

Stabilizing Heegaard splittings of high-distance knots

GEORGE MOSSESIAN

Suppose K is a knot in S^3 with bridge number n and bridge distance greater than $2n$. We show that there are at most $\binom{2n}{n}$ distinct minimal-genus Heegaard splittings of $S^3 \setminus \eta(K)$. These splittings can be divided into two families. Two splittings from the same family become equivalent after at most one stabilization. If K has bridge distance at least $4n$, then two splittings from different families become equivalent only after $n - 1$ stabilizations. Furthermore, we construct representatives of the isotopy classes of the minimal tunnel systems for K corresponding to these Heegaard surfaces.

[57M25](#); [57M27](#)

1 Introduction

In 1933, Reidemeister [11] and Singer [13] independently showed that any two Heegaard splittings of a 3-manifold become equivalent after some finite number of stabilizations to each. No upper bound on the genus of this common stabilization was known until 1996, when Rubinstein and Scharlemann [12] showed that, for a non-Haken manifold, there is an upper bound which is linear in the genera of the two respective Heegaard splittings. Later, they extended these results to a quadratic bound in the case of a Haken manifold. In 2011, Johnson [6] showed that for any 3-manifold, there is a linear upper bound $3p/2 + 2q - 1$, where $p \geq q$ are the genera for the Heegaard surfaces. Examples of Heegaard splittings which required more than one stabilization to the larger-genus surface were not known until Hass, Thompson and Thurston [2] constructed examples with stable genus $p + q$ in 2009. Later that year, Bachman [1] gave examples with stable genus $p + q - 1$, and Johnson and Tomova [7] gave a combinatorial description of the geometric construction given by Hass, Thompson and Thurston.

One goal of this paper is to show that for most n -bridge knots, there are many pairs of genus- n Heegaard splittings of the knot exterior which have stable genus $2n - 1$. This is accomplished by means of the same machinery that Rubinstein and Scharlemann developed and Johnson refined for their respective upper bounds, the Rubinstein–Scharlemann graphic.

It was not previously known whether a knot complement can have Heegaard surfaces which require more than one stabilization to become equivalent. [Theorem 2.1](#) shows that such knots do exist. Though we do not prove it here, it is easy to see that the knots we consider are all hyperbolic. The 3-manifolds with stable genus $p + q$ constructed by Hass, Thompson and Thurston in [\[2\]](#) are negatively curved, but not hyperbolic. Their methods can be used to construct a hyperbolic 3-manifold with two Heegaard surfaces with stable genus $p + q - 4$, so [Theorem 2.1](#) extends this bound, as well. The geometry of high-distance knot complements and tunnel systems is a topic for future investigation.

The starting point of this work is a theorem of Tomova which states that for a knot with sufficiently high distance, the minimal bridge sphere is unique up to isotopy. This allows us to give a combinatorial description of all of the minimal-genus Heegaard splittings of the complement of such a knot. In turn, we use this description to come up with a list of all the tunnels belonging to any minimal tunnel system for such a knot. It turns out that this description precisely generalizes Kobayashi's [\[8\]](#) classification of unknotting tunnels for 2-bridge knots.

This work suggests some questions for further investigation. In particular, is it possible to extend this construction to closed manifolds? How can this construction be extended to positive-genus bridge surfaces in arbitrary manifolds? And, can a criterion be found for determining whether two Heegaard splittings of a knot complement with low-genus common stabilization are isotopic, as Morimoto and Sakuma [\[10\]](#) did for 2-bridge knots?

In [Section 2](#), we will define all the terms necessary for stating the main theorem. In [Section 3](#), we construct these Heegaard splittings. It follows from work of Tomova [\[14\]](#) that there are no others. In [Section 4](#), we show how they divide into two families, and we show that two splittings from the same family become equivalent after one stabilization. [Section 5](#) is devoted to recalling the machinery of Johnson developed in [\[4\]](#) and [\[7\]](#), and proving that it can be used in this situation. In [Section 6](#), this machinery is used to place a lower bound on the stable genus of two splittings from different families. Finally, in [Section 7](#), we show how to construct a tunnel system which corresponds to one of the Heegaard splittings constructed in the proof of [Theorem 2.1](#), so that any minimal tunnel system for K is equivalent to one of these tunnel systems.

Acknowledgments I would like to thank Abby Thompson for many discussions and suggestions without which this work would not have been possible, Jesse Johnson for the initial conversation which led to this work, and the anonymous referee who read the paper very thoroughly and suggested a very large number of improvements.

2 Preliminaries

Let M be a closed, compact, connected, orientable 3-manifold. A *Heegaard surface* for M is a closed, orientable, connected surface Σ embedded in M so that $\overline{M \setminus \Sigma}$ is a pair of handlebodies H_1, H_2 . The triple (Σ, H_1, H_2) is a *Heegaard splitting* of M . The minimal genus of all such surfaces Σ is the *Heegaard genus* of M . Two Heegaard splittings (Σ, H_1, H_2) and (Σ', H'_1, H'_2) are *isotopic* or *equivalent* if there is an ambient isotopy of M taking Σ to Σ' and H_i to H'_i . Otherwise, they are *distinct*. For a manifold with boundary, there is a natural generalization of this idea, discussed below.

Given a handlebody H , let $\Gamma \subset H$ be a graph which is a deformation retract of H . Then Γ is a *spine* of H . Removing a regular neighborhood $N \subset \text{Int } H$ of some (possibly empty, possibly disconnected) subgraph of Γ contained in H results in a *compression body*. However, we assume that no component of the boundary of this regular neighborhood is spherical. The spine of this compression body is the remainder of Γ , union the boundary of the regular neighborhood which was removed. If M is a compact orientable manifold with boundary, then ∂M is a closed orientable surface, and thus a Heegaard splitting of M can be defined as a triple (Σ, W_1, W_2) , where the W_i are compression bodies.

If $K \subset M$ is a knot, a *bridge surface* for K is a Heegaard surface Σ for M such that K intersects Σ transversely and $K \cap H^\pm$ is a boundary-parallel collection of arcs in each handlebody, called *bridge arcs*. It is possible that $K \cap \Sigma = \emptyset$, in which case we require that K can be isotoped to lie in Σ . Equivalently, $K \cap H^\pm$ is a collection of arcs γ_i such that there is a collection of arcs $(\alpha_i, \partial\alpha_i) \subset (\Sigma, \partial\gamma_i)$ which, together with the γ_i , cobound disks D_i which, within each handlebody, are pairwise disjoint. These are called *bridge disks*. If the genus of Σ is g and the number of points of $\Sigma \cap K$ is $2n$, then Σ is a (g, n) -bridge surface. If n is minimal over all surfaces of some fixed genus g , then K is said to be a (g, n) -bridge knot. If $g = 0$, this is reduced to simply saying that K is n -bridge, and Σ is a *bridge sphere*.

If a simple closed curve γ on a Heegaard or bridge surface Σ does not bound a disk in Σ or a disk punctured once by $K \cap \Sigma$ in Σ , we say that γ is *essential*. We define the *curve complex* $C(\Sigma)$, first introduced by Hempel in [3], of the surface Σ as follows: let vertices correspond to isotopy classes of essential simple closed curves on Σ , and let any two vertices corresponding to curves which have disjoint representatives on Σ be connected by an edge. To define the distance of the splitting (Σ, H_1, H_2) , let K_1 and K_2 be the collections of vertices corresponding to curves that bound disks to either side of Σ , which are disjoint from K in the case of a bridge splitting. Then the minimal length edge path in $C(\Sigma)$ between a vertex of K_1 and a vertex of K_2 is called the *distance* of Σ , and is denoted $d(\Sigma)$. If Σ is a bridge sphere for a knot $K \subset S^3$ which

realizes the minimal distance of all bridge spheres for K , then we say that the knot has distance $d(\Sigma)$, which is denoted as $d(K)$.

Let (Σ, W_1, W_2) be a genus- g Heegaard splitting of M . If γ is a properly embedded arc in, say, W_2 , which is parallel into Σ , let $W'_1 = W_1 \cup N(\gamma)$ and $W'_2 = \overline{W_2} \setminus N(\gamma)$, with $\Sigma' = \partial W'_2 = \partial W'_1$. Then W'_2 and W'_1 are still compression bodies, and (Σ', W'_1, W'_2) is a genus- $(g+1)$ Heegaard splitting for M which is a *stabilization* of (Σ, W_1, W_2) . If Σ was a (g, n) -bridge surface for a knot K , and γ a bridge arc, then Σ' is a *meridional stabilization* of Σ and is a $(g+1, n-1)$ -bridge surface. Reidemeister [11] and Singer [13] showed independently that any two inequivalent Heegaard surfaces for M have stabilizations which are equivalent. The minimal genus of this common stabilization is called the *stable genus* of the two surfaces.

Let $X_K = S^3 \setminus \eta(K)$. In this paper, we classify the minimal-genus Heegaard splittings and their common stabilizations for X_K when K is a high-distance knot:

Theorem 2.1 *Let $K \subset S^3$ be an n -bridge knot with $n \geq 3$ and bridge distance $d > 2n$. Then the Heegaard genus of X_K is n , and there are at most $\binom{2n}{n}$ distinct minimal-genus Heegaard splittings of X_K . If two Heegaard surfaces have K on the same side, then they have a stable genus of at most $n + 1$. If two Heegaard surfaces have K on opposite sides, then their stable genus is at least $\min(2n - 1, \frac{1}{2}d)$.*

Corollary 2.2 *Let $K \subset S^3$ be an n -bridge knot with $n \geq 3$ and bridge distance $d > 2n$. Then the tunnel number of K is $n - 1$, and there are at most $\binom{2n}{n}$ distinct tunnel systems for K , up to isotopy and edge slides.*

This extends Kobayashi's classification of unknotting tunnels for 2-bridge knots in [8]: the $\binom{4}{2} = 6$ unknotting tunnels classified there correspond to genus-2 Heegaard splittings of $S^3 \setminus K$, all of which become equivalent after at most one stabilization.

3 Constructing Heegaard splittings

The following is a special case of a theorem of Tomova:

Theorem 3.1 [14, Theorem 10.3] *Let K be a nontrivial knot in S^3 with bridge number $n \geq 3$. Let S be a bridge sphere for K . If Σ is a bridge surface for K , or a Heegaard surface for X_K , then either Σ is obtained from S by a sequence of meridional stabilizations and stabilizations, or $d(S) \leq 2 - \chi(\Sigma - K)$.*




Figure 1: Diagram of proof of Lemma 3.2

Thus, if K is n -bridge and $d(K) > 2n$, then every nonstabilized (g, n) -bridge surface for K , and every genus- n Heegaard surface for X_K , is constructed by meridionally stabilizing the unique bridge sphere S . This in turn implies that no sequence of meridional stabilizations of S yields a surface which is stabilized. This also implies that the bridge sphere S which realizes the bridge distance is unique up to isotopy.

For the purposes of establishing some notation, we prove the following obvious lemma:

Lemma 3.2 *If Σ' is obtained from a (g, n) -bridge surface Σ by a single meridional stabilization, then Σ' is a $(g + 1, n - 1)$ -bridge surface.*

Proof Note that, in order to be able to perform a meridional stabilization, n must be greater than 0.

Let V and W be the handlebodies on either side of Σ and γ an arc of $K \cap V$. Let $W' = W \cup N(\gamma)$ and $V' = \overline{V} \setminus N(\gamma)$, so $V' \cap W' = \Sigma'$. Since γ can be isotoped into ∂V , there is a disk $D_\gamma \subset V$ such that $\partial D_\gamma = \gamma \cup \alpha$ with $\alpha \subset \partial V$. Let $D'_\gamma = \overline{D_\gamma} \setminus N(\gamma)$. If $\varphi: N(\gamma) \rightarrow I \times D^2$ is the obvious homeomorphism sending γ to $I \times \{0, 0\}$, let $D = \varphi^{-1}(\{\frac{1}{2}\} \times D^2)$ be a meridian disk of $N(\gamma)$. Since $D'_\gamma \cap N(\gamma)$ is an arc parallel to γ , we have that $D'_\gamma \cap D$ is a single point, and therefore Σ' is a stabilization of Σ , and is a Heegaard splitting of S^3 . The genus of Σ' is $g + 1$, and the number of arcs of K in W' and V' each is $n - 1$. See Figure 1. \square

If Σ is an n -bridge sphere, a sequence of n meridional stabilizations results in an $(n, 0)$ -bridge surface Σ' . If the two handlebodies on either side of Σ' are V' and W' , one of them contains the knot K , say V' . Then V' has n compression disks D_i (the meridian disks D of $N(\gamma)$ in the proof of Lemma 3.2), which each intersect K in exactly one point. Thus K is parallel to a subset of a spine of V' , and therefore, Σ' is also a Heegaard surface for the knot complement X_K , with handlebody W' and compression body $\overline{V' \setminus N(K)}$.

The isotopy class of Σ' may depend on the choice of meridional stabilizations, or the order in which they are performed. For example, Figure 2 shows three possible




Figure 2: Three possible splittings of a 2-bridge knot complement

splittings of a 2-bridge knot complement resulting from meridional stabilizations of a bridge sphere. Though [Theorem 2.1](#) concerns knots with bridge number 3 or greater, [Figure 2](#) presents heuristic pictures of 2-bridge knots for simplicity. As we will see in [Section 4](#), for a 3-bridge knot with distance 6 or greater, the constructions in the left and middle of [Figure 2](#) do not represent the same surface because K is on opposite sides of the resulting surface. For a figure-8 knot, these two Heegaard splittings correspond to the two unique unknotting tunnels. The surface in [Figure 2](#) (right) may or may not be isotopic to the surface of [Figure 2](#) (left). For the figure-8 knot, they are indeed isotopic, but the question of when they are distinct remains open for knots with bridge number 2 or greater. For 2-bridge knots, it is answered by Morimoto and Sakuma in [\[10\]](#).

Before proving the next lemma, we note some definitions. Let S be an n -bridge sphere for K , so that $K \cap S = \{x_1, x_2, \dots, x_{2n}\}$, labeled by picking a point on K , following K along a fixed orientation, and labeling the points of $S \cap K$ in order of increasing index. The Heegaard surface Σ which was constructed by meridionally stabilizing S may be viewed as $(S \setminus \bigcup_{j=1}^{2n} N(x_j)) \cup (\bigcup_{i=1}^n A_i)$, where the neighborhoods $N(x)$ are pairwise disjoint and the A_i are annuli with $A_i \cap S = \partial A_i$ for all i . Each such annulus A_i is the image of an embedding

$$\varphi_i: (S^1 \times I, S^1 \times \{0\}, S^1 \times \{1\}) \rightarrow (S^3 \setminus K, \partial N(x_i), \partial N(x_j)),$$

and the images of all the φ_i are pairwise disjoint.

Definition 3.3 A *tubed Heegaard surface* is a Heegaard surface for X_K consisting of the $2n$ -punctured bridge sphere together with n annuli as shown above.

Note again that [Theorem 3.1](#) implies that if K is n -bridge and $d(K) > 2n$, then every minimal-genus Heegaard surface of X_K is a tubed Heegaard surface.

Definition 3.4 Let $\gamma, \alpha, D_\gamma, D'_\gamma$ be as in [Lemma 3.2](#), let $z \in S^1$ and choose φ_i so that $\varphi_i(\{z\} \times I) = D_\gamma \cap A_i$. Then $D_\gamma \setminus D'_\gamma$ is a square with sides $\gamma, \partial D'_\gamma \setminus \alpha$ and the two components of $\alpha \cap N(\gamma)$. Given an orientation of K and orienting I from 0 to 1, the annulus A_i runs along K if the arcs γ and $\varphi(\{z\} \times I)$ with the orientations induced by I are parallel in the square.

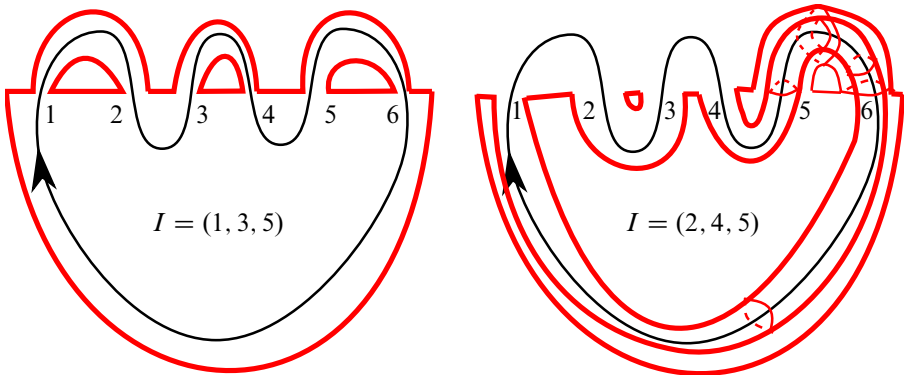


Figure 3: Diagram of Definition 3.6

The point of Definition 3.4 is just to have a unique way of identifying the annuli by one of the endpoints, as follows:

Definition 3.5 If A_i is an annulus that runs along K with $\varphi_i(S^1 \times \{0\}) \cap S = \partial N(x_i)$ and $\varphi_i(S^1 \times \{1\}) \cap S = \partial N(x_j)$, then A_i is said to have its *left foot* at x_i and *right foot* at x_j . The difference $j - i \pmod{2n}$ is the *length* of the annulus.

We take the convention that the annulus A_i has left foot at x_i , and $r(i)$ denotes the index of the right foot of the annulus.

Definition 3.6 Let K be an n -bridge knot and S a bridge sphere with $S \cap K = \{x_1, \dots, x_{2n}\}$ labeled as above. Let Σ be a tubed Heegaard surface with

$$\Sigma = \overline{\left(S \setminus \bigcup_{i=1}^{2n} N(x_i) \right)} \cup \left(\bigcup_{i=1}^n A_i \right),$$

such that each A_i runs along K . Let $I_0 = \{\alpha_i\}_{i=1}^{2n}$ be the set of indices from 1 to $2n$ that are left feet of the annuli. The set I_0 is the *index* of the tubed Heegaard surface Σ for the fixed orientation of K and labeling of $S \cap K$. See Figure 3.

Since every tubed Heegaard surface has an index, this definition gives the upper bound of $\binom{2n}{n}$ on the number of tubed Heegaard surfaces of K . Now we show that every choice of n numbers is the index of some tubed Heegaard surface:

Lemma 3.7 Let I_0 be any n -element subset of $\{1, 2, \dots, 2n\}$. This I_0 can be realized as the index of a tubed Heegaard surface.

Proof The desired surface is constructed by attaching the annuli to $S \setminus \bigcup_i N(x_i)$ in order of increasing length, which is equivalent to performing meridional stabilizations in the correct order. Define $J_0 = \emptyset$. For each $1 \leq k \leq n$, let

$$J_k = \left\{ \lambda \in I_0 \setminus \left(\bigcup_{\ell=0}^{k-1} J_\ell \right) \mid \lambda + (2k - 1) \pmod{2n} \notin I_{k-1} \right\}$$

and

$$I_k = I_{k-1} \cup \{ \lambda + (2k - 1) \pmod{2n} \mid \lambda \in J_k \}.$$

For each $\lambda \in J_k$, attach an annulus with left foot at x_λ and right foot at $x_{\lambda+2k-1}$, where the indices are all computed mod $2n$.

Thus, at each k , we have that I_k is the set I_0 together with the indices of points which are right feet of annuli that have already been attached, that is, indices which are no longer available for attaching annuli, and J_k is the set of indices of left feet of annuli of length $2k - 1$. By attaching the annuli shortest first, we guarantee that all the annuli can be made pairwise disjoint, since longer ones simply pass between the shorter annuli and the knot. After gluing the annuli of length $2n - 1$ with left feet in J_n , the resulting tubed Heegaard surface has index I_0 by construction. Thus every such set is the index of some tubed Heegaard surface, proving the lemma. \square

Lemma 3.7 together with the uniqueness of the n -bridge sphere given by **Theorem 3.1** imply the following corollary:

Corollary 3.8 *If K is an n -bridge knot with bridge distance at least $2n$, there are at most $\binom{2n}{n}$ distinct minimal Heegaard splittings of X_K .*

4 Common stabilizations of same-side tubed Heegaard splittings

The set of all tubed Heegaard surfaces for K can be divided into two families, depending on which side of the initial bridge sphere is adjacent to K in the final splitting. In this section, we prove that if two tubed Heegaard surfaces lie on the same side of K , they become equivalent after just one stabilization.

Definition 4.1 Let Σ be a surface in S_3 , and \mathcal{A} an annulus properly embedded in $S^3 \setminus \Sigma$. Let Σ' be obtained from Σ by cutting Σ along $\partial\mathcal{A}$ and gluing in two copies of \mathcal{A} . We say that Σ' is obtained from Σ by an *annulus compression* along the *compressing annulus* \mathcal{A} . See **Figure 4**.




Figure 4: Diagram of proof of Lemma 4.3

Definition 4.2 Let \mathcal{A} be a compressing annulus for a tubed Heegaard surface such that $\partial\mathcal{A}$ consists of essential loops contained in two annuli of Σ , namely $\gamma_i \subset A_i$ and $\gamma_j \subset A_j$. Let $\psi: S^1 \times I \rightarrow \mathcal{A} \hookrightarrow S^3$ be the homeomorphism with image \mathcal{A} , with $\psi(S^1 \times \{0\}) = \gamma_i$ and $\psi(S^1 \times \{1\}) = \gamma_j$. With the orientation on I from 0 to 1, if the induced orientation on $\psi(\{z\} \times I)$ for some $z \in S^1$ is parallel to the fixed orientation of K , then \mathcal{A} runs along K and is denoted as $[ij]$. Furthermore, to allow for the case $i = j$ without ambiguity, we require that $[ij]$ have $\gamma_i = \partial N(x_{r(i)})$ and $\gamma_j = \partial N(x_j)$. Note that the compressing annulus $[ij]$ is therefore different from the compressing annulus $[ji]$.

It is easy to construct the annulus $[ij]$ for any distinct i and j by taking a neighborhood of K which is small enough to be disjoint from Σ , and taking $[ij]$ to be an annular subset of the boundary of this neighborhood, running from x_i to x_j , together with a collar at each end to connect this subset annulus to Σ .

Lemma 4.3 If Σ is a tubed Heegaard surface for K and \mathcal{A} a compressing annulus $[ij]$ with $i \neq j$, then an annulus compression along \mathcal{A} is equivalent to a single stabilization and destabilization of Σ . If the index of Σ is I_0 , then after such an annulus compression the index of the resulting tubed Heegaard surface is $\sigma(I_0)$, where $\sigma \in S_{2n}$ is the 2-cycle $(r(i) j)$.

Proof Let ψ be as in Definition 4.2, and let $\beta = \psi(\{z\} \times I)$ be a spanning arc in \mathcal{A} connecting γ_i and γ_j . Let $N(\beta)$ be a small enough neighborhood of β in S^3 such that $N(\beta) \cap \Sigma$ is just two disks D_i and D_j in Σ which are neighborhoods of $\partial\beta$. Since \mathcal{A} runs along K , we have that β is parallel into Σ , and therefore there is a disk Δ properly embedded in $S^3 \setminus (\Sigma \cup \mathcal{A})$ which is cobounded by β and a copy of β which has been pushed into Σ . The disk Δ intersects a meridian disk of $N(\beta)$ in exactly one point, so $\Sigma' = (\Sigma - D_i - D_j) \cup \partial N(\beta)$ is a stabilization of Σ .

Consider $D = \mathcal{A} \setminus N(\beta)$. Its boundary ∂D consists of $\gamma_i \setminus N(\beta)$, $\gamma_j \setminus N(\beta)$, and two spanning arcs on $N(\beta)$. Since both A_i and A_j are the results of meridional

stabilization, they both have the disks D'_j which are subsets of bridge disks of K as in the proof of Lemma 3.2, and since $i \neq j$, each of these disks only intersects D in one point. Therefore, compressing along D is a destabilization of Σ' , and the end result is identical to performing an annulus compression along A .

The annulus compression $[ij]$ has the effect of connecting the left foot of A_i to the right foot of A_j , and the right foot of A_i is now the left foot of an annulus that ends at the left foot of A_j . Therefore, the index of the new tubed Heegaard surface is I_0 with $r(i)$ instead of j . See Figure 4. □

Observe that if $i = j$, the surface after compressing along $[ij]$ is disconnected, with one boundary-parallel component. For this reason, we will always assume that $i \neq j$ unless stated otherwise.

Definition 4.4 Given a bridge sphere for K , label one side $+$ and the other $-$. Let Σ be a tubed Heegaard surface derived from this bridge sphere. If K is adjacent to the minus ($-$) side of the portion of the bridge sphere outside the $N(x_i)$, we say K is *below* Σ . Otherwise, K is *above* Σ .

It is also possible to think of “above” and “below” in terms of Morse functions: let $K \subset S^3$, and let $f: S^3 \rightarrow [-1, 1]$ be a Morse function on S^3 such that $f^{-1}(t)$ is a 2-sphere for $t \in (-1, 1)$, $f|_K$ is also Morse with image $[-\frac{1}{2}, \frac{1}{2}]$, and $S = f^{-1}(0)$ is a bridge sphere for K . Let Σ be a tubed Heegaard surface for K from S such that $f|_\Sigma$ has image $[-\frac{3}{4}, \frac{3}{4}]$. If $f^{-1}(-1)$ is a point in the handlebody containing K , we say that K is *below* Σ . If, on the contrary, $f^{-1}(1)$ is contained in the handlebody containing K , then K is *above* Σ .

There are two obvious tubed Heegaard surfaces for any knot in bridge position: those with indices $(1, 3, \dots, 2n - 1)$ and $(2, 4, \dots, 2n)$. We take the convention that in the first, K is below Σ , and in the second, K is above Σ . Thus these will be referred to as Σ_b and Σ_a respectively.

If there is a genus- g tubed Heegaard surface Σ with index I such that $i, j \in I$, compressing along the annulus $[ij]$ yields another splitting Σ' with index I' . Considering the annulus compression as a stabilization followed by a destabilization as in the proof of Lemma 4.3, call the intermediate genus- $(g+1)$ surface Σ_1^+ . Another annulus compression along $[i'j']$ with $i', j' \in I'$ results in another surface Σ'' and is equivalent to another stabilization-destabilization pair with intermediate genus- $(g+1)$ surface Σ_2^+ . However, stabilizations are unique, so Σ_1^+ is, in fact, equivalent as a Heegaard surface to Σ_2^+ ; thus, after a sequence of handle slides, there is a single destabilization of Σ_1^+ which results in the surface Σ'' directly. Therefore, extending this argument inductively, we have proved the following:

Lemma 4.5 *Any two genus- g tubed Heegaard surfaces related by a sequence of $i \neq j$ annulus compressions have a common stabilization of genus $g + 1$.*

We also observe the following:

Lemma 4.6 *Two tubed Heegaard surfaces have K on the same side if and only if they are related by a sequence of annulus compressions of the $[ij]$ type.*

Proof First, we show that an annulus compression does not change on which side of the surface the knot lies. Pick a point on K , and consider a path γ transverse to Σ from this point to $f^{-1}(-1)$, where f is the fixed Morse function defined above. Let $\nu = |\gamma \cap \Sigma|$. We see that $\nu = 0 \pmod 2$ if and only if K is below Σ . Let Σ' be the result of performing the annulus compression along $[ij]$ on Σ , and let $\nu' = |\gamma \cap \Sigma'|$. If $|\gamma \cap [ij]| = m$, then $\nu' = \nu + 2m$, and so K is on the same side of Σ' as of Σ .

Conversely, suppose without loss of generality that Σ_1 and Σ_2 are two tubed Heegaard surfaces with K above both. Both surfaces are constructed from the bridge sphere S and $\gamma_1, \dots, \gamma_{2n}$ consecutive bridge arcs of K with respect to S . We let p_i (resp. q_i) be the number of annuli running along γ_i in Σ_1 (resp. Σ_2). Since K is above both Σ_1 and Σ_2 , we have $p_i - q_i = 0 \pmod 2$ for all i . Also note that $p_{i+1} = p_i \pm 1$, and similarly for q_i , so $\Delta = (p_{i+1} - q_{i+1}) - (p_i - q_i)$ is either 0 or ± 2 . Therefore, we can construct a collection of compressing annuli that will transform one surface into the other: if $\Delta = 0$, then the number of compressing annuli over γ_i is equal to the number passing over γ_{i+1} , and so no compressing annulus has a foot at $\partial N(\gamma_i \cap \gamma_{i+1})$. Otherwise, a compressing annulus has either a left or right foot at $\partial N(\gamma_i \cap \gamma_{i+1})$ for $\Delta = -2$ or 2, respectively. This gives a way to construct a sequence of compressing annuli which takes Σ_1 to Σ_2 . \square

Thus we have proved the main result of this section:

Lemma 4.7 *Given two minimal-genus tubed Heegaard surfaces for X_K where K is an n -bridge knot with bridge distance at least $2n$, if K either lies above or lies below both surfaces, then they become equivalent after one stabilization.*

5 Sweep-outs and the graphic

In this section, we recall the machinery developed in [4] and used in [7] to construct Heegaard splittings with large stable genus, and modify it for application in this situation.

A (g, n) -bridge surface S in (S^3, K) , with $n > 0$, splits the 3-sphere into two handlebodies H^\pm , each containing a boundary-parallel collection τ^\pm of n properly embedded arcs. If we restrict S to X_K , it splits X_K into two handlebodies with neighborhoods of the n arcs removed (we note that this is homeomorphic to a genus- $(g+n)$ handlebody). To begin generalizing this machinery to punctured handlebodies, we define a *spine* for $H^+ \setminus \eta(\tau^+)$ (similarly for $H^- \setminus \eta(\tau^-)$) to be a graph embedded in $H^+ \setminus \eta(\tau^+)$ which contains the spine of H^+ , together with an edge along a meridian circle on the boundary of each arc of bridge, and a vertical edge from a vertex of the spine of H^+ to each of the meridian circles. Note that a regular neighborhood of the spine on either side of S in X_K has boundary which is equivalent to $S - \eta(K)$, and that $H^\pm \setminus \eta(\tau^\pm)$ deformation retracts onto its spine. In the case that $g = 0$ and H^\pm are 3-balls, the spine of H^+ should be chosen to be a single vertex. We will call this the *central vertex*.

Let $\partial_- W$ be some collection of closed positive-genus surfaces. The *compression body* W is formed by attaching at least enough 1-handles to $\partial_- W \times I$ so that the result is connected. Then a *spine of W* is a deformation retract of W , best constructed by taking the collection $\partial_- W \times \{\frac{1}{2}\}$ together with arcs that are cores of the attached 1-handles. If W is a handlebody, then $\partial_- W = \emptyset$ and the spine is just a graph of which W is a regular neighborhood. For a Heegaard splitting (Σ, W_1, W_2) of some manifold M , where either W_i could be a compression body, Σ is equivalent to the positive boundary $\partial_+ W_i$ of a regular neighborhood in X_K of either the spine of W_1 or that of W_2 , and $\partial M = \partial_- W_1 \cup \partial_- W_2$.

A *splitting surface* in X_K is either a bridge surface for (S^3, K) restricted to X_K , or a Heegaard surface for X_K . A *sweep-out* for X_K is a map $f: S^3 \setminus \eta(K) \rightarrow [-1, 1]$ such that $f^{-1}(t)$ is equivalent to some splitting surface Σ for each $t \in (-1, 1)$, and $f^{-1}(\pm 1)$ is the two spines to either side of Σ . We say that f is a *sweep-out given by Σ* .

Let f be a sweep-out of X_K given by a tubed Heegaard surface Σ , and g a sweep-out of X_K given by the bridge sphere S . Given the map $f \times g: X_K \rightarrow [-1, 1]^2 = [-1, 1] \times [-1, 1]$, the *graphic* of $f \times g$ is the set of points $(s, t) \subset [-1, 1]^2$ for which $f^{-1}(s) = \Sigma_s$ is not transverse to $g^{-1}(t) = S_t$. For a generic $f \times g$, the graphic Γ is a 2- and 4-valent graph subset of $[-1, 1]^2$, except in one situation noted below, with a finite number of points along any vertical or horizontal line $\{s\} \times [-1, 1]$ or $[-1, 1] \times \{t\}$, and no two vertices sharing an s or t coordinate. See [4] for details. We note that at some subset of the four corners of the square, the spine corresponding to Σ is not transverse to the spine corresponding to S on the boundary of X_K ; the edges of the spine corresponding to S which are meridian circles lie on $\partial_- W_i$. Generically, the closure of these points does not intersect the interior of the square, so their existence will not affect our arguments; we thus ignore these points.

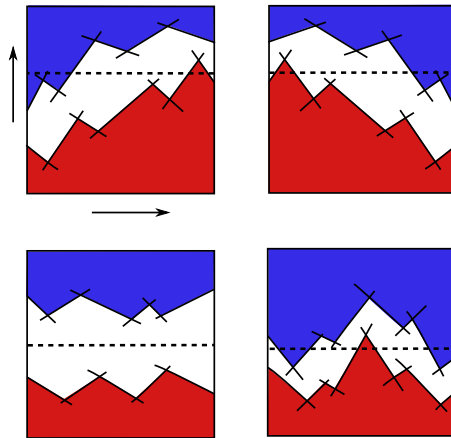


Figure 5: [4, Figure 2] Red and blue regions are where S_t is mostly above and below Σ_s , respectively. Clockwise from top left: spanning positively, spanning negatively, spanning with both signs, splitting

Let $\Theta_s^- = f^{-1}([-1, s])$ and $\Theta_s^+ = f^{-1}([s, 1])$. We say that S_t is *mostly above* Σ_s if each component of $S_t \cap \Theta_s^-$ is contained in a disk subset of S_t or in an annulus of S_t which has one boundary component in ∂X_K . Similarly, S_t is *mostly below* Σ_s if each component of $S_t \cap \Theta_s^+$ is contained in a disk subset of S_t or in an annulus of S_t which has one boundary component in ∂X_K . If there exist t_0, t_1 and s such that S_{t_0} is mostly below Σ_s and S_{t_1} is mostly above Σ_s , we say that Σ *spans* S . If $t_0 < t_1$, then Σ spans S *positively*, and if $t_1 < t_0$, then Σ spans S *negatively*. Note that it is possible for Σ to span S both positively and negatively.

If there is some value of s so that $\{s\} \times [-1, 1]$ contains no vertices of Γ and no points where S_t is either mostly above or mostly below Σ_s for all values $t \in [-1, 1]$, then we say that Σ *splits* S . Note that in this case, for this value of t , every $S_t \cap \Sigma_s$ contains at least one loop that neither bounds a disk in S_t nor is parallel to ∂S_t . A curve with these properties is called *essential* in S_t . See Figure 5.

If an essential curve on Σ bounds an embedded disk D in $X_K \setminus \Sigma$, this disk is a *compressing disk* for Σ . If the curve bounds an once-punctured disk whose other boundary component is a meridian of ∂X_K , this once-punctured disk is a *cut disk*. A *compression* of a connected surface Σ along the compressing disk D is a surface equivalent to the set of boundary components of $N(\Sigma \cup D)$ which are not isotopic to Σ . A *cut compression* of a connected surface Σ along the cut disk D is the set of boundary components of $N(\Sigma \cup D)$ which meet $N(D)$. A *c-compression* along a *c-disk* is either a compression along a compressing disk or a cut compression along a cut disk.

The following lemma is a version of [7, Theorem 3.1] and follows the same line of proof. We write it out in full to show how it adapts to the current situation.

Lemma 5.1 *Let Σ and S be two splitting surfaces for X_K . Suppose Σ spans S both positively and negatively. Then there is a sequence of compressions and cut compressions of Σ after which there are two components of the compressed surface which are parallel to S .*

Proof Let t_-, t_0, t_+ and s be the values such that S_{t_-} and S_{t_+} are mostly below Σ_s , and S_{t_0} is mostly above Σ_s . Every curve of $S_{t_-} \cap \Sigma_s$ is inessential in S_{t_-} . First, we claim that if for some t , an innermost curve $\gamma \in S_t \cap \Sigma_s$ is also inessential in Σ_s , then γ can be removed by isotopy of the surfaces.

Case 1 If γ is contained in a disk subset of S_t , then by the Jordan curve theorem, γ must itself bound a disk in S_t . If γ also bounds a disk in Σ_s , then since X_K is irreducible, the two disks cobound a ball across which Σ_s can be isotoped. Note that in this case, γ cannot bound an annulus in Σ_s , since this annulus together with the disk would form a sphere in S^3 punctured exactly once by K , but since S^3 contains no copies of $S_1 \times S_2$, this is a contradiction.

Case 2 If γ is contained in an annulus subset of S_t , it bounds either a disk or an annulus in S_t . If it bounds a disk, we revert to case 1. If it bounds an annulus, by the same argument as in case 1 we note that γ must also bound an annulus in Σ_s . If we replace $N(K)$ (ie perform a trivial Dehn filling), γ bounds two disks in S^3 , so the curve of intersection in X_K can be removed by sliding the two boundary components of the annuli in ∂X_K past each other.

After removing all curves of intersection which are inessential in both surfaces, we consider a curve $\gamma \in S_{t_-} \cap \Sigma_s$ which is innermost in S_{t_-} and essential in Σ_s . Therefore this curve bounds a c -disk for Σ_s , along which we compress Σ_s . We repeat this process with the compressed surface until there are no more curves of intersection of S_{t_-} and the resulting compressed, possibly disconnected surface, which we will call Σ'_s . Note that now S_{t_-} lies entirely to one side of Σ'_s .

We repeat the same process with all curves of intersection between Σ'_s and S_{t_+} , and we call the resulting compressed surface Σ''_s . Observe that now S_{t_-} and S_{t_+} are both on the same side of Σ''_s as they were both mostly below Σ_s . Finally, we repeat the process with all curves of intersection between S_{t_0} and Σ''_s , resulting in a surface Σ'''_s now disjoint from all three surfaces S_{t_-} , S_{t_+} and S_{t_0} , with S_{t_-} and S_{t_+} to one side of Σ'''_s and S_{t_0} on the other side of Σ'''_s .

We maximally compress and ∂ -compress Σ'''_s in the complement of $S_{t_-} \cup S_{t_0} \cup S_{t_+}$, and call the resulting surface F . This is an incompressible surface properly embedded

in a manifold homeomorphic to $S \times [0, 2]$, a $2n$ -punctured sphere cross interval. It separates $S \times \{0\}$ from $S \times \{1\}$, and separates $S \times \{1\}$ from $S \times \{2\}$. Therefore, F has at least two incompressible, ∂ -incompressible components, each inside a copy of $S \times I$ separating $S \times \{0\}$ from $S \times \{1\}$. It follows from [15, Proposition 3.1] that each of these two components of F is homeomorphic to S . \square

Corollary 5.2 *Let K be an n -bridge knot. If Σ is a Heegaard splitting for X_K which spans a splitting sphere S both positively and negatively, then the genus of Σ is at least $2n - 1$.*

Proof By Lemma 5.1, there is a sequence of compressions and cut compressions of Σ after which the resulting collection of surfaces contains two copies of S . Since $\chi(\Sigma)$ is nondecreasing under compressions and cut compressions, we have $\chi(\Sigma) \leq 2\chi(S)$, which implies that $2 - 2g \leq 4 - 4n$, yielding $g \geq 2n - 1$. \square

The following lemma is a version of [7, Theorem 4.2]. The proof follows identical arguments as in [4] and [7], so we only sketch details here and direct the reader there for details. It is generalized to splitting surfaces in the same way as Lemma 5.1 above.

Lemma 5.3 *If S and Σ are splitting surfaces for X_K , $\chi(\Sigma) < 0$, and Σ splits S , then $d(S) \leq 2 - \chi(\Sigma)$.*

Sketch of Proof. Let s be such that $\{s\} \times [-1, 1]$ is disjoint from any vertices of Γ , and S_t is neither mostly above nor below Σ_s for any $t \in (-1, 1)$. For every t , the intersection $\Sigma_s \cap S_t$ contains at least one curve essential in S_t . Let g and f be the sweep-outs given by S and Σ , respectively. Let $[\alpha, \beta]$ be the largest interval such that for no $t \in (\alpha, \beta)$ is any curve of $\Sigma_s \cap S_t$ both essential in S_t and bounding a disk or cut disk in either compression body to either side of Σ_s . (If no such essential interval exists, $d(S) \leq 1$ and we are done).

Consider $\Sigma' = \Sigma \cap g^{-1}([\alpha + \epsilon, \beta - \epsilon])$ for a small $\epsilon > 0$. The projection $\pi: \Sigma' \rightarrow S$ which factors through inclusion into $S \times I$ will send isotopy classes of curves in Σ' to isotopy classes of curves in S , and it will send curves bounding pairs of pants in Σ' to disjoint curves in S . Therefore, the induced map $\pi^*: \Sigma' \rightarrow \mathcal{C}(S) \cup \{0\}$ projects a pants decomposition of Σ' to a path in $\mathcal{C}(S)$, which has length at most $-\chi(\Sigma')$. Since $[\alpha, \beta]$ was chosen to be as large as possible, one more edge at each end of this path connects it to curves which bound disks in the handlebodies to either side of Σ , and $d(S) \leq 2 - \chi(\Sigma') \leq 2 - \chi(\Sigma)$. \square

Furthermore, we note that the stabilization of Σ can be done in such a way as to preserve the spanning properties of the graphic.

Lemma 5.4 *If Σ and S are splitting surfaces such that Σ spans S positively (resp. negatively), then a stabilization of Σ also spans S positively (resp. negatively).*

Proof Let f and g be sweep-outs respectively given by Σ and S such that Σ spans S positively. Let t_1, t_2 and s be the values such that S_{t_1} is mostly below Σ_s and S_{t_2} is mostly above Σ_s . Let Σ' be a stabilization of Σ_s , and let B be a ball which contains the canceling pair of handles added to Σ_s and intersects Σ_s in a disk, contained in $g^{-1}((t_1, t_2))$. Let f' be a sweep-out of Σ' so that $f'^{-1}(s) \setminus B$ is setwise equal to $f^{-1}(s) \setminus B$. Since S_{t_1} is mostly below Σ_s and disjoint from B , it is also mostly below Σ'_s , and likewise S_{t_2} is mostly above Σ'_s . Therefore, Σ' spans S with the same sign that Σ does. \square

Finally, the following lemma is proved by Johnson [5, Lemma 9] for graphics of Morse functions. We reproduce the proof, modifying it for sweep-outs and introducing some minor corrections and clarifications.

Lemma 5.5 *Let f_0 and f_1 be two sweep-outs given by equivalent Heegaard surfaces Σ and Σ' , and let g be a sweep-out given by the splitting surface S such that $f_0 \times g$ and $f_1 \times g$ are both generic. Then there is a family of sweep-outs f_r , for $r \in [0, 1]$, such that $f_r \times g$ is generic for all but finitely many r . At the nongeneric values, the graphic has at most two valence-two or -four vertices at the same level, or one valence-six vertex.*

The idea of the proof is as follows: We approximate $\{f_r\}$ with a piecewise linear path in $C^\infty(X_K, \mathbb{R})$ and consider consecutive vertices of this path, φ_0 and φ_1 . Then we consider the image in $[-1, 1]^3 = \{(s_0, s_1, t)\}$ of points in M where the surfaces given by $\varphi_0^{-1}(s_0)$, $\varphi_1^{-1}(s_1)$ and $g^{-1}(t)$ are not transverse, a “3D graphic” of sorts. Then we examine the projection of this image onto a plane in the cube which rotates about the t -axis from the $\varphi_0 \times g$ plane to the $\varphi_1 \times g$ plane, and classify the types of degeneracies which can occur. It can probably be proved for splitting surfaces in general with some attention given to singularities on the boundary, but we only need it for Σ and Σ' as Heegaard surfaces.

Proof If Σ and Σ' are equivalent, there is an ambient isotopy of M , say ϕ_r , which takes Σ to Σ' and takes the spines of Σ to the spines of Σ' . Let $f_r = f \circ \phi_r$ be the sweep-outs given by the surfaces of this one-parameter family. Because Morse functions are dense in $C^\infty(M, \mathbb{R})$, each f_r can be approximated arbitrarily closely by a Morse function f'_r in the C^∞ topology, meaning that for any $\epsilon > 0$, there exists a 1-parameter family of Morse functions f'_r so that for $t \in [-1 + \epsilon, 1 - \epsilon]$,

$f_r'^{-1}(t) = f_r^{-1}(t)$, so the graphic for $f_r' \times g$ is identical to that of $f_r \times g$ except on an arbitrarily small neighborhood of $\{-1, 1\} \times I$. In addition, we approximate g with a Morse function g' so that the graphic of $f_r' \times g'$ is different from that of $f_r \times g$ only in a small neighborhood of $\partial([-1, 1]^2)$.

Since the spines of the two splitting surfaces are disjoint except where they lie on ∂X_K , the difference between the graphics of $f_r' \times g'$ and $f_r \times g$ is the sequence of births and deaths that occur within ϵ of the boundary of the square, none of which share either s or t coordinates. By pulling these cusps back to the boundary of the square, we recover the graphic of $f_r \times g$.

Now, every f_r' is Morse, so there is an open, convex neighborhood N_r of f_r' in $C^\infty(X_K, \mathbb{R})$ such that every $\varphi \in N_r$ is isotopic to f_r' . We can cover $\{f_r'\}$ by such neighborhoods, and since $\{f_r'\} \cong [0, 1]$ is compact, there is a finite cover by such neighborhoods, so the path $\{f_r'\}$ can be replaced by a piecewise-linear path with each line segment contained in one convex neighborhood, and consecutive vertices $\varphi_0, \dots, \varphi_n \in \{f_r'\}$. For $\alpha \in [0, 1]$, let $\beta = \alpha/(1 - \alpha)$ and let $\pi_\beta: \mathbb{R}^2 \times \mathbb{R}^1 \rightarrow \mathbb{R}^1 \times \mathbb{R}^1$ be the projection of the \mathbb{R}^2 factor onto a line through the origin which has slope β . The graphic of the function $((1 - \alpha)\varphi_0 + \alpha\varphi_1) \times g'$ is the graphic of $\pi_\beta \circ (\varphi_0 \times \varphi_1 \times g')$: $X_K \rightarrow \mathbb{R}^1 \times \mathbb{R}^1$.

The maps can all be chosen so that $\varphi_i, \varphi_{i+1}, g', \varphi_i \times g', \varphi_{i+1} \times g'$ and $\varphi_i \times \varphi_{i+1}$ are all stable, meaning each has an open neighborhood of isotopic Morse functions in the appropriate vector space. Since a projection is a continuous map, the preimage of each of these open neighborhoods in $C^\infty(X_K, \mathbb{R}^3)$ under projection is an open neighborhood of $\varphi_i \times \varphi_{i+1} \times g'$, and their intersection is also open. Since stable functions between 3-manifolds are dense in the C^∞ topology, as shown by Mather [9], we can ensure that $\varphi_i \times \varphi_{i+1} \times g'$ will be stable. Then, by Mather's classification of singularities of stable maps between 3-manifolds [9], the discriminant set of $\varphi_i \times \varphi_{i+1} \times g'$ is a compact 2-submanifold $S \subset X_K$, and its image is an immersed 2-manifold with cusps.

Let $F = \varphi_i \times \varphi_{i+1} \times g'$. For $p \in S$, if $F(p)$ is not a cusp, then there is a map $T_p S \rightarrow T_{F(p)} F(S)$. If this plane is parallel to the $s_0 - s_1$ plane, this means that the tangency between $\varphi_i^{-1}(F(p))$ and $\varphi_{i+1}^{-1}(F(p))$ is preserved under two degrees of freedom independent of the parameter of g' , which means p is at a vertex of the graphic of $\varphi_i \times \varphi_{i+1}$. As this graphic is generic, there are only finitely many such points. When $T_{F(p)} F(S)$ is not parallel to the $s_0 - s_1$ plane, the intersection of $T_{F(p)} F(S)$ with the plane $t = g'(p)$ determines a slope in that plane. For each $p \in S$ where $F(p)$ is not a cusp and $T_{F(p)} F(S)$ is not parallel to $t = g'(p)$, let $\sigma(p)$ be this slope. Observe that, though the number of noncusp points where the slope is not well-defined is finite, the intermediate value theorem implies that any connected level curve of


$\sigma^{-1}(\beta)$ must contain a nonzero, even number of points p at which $T_{F(p)}F(S)$ is parallel to $t = g'(p)$.

We may perturb F slightly to ensure that σ is a Morse function on S away from the finite number of points and curves (corresponding to the vertices of the graphic $\varphi_i \times \varphi_{i+1}$ and cusps of $F(S)$). This allows us to locally identify a patch of $F(S)$ with the graph of some function $\gamma: \mathbb{R}^2 \rightarrow \mathbb{R}$. If $x = \gamma(y, z)$, and for some fixed $z = g'(p)$, $dx/dy = \sigma(F^{-1}(x, y, z))$, we have that there is some smooth $\gamma': \mathbb{R} \rightarrow \mathbb{R}$ so that

$$\gamma(y, z) = \gamma'(z) + \int_{(x,0,z)}^{(x,y,z)} \sigma \circ F^{-1} dy.$$

The discriminant set of $\pi_\beta \circ F$ is the image under F of the closure of the set of points $\sigma^{-1}(\beta)$ in S .

Since σ is Morse, all but finitely many β are regular values. A preimage of a regular β value is a 1-dimensional submanifold of S . When β is passing through a regular value, points in the interior of $\sigma^{-1}(\beta)$ map to the interior of an edge of the graphic of $\pi_\beta \circ F$. If $p \in \overline{\sigma^{-1}(\beta)} \setminus \sigma^{-1}(\beta)$, then it is either a point at which $T_{F(p)}F(S)$ is parallel to $t = g(p)$, or $F(p)$ belongs to a cusp of $F(S)$. In the first case, the intermediate value theorem implies that such points must come in pairs. If no such points are introduced or removed at this level set β , then such points get mapped to the interior of an edge as well. If a pair of such points is introduced or removed, the graphic of $\pi_\beta \circ F$ gains or loses a pair of cusps connected by an edge in a “dovetail” type move. In the case that $F(p)$ is a cusp, p gets sent to a cusp of the graphic of $\pi_\beta \circ F$. Generically, the preimage of cusp points of $F(S)$ is transverse to level sets of σ . When β is regular and $\sigma^{-1}(\beta)$ is tangent to an arc of a cusp, a pair of cusps is introduced in the graphic of $\pi_\beta \circ F$ in either another dovetail move or an eye, which can be pictured as follows:



When β passes through a critical value of σ , the index of the critical point is either 0, 1 or 2. If it is 0 or 2, then there must be a pair of tangencies to $\mathbb{R}^2 \times \{g'(p)\}$ immediately after (resp. before) this value of β , which do not exist before (resp. after) this value of β . In the graphic of $\pi_\beta \circ F$ these two corresponding cusps cancel each other out in an “eye,” a process during which two vertices occupy the same level. If the critical point has index 1, a pair of points p_i with tangent planes $T_{F(p_i)}F(S)$ which

are parallel to $g'(p)$ cancel and another pair is born, leading to an exchange of the order of vertices of the graphic during which there are two valence-two vertices at the same level.

Thus there are only finitely many values of β (and so, of α as well) at which $((1 - \alpha)\varphi_i + \alpha\varphi_{i+1}) \times g'$ has a critical point which fails to be stable. In between these critical points, the graphic changes by some homotopy of the image of the discriminant set, which can be done generically except at a finite number of points introducing a triple point, self-tangency or double cusp points. \square

6 Common stabilizations of opposite-sided tubed Heegaard surfaces

To use the above results, we must understand how tubed Heegaard surfaces span a bridge sphere S . We begin by showing that, for the knots and surfaces under consideration here, every graphic must either split or span.

Lemma 6.1 *Let f and g be sweep-outs for a Heegaard surface Σ and a minimal bridge sphere S , respectively. If the bridge number of K is at least 3, the graphic must either split or span.*

Proof If the graphic neither splits nor spans, then the graphic of $f \times g$ must contain a vertex at the coordinates (s, t) such that S_t is mostly above $\Sigma_{s-\epsilon}$ and mostly below $\Sigma_{s+\epsilon}$ for a small ϵ . Then $f|_{S_{t+\epsilon}}$ is a Morse function on $S_{t+\epsilon}$. If (s, t) is a valence-4 vertex, this Morse function has exactly two index-1 critical points, and the rest are index 0 or 2. Then S admits a decomposition into a collection of disks and once-punctured disks, which have Euler characteristic at least 0, connected by two bands of Euler characteristic -1 , so $\chi(S) \geq -2$. If (s, t) is a valence-6 vertex, then a similar argument shows that $\chi(S) \geq -3$. But if $n \geq 3$, we have that $\chi(S) \leq -4$, a contradiction. \square

Lemma 6.2 *If K lies below (resp. above) Σ , then there are sweep-outs f for Σ and g for S such that Σ spans S positively (resp. negatively).*

Proof First, note that if $d(S) > 2n$ and Σ is a genus- n surface, Lemmas 5.1, 5.3 and 6.1 guarantee that Σ spans S either positively or negatively, but not with both signs. We take the convention that the sweep-outs are oriented so that $\partial M \subset f^{-1}(-1)$, and $g^{-1}(-1)$ is the spine on the minus side of S .

If K lies below Σ , then for any value of $s \in (-1, 1)$, there is a spine of S on the minus side of Σ_s which is entirely disjoint from Σ_s , and therefore, there is a value t_- close to -1 and a sweep-out g so that $S_{t_-} \subset f^{-1}([-1, s])$, so S_{t_-} is mostly below (in fact, entirely below) Σ_s . Fix one such s . Since the graphic must span either positively or negatively, there must be another $t_+ > t_-$ such that S_{t_+} is mostly above Σ_s , and so Σ spans S positively. One way to construct this pair (s, t_+) explicitly is to take s very close to -1 , so that Σ_s is very close to ∂M , and then for some t_+ , it is clear that $S_{t_+} \cap f^{-1}([-1, s])$ is a collection of once-punctured disks in S_{t_+} .

A symmetric argument shows that if K lies above Σ , then there are sweep-outs so that Σ spans S negatively. □

Furthermore, every pair of sweep-outs f and g for Σ and S define a graphic in which Σ spans S with the same sign, independent of the choice of sweep-outs.

Lemma 6.3 *Let K be an n -bridge knot with $n \geq 3$, $d(K) > 2n$, and Σ a Heegaard surface for X_K . If there are sweep-outs f and f' corresponding to Σ such that f spans g positively and f' spans g negatively, then during any isotopy f_r with $f_0 = f$ and $f_1 = f'$, there is some generic r such that f_r must either split g or span g with both signs.*

Proof Since f_0 and f_1 are given by equivalent surfaces, there is a one-parameter family of sweep-outs $\{f_r\}$, and by Lemma 5.5, we know that the graphic $f_r \times g$ is generic except at finitely many values of $r \in (0, 1)$, where it has two valence-two or valence-four vertices occupying the same level, or a single valence-6 vertex.

For contradiction, assume that f_r neither spans nor splits g with both signs away from the nongeneric r . Therefore, f_r splits g either positively or negatively, but not both. Then there is some nongeneric r_0 such that $f_{r_0-\epsilon}$ splits g positively and $f_{r_0+\epsilon}$ spans g negatively. At this r_0 , the regions where S is mostly above and mostly below Σ must exchange, so there is a vertex (s, t) in the graphic of $f_{r_0} \times g$ so that S_t is, say, mostly above $\Sigma_{s-\epsilon}$ and mostly below $\Sigma_{s+\epsilon}$. The proof of Lemma 6.1 shows that this cannot happen.

Therefore, for some generic r_0 , either f_{r_0} spans g both positively and negatively, or f_{r_0} splits g . □

Corollary 6.4 *If Σ is a tubed Heegaard surface with K below (resp. above) Σ , $b(K) \geq 3$ and $d(K) > 2b(K)$, then for any sweep-outs f of Σ and g of S with $\partial M \subset f^{-1}(-1)$ and $g^{-1}(-1)$ on the minus side of S , we have that f spans g positively (resp. negatively).*

Proof If there are two graphics, one in which f spans g positively, and the other in which f spans g negatively, by Lemma 6.3, there is an isotopy between them during which the graphic will either split or span with both signs. But $g(\Sigma) = n$ and $n < 2n - 1$, so by Corollary 5.2, the graphic cannot span with both signs; also, $n < \frac{1}{2}d(S)$, so by Lemma 5.3, the graphic cannot split. \square

Corollary 6.5 *Let K be an n -bridge knot with $n \geq 3$ and $d(K) > 2n$. Let Σ be a tubed Heegaard surface such that K lies below Σ , and Σ' a tubed Heegaard surface such that K lies above Σ' . Then a common stabilization of Σ and Σ' has genus at least $\min\{\frac{1}{2}d(S), 2n - 1\}$.*

Proof Let f_0'' be the sweep-out given by Σ'' as a stabilization of Σ , so by Lemmas 6.2 and 5.4 f_0'' spans g positively. Likewise, let f_1'' be the sweep-out given by Σ'' as a stabilization of Σ' , so f_1'' spans g negatively. By Lemma 6.3, at some point during the isotopy between these, the graphic either spans with both signs or splits. The conclusion follows. \square

Theorem 2.1 now directly follows from Corollary 3.8, Lemma 4.5 and Corollary 6.5.

Corollary 6.6 *If $d(S) \geq 4n$, then $g(\Sigma'') = 2n - 1$.*

Proof Σ corresponds to some system of $n - 1$ tunnels for K , and Σ' corresponds to another such system. We take a neighborhood of K union all $2n - 2$ tunnels; the boundary of the result is a genus- $(2n - 1)$ Heegaard surface which is a stabilization of both Σ and Σ' . \square

7 The tunnel systems

Any Heegaard splitting of a knot complement determines a tunnel system for that knot. We now show how to construct such a tunnel system for any tubed Heegaard surface, and we state the generalization of Kobayashi's classification of unknotting tunnels for 2-bridge knots in [8]. A tubed Heegaard splitting splits the knot complement into one handlebody and one compression body. By choosing a nonseparating system of $n - 1$ compression disks for the compression body, the cocores of these 2-handles can be extended to a tunnel system for K . In this section, we give a canonical way to make this choice.

Definition 7.1 Let A_i be an annulus of a tubed Heegaard splitting which runs along K and has left foot at x_i . If there is no annulus A_j with left foot at x_j running along K with length $\ell > (i - j) \pmod{2n}$, then A_i is *adjacent to K* .

In other words, A_i is adjacent to K if no other annulus runs between A_i and K . If $\ell \leq (i - j) \pmod{2n}$, then x_j and $x_{r(j)}$ are separated by x_i in the direction of the orientation of the annulus. Otherwise, any annulus with left foot between i and $r(i)$ also has its right foot in this range. Let all such annuli be *surrounding* A_i .

Definition 7.2 Let Σ be a tubed Heegaard surface. An annulus of Σ which is adjacent to K , together with all the annuli which surround this one, is called a *chunk* of Σ . The *size* of the chunk is the total number of annuli. The annulus which is adjacent to K is the *defining annulus* for the chunk.

The splittings Σ_a and Σ_b have n chunks all of size 1, the 3–bridge splitting with index $(1, 2, 3)$ has a single chunk of size 3, and the 7–bridge splitting with index $(1, 3, 5, 6, 7, 11, 12)$ has four chunks of sizes 1, 1, 3 and 2. We now show how to associate $k - 1$ tunnels to each chunk of size k .

Assume without loss of generality that K is below Σ . Annuli of Σ which have their left feet at odd values of i will be said to go *up*, while those with even-numbered left-foot indices will be said to go *down*. Since K is below Σ , any annulus which is adjacent to K must go up. Let there be a chunk of size k with defining annulus A_i . Let γ_i be the bridge arc of K between x_i and x_{i+1} . The $k - 1$ tunnels associated to this chunk, $\tau_1^i, \dots, \tau_{k-1}^i$, will all have $\partial\tau_j^i \subset \gamma_{i-1}$.

Let A_j be an annulus of this chunk which goes down. If A_j has length 1, let D be the bridge disk associated to γ_j . Let W_1 be the compression body which contains K . The component of $D \cap W_1$ which meets $S - \eta(K)$ is a nonseparating compression disk for Σ contained in W_1 . A tunnel for K which is dual to this compression disk is an arc with endpoints on γ_{i-1} which intersects this disk in one point and can be leveled into Σ . Such an arc is shown in Figure 6.




Figure 6: A tunnel (blue) which is dual to a compression disk (green) which corresponds to a down-annulus of length 1 when K is below Σ

If A_j is a down-annulus with length greater than 1, then we construct a compression disk for Σ on the same side as K by banding together the two bridge disks associated to γ_j and $\gamma_{r(j)-1}$ with a strip which follows A_j , as shown in Figure 7.




Figure 7: Banding together two bridge disks along a down-annulus of length 3 to make a compression disk for Σ whose dual is a tunnel

If A_k is an up-annulus which is *not* the defining annulus A_i , we make use of a nonseparating compression disk for Σ whose boundary separates the ends of A_k . We choose this disk so that $\partial D \subset S$ separates the punctures $\{x_{r(k)}, x_{r(k)+1}, \dots, x_{r(i)-1}\}$ from the others. The arc which starts at γ_{i-1} , pierces this disk once, goes through A_k and returns to γ_{i-1} , is parallel into Σ ; see Figure 8.




Figure 8: Constructing a tunnel corresponding to an up-tube in a chunk when K is below Σ

In this way, assigning tunnels dual to banded bridge disks for down-tubes and dual to nonseparating compression disks for up-tubes, every chunk of size k contributes $k - 1$ tunnels, since all annuli but the defining annulus contributes exactly one tunnel. The remainder of the tunnels connect adjacent chunks: if there is a chunk defined by the annulus A_i , then there is a tunnel τ_i with $\partial_- \tau_i \in \gamma_{i-1}$ and $\partial_+ \tau_i \in \gamma_{r(i)}$. However, if we connect all adjacent chunks together, there will be n tunnels total. One of these is dual to a compression disk for Σ in S^3 which intersects K in one point, so we simply choose which chunk to not connect to its neighbor, thereby omitting one redundant tunnel. So that this choice is consistent for all partitions of $\{1, \dots, 2n\}$ into chunks,

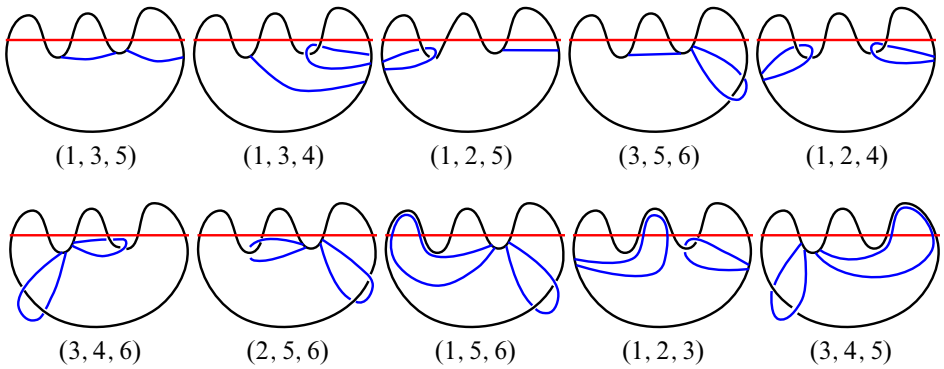


Figure 9: 10 of the 20 tunnels of a 3-bridge knot, and their indices

we omit the tunnel τ_i which connects γ_{i-1} to $\gamma_{r(i)}$ where i is the smallest index over all defining annuli for Σ .

In constructing these tunnels, several choices were made: the decision to use the structure of the chunks is a choice of a system of compression disks. Within this choice, the structure of each chunk as being given by up- and down-tunnels is once again a choice of compression disks. Since any two complete systems of compression disks for a given handlebody are related by a sequence of handle slides, the tunnel systems corresponding to these choices are equivalent after a sequence of edge slides. The choice of endpoints for the tunnels is also trivial modulo edge slides of the tunnels. Finally, the choice of which redundant chunk-connector tunnel is to be omitted is arbitrary, but also corresponds to a choice of system of compressing disks. Thus, any minimal tunnel system for K is equivalent to a minimal tunnel system for K constructed in this way, after a sequence of isotopies and edge slides.

Figure 9 shows how we construct the $\frac{1}{2}\binom{6}{3} = 10$ two-tunnel tunnel systems of a 3-bridge knot which correspond to a Heegaard surface which is above K . The horizontal line represents a bridge sphere, and the braids are omitted for simplicity. The other ten tunnel systems are symmetric to the given ones, on the other side of the bridge sphere.

References

- [1] **D Bachman**, *Heegaard splittings of sufficiently complicated 3-manifolds, I: Stabilization*, preprint (2009) [arXiv](#)
- [2] **J Hass**, **A Thompson**, **W Thurston**, *Stabilization of Heegaard splittings*, *Geom. Topol.* 13 (2009) 2029–2050 [MR](#)
- [3] **J Hempel**, *3-manifolds as viewed from the curve complex*, *Topology* 40 (2001) 631–657 [MR](#)

- [4] **J Johnson**, *Bounding the stable genera of Heegaard splittings from below*, J. Topol. 3 (2010) 668–690 [MR](#)
- [5] **J Johnson**, *Automorphisms of the three-torus preserving a genus-three Heegaard splitting*, Pacific J. Math. 253 (2011) 75–94 [MR](#)
- [6] **J Johnson**, *An upper bound on common stabilizations of Heegaard splittings*, preprint (2011) [arXiv](#)
- [7] **J Johnson**, **M Tomova**, *Flipping bridge surfaces and bounds on the stable bridge number*, Algebr. Geom. Topol. 11 (2011) 1987–2005 [MR](#)
- [8] **T Kobayashi**, *Classification of unknotting tunnels for two bridge knots*, from “Proceedings of the Kirbyfest” (J Hass, M Scharlemann, editors), Geom. Topol. Monogr. 2, Geom. Topol. Publ., Coventry (1999) 259–290 [MR](#)
- [9] **JN Mather**, *Stability of C^∞ mappings, V: Transversality*, Advances in Math. 4 (1970) 301–336 [MR](#)
- [10] **K Morimoto**, **M Sakuma**, *On unknotting tunnels for knots*, Math. Ann. 289 (1991) 143–167 [MR](#)
- [11] **K Reidemeister**, *Zur dreidimensionalen Topologie*, Abh. Math. Sem. Univ. Hamburg 9 (1933) 189–194 [MR](#)
- [12] **H Rubinstein**, **M Scharlemann**, *Comparing Heegaard splittings of non-Haken 3-manifolds*, Topology 35 (1996) 1005–1026 [MR](#)
- [13] **J Singer**, *Three-dimensional manifolds and their Heegaard diagrams*, Trans. Amer. Math. Soc. 35 (1933) 88–111 [MR](#)
- [14] **M Tomova**, *Multiple bridge surfaces restrict knot distance*, Algebr. Geom. Topol. 7 (2007) 957–1006 [MR](#)
- [15] **F Waldhausen**, *On irreducible 3-manifolds which are sufficiently large*, Ann. of Math. 87 (1968) 56–88 [MR](#)

Department of Mathematics, University of California, Davis,
1 Shields Avenue, Davis, CA 95616, United States

gmooss@math.ucdavis.edu

<http://math.ucdavis.edu/~gmooss>

Received: 8 September 2015 Revised: 15 February 2016

