Under these circumstances, note that a cycle in G_0 that has no vertices or edges in its interior, and which moves clockwise around its interior (we call it a *clockwise face*), must have corners that are always on the w -side of w -edges. In particular, if such a cycle can be found, then its interior would, before the ∂ -compression that changed F to A , correspond to a compressing disk E for F . (We know that ∂E would be essential in F , since it crosses a proper arc in F , namely the one ∂ -compressed to ω , always in the same direction.) Such a compression, of course, violates our assumption that F is incompressible.

Symmetrically, in a component of G for which the w -side of any w -end is to the left of the oriented end, there could be no counterclockwise cycle whose interior is empty (i.e., no counterclockwise face).

Although it might not be easy to see if a given component of the graph G in P is “right-handed” or “left-handed” in this sense, it is possible to use the extreme regularity of G (guaranteed by the fact that all edges in G are parallel in A) to identify circumstances in which there are *both* clockwise and counterclockwise cycles in the same component of G with no vertices in the interior of either. That, then,

forces one or the other to define a compression of F , a contradiction. For example, we have

Lemma 7.10. *No two faces of G can be adjacent and have boundaries that are cycles.*

Proof. Since the cycles are adjacent, one is clockwise and one is counterclockwise. \square

Definition 7.11. A disk component of $P - G$ will be called a *face*. If the boundary of the face is a cycle, we call it a *face cycle*. A clockwise (resp. counterclockwise) face cycle will be called a *clockwise* (resp. *counterclockwise*) *face*. A face incident only to a -vertices (resp. b -vertices) will be called an *a -face* (resp. *b -face*).

Lemma 7.12. *Any a -face is a face cycle. If $p = 1$, any b -face is a face cycle.*

Proof. At any a -vertex exactly one end of an edge is not a w -end, and it points into the vertex. Hence there cannot be two adjacent ends of edges pointing out, as there would be in an a -face that is not a cycle. \square

Lemma 7.13. *No two a -faces can be adjacent. If $p = 1$, no two b -faces can be adjacent.*

Proof. Combine Lemmas 7.10 and 7.12 \square

It is natural to seek features of the graph which guarantee the existence of cycles. The following lemma suggests a possibility.

Lemma 7.14. *No distinct b -vertices can have edges pointing to (resp. from) the same vertex.*

Proof. Suppose α_1, α_2 are two edges in G with their heads, say, at the same vertex of G . Then the ends of the α_i in $\partial_- A$ are some multiple of $l(b)$ apart. (Recall that $l(a) = pl(b)$.) The ends of the α_i in $\partial_+ A$ are the same distance apart. So if both ends on $\partial_+ A$ are at b -labels, they must be at the same b -label. \square

It seems from this lemma that bigons may be prevalent. To be precise, define a *bigon* in G to be a pair of edges, each running between the same pair of vertices. A *parallel bigon* will be a bigon in which both edges of the bigon are oriented in the same direction. An *anti-parallel bigon* will be one in which the edges are oriented in the opposite direction, forming a cycle in G of length two. In the case of a parallel bigon we will denote the vertex from which the edges of the bigon point out by v_+ and the vertex into which the edges point by v_- .

Lemma 7.15. *Suppose there is a parallel bigon in G , and let B be a disk in P that it bounds. Suppose in B there is an oriented path from v_- to v_+ . Then in the interior of B there is a cycle that is disjoint from both v_{\pm} .*

Proof. With no loss we may assume that B contains no other parallel bigon, else we would focus on an innermost one. Since w contains both letters a and \bar{b} , any vertex is incident to an edge pointing out and an edge pointing in; no vertex is a sink or source. Hence any component of G contains a cycle, so we may as well assume that every vertex in the disk belongs to the same component G_0 of G as the bigon. Suppose, with no loss (as explained above), G_0 contains no clockwise face.




FIGURE 26.

If there were an oriented path that runs from v_- to v_+ inside B , then the closed disk would contain both a clockwise and a counterclockwise cycle. Consider an innermost clockwise cycle (perhaps passing more than once through the same vertex, but not crossing at such a vertex) and the disk B' that it bounds. Suppose B' contains a vertex. That vertex must be part of an oriented path in the interior of B' . If that path forms a cycle completely in the interior of B' , we are done. If not (e.g., the ends of the path are at the same vertex of the cycle $\partial B'$), the path would cut off a clockwise cycle that would be even further in, a contradiction. (See Figure 26.) So there is no vertex in the interior of B' . Similarly, if there were an edge in the interior of B' , there would be a further in clockwise cycle. We conclude that B' would have to be a clockwise face, which is impossible. \square

Lemma 7.16. *Suppose there is a parallel bigon in G , and let B be a disk in P that it bounds. Then in the closure of B (i.e., including the vertices v_{\pm}) there is a cycle that includes at most one of v_{\pm} .*

Proof. As in the proof of Lemma 7.15, we can assume, with no loss, that B contains no other parallel bigon, that no vertex is a sink or source, that every vertex in the disk belongs to the same component G_0 of G as the bigon, and that G_0 contains no clockwise face.

If any end of G in B points out from v_- , then it is part of an oriented path (since no vertex is a source or sink). If the path ends in v_+ , we are done by Lemma 7.15. If not, it must contain a cycle not incident to v_+ , and we are done.

If the only ends of G in the bigon that are incident to v_- point into v_- , then there can be no edges other than the bigon itself, since even an a -vertex can have at most two adjacent ends pointing in. (Recall that w -ends alternate between pointing in and pointing out.) So we may as well assume that no end of G in the interior of B is incident to v_- . In that case, any end lying in the interior of B and incident to v_+ must be part of a cycle in the bigon incident only to v_+ , and we are done.

The possibility remains that the interior of the bigon is empty. In that case, at least one end of the edges of the bigon at each vertex is not a w -end (since w -ends alternate between pointing in and pointing out) so v_- is an a -vertex and v_+ is a b -vertex. If neither end at v_+ is a w -end, then, considering how non- w -ends arise, necessarily $p \geq 2$ and the edges of the bigon, when viewed in A , are some $kl(a)/p$, $1 \leq k < p$, apart. But then they can't have their other end at the same a -vertex, since two ends in $\partial_- A$ with the same a -label are at least $l(a)$ apart in A . Hence we conclude that exactly one end of the bigon at each of v_+ and v_- is




FIGURE 27.

a w -end. Necessarily their w -side is the same side and not the side in the bigon. Hence exactly one edge α_l in the bigon (the left one, say) has both of its ends w -ends, and *neither* of the ends of the other edge α_r are w -ends. In particular, α_r connects, in the annulus A , a point in $\partial_- A$ corresponding to a point in the last letter of wa , to a point in $\partial_+ A$ that lies in the final syllable \bar{b}^p of $w\bar{b}^p$.

Now $l(a) = pl(b)$; suppose the end of α_r in $\partial_+ A$ is the i th end in the final syllable \bar{b}^p of $w\bar{b}^p$ and the end in $\partial_- A$ is the j th end in the final letter a of wa . Suppose for concreteness that $i \geq j$. Now consider α_l . Since its ends are the same as those of α_r , the end of α_l in $\partial_- A$ is the j th end in some letter a in w . But the length of the terminal segment of w is of course the same in $\partial_+ A$ as in $\partial_- A$ and so the end of α_l in $\partial_+ A$ is the point exactly $i - j$ later than the end is in $\partial_- A$. That is, it is still part of the same letter a of w . (See Figure 27.) This is a contradiction, since the other end of α_l is at a b -vertex. If $i \leq j$, we get the same contradiction, using the initial segment of w instead of the terminal segment. \square

With a little determination, we have more:

Lemma 7.17. *The interior of any parallel bigon contains a b -vertex.*

Proof. With no loss we can assume that the parallel bigon is innermost. Suppose the disk B contains only a -vertices in its interior. By Lemma 7.16, B contains a cycle that passes through some vertices in its interior. Since the valence of each vertex is greater than two, there are adjacent faces in B . Then Lemma 7.13 shows that v_\pm can't both be a -vertices. Similarly, all vertices in B are in the same component of G as the bigon.

Let Γ denote the subgraph of G , lying in the interior of B , obtained by deleting all edges incident to v_\pm .

Claim. Γ contains a cycle.

Otherwise consider an oriented path from a source vertex a_+ to a sink vertex a_- . All edges (of G) pointing into a_+ must have their other ends at v_\pm and not the same vertex, since this would exhibit a further in parallel bigon. There must be at least two such edges, since by Lemma 7.9, w contains at least one occurrence of a , so the word wa contains at least two. This implies that there are exactly two edges pointing into a_+ , and one edge comes from each of v_\pm . Then Lemma 7.14 implies that v_\pm can't both be b -vertices. Suppose that v_+ were an a -vertex. We have already identified three edges pointing out from v_+ : the edges of the bigon and an edge to a_+ . Then three edges point out from a_- , and at least two would have to go to the same vertex in the pair v_\pm , creating a further in bigon. So v_+ is




FIGURE 28.

a b -vertex and v_- is an a -vertex. Now consider a_- . If an edge in G pointing out from a_- goes to v_- , then v_- , hence every a -vertex, has three edges pointing in. This would force a_+ to be part of a further in bigon, a contradiction. So the edge pointing out from a_- goes to v_+ and this edge, together with the path from a_+ to a_- , together with the edge from v_- pointing into a_+ , gives an oriented path from v_- to v_+ . Then Lemma 7.15 provides a cycle in Γ , as claimed. \square

Having established the claim, we continue with the proof of Lemma 7.17. Note that, once we have a cycle in Γ , we know that Γ contains a face, hence, following Lemma 7.12, an a -face cycle.

Case 1: w contains only one occurrence of a .

In this case each a -vertex has valence three – two ends of edges pointing in and one pointing out. (So, for example, v_+ must be a b -vertex.) We have seen that Γ contains an a -face cycle σ . Now notice that there are only two possible “corners” of cycles at each a -vertex, since the valence is three; one set of corners occurs only in clockwise cycles and one in counterclockwise cycles. As a result, all the corners of σ come from a single occurrence of the letter a in ∂_-A and a single occurrence of the letter a in ∂_+A . We claim this is impossible: Of all the vertices of σ , let a_l be the one that is first encountered when passing along the oriented edge e^- , so, in any occurrence of the letter a in ∂A , the label corresponding to a_l lies most to the left among all labels coming from vertices of σ . Similarly, define a_r to be the last vertex of σ that is encountered along the oriented edge e^- . We repeat: the edges pointing out from a_l and from a_r , when viewed in A , leave from the same a interval a_+ in ∂_+A and end in the same a interval a_- in ∂_-A . But, by definition of a_r and a_l , the edge pointing out from a_r in $a_+ \subset \partial_+A$, goes to a label to the left of a_r in $a_- \subset \partial_-A$ whereas the edge pointing out from the label a_l in a_+ goes to a label to the right of a_l in a_- . This presents a clear contradiction: there are more ends of edges between the two edges in ∂_+A than there are in ∂_-A . See Figure 28.

Case 2: w contains three or more occurrences of a .

In this case we will show that there are two adjacent a -faces in Γ , contradicting Lemma 7.13. Let Γ_0 be a component of Γ and let k denote the number of vertices in Γ_0 . If v_{\pm} are both b -vertices, then at each a -vertex in B there can be at most two edges also incident to one of the v_{\pm} , one pointing in and one pointing out, since two different b -vertices can't have edges pointing toward (or away from) the same a -vertex (Lemma 7.14), and if one b -vertex had two edges pointing toward (or away from) the same a -vertex, it would be a further in bigon. On the other hand,

if either vertex (say v_-) were an a -vertex, then at most two edges could be incident to both Γ_0 and v_- , for otherwise G would have adjacent a -faces. So clearly $k \geq 2$ (the valence of each vertex is at least 7) and at most $2k + 2 \leq 3k$ edges connect Γ_0 to v_{\pm} in G . It follows that Γ_0 contains at least $7k/2 - 3k/2 = 2k$ edges. The proof that this would provide two adjacent a -faces now follows from this simple observation:

Claim. *If Γ is a connected graph in the plane with k vertices and at least $2k - 1$ edges, then either Γ contains a trivial loop, or two faces of Γ (i.e., compact complementary components, necessarily disks) are adjacent.*

Here's the proof of the claim. Let $e \geq 2k - 1$ be the number of edges. If any face is a monogon, we are done. So suppose every face has at least two edges. Either some edge is incident to two faces, and we are done, or the number of faces $f \leq e/2$. Then consider the Euler characteristic of Γ and all its faces: $1 = v - e + f \leq v - e/2 \leq k - (k - 1/2) = 1/2$, a contradiction. \square

Case 3: w contains exactly two occurrences of a .

Suppose first that v_{\pm} are both b -vertices. As usual, for each of the vertices in the interior of B , there is at most one edge pointing from v_{\pm} to the interior vertex, by Lemma 7.14 and our assumption that B is an innermost bigon. Now consider a component Γ_0 of Γ with, say, k vertices. We have just shown that there are at least 2 edges in Γ pointing into each a -vertex (for at most one edge pointing into the a -vertex comes from a b -vertex). Hence there are at least $2k$ edges in Γ_0 . The proof now follows as in the previous case.

Suppose finally that one of v_{\pm} is a b -vertex v_b and one is an a -vertex v_a . If, for some component Γ_0 of Γ there are no edges (in G) running from Γ_0 to the b -vertex, we are done much as before. Similarly, if there is at most one edge α in G pointing from v_a to Γ_0 , we are done: In Γ_0 there are still at least 2 edges pointing into every vertex, except for the single vertex at the end of α . Hence there are at least $2k - 1$ edges in Γ_0 , and we can still apply the combinatorial claim above. If there are at least two edges, say α_1 and α_2 , in G pointing from v_a to Γ_0 , then, since no adjacent ends at an a -vertex point out, there is another end of an edge α_3 between the ends of the α_i at v_a . If α_3 also goes from v_a to Γ_0 , then on either side of it are adjacent a -faces, contradicting Lemma 7.13. If instead it goes to another component of Γ , then that component is cut off from v_b by $\Gamma \cup \alpha_1 \cup \alpha_2$, so no edge in G connects it to v_b , a case we have already established. \square

Lemma 7.16 immediately eliminates the possibility that w is a long word. Explicitly, we have:

Lemma 7.18. *The letter \bar{b} occurs at most once in w .*

Proof. Suppose α and α' are two edges in G that point out from the same b -vertex. Then (echoing the argument of Lemma 7.14) the distance between α and α' as measured along either of $\partial_{\pm}A$ is some multiple of $l(b) = l(a)/p$. In particular, there are at most p candidates for a -vertices the other ends of α and α' might be incident to, plus a b -vertex. If the letter \bar{b} occurs more than once, then in ∂_+A there are at least $p + 2$ occurrences of each b -label. And, for each b -label, we have just argued that there are at most $p + 1$ possible labels in ∂_-A to which they can point, p of them a -labels and one a b -label. Hence at least two of the edges point to the same label. This shows that every b -vertex is part of a parallel bigon. An




FIGURE 29.

innermost parallel bigon then would have to contain only a -vertices, contradicting Lemma 7.17. \square

Lemma 7.19. *If \bar{b} occurs in w , then the letter a occurs at most once.*

Proof. Following Lemmas 7.9 and 7.18, we can restrict to the case in which \bar{b} occurs exactly once in w . We will assume that a occurs $m \geq 2$ times and derive a contradiction. The structure of the proof depends on whether $p = 1$ or $p \geq 2$.

If $p = 1$, then any two occurrences of the same a -label or the same b -label in $\partial_{\pm}A$ occur a multiple of $l(a) = l(b)$ apart. It follows that at least two edges pointing into any given a -vertex in G have their other ends at the same vertex. In other words, each a -vertex is contained in some parallel bigon. Thus an innermost parallel bigon contains only b -vertices. Any b -vertex has at least one edge pointing out that goes to an a -vertex (since there is only one occurrence of the letter \bar{b} in ∂_-A but two in ∂_+A), and edges that point from b -vertices to a given a -vertex must all come from the same b -vertex by Lemma 7.14. It follows that there are at most two b -vertices in the interior of the bigon. If there were only one, then the two edges coming out from it can't go to the same vertex, for that would form another parallel bigon, so one goes to each of the vertices forming the bigon. This would force v_+ to be a b -vertex and v_- to be an a -vertex. Furthermore, b -vertices are of valence 3, and the edge pointing into the interior b -vertex has nowhere to come from but v_- . The two adjacent face cycles contradict Lemma 7.10. See Figure 29i.

So there are exactly two b -vertices inside the bigon. As noted above, at least one edge pointing out from each of these b -vertices must go to an a -vertex, and edges can't point from different b -vertices to the same a -vertex. It follows that both v_{\pm} are a -vertices. If both of the edges pointing out from a b -vertex go to a -vertices, then, since $p = 1$, they would in fact go to the same a -vertex, contradicting our assumption that the parallel bigon is innermost. So each b -vertex has an edge pointing from it to the other b -vertex. In other words, the two b -vertices in the bigon are the vertices of a 2-cycle, necessarily a b -face cycle. (See Figure 29ii.) But this leads to the same contradiction as in the proof of Case 1 of Lemma 7.17.

If $p \geq 2$, then we know from Corollary 7.8 that w either begins or ends in \bar{b} , so the words corresponding to $\partial_{\pm}A$ are (up to cyclic rotation) $a^m\bar{b}^{p+1}$ and $a^{m+1}\bar{b}$ respectively. Here $l(a) = pl(b)$ as usual. Suppose first that no edge in G runs from

one b -vertex to another b -vertex. Then the $p + 1$ occurrences of any b -vertex in $\partial_+ A$ have their other ends in only the p possible a -vertices, so there is a parallel bigon at each b -vertex. An innermost parallel bigon then contains only a -vertices. This contradicts Lemma 7.17.

Suppose then that some edge β in G has both ends at b -vertices. Consider the distance in A between copies of the same a -label, counting distance (i.e., intersection with P) along the arcs of ∂A that *don't* intersect β . Measured on this side, the distance between any two copies of the same a -label in either component of ∂A is a multiple of $l(a)$. It follows that the $m + 1 \geq 3$ copies of the same label in $\partial_- A$ have their other ends at at most two labels in $\partial_+ A$, one an a -label and one a b -label. Thus in this case every a -vertex is part of a parallel bigon. Then an innermost bigon contains only b -vertices. The proof now follows as for the case $p = 1$ but is easier, since any b -vertex has $p + 1 \geq 3$ edges pointing out. \square

At this point there are only two remaining words to consider: $w = \bar{a}\bar{b}$ and $w = \bar{b}a$. Eliminating these two requires a bit more detailed argument.

Lemma 7.20. $w \neq \bar{a}\bar{b}$ or $\bar{b}a$.

Proof. The cases are symmetric, so without loss of generality suppose $w = \bar{a}\bar{b}$. Then $\partial_- A$ is represented by the word $\bar{a}\bar{b}a$ and $\partial_+ A$ is represented by the word $\bar{a}\bar{b}^{p+1}$. Let $n = l(a) = pl(b)$, $p \geq 1$. The type of contradiction depends on how the copies of \bar{b} in ∂A are aligned with each other. There are three cases. See Figure 30, where the orientation of w is meant to be clockwise around A .

Case 1: The \bar{b} -segment in $\partial_- A$ lies completely opposite a subsegment of the \bar{b}^{p+1} segment in $\partial_+ A$ (Figure 30i).

Then symmetrically, the a -segment in $\partial_+ A$ lies completely opposite a subsegment of the a^2 segment in $\partial_- A$. Consider the collection of edges incident to the copy of the a segment in $\partial_+ A$. Among those edges, every a -label occurs exactly once in $\partial_+ A$ and once in $\partial_- A$. It follows that these edges, when viewed in $G \subset P$, form a collection of a -cycles containing every a -vertex. Similarly there is a collection of b -cycles containing every b -vertex. An innermost pure a - or pure b -cycle can then contain no vertices in its interior and so must be an a -face or a b -face. Either way, the argument presented in Case 1 of Lemma 7.17 presents a contradiction.

Case 2: The \bar{b} -segment in $\partial_- A$ lies completely opposite a subsegment of the a segment in $\partial_+ A$ (Figure 30ii).

Then dually the entire \bar{b}^{p+1} segment in $\partial_+ A$ lies completely opposite a subsegment of the a^2 segment in $\partial_- A$. This means that for each b -vertex, each of the $p + 1$ edges in A with that label in $\partial_+ A$ can go to at most p different a -labels in $\partial_- A$. It follows that every b -vertex is part of a parallel bigon. An innermost one can contain only a -vertices, contradicting Lemma 7.17.

Case 3: The \bar{b} segment in $\partial_- A$ lies partly opposite an end of the a segment in $\partial_+ A$ and partly opposite an end of the \bar{b}^p segment (Figure 30iii).

The argument in this case is a kind of degenerate variant of the argument in Lemma 7.16. Suppose, with no loss and as shown in the figure, that part of the \bar{b} segment in $\partial_- A$ is opposite the beginning end of \bar{b}^p , overlapping say on $j < l(b)$ edges. ($j = 3$, $l(b) = 5$ in the figure). Now consider any of the last $n - j$ labels in the first occurrence of \bar{b} in $\partial_+ A$ and the corresponding label in the last occurrence of \bar{b} .




FIGURE 30.

The distance between them is $l(a)$ in $\partial_+ A$. It follows that each of the corresponding $(n - j)$ b -vertices is part of a parallel bigon in P , each with an a -vertex for v_- .

Consider an innermost parallel bigon in P and the disk B that it bounds. The interior of B must contain b -vertices, by Lemma 7.17, and it must also contain a -vertices, by the argument of Case 1 of the proof of that lemma. We claim that there are an oriented edge pointing from v_- into B and an oriented edge pointing out from B into v_+ . To see the former, consider the a -vertex a_l which, among all a -vertices lying in B , is the first encountered by e^- . Then the corresponding label in the second copy of a in $a^2 \subset \partial_- A$ lies across from the label of an earlier vertex

in e^- ; hence the label is that of v_- , since there is no alternative. So the edge in P between them must connect a_l to v_- , pointing toward a_l . A symmetric argument, using the last b -vertex b_r in the interior of B encountered by e^+ , shows that there is an edge pointing from b_r into v_+ .

Since each a -vertex has valence 3, we have now accounted for all edges incident to v_- . In particular, one of the two edges of the bigon has a w -end at v_- with the w -side lying within the bigon. What's important here is not that one of the edges of the bigon has a w -end at v_- – that fact can be seen simply because one of the ends of the edges of the bigon, viewed in A , lies in the second occurrence of a in a^2 (see Figure 30iii), hence in the $a\bar{b}$ section of the word $wa = a\bar{b}a$. What is important is that the w -side of this edge lies in the interior of the bigon. But examining the figure again, we see that the *other* edge of the bigon lies in the first occurrence of \bar{b} in \bar{b}^{p+1} , hence in the w section of the word $w\bar{b}^p = a\bar{b}b^p$. That is, the *other* edge has a w end at v_+ . So the w -side of that edge lies *outside* of B . On the other hand, we've shown that some edge at v_+ points into v_+ from the interior of B and, since any b -vertex has only one edge pointing into it, that edge must be adjacent to the (only) w -corner at v_+ , so that corner must be *inside* B . This contradiction proves the lemma, hence the theorem. \square

Having eliminated every possible word for w , we deduce that no dividing sphere can intersect A only in essential arcs, completing the proof of Theorem 7.5 \square

8. THE GODA-TERAGAITO CONJECTURE

Theorem 8.1. *Suppose K is a tunnel number one knot of genus one and γ is an unknotting tunnel. Then either there is a genus one Seifert surface F for K that contains γ , or γ can be slid and isotoped until it is an unknotted loop.*

Proof. According to Proposition 5.7, if neither of the outcomes above occurs, then there is an appropriate Θ -graph for (K, F) , thinly presenting it, say, as a (p, q) quasi-cable. Let θ be a thinnest appropriate Θ -graph for (K, F) and, among all such possibilities, choose one with p maximal. According to Propositions 5.7 and 5.9, $q = 1$ and θ is in bridge position.

We claim that if $p \geq 2$ the cycle $e^\perp \cup e^+$ is unknotted and, if $p = 1$, one of the two cycles $e^\perp \cup e^\pm$ is unknotted. This follows immediately from Corollary 5.13 unless a dividing sphere is a critical sphere that is disjoint from some edge. Consider the possibilities for such an edge: If the disjoint edge is e^\perp , then the claim is established by Theorem 7.5. If e^- is the disjoint edge, then, since θ has been chosen to have maximal p , it follows from Proposition 7.4 that $p = 1$ and one of $e^\perp \cup e^\pm$ is unknotted, establishing the claim. Similarly, if e^+ is the disjoint edge, then it follows from Proposition 7.3 that the wave is based at μ^- . Then Theorem 7.1 establishes the claim. So the claim is established in all cases.

Now let L be the unknotted solid torus neighborhood of $e^+ \cup e^\perp$ in H . Since $q = 1$, we can apply the “vacuum cleaner trick”: slide the ends of the 1-handle corresponding to e^- along the arc $K \cap L$ until K has been made disjoint from a meridian of L . At that point, L has become a tunnel for K and remains unknotted. \square

Corollary 8.2 (Goda-Teragaito Conjecture). *Suppose K is a tunnel number one knot of genus one that is not a satellite knot. Then K is 2-bridge.*

Proof. Let γ be an unknotting tunnel for K . If γ can be slid and isotoped to lie on a genus one Seifert surface F , then K is necessarily 2-bridge (see Corollary 5.4 of [ST1]). If not, then according to Theorem 8.1 γ can be slid and isotoped until it is an unknotted loop. The following argument (shown to me by Abby Thompson) shows that then K is 1-bridge on an unknotted torus. Let W denote the solid torus neighborhood of the loop, containing a short, ∂ -parallel arc of K . Let K_- denote the arc of K that lies outside of W . Since $\eta(K \cup \gamma)$ is an unknotted handlebody, it follows that the 1-handle with K_- at its core constitutes a genus two Heegaard splitting of the solid torus $S^3 - W$. Any non-trivial splitting of a handlebody (e.g., of $S^3 - W$) is stabilized [ST3], so in fact K_- is also parallel to ∂W . This shows that K is 1-bridge with respect to the unknotted torus ∂W .

Matsuda [Ma] has proven the statement for this class of knots. □

REFERENCES

- [BZ] G. Burde and H. Zieschang, *Knots*, de Gruyter Studies in Mathematics **5**, Walter de Gruyter, Berlin - New York (1985). MR **87b**:57004
- [EM] M. Eudave-Muñoz, *On non-simple 3-manifolds and 2-handle addition*, Topology Appl. **55** (1994) 131-152. MR **95e**:57029
- [EU] M. Eudave-Muñoz and Yoshiaki Uchida, *Non-simple links with tunnel number one*, Proc. AMS, **124** (1996) 1567-1575. MR **96d**:57007
- [G] D. Gabai, *Foliations and the topology of 3-manifolds. III*, J. Differential Geom. **26** (1987), 479-536. MR **89a**:57014b
- [GST] H. Goda, M. Scharlemann and A. Thompson, *Levelling an unknotting tunnel*, Geometry and Topology **4** (2000) 243-275. MR **2002h**:57011
- [GT] H. Goda, M. Teragaito, *Tunnel number one genus one non-simple knots*, Tokyo J. Math. **22** (1999) 99-103. MR **2000j**:57011
- [GL] C. Gordon and R. Litherland, *Incompressible planar surfaces in 3-manifolds*, Top. Appl. **18** (1984), 121-144. MR **86e**:57013
- [Ma] H. Matsuda, *Genus one knots which admit (1, 1)-decompositions*, Proc. AMS. **130** (2001), 2155-2163.
- [Mo] K. Morimoto, *Planar surfaces in a handlebody and a theorem of Gordon-Reid*, Proc. Knots '96, ed. S. Suzuki, World Sci. Publ. Co., Singapore (1997), 127-146. MR **99k**:57040
- [MS] K. Morimoto, M. Sakuma, *On unknotting tunnels for knots*, Math. Ann. **289** (1991) 143-167. MR **92e**:57015
- [OZ] R. P. Osborne and H. Zieschang, *Primitives in the free group on two generators*, Inv. Math. **63** (1981) 17-24. MR **82i**:20042
- [Sc] M. Scharlemann, *Outermost forks and a theorem of Jaco*, Combinatorial methods in topology and algebraic geometry (Rochester 1982), Contemp. Math. **44** (1985) 189-193. MR **87e**:57021
- [Sc2] M. Scharlemann, *The Goda-Teragaito Conjecture: an overview*, Surikaiseikikenkyusho Kokyuroku **1229** (2001) 87-102 or <http://front.math.ucdavis.edu/math.GT/0108079>.
- [ST1] M. Scharlemann and A. Thompson, *Unknotting tunnels and Seifert surfaces*, preprint (<http://front.math.ucdavis.edu/math.GT/0010212>).
- [ST2] M. Scharlemann and A. Thompson, *Thinning genus two Heegaard spines in S^3* , to appear.
- [ST3] M. Scharlemann and A. Thompson, *Thin position and Heegaard splittings of the 3-sphere*, Jour. Diff. Geom. **39** (1994), 343-357. MR **95a**:57026
- [T] A. Thompson, *Thin position and bridge number for knots in the 3-sphere*, Topology **36** (1997), 505-507. MR **97m**:57013

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, SANTA BARBARA, CALIFORNIA 93106

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