

An infinite family of links with critical bridge spheres

DANIEL RODMAN

A closed orientable splitting surface in an oriented 3–manifold is a topologically minimal surface of index n if its associated disk complex is $(n-2)$ –connected but not $(n-1)$ –connected. A critical surface is a topologically minimal surface of index 2. In this paper, we use an equivalent combinatorial definition of critical surfaces to construct the first known critical bridge spheres for nontrivial links.

57M25

1 Introduction

In the 1960s Haken developed a framework for studying manifolds that contain incompressible surfaces. He demonstrated that for such manifolds, one can reduce the proofs of many theorems to induction arguments by using “hierarchies”, consecutively cutting a manifold into pieces along incompressible surfaces until a collection of balls is obtained. This powerful approach clearly demonstrated the utility of incompressible surfaces in the study of low-dimensional topology.

In 1987 Casson and Gordon [4] introduced the idea of strongly irreducible surfaces. Unlike incompressible surfaces, which have no compressing disks, strongly irreducible surfaces have potentially many compressing disks to both sides, but any two compressing disks on opposite sides necessarily intersect. Somewhat surprisingly, it turns out that many theorems that are easy to prove for manifolds with incompressible surfaces also hold true for manifolds with strongly irreducible surfaces, although the proofs can be somewhat more involved. Casson and Gordon’s watershed result in the theory of strongly irreducible surfaces is that if a closed, connected, orientable 3–manifold’s minimal-genus Heegaard splitting is not strongly irreducible, then the manifold must contain an essential surface [4].

In 2002, Bachman [1] introduced the notion of a critical surface, which he followed in 2009 [2] with the introduction of the more general concept of a topologically minimal surface. A closed orientable splitting surface in an oriented 3–manifold M

is a topologically minimal surface of index n if its associated disk complex is $(n-2)$ -connected but not $(n-1)$ -connected. This latter definition provides the framework into which incompressible, strongly irreducible, and critical surfaces all fit: they are topologically minimal surfaces of indices 0, 1 and 2 respectively. Just like the incompressible and strongly irreducible surfaces before them, critical surfaces and topologically minimal surfaces in general have been used to prove long-standing conjectures that had remained unresolved for many years. For example, Bachman [3] used these surfaces to prove the Gordon conjecture¹ and to provide counterexamples to the stabilization conjecture.²

Most of the work done thus far regarding topologically minimal surfaces deals specifically with surfaces which are Heegaard splittings, two-sided surfaces that split a manifold into two compression bodies. Less is known about topologically minimal bridge surfaces, which are a natural type of surface to consider when the 3-manifold is a link complement. Lee has shown in [8] that all bridge spheres for the unknot with any number of bridges in S^3 are topologically minimal, and his results provide upper bounds for the indices of such bridge spheres. In particular, he concludes that the bridge sphere for an unknot in a 3-bridge position has index exactly 2 (ie it is critical).

In this paper we provide the first known examples of critical bridge spheres for nontrivial links. Our construction is inspired by recent work of Johnson and Moriah [6] in which they construct links with bridge surfaces of arbitrarily high distance. The central result of this paper is the following:

Theorem 6.9 *There is an infinite family of nontrivial links with critical bridge spheres.*

In Section 2 we go over some of the foundational topological definitions upon which this paper relies, including an equivalent, combinatorial definition of a critical surface. Then in Section 2 we describe what it means for a link to be in a plat position. Following that, in Section 3 we embed a link L in S^3 in a plat position with bridge sphere F and discuss some of the specific details of the embedding as well as build some of the tools (certain arcs, loops, disks and projection maps) which we will use throughout the rest of the paper. Sections 2 and 3 should be considered setup for the rest of the paper, and this is essentially the same setup as in Johnson and Moriah's paper [6]. In particular, we make use of Johnson and Moriah's plat links and projection maps. In Section 4 we

¹Gordon conjecture: if the Heegaard splitting $U \cup_H V$ is a connected sum of two Heegaard splittings, and H is stabilized, then one of its summands is stabilized.

²Stabilization conjecture: given any pair of Heegaard splittings, stabilizing the higher-genus splitting once results in a stabilization of the other splitting.

develop a link-diagrammatic way to visualize boundary loops of compressing disks for F . Theorem 6.9 is proved in Sections 5 and 6.

2 Definitions

Suppose Σ is a compact orientable surface embedded in a compact orientable 3-manifold M , and D is a disk embedded in M (not necessarily properly). We say that D is a *compressing disk* for Σ if $D \cap \Sigma = \partial D$ and ∂D neither bounds a disk in Σ nor is parallel to a boundary component of Σ .

A compact, orientable surface Σ embedded in a connected 3-manifold M is a *splitting surface* if $M \setminus \Sigma$ has two components. We associate a simplicial complex Γ called the *disk complex* to (M, Σ) in the following way: Vertices of Γ are isotopy classes of compressing disks for Σ . A set of $m + 1$ vertices will be filled in with an m -simplex if and only if the corresponding isotopy classes of compressing disks have pairwise disjoint representatives. For $n \geq 0$, Σ is called *topologically minimal of index n* if Γ is $(n-2)$ -connected but not $(n-1)$ -connected. (Note that every space is (-2) -connected, and a space is (-1) -connected if and only if it is nonempty.) Σ is called a *critical surface* if it is a topologically minimal surface of index 2. In [2], Bachman gives an alternative, combinatorial definition for critical surfaces and proves the two definitions are equivalent. This second definition, given below, is the one we will use in this paper.

A *critical surface* is a splitting surface $\Sigma \subset M$ whose isotopy classes of compressing disks can be partitioned into $\mathcal{C}_1 \sqcup \mathcal{C}_2$ in a way that satisfies the following two conditions:

- (1) Whenever $[C] \in \mathcal{C}_1$ and $[D] \in \mathcal{C}_2$ for C and D on opposite sides of Σ , $\partial C \cap \partial D$ must be nonempty.
- (2) For $i = 1, 2$, there exists a pair of disjoint compressing disks, one on either side of Σ , where each disk belongs to an isotopy class in \mathcal{C}_i .

Any link $L \subset S^3$ can be isotoped so that all of its maxima lie above all of its minima (with respect to the standard height function on S^3). After such an isotopy, L is said to be in *bridge position*. Let $\eta(L)$ be an open regular neighborhood of L , and let F be a level sphere in S^3 which separates all of the maxima of L from all of the minima. Then the surface $F' = F \setminus \eta(L) \subset M = S^3 \setminus \eta(L)$ (a sphere with a finite number of open disks removed) is called a *bridge sphere* for L . See Figure 2.

In our context, studying the bridge sphere $F' \subset S^3 \setminus \eta(L)$ is equivalent to studying the “sphere $F \subset S^3$ with marked points”, where the *marked points* are the points of $L \cap F$.




Figure 1: If a bridge disk on one side of F intersects a bridge disk below F in a single point of the link, and are disjoint otherwise, then the link is perturbed.




Figure 2: A bridge sphere and a compressing disk

A disk in S^3 is considered a compressing disk for the sphere F with marked points if and only if it is a compressing disk for F' in M . This implies that a compressing disk for F cannot intersect L , and two compressing disks for F are considered to be in the same isotopy class if and only if they are isotopic in $S^3/\eta(L)$.

A *bridge arc* is a component of $L \setminus F$, and each of these bridge arcs has exactly one critical point, so each is parallel into F . An isotopy of a bridge arc into F sweeps out a disk, called a *bridge disk*. (Bridge disks are not unique, even up to isotopy.)

A link is *perturbed* if there is a bridge disk above F and another bridge disk below F which intersect in a point of L , and which are otherwise disjoint. See Figure 1. We will say the link is *perturbed at* a bridge arc α if α corresponds to either of these two bridge disks. If a link L is perturbed, then there is an isotopy of L through the bridge disks which reduces the number of critical points of L by one maximum and one minimum.

A *cap* for F is a compressing disk C such that ∂C bounds a disk in F that contains exactly two marked points. The existence of caps is guaranteed by the existence of bridge disks. See Figure 3.

The notion of boundary-compression along boundary-compressing disks will also be useful. Suppose Σ is a surface with boundary properly embedded in a 3-manifold M .




Figure 3: The two disks depicted are both caps for the leftmost bridge arc. In each figure, the disk in F with two marked points is shaded.




Figure 4: A boundary-compressing disk is shaded dark gray.

A *boundary-compressing disk* (See Figure 4 for an example) is defined to be an embedded disk D with the following properties:

- $D \cap \partial M$ is an arc β .
- $D \cap \Sigma$ is an arc $\tilde{\beta}$.³
- $\partial D = \beta \cup \tilde{\beta}$.
- $\beta \cap \tilde{\beta} = \partial\beta = \partial\tilde{\beta}$.

Next we define a half-twist. Let U be a twice punctured (marked) disk. The half-twist can be simply understood to be the homeomorphism $U \rightarrow U$ depicted in Figure 5. To be technical, the *half-twist* H of U is the homeomorphism $H: U \rightarrow U$ defined as follows: View U as a union $A \cup_{\phi} D$ of the closed unit disk $D \subset \mathbb{C}$ with an outer annulus A whose intersection is an embedded circle containing the two marked points. See Figure 5. Give D the coordinates $re^{i\theta}$, and identify A with $S^1 \times I$, with the coordinates (θ, t) (here A is glued to D by the map ϕ taking $(\theta, 0)$ in one component of ∂A to $e^{i\theta}$ in ∂D). Now define

$$H(p) = \begin{cases} (\theta + \pi t, t) & \text{for } p = (\theta, t) \in A, \\ re^{i(\theta+\pi)} & \text{for } p = re^{i\theta} \in D. \end{cases}$$

³Our definition here is looser than the standard definition since we do not require $\tilde{\beta}$ to be an *essential* arc in Σ . We drop this condition because we will want to perform boundary-compressions on disks, which have no essential arcs.




Figure 5: Half-twist

For any punctured (marked) surface S with more than one puncture (marked point), if $l \subset S$ is a loop which bounds an embedded twice-punctured (marked) disk $U_l \subset S$, we can identify U_l with U and define the *half-twist* H_l about l to be the homeomorphism $H_l: S \rightarrow S$ given by

$$H_l(p) = \begin{cases} H(p) & \text{for } p \in U, \\ p & \text{for } p \notin U. \end{cases}$$

Now we will define plat position for a link. We start by describing plat position for a braid. Fix integers $h \geq 2$ and $b \geq 3$. We begin by constructing a set of vertical cylinders in \mathbb{R}^3 which will provide a frame for our braid. See Figure 6. In the xy -plane, define $c^1, c^2, \dots, c^{2b-1}$ to be the circles of radius $\frac{1}{2}$ such that c^j has center $(j + \frac{1}{2}, 0, 0)$. Now define Cyl^j to be the cylinder obtained by crossing c^j with the z -interval $[1, h]$.

Next we define the *twist regions*, subcylinders of $\{\text{Cyl}^j\}$ depicted in Figure 6. Let $i \in \{2, 3, \dots, h\}$. For odd i , let j range from 1 to b ; for even i , let j range from 1 to $b - 1$. Then define

$$\text{Tw}_i^j = \begin{cases} c^{2j-1} \times [i - \frac{3}{4}, i - \frac{1}{4}] & \text{for } i \text{ odd,} \\ c^{2j} \times [i - \frac{3}{4}, i - \frac{1}{4}] & \text{for } i \text{ even.} \end{cases}$$

In this paper, superscripts will usually denote an object’s horizontal position (ie in the x -direction), and subscripts will usually denote an object’s vertical position (ie in the z -direction). This is reminiscent of typical matrix notation: an entry a_i^j of a matrix is in the i^{th} position vertically and the j^{th} position horizontally. Here, however, unlike in a matrix, an array of objects is enumerated from bottom to top instead of from top to bottom.

Now we define a (nonclosed) braid B to be in (h, b) -plat position if it satisfies the following conditions:

- (1) B is a $2b$ -strand braid embedded in $\bigcup \text{Cyl}^j$ whose bottom endpoints’ coordinates are $(j, 0, 1)$ for $j = 1, 2, \dots, 2b$ and whose top endpoints’ coordinates are $(j, 0, h)$ for $j = 1, 2, \dots, 2b$.

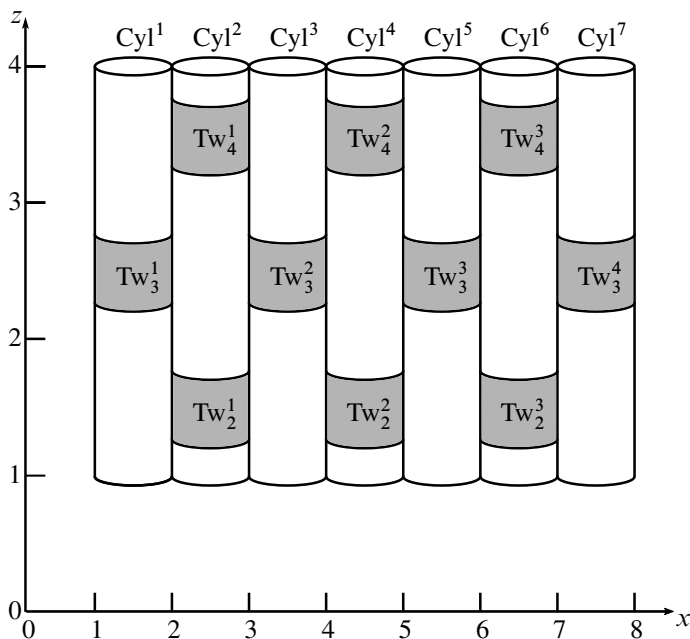


Figure 6: The cylinders which form a frame for plat links (here $h = b = 4$)

- (2) Arcs of B are partitioned into two types: vertical arcs which lie in the intersection of $\bigcup \text{Cyl}^j$ with the xz -plane, and twisting arcs which are not contained in the xz -plane but are properly embedded in the twist regions.
- (3) Every twisting arc is strictly increasing in a twist region, and as it ascends, it moves either strictly clockwise around the cylinder or strictly counterclockwise.

These conditions guarantee that each twist region contains exactly two arcs of B . Observe that the endpoints of every twisting arc must lie in the xz -plane. Since twisting arcs move strictly clockwise or counterclockwise, they intersect the xz -plane minimally. That is, there is no isotopy of a twisting arc in a twist region, relative to its endpoints, which would decrease the intersection of that twisting arc and the xz -plane. Looking down at the link from above, if the twisting arc moves counterclockwise as it ascends, we define t to be -1 plus the number of times the twisting arc intersects the xz -plane. Similarly, if the twisting arc moves clockwise as it ascends, we define $-t$ to be -1 plus the number of times the twisting arc intersects the xz -plane. In other words, to each twisting arc we assign an integer t such that the twisting arc travels around the twist region through an angle of $t\pi$. Notice that the other twisting arc in the same twist region must twist through an equal angle $t\pi$ (or else the two arcs would




Figure 7: Closing a braid in (h, b) -plat position, for (left) h even, and (right) h odd. In this paper, we only consider $h = 4$.

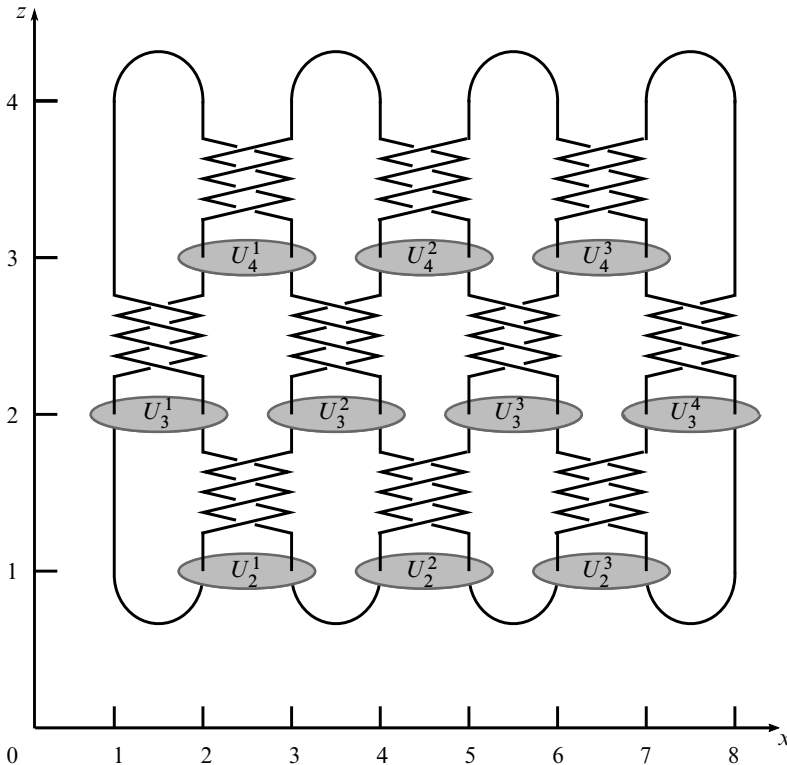


Figure 8: A 4-twisted $(4, 4)$ -plat link with $t_1^j = t_3^j = 4$ for all j and $t_2^j = -4$ for all j . Also pictured are the U -disks defined in Section 3.

intersect). Thus we can define the number t to be the *twist number* corresponding to that twist region. We will use t_i^j to denote the twist number for Tw_i^j .

There is a standard way to create a link from a braid in plat position. See Figure 7. Along the bottom of B , we have a row of $2b$ endpoints at a height of $z = 1$. For each odd j between 1 and $2b$, connect the endpoints $(j, 0, 1)$ and $(j + 1, 0, 1)$ with a lower semicircle in the xz -plane. The way we connect the top endpoints of B depends




Figure 9: Point of view

on the parity of h . Notice that if h is even, then the highest row of twist regions consists of $b - 1$ twist regions. Along the top of B we have a row of $2b$ endpoints at a height of $z = h$. For each odd j between 1 and $2b$, connect the endpoints $(j, 0, h)$ and $(j + 1, 0, h)$ with an upper semicircle in the xz -plane. If h is odd, there are b twist regions at the top. For each even j between 2 and $2b - 2$, connect the endpoints $(j, 0, h)$ and $(j + 1, 0, h)$ with an upper semicircle in the xz -plane. Then connect the points $(j, 0, 1)$ and $(j, 0, 2b)$ with a larger upper semicircle in the xy -plane. See Figure 7. Thus we obtain a link from B .

Any link constructed in this way from a braid in (h, b) -plat position will be said to be a link in (h, b) -plat position. Observe that such a link is in a bridge position with bridge number b . For $n \in \mathbb{N}$, a link in (h, b) -plat position is called n -twisted if $|t_i^j| \geq n$ for all i, j . See Figure 8 for an example of a 4-twisted link in $(4, 4)$ -plat position.

3 Setting

Let $f: S^3 \rightarrow [-\infty, \infty]$ be a height function on S^3 , and let Λ be a strictly increasing arc in S^3 whose endpoints are the two critical points of S^3 . Then $S^3 \setminus \Lambda$ is homeomorphic to \mathbb{R}^3 . Let $\phi: \mathbb{R}^3 \rightarrow S^3 \setminus \Lambda$ be a homeomorphism that respects the height function f on S^3 and the standard height function on \mathbb{R}^3 . That is, if P_i is the level plane in \mathbb{R}^3 of height i , and F_i is the level sphere in S^3 of height i , then $\phi(P_i)$ is F_i minus the point $F_i \cap \Lambda$. Throughout the rest of the paper, we adopt the point of view of someone standing on the $-y$ side of the xz -plane. This gives meaning to words like “left”, “right”, “up”, “down”, “horizontal” and “vertical”. See Figure 9.

The link L Let L be a 2-twisted link in $(4, 4)$ -plat position with bridge sphere $F \simeq F_1 = f^{-1}(1) \subset S^3$. Suppose L has twist numbers $\{t_i^j\}$ such that for all j , the




Figure 10: The bridge disks D^k above F



Figure 11: The bridge arcs below F . Of these, we care especially about the rightmost bridge arc $\alpha^{4'}$, the straight line segment β' between its endpoints, and the bridge disk $D^{4'}$ cobounded by $\alpha^{4'}$ and β' .

twist numbers t_2^j, t_4^j are positive, and t_3^j is negative. We will work with this link L for the rest of the paper.

Consider the diagram $D(L)$ for L obtained by projection to the xz -plane. Our choice of signs for the twist numbers $\{t_i^j\}$ makes $D(L)$ an alternating diagram, and we know that L is a split link if and only if $D(L)$ is a split diagram (see Theorem 4.2 of [9]). Since $D(L)$ is clearly a nonsplit diagram, it follows that L is a nonsplit link.

The bridge disks Refer to Figure 10. Above F_4 in the xz -plane lie the four upper bridge arcs. We will name these arcs consecutively from left to right $\alpha^1, \alpha^2, \alpha^3, \alpha^4$. Vertical projection of the bridge arcs into F gives us four straight line segments at level 4, which we will name consecutively from left to right $\beta^1, \beta^2, \beta^3, \beta^4$. We will refer to these as β -arcs. Observe that α^j and β^j form a loop. Let D^j be the disk in the xz -plane bounded by this loop. Notice each D^j is a bridge disk. Define γ^j to be the straight line segment connecting the points $(2j, 0, 4)$ and $(2j + 1, 0, 4)$. We will refer to these arcs as γ -arcs.

Below F_1 in the xz -plane lie the four lower bridge arcs, which we will call consecutively from left to right $\alpha^{1'}, \alpha^{2'}, \alpha^{3'}, \alpha^{4'}$. We will also define β' to be the straight line segment between the endpoints of $\alpha^{4'}$. See Figure 11. The union $\beta' \cup \alpha^{4'}$ is a loop bounding a bridge disk in the xz -plane, which we will call $D^{4'}$.

l -loops and U -disks For $i = 1, 3$, define l_i^j to be the circle of radius $\frac{3}{4}$ in the horizontal plane $z = i - 1$ centered at the point $(2j + \frac{1}{2}, 0, i - 1)$. For $i = 2$, define l_2^j to be the circle of radius $\frac{3}{4}$ in the horizontal plane $z = 2$ centered at the point

$(2j - \frac{1}{2}, 0, i - 1)$. Then each l_i^j bounds a twice-punctured disk which we call U_i^j . Each disk U_i^j lies below a corresponding twist region Tw_i^j . See Figure 8.

σ -projection It will be convenient to be able to talk about simple vertical projections. If P_i denotes the horizontal plane in \mathbb{R}^3 of height i , let $\sigma_i: \mathbb{R}^3 \rightarrow P_i$ be the vertical projection map defined by $(x, y, z) \mapsto (x, y, i)$.

τ -projection We will also need a more subtle type of projection which respects L . We closely follow Johnson and Moriah [6]. Note that L intersects each F_i in the same number of points and these points vary continuously as i varies from 1 to h . This can be thought of as an isotopy of these $2b$ marked points in F which extends to an ambient isotopy of F_h . To be precise, there is a projection map $\tau_i: F \times [1, h] \rightarrow F_i$ for each $i \in [1, h]$ that sends each arc component of the plat braid to a point $(j, 0, i)$ for some $j = 1, 2, \dots, 2b$ and defines a homeomorphism $F_{i'} \rightarrow F_i$ for each $i' \in [1, h]$. These homeomorphisms are canonical up to isotopy fixing the points $L \cap F_i$, and the induced homeomorphism $F_i \rightarrow F_i$ is the identity. Further, for each $i = 2, 3, 4$, the homeomorphism induced by $\tau_i: F_{i-1} \rightarrow F_i$ is a composition of half-twists about the l -loops, which can be expressed as the product $\prod_j H_{l_j^i}^{t_j^i}$. (Note that since the U -disks at any given level are pairwise disjoint, the corresponding half-twists all commute with each other.) Therefore the homeomorphism $g: F_0 \rightarrow F_4$ induced by τ_4 can be expressed as

$$(*) \quad g = \left(\prod_j H_{l_4^j}^{t_4^j} \right) \left(\prod_j H_{l_3^j}^{t_3^j} \right) \left(\prod_j H_{l_2^j}^{t_2^j} \right).$$

Disk partition Our goal is to show that F is critical, so we need to exhibit a partition $\mathcal{C}_1 \sqcup \mathcal{C}_2$ of the isotopy classes of the compressing disks for F . (Recall the definition of “critical” from the beginning of Section 2.) Above F , let B be the frontier of a regular neighborhood of D^1 . Below F , let B' be the frontier of a regular neighborhood of D^4 . B and B' are depicted in Figure 12. Let $[B], [B'] \in \mathcal{C}_2$, and let all of the other isotopy classes of compressing disks be in \mathcal{C}_1 . It will be convenient to refer to a compressing disk C as *red* if $[C] \in \mathcal{C}_1$ and *blue* if $[C] \in \mathcal{C}_2$.

4 The labyrinth

We are interested in the image of $\partial B'$ under g , the homeomorphism induced by τ_4 . As τ_4 takes $\partial B'$ up from level 1 to level 4, $\partial B'$ undergoes a series of twists, following the strands of the link. As explained in Section 3, g can be expressed as a product of




Figure 12: The two blue compressing disks, B and B'

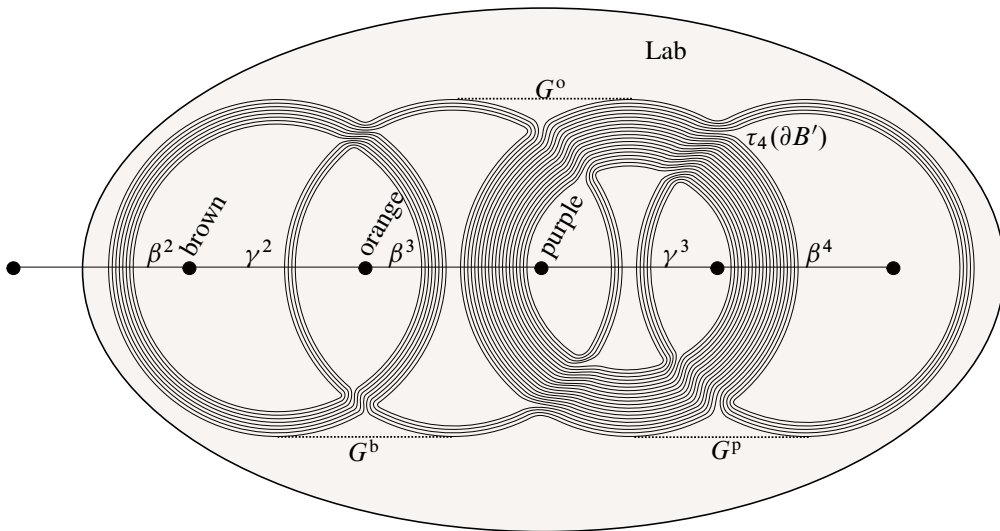


Figure 13: The thin spiraling line is $\tau_4(\partial B')$ in the case when $t_1^j = t_3^j = 2$ and $t_2^j = -2$ for all j . Also depicted is the disk Lab which contains $\tau_4(\partial B')$, and the three gates, G^b , G^o , and G^p , as well as the three colored points. In general, the five marked points in Lab may be permuted differently when different twist numbers are chosen.

half-twists around the l -loops via the equation (*). In Figure 13 we show exactly what $\partial B'$ looks like under g in a case where $|t_i^j| = 2$ for all j .




Figure 14: How to view a half-twist about an l -loop as a link diagram




Figure 15: Rule for smoothing a crossing

Of course, we want to understand what $\tau_4(\partial B')$ will look like in general for any set of twist numbers $\{t_i^j\}$. To do so, we will devise a convenient pictorial notation for loops in F which undergo half-twists about l -loops. Suppose λ is a loop in F with one arc of intersection with U_i^j which passes between the two marked points, as in the first picture in Figure 14. The second and third pictures in Figure 14 are homotopic pictures of $H_{l_i^j}(\lambda)$. Observe that the third picture is reminiscent of knot diagram smoothing. In fact, using the smoothing rule in Figure 15, we can “unsmooth” $H_{l_i^j}(\lambda)$ into the fourth picture, which consists of a link diagram on F with two components: (1) the original λ and (2) l_i^j . This suggests a convenient algorithm to draw any such loop λ after a (positive) half-twist about an l -loop: We start by drawing λ . Next we draw the l -loop so that it passes *under* λ . Last we smooth the crossings using the rule in Figure 15. If we want to perform a negative half-twist instead, the procedure is the same, except that we draw the l -loop crossing *over* λ .

We can generalize this process to the situation where there are n strands of λ passing through U_i^j . In Figure 16, the top row is an example in which $n = 3$, and the bottom row depicts the general case. As in Figure 14, the second and third pictures of each row in Figure 16 are homotopic pictures of $H_{l_i^j}(\lambda)$, and we can “unsmooth” $H_{l_i^j}(\lambda)$ into the fourth picture, which consists of a link diagram on F . This time, the link diagram contains $n + 1$ components: λ and n parallel copies of l_i^j . Thus our generalized procedure for drawing a half-twist of λ about l_i^j is the following: First determine the number n of times that λ passes between the two punctures of U_i^j , then draw n parallel, disjoint copies of l_i^j so that at each crossing, l_i^j passes *under* λ for a positive




Figure 16: How to view a half-twist about an l -loop as a link diagram. In the top row of pictures, we show specifically what the situation looks like when $n = 3$. The bottom row shows that the whole process can be done with an arbitrary number n of strands through U_i^j . In all the following figures that depict half-twists, when we label the number of parallel strands present, as in the bottom row here, we will box the labels with squares.

half-twist, or *over* λ for a negative half-twist, and then smooth the crossings using the rule from Figure 15.

Note that to perform $|t_i^j|$ consecutive half-twists around l_i^j , we simply perform this process $|t_i^j|$ times, taking care that every time we add new parallel copies of l_i^j , they are nested inside the ones previously drawn. Thus if λ is half-twisted around l_i^j a total of $|t_i^j|$ times, and λ passes between the punctures of U_i^j a total of n times, then to build the link diagram depicting $H_{l_i^j}^{|t_i^j|}(\lambda)$, we will add a total of $n|t_i^j|$ parallel copies of l_i^j (all passing under λ if $t_i^j > 0$ or over λ if $t_i^j < 0$). Thus after any sequence of half-twists about l -loops, we obtain a link diagram representing λ . Instead of drawing all of the parallel copies of each l -loop, we will draw only one of each and label it with the number \boxed{n} of strands it represents.

In the resulting link diagram, we will have crossings involving two strands which are labeled with different numbers. Figure 17 depicts an example in which we show how to recover what λ actually looks like near such a crossing, where a strand marked with a $\boxed{7}$ crosses a strand marked with a $\boxed{3}$. Figure 18 shows the general case.

Our purpose in developing this link-diagrammatic representation of a twisted loop is to understand what $\partial B'$ looks like after performing (in order) half-twists about all of the l -loops for any set of twist numbers $\{t_i^j\}$. With this method in hand, this is a

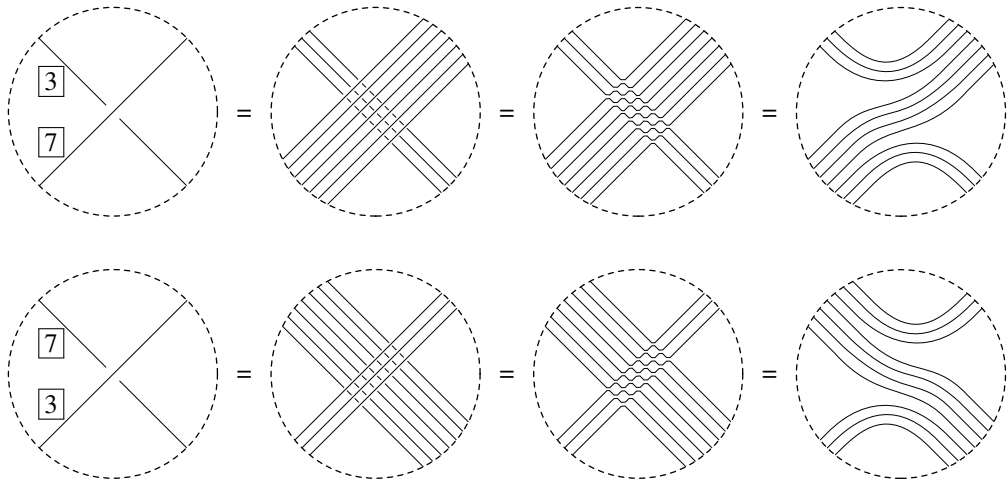


Figure 17: After twisting a loop λ about l -loops, we obtain a link diagram on F . Here is an example of a crossing in such a diagram: In this crossing, one strand represents three parallel strands, and the other represents seven parallel strands. Consider the top row of pictures, where the over-strand has a larger number. The first picture is a neighborhood of the crossing of the link diagram we have obtained. The second picture explicitly shows all of the parallel strands at the crossing. The third picture shows the result after smoothing all 21 crossings. The fourth picture is a smoothed version of the third. All four pictures represent the same thing. The bottom row is similar, but this time the under-strand has the larger number.




Figure 18: How to interpret a crossing in the link diagram of N strands with n strands, where $N > n$




Figure 19: The components of the link diagram that are newly added as a result of having just moved up a level are highlighted by being drawn with dotted lines. The reader may find it helpful to compare this figure to Figure 12, which shows the link diagram’s location with respect to L .

straightforward task. Refer to Figure 19 as we describe this process. We start with the simple circle $\partial B'$ in F at level 1, and then push it up to level 2 (ie we apply τ_2). The boundary $\partial B'$ had to pass through the twist region Tw_2^3 , so this means it underwent t_2^3




Figure 20: The link diagram representing $\tau_4(\partial B')$ (with values of the N^i and inequalities between them as given below)

half-twists about l_2^3 . Before twisting (ie at level 1), $\partial B'$ passed one time (ie $n = 1$) between the punctures of U_2^3 ; therefore what $\partial B'$ looks like at level 2 is a link diagram consisting of the original circle $\partial B'$, plus $1 \times t_2^3$ parallel copies of l_i^j . Since we specified in Section 3 that t_2^j is positive for all j , each copy of l_i^j will be drawn to pass under $\partial B'$. Now the original $\partial B'$ and the t_2^3 copies of l_i^j together with the crossing information make a link diagram which represents the image of $\partial B'$ under τ_2 . Since the value of t_2^3 is arbitrary, we cannot smooth the crossings to see exactly what $\tau_2(\partial B')$ looks like. But that is not a problem; we can still push this loop up from level 2 to level 3. As we do, we add copies of l_3^3 and l_3^4 to the link diagram. This time, since $t_3^j < 0$ for all j , these l -loops will all be drawn with overcrossings, which gives us a link diagram representing $\tau_3(\partial B')$. Finally we push this up to level 4, adding copies of l_4^2 and l_4^3 (with undercrossings) to the link diagram. In this way we see that $\tau_4(\partial B')$ can be represented by the diagram of the four-component unlink in Figure 20.

Each circle in the diagram in Figure 20 is marked with a number N^i indicating how many parallel copies of that circle are present, where

$$\begin{aligned}
 N^1 &= t_2^3 |t_3^3| t_4^2, \\
 N^2 &= t_2^3 |t_3^3|, \\
 N^3 &= t_2^3 + t_4^3 (1 + t_2^3 |t_3^3| + t_2^3 |t_3^4|), \\
 N^4 &= 1 + t_2^3 |t_3^4|.
 \end{aligned}$$

Recall that $t_2^j, t_4^j \geq 2$ and $t_3^j \leq -2$ for all j . No matter what values the twist numbers $\{t_i^j\}$ take on, provided they follow this rule, it can be shown that $N^i > 0$ for each i , and these three inequalities hold true:

$$N^1 > N^2, \quad N^2 < N^3, \quad N^3 > N^4.$$




Figure 21: Proposition 4.2 asserts that Lab is isotopic to this picture. The boundary $\partial B'$ and the three gates cut Lab into five components: three once-punctured disks, an annulus, and a twice-punctured disk.

This means that at each of the six crossings, the strand labeled with the higher number passes under the strand labeled with the lower number, so all six crossings look like the bottom picture in Figure 18.

Proposition 4.1 *The position of $\partial B'$ described by the link diagram in Figure 20 is minimal with respect to $\beta^2 \cup \beta^3 \cup \beta^4$.*

Proof The bigon criterion in [5] tells us that as long as $\partial B'$ cobounds no bigons with $\beta^2 \cup \beta^3 \cup \beta^4$, then they are in minimal position. Consider $\partial B'$ as being cut into component strands by $\beta^2 \cup \beta^3 \cup \beta^4$. To bound a bigon, one of these strands would have to have both endpoints on a single β -arc. However, each strand has endpoints on distinct β -arcs. \square

Let Lab be the disk in F_4 which contains $\tau_4(\partial B')$ and whose complement is a regular neighborhood of $\beta^1 \cup \gamma^1$. (Lab is depicted in Figures 13 and 20, and also, Lab is isotopic in F to the top gray rectangle in Figure 12.) We will refer to Lab as the *labyrinth*.

Define G^b , G^o , and G^p to be the dotted arcs depicted in Figures 13 and 20. We will call these the *brown*, *orange*, and *purple gates* of the labyrinth, respectively. Each gate is an arc in Lab with endpoints on $\partial B'$.

Proposition 4.2 *$\partial B'$ and the three gates cut Lab into five components: three once-punctured disks, an annulus, and a twice-punctured disk.*

Proof Since each crossing in our link diagram in Figure 20 looks like the bottom picture of Figure 18, it is apparent that there is an annulus component, which we will




Figure 22: In Case 1, α^2, α^3 and α^4 appear on both sides of R .

call A . A has $\partial(\text{Lab})$ as one boundary component, and the other component of ∂A alternates between the three gates and three subarcs of $\partial B'$. Observe that A is not punctured. The boundary $\partial B'$ cuts off the twice-punctured disk U_2^3 from Lab by definition, and clearly none of the gates are in this disk. $\text{Lab} \setminus U_2^3$ is a thrice-punctured annulus in which the three gates are properly embedded but not nested: if any subset of the gates were nested, then at least one of them would not be a boundary component of A . None of the gates can be parallel to $\partial B'$ because that would either force A to be punctured or force the other gates to be nested. Thus the only possible configuration is (isotopic to) Figure 21. □

We will refer to the three once-punctured disks in Proposition 4.2 as the *brown, orange, and purple punctured disks*, according to the color of the corresponding gate, and we will call the marked point contained therein a *brown, orange, or purple marked point*. In the brown disk, there is a unique arc (up to isotopy) connecting the brown point to G^b , which we will call the *brown escape route*. We define the *orange and purple escape routes* similarly. (The three escape routes are depicted in the second picture of Figure 32.) We can think of the colored points as escaping from a maze whose walls are $(\partial B')$, and the escape routes are the paths they take to get to the exits (the gates).

5 Red disks enter the labyrinth

Refer to Figure 10.

Lemma 5.1 *If R is a red disk above F , then R must intersect γ^2 or γ^3 .*

Proof F cuts S^3 into two 3-balls: M_+ above F and M_- below F . R divides M_+ into two 3-balls which we will call M_+^1 and M_+^2 and refer to as the two sides of R .

Case 1 (α^2, α^3 and α^4 are not all on the same side of R) Then for $k = 2$ or $k = 3$, α^k and α^{k+1} are on opposite sides of R . Since γ^k connects an endpoint of α^k and an endpoint of α^{k+1} , γ^k must intersect R .




Figure 23: In Case 2, $\alpha^1 \subset M_+^1$ and $\alpha^2, \alpha^3, \alpha^4 \subset M_+^2$.




Figure 24: In Subcase 2.2, $R \cap D^k$ is nonempty. Possible arcs of intersection are depicted on D^k .

Case 2 (all of the bridge arcs α^2, α^3 and α^4 are on the same side of R) Without loss of generality, say they are in M_+^2 . See Figure 23. Notice that α^1 cannot also be in M_+^2 because that would imply ∂R is null-homotopic in F , contradicting the fact that R is a compressing disk. Therefore, in Case 2, R separates α^1 from the other bridge arcs, which implies R is a cap for α^1 . Let $\mathcal{D} = D^2 \cup D^3 \cup D^4$, and isotope R to minimize $\#|R \cap \mathcal{D}|$.

Subcase 2.1 (R is disjoint from \mathcal{D}) Assume (for contradiction) that R is also disjoint from γ^2 and γ^3 . Then we have a single straight arc $\xi = \beta^2 \cup \gamma^2 \cup \beta^3 \cup \gamma^3 \cup \beta^4$ disjoint from ∂R because $\beta^j \subset D^j \subset \mathcal{D}$. So ∂R cuts F into two disks, one of which contains ξ , and the other of which contains $\partial\beta^1$. But of course, there is only one such loop in F , which is ∂B . Thus $R \simeq B$, so the red disk R is blue, a contradiction. We conclude that in this subcase, ∂R must intersect γ^2 or γ^3 .

Subcase 2.2 (R is not disjoint from \mathcal{D}) For some $D^k \subset \mathcal{D}$, $R \cap D^k$ is nonempty. $R \cap D^k$ cannot contain loop intersections because $\#|R \cap \mathcal{D}|$ is minimal, and $S^3 \setminus \eta(L)$ is irreducible. $R \cap D^k$ cannot contain an arc with either endpoint on α^k because that would imply that R intersects L . Thus the components of intersection in D^k must all be arcs with both endpoints on β^k . See Figure 24.




Figure 25: In Subcase 2.2, $R \cap D^k$ is nonempty. Here we see R and the outermost disk \tilde{D} .

Take an arc of intersection outermost on D^k and call it $\tilde{\alpha}$. Let $\tilde{\beta}$ be the subarc of β^k that shares endpoints with $\tilde{\alpha}$. Then let \tilde{D} be the subdisk of D^k cobounded by $\tilde{\alpha}$ and $\tilde{\beta}$. See Figures 24 and 25. $\tilde{D} \subset M_+^j$ for either $j = 1$ or $j = 2$. Whichever is the case, \tilde{D} cuts M_+^j into two balls which we call the sides of \tilde{D} . Recall that we assume $\alpha^2, \alpha^3, \alpha^4 \subset M_+^2$, so therefore $\alpha^1 \subset M_+^1$. \tilde{D} cannot be in M_+^1 with α^1 . If it were, then since α^1 cannot intersect R or \tilde{D} , α^1 would have to be completely contained on one side or the other of \tilde{D} . Then the other side of \tilde{D} would be an empty 3–ball through which we can isotope R , removing at least one component of intersection with D^k , which is a contradiction since $\#|R \cap \mathcal{D}|$ is already minimal. Thus \tilde{D} is in M_+^2 with α^2, α^3 and α^4 .

If α^2, α^3 and α^4 were on the same side of \tilde{D} , that would mean that the other side of \tilde{D} was an empty 3–ball through which we could isotope R , removing the arc of intersection $\tilde{\alpha} \subset R \cap \mathcal{D}$, contradicting minimality. Thus it must be the case that there are some α –arcs in either side of \tilde{D} .

Then for $k = 2$ or $k = 3$, α^k and α^{k+1} are on opposite sides of \tilde{D} . An endpoint of α^k and an endpoint of α^{k+1} are connected by the arc γ^k in F . None of the γ –arcs have interiors that intersect \mathcal{D} , so γ^k must pass through ∂R in order to connect α^k to α^{k+1} . Thus we have proved that in Subcase 2.2, $\partial R \cap (\gamma^2 \cup \gamma^3) \neq \emptyset$.

Now we have proved both subcases, concluding Case 2 and finishing the proof of Lemma 5.1. □

Corollary 5.2 *Every red disk intersects at least one of $\tau_4(\partial B')$, G^b , G^o and G^p .*

Proof Recall that by Proposition 4.2, $\text{Lab} \setminus (\tau_4(\partial B') \cup G^b \cup G^o \cup G^p)$ has five components: three (colored) once-punctured disks, one twice-punctured disk (U_2^3), and

an annulus. See Figure 21. Let R be a red compressing disk above F . The boundary of a compressing disk cannot bound a once-punctured disk in F . Thus ∂R cannot be contained in one of the three colored once-punctured disks. If ∂R lies in the twice-punctured disk, then ∂R must be isotopic to $\tau_4(\partial B')$. But then after the isotopy, $R \cup B'$ would be a splitting sphere for L , a contradiction since L is nonsplit. We see in Figures 13 and 20 that γ^2 and γ^3 are disjoint from the annulus component of Lab , so if ∂R is contained in the annulus, ∂R would fail to intersect γ^2 or γ^3 , contradicting Lemma 5.1. We conclude ∂R cannot lie in a single component of $\text{Lab} \setminus (\tau_4(\partial B') \cup G^b \cup G^o \cup G^p)$; therefore it must intersect $\tau_4(\partial B') \cup G^b \cup G^o \cup G^p$. \square

6 All red disks above F intersect all blue disks below

Recall the disks and arcs defined in Section 3 and pictured in Figures 10 and 11. Let $\mathcal{D} = D^2 \cup D^3 \cup D^4$ as in Section 5. Let \mathcal{R} be the set of red compressing disks above F which are disjoint from β' . Throughout Section 6, whenever \mathcal{R} is nonempty, assume that R is a disk in \mathcal{R} such that $|R \cap \mathcal{D}| \leq |R' \cap \mathcal{D}|$ for all $R' \in \mathcal{R}$, and assume R is in minimal position with respect to \mathcal{D} and with respect to the gates. β' cuts β^i into multiple subarcs, which we call *lanes*.

Lemma 6.1 *Assume \mathcal{R} is nonempty. If two points of $R \cap \beta^i$ lie in the same lane, then they cannot be the endpoints of a common arc of $R \cap D^i$.*

Proof Suppose there exists an arc of $R \cap D^i$ whose endpoints lie in the same lane. We may assume the arc is outermost on D^i . It cuts off a small subdisk $\tilde{D}^i \subset D^i$ which does not intersect β' . Boundary-compress R along \tilde{D}^i . The result is two disks, R^1, R^2 whose boundaries lie in F , and which are disjoint from β' . Suppose R^1 is trivial. Then it represents an isotopy through which we can move R to decrease its intersection with the bridge disks. But that implies R was not in minimal position with respect to the bridge disks, a contradiction. Thus R^1 is not trivial, and similarly, neither is R^2 .

This means both R^1 and R^2 are compressing disks. By construction, neither of them intersects β' , and $|R^i \cap \mathcal{D}| < |R \cap \mathcal{D}|$ for both i . Thus $R^i \notin \mathcal{R}$ by our choice of R as a minimal representative from \mathcal{R} . Then by the definition of \mathcal{R} , both R^i must be blue. There is only one blue disk above F , which is B , so R^1 and R^2 must be parallel copies of B . This implies R is a band sum of parallel disks, but any band sum of parallel disks is trivial, so R is a trivial disk, which is a contradiction. \square




Figure 26: Here we depict the brown, orange, and purple tracks in ∂R with their respective sequences of numbered points for the example depicted in Figure 13 (in which $t_1^j = t_3^j = 2$ and $t_2^j = -2$ for all j). We will not prove this figure, nor does the paper depend on it. The reader can check it by taking a pencil (and perhaps a magnifying glass) to Figure 13 and tracing out a path that winds in a gate, around a marked point inside the labyrinth, and then back out.

By definition, $R \in \mathcal{R}$ implies $R \cap \beta' = \emptyset$, so R can also be made disjoint from $\partial B'$ since $\partial B'$ is the boundary of a regular neighborhood of β' . It follows from Corollary 5.2 that R must intersect at least one gate. Figure 21 makes clear something not at all obvious in Figures 13 and 20, which is that if an arc of R “enters” the labyrinth through a particular gate, it must subsequently “exit” the labyrinth from the same gate. We will call the components of the intersection of ∂R and the brown disk *brown tracks*, and we will define *orange tracks* and *purple tracks* similarly. Observe that these tracks are pairwise disjoint, and each track has endpoints on the gate of the corresponding color. We label the points of $\partial R \cap D^2$, $\partial R \cap D^3$, and $\partial R \cap D^4$ *2–points*, *3–points*, and *4–points*, respectively, and we will collectively refer to these as *numbered points*. Likewise, we will label the arcs of $R \cap D^2$, $R \cap D^3$, and $R \cap D^4$ *2–arcs*, *3–arcs*, and *4–arcs*, respectively, and we will collectively call them *numbered arcs*. Observe that the endpoints of a j –arc are j –points. The next result follows from these definitions.

Lemma 6.2 *If \mathcal{R} is nonempty, at least one endpoint of each 2–arc must lie in a track, and all 3– and 4–points lie in tracks.*

Proof The leftmost lane of β^2 is not completely contained in Lab, but all the other lanes of β^2 , as well as all the lanes of β^3 and β^4 , are contained in the union of the three colored disks and U_2^3 . The boundary ∂R is disjoint from U_2^3 , so it can only intersect these lanes inside the colored disks. Since every arc component of ∂R inside a colored disk is, by definition, a track, we conclude that aside from the leftmost lane of β^2 , ∂R only intersects the lanes of β^2 , β^3 and β^4 in tracks. Thus all 3– and 4–points lie in tracks. If both endpoints of a 2–arc lie outside of the tracks, then those endpoints must lie in the leftmost lane of β^2 , contradicting Lemma 6.1. Therefore at least one endpoint of each 2–arc must lie in a track. □




Figure 27: Each colored track contains a sequence of numbered points. No matter the twist numbers $\{t_i^j\}$, we can say what the outermost and second-outermost points will be for each type of track.

Corresponding to each track color is a particular sequence of numbered points. Each of these three sequences is symmetric in the sense that the outermost numbered points in the track have the same number, the second-outermost numbered points match each other, the third-outermost numbered points match, etc. Figure 26 depicts the sequence of numbered points along each colored track in the case depicted in Figure 13.

Observe Figure 28, which depicts the possibilities for which directions tracks can go after they intersect the β -arcs. We see that the outermost numbered points of a brown track are 3-points because the first and last β -arc that the track will intersect is β^3 . After intersecting β^3 , a brown track will loop around and intersect β^2 , so the second-outermost numbered points of a brown track are 2-points. Similarly, the outermost numbered points of an orange track are 3-points, and the second-outermost numbered points of an orange track are 4-points. Finally, the outermost numbered points of a purple track are 4-points, and the second-outermost numbered points of a purple track are 3-points. This information is summarized in Figure 27. Note that this information about the outermost and second-outermost numbered points does not depend on the set $\{t_i^j\}$ of twist numbers.

Lemma 6.3 *Assume \mathcal{R} is nonempty. No numbered arc can have endpoints which lie on the same track.*

Proof Since ∂R must intersect at least one of the gates, it follows that there exists at least one track in ∂R . Refer again to Figure 28: In the top picture, we see that after intersecting β^3 , a track will always either intersect β^2 or β^4 or leave the labyrinth. In the bottom picture, we see that after intersecting β^2 or β^4 , a track will always either intersect β^3 or leave the labyrinth. This means that in each sequence of numbered points in a track, every other numbered point is a 3-point.




Figure 28: If λ is a track, then λ contains a sequence of numbered points. The top picture shows that each 3–point on λ either has a 2–point or a 4–point to both sides, or it has a 2–point or a 4–point to one side and exits through a gate to the other side. The bottom picture shows that each 2–point on λ has a 3–point to either side, and that each 4–point either has a 3–point to either side, or it has a 3–point to one side and exits through a gate to the other side.

Suppose two numbered points from the same track are connected in R by a numbered arc λ . Since the two endpoints of λ must have the same number, and since every other numbered point in the track is a 3–point, there must be an odd number of numbered points between the endpoints of λ . But this leads to a contradiction: numbered arcs never intersect each other, so it is impossible to pair up the numbered points between the endpoints of λ with disjoint numbered arcs. \square

Lemma 6.4 *If \mathcal{R} is nonempty, ∂R contains tracks of all three colors.*

Proof Consider a numbered arc λ which is outermost in R . Being outermost guarantees that the endpoints of λ are adjacent same-numbered points on ∂R . Since numbered points in each track alternate between 3–points and 2– or 4–points, there are only three ways to have adjacent same-numbered points on ∂R :




Figure 29: Case 1: The outermost 4-arc ζ connects 4-points in adjacent brown tracks.

- (1) The endpoints of λ are consecutive 2-points which do not lie in tracks, which would contradict Lemma 6.2.
- (2) The endpoints of λ are a 2-point off of a track, and an outermost numbered point on a track, which is impossible since no outermost point on any track is a 2-point.
- (3) The endpoints of λ are outermost same-numbered points of adjacent tracks.

Suppose ∂R contains no brown tracks. Since the outermost numbered points of orange tracks are 3-points, and the outermost numbered points of purple tracks are 4-points, λ must have endpoints which are adjacent outermost numbered points of two adjacent same-colored tracks, and we reach a contradiction by Lemma 6.1 since these two endpoints lie in the same lane. Therefore ∂R must contain at least one brown track. By a similar argument, ∂R contains at least one orange track.

Suppose ∂R contains no purple tracks. Consider a numbered 4-arc ζ of $R \cap D^4$ which is outermost in R (ie outermost among the set of 4-arcs). The arc ζ cuts ∂R into two components. Since ζ is outermost among the 4-arcs, one of these components contains no 4-points. Call this piece the *primary* piece. There are three possibilities: ζ connects two 4-points in two brown tracks, it connects two 4-points in two orange tracks, or it connects a 4-point of a brown track and a 4-point of an adjacent orange track. We will derive a contradiction in each case.

Case 1 (ζ connects two 4-points in two brown tracks)

For this to be possible, the set $\{t_i^j\}$ of twist numbers must be such that brown tracks contain 4-points, which is only true for some sets of twist numbers. In contrast, orange and purple tracks necessarily contain 4-points, no matter the twist numbers. This means that since the primary piece of ∂R cannot contain any 4-points, it cannot contain any complete track of any color. Then the 4-points which are the endpoints of ζ must lie in adjacent brown tracks. See Figure 29. Consider all of the numbered points in the




Figure 30: Case 2: The outermost 4–arc ζ connects 4–points in adjacent orange tracks. If brown tracks do not contain 4–points, then there may be any number of brown tracks between these two orange tracks (and there may be any number of 2–points between each pair of tracks).

primary piece of ∂R . They must be paired with numbered arcs, and there must be an outermost such arc λ , which connects same-numbered points adjacent on ∂R . If there are no 2–points between the brown tracks in the primary piece of ∂R , then the only pair of adjacent same-numbered points in the primary part of ∂R are an outermost 3–point of one brown track, and an outermost 3–point of the other brown track. But these points are the same lane of β^3 , which contradicts Lemma 6.1. If there is exactly one 2–point between the brown tracks in the primary piece of ∂R , then there does not exist a pair of adjacent same-numbered points in the primary part of ∂R to be the endpoints of λ . If there is more than one 2–point between the brown tracks in the primary piece of ∂R , then the only pair(s) of adjacent same-numbered points in the primary part of ∂R are pairs of these 2–points, but they cannot be the endpoints of λ by Lemma 6.2. We conclude Case 1 is not possible.

Case 2 (ζ connects two 4–points in two orange tracks)

In this case, brown tracks may or may not contain 4–points, depending on the twist numbers; if they do not, then the primary piece of ∂R may contain any number of brown tracks. If there are zero brown tracks in the primary piece of ∂R , then the proof is similar to Case 1. Suppose the primary piece of ∂R contains at least one brown track. Then since the second-outermost numbered points of brown tracks are 2–points, there are at least two 2–points in the primary piece of ∂R . Since numbered arcs cannot intersect each other, they cannot intersect ζ in particular, so all of the 2–points in the primary piece must be paired with each other by a set S of 2–arcs. Consider an outermost 2–arc λ in S (ie outermost in R with respect to the other 2–arcs of S). The endpoints of λ are a pair of 2–points with no other 2–points between them. There are four ways this can happen (depicted in Figure 30), though each one leads to a contradiction:




Figure 31: Case 3: Adjacent brown and orange tracks with 4–points connected by the outermost 4–arc ζ .

- (1) The pair of 2–points lies in a single brown track, contradicting Lemma 6.3.
- (2) Both 2–points are off of tracks, contradicting Lemma 6.1.
- (3) The pair consists of a second-outermost numbered point in a brown track and a 2–point not contained in a track. But then λ would separate a single 3–point from all the other numbered points of ∂R , so the numbered arc corresponding to that 3–point would have to intersect λ .
- (4) The pair consists of second-outermost numbered points in adjacent brown tracks, points which lie in the same lane, again contradicting Lemma 6.1.

Therefore Case 2 is impossible.

Case 3 (ζ connects a 4–point of a brown track and a 4–point of an adjacent orange track)

See Figure 31. Again, the numbered points in the primary piece of ∂R must all be connected in pairs. Recall that in every track the numbered points alternate between 3–points and 2– or 4–points, so in the intersection of the primary piece and the brown track, there are at least two 3–points, one on either side of the second-outermost 2–point. In the orange track, the second-outermost numbered point is a 4–point, so it must be an endpoint of ζ , so there is exactly one other numbered point in the intersection of the primary piece and the orange track: an outermost 3–point. Since the primary piece of the brown track contains strictly more than one 3–point, and the primary piece of the orange track contains exactly one 3–point, it is impossible to pair up the 3–points in such a way that no pair lies in the same track, so Case 3 contradicts Lemma 6.3 and is therefore impossible. We conclude ∂R must contain purple tracks, finishing the proof of Lemma 6.4. □

Lemma 6.5 *Assume \mathcal{R} is nonempty. Let $b \subset \beta^i$ be a lane whose endpoints are interior points of β' . Then a track must intersect b .*

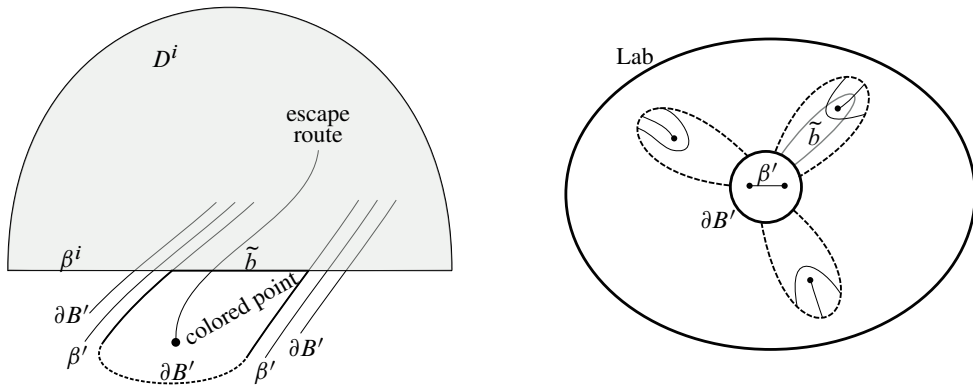


Figure 32: On the left, \tilde{b} is the bold horizontal segment of β^i . On the right, \tilde{b} is the gray arc. Also depicted on the right are the three escape routes, and the three corresponding tracks.

Proof Recall that $\partial B'$ and the three gates cut Lab into five components, one of which is the annulus A , and observe that no lane of β^2, β^3 or β^4 intersects A except for the leftmost lane of β^2 . By definition, b cannot be that leftmost lane, so b is properly contained in the union of U_2^3 and the three colored disks. Let $\tilde{b} = b \setminus (U_2^3)^\circ$ as in the first picture of Figure 32; \tilde{b} is an arc properly contained in a colored disk, with endpoints on ∂U_2^3 . \tilde{b} cannot cobound a bigon with $\partial B'$ because that would contradict Proposition 4.1, so \tilde{b} must be an arc cutting out a punctured disk from a colored disk, as in the second picture of Figure 32. Suppose the colored disk in question is brown. By definition, the brown escape route does not intersect $\partial B'$, so the only way for it to connect the brown point with the brown gate is to intersect \tilde{b} . Thus if \tilde{b} is contained in the brown disk, it must intersect the brown escape route. Every brown track is a frontier in the brown disk of the brown escape route, so every brown track must also intersect \tilde{b} . Since the color was arbitrary, the lemma is proved. \square

Proposition 6.6 *If L is not perturbed at $\alpha^{4'}$, then \mathcal{R} is empty.*

Proof Assume \mathcal{R} is nonempty. Consider a numbered arc λ which for some i is outermost in D^i . Let p and p' be the endpoints of λ , and let $\tilde{\beta}^i$ be the segment of β^i between p and p' , as in Figure 33. The points p and p' cannot lie in a single lane by Lemma 6.1. If the two lanes containing p and p' have at least one lane between them (ie if the lanes are not adjacent), then by Lemma 6.5, there must be at least one numbered point between p and p' . However, λ is outermost in D^i , so that cannot




Figure 33: λ straddles exactly one strand of β'




Figure 34: In Case 1, the bridge arcs are all disjoint from R^1 and R^2 . One of the disks F^{R^1} or F^{R^2} must contain exactly one pair of partners, the endpoints of a bridge arc α^j . Then L is perturbed at α^j .

happen. Therefore the lanes containing p and p' must be adjacent, so $\tilde{\beta}^i$ contains exactly one point of β' .

The arc λ cuts off a small disk \tilde{D}^i from D^i . We boundary-compress R along \tilde{D}^i , which results in two new disks, R^1 and R^2 , properly embedded in M_+ . Since each intersects β' exactly once, ∂R^1 and ∂R^2 cut F into an annulus and two disks, and each of the two disks contains an endpoint of β' . Label these three regions F^{R^1} , F^{R^2} and F^A . The disk R^i cannot be trivial since F is punctured to either side of ∂R^i . Therefore both R^i are compressing disks. Notice that R^1 and R^2 cannot be parallel because the band sum dual to the boundary compression just performed would recover R , but any band sum of parallel disks is a trivial disk.

Define a pair of marked points on F to be *partners* if they are endpoints of a common bridge arc above F . Let the endpoints of β' be q^1 and q^2 .




Figure 35: A disjoint pair of blue disks (B and B') and a disjoint pair of red disks (B'' and B''') which fulfill condition (2) of the definition of a critical surface

Case 1 For all four pairs of partners, it is the case that both points are contained in the same region: F^{R^1} , F^{R^2} or F^A . Since an endpoint of β' lies in each of the disks, a pair of partners lies in each of the disks. One pair of partners must be in F^A ; otherwise, R^1 and R^2 would be parallel compressing disks. This accounts for three of the four pairs of marked points. The fourth pair may be in any of the three regions F^{R^1} , F^{R^2} or F^A . This means that F^{R^1} or F^{R^2} must contain exactly one pair of partners. See Figure 34. Suppose these partners are the endpoints of α^j for some j . Then there is some bridge disk \hat{D} for α^j (which is not necessarily isotopic to D^j) which is disjoint from R^i . If $\hat{d} = F \cap \hat{D}$, then \hat{d} and β' can be made disjoint except for their shared endpoint. Then $\hat{d} \cup \beta'$ is an embedded arc. This means that \hat{D} and $D^{4'}$ are bridge disks above and below F which intersect in a single point of the link, so L is perturbed at $\alpha^{4'}$. This contradicts the hypothesis of the proposition, so Case 1 is not possible.

Case 2 There is a pair of partners which lie in different regions F^{R^1} , F^{R^2} and F^A . But then the bridge arc connecting them must intersect R^1 or R^2 , a contradiction since compressing disks are disjoint from the link. Having arrived at a contradiction, we conclude R cannot exist, so \mathcal{R} is empty, and the proposition is proved. \square

Define \mathcal{R}' to be the set of red compressing disks below F disjoint from B . The next result follows directly from Proposition 6.6 and rotational symmetry of L .

Corollary 6.7 *If L is not perturbed at α^1 , then \mathcal{R}' is empty.*

Theorem 6.8 *Let L be a 2-twisted link in $(4, 4)$ -plat position with bridge sphere $F \simeq F_4 = f^{-1}(4) \subset S^3$. Suppose L has twist numbers $\{t_i^j\}$ such that for all j , the twist numbers t_2^j, t_4^j are positive and t_3^j is negative. If L is perturbed at neither α^1 nor $\alpha^{4'}$, then F is a critical bridge sphere.*




Figure 36: Links L in the family \mathcal{L} . Since both t_2^2 and t_4^2 are even, L will have at least two components, L^1 (drawn with dotted arcs) and L^2 (drawn with solid arcs).

Proof By Proposition 6.6 and Corollary 6.7 we conclude that \mathcal{R} and \mathcal{R}' are empty. In other words, all red disks above F intersect all blue disks below F , and vice versa, which fulfills condition (1) of the definition of a critical surface. Observe that the link diagram representing $\partial B'$ (in Section 4) is disjoint from ∂B , so B and B' are disjoint. This gives us a pair of disjoint blue disks, one above F and one below. Consider B'' and B''' depicted in Figure 35. (Explicitly, these are the frontiers of regular neighborhoods of D^4 and $D^{1'}$ above and below F , respectively.) By a symmetric argument, these two compressing disks are a pair of disjoint red disks, one above F and one below. These two pairs of disjoint disks fulfill the condition (2) of the definition of a critical surface. □

We finally come to the proof of Theorem 6.9.

Theorem 6.9 *There is an infinite family of nontrivial links with critical bridge spheres.*

Proof All we must do to prove Theorem 6.9 is to demonstrate that there exists an infinite family of links which satisfy the hypothesis of Theorem 6.8. Let \mathcal{L} be the set of 2-twisted links in $(4, 4)$ -plat position with twist numbers $\{t_i^j\}$, such that all twist numbers in the top and bottom rows are positive, and all twist numbers in the middle row are negative, and such that t_2^2 and t_4^2 are both even. The parity requirement guarantees that if $L \in \mathcal{L}$, then L is a union of two different links, say L^1 and L^2 . See Figure 36. Consider L^1 ; it is a link in a 2-bridge position (with respect to F), so it can have at most two components.

Case 1 (L^1 has one component, ie it is a knot) Let $D(L^1)$ be the diagram for L^1 obtained by projection to the xz -plane. This diagram is alternating (for the same reason $D(L)$ was shown to be alternating in Section 3). Observe that $D(L^1)$ is a reduced diagram. (In other words, as a graph, it has no cut vertex.) One of the

famous Tait conjectures states that a reduced alternating diagram for a knot realizes the knot's minimal crossing number [9]. (This was proved in 1987 by Kauffman [7] and Murasugi [10].) Therefore as $D(L^1)$ has at least eight crossings, L^1 is not the unknot, so it must be a 2-bridge knot. Since L^1 is in bridge position, it is by definition not perturbed at any of its bridge arcs, one of which is α^1 . We conclude that L^1 is not perturbed at α^1 . Since L^1 and L^2 are not connected to each other, α^1 cannot share an endpoint with either of the two bridge arcs below F contained in L^2 . Therefore we conclude that L is not perturbed at α^1 .

Case 2 (L^1 has two components) Then each component is individually an unknot in bridge position. Unknots are not perturbed, so in particular, L^1 is not perturbed at α^1 . From there, it follows as in Case 1 that L is not perturbed at α^1 either.

By a similar argument, L is not perturbed at $\alpha^{4'}$ either. Therefore L satisfies Theorem 6.8 and its bridge sphere F is critical. This proves \mathcal{L} is an infinite family of links satisfying Theorem 6.9. \square

References

- [1] **D Bachman**, *Critical Heegaard surfaces*, Trans. Amer. Math. Soc. 354 (2002) 4015–4042 MR
- [2] **D Bachman**, *Topological index theory for surfaces in 3-manifolds*, Geom. Topol. 14 (2010) 585–609 MR
- [3] **D Bachman**, *Stabilizing and destabilizing Heegaard splittings of sufficiently complicated 3-manifolds*, Math. Ann. 355 (2013) 697–728 MR
- [4] **A J Casson**, **C M Gordon**, *Reducing Heegaard splittings*, Topology Appl. 27 (1987) 275–283 MR
- [5] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Math. Ser. 49, Princeton Univ. Press (2012) MR
- [6] **J Johnson**, **Y Moriah**, *Bridge distance and plat projections*, Algebr. Geom. Topol. 16 (2016) 3361–3384 MR
- [7] **L H Kauffman**, *State models and the Jones polynomial*, Topology 26 (1987) 395–407 MR
- [8] **J H Lee**, *Bridge spheres for the unknot are topologically minimal*, Pacific J. Math. 282 (2016) 437–443 MR
- [9] **W B R Lickorish**, *An introduction to knot theory*, Graduate Texts in Mathematics 175, Springer (1997) MR
- [10] **K Murasugi**, *Jones polynomials and classical conjectures in knot theory*, Topology 26 (1987) 187–194 MR

*Mathematics Department, Taylor University
Upland, IN, United States*

daniel_rodman@taylor.edu

Received: 17 February 2016 Revised: 6 May 2017