Algebraic & Geometric Topology 13 (2013) 1125–1141

Bridge number and tangle products

RYAN BLAIR

We show that essential punctured spheres in the complement of links with distance three or greater bridge spheres have bounded complexity. We define the operation of tangle product, a generalization of both connected sum and Conway product. Finally, we use the bounded complexity of essential punctured spheres to show that the bridge number of a tangle product is at least the sum of the bridge numbers of the two factor links up to a constant error.

57M25, 57M27, 57M50

1 Introduction

Bridge number is a classical link invariant originally introduced by Schubert as a tool to study companion tori. In [4], Schubert proves the remarkable fact that given a composite knot K with summands K_1 and K_2, then the following equality holds: $\beta(K) = \beta(K_1) + \beta(K_2) - 1$, where $\beta(L)$ denotes the bridge number of a link L. Schultens gives a modern proof of this equality in [5].

Connected sum is a classical and intentionally restrictive method of amalgamating two links in S^3 together to create a new link in S^3. Tangle products are the natural generalization of this amalgamation operation. Roughly speaking, to form an n–strand tangle product of links K_1 and K_2 remove an n–strand rational tangle from the 3–sphere containing K_1 and the three sphere containing K_2. Now, glue the resulting tangles together via some homeomorphism of the $2n$–punctured sphere, S. The result is a tangle product, denoted $K_1 *_S K_2$. For a rigorous definition see Section 4. In particular, connected sums are 1–strand tangle products and Conway products are 2–strand tangle products. Conway products were studied in [3] where Scharlemann and Tomova produced Conway products which respected multiple bridge surfaces. How bridge number behaves with respect to Conway products was studied by the author in [2]. The goal of this paper is to generalize Schubert's equality for bridge number to the operation of tangle product.

Because of their generality and the choices involved, tangle products are exceptionally poorly behaved. For example, for any two n bridge knots K_1 and K_2 there exists

Published: 17 April 2013 DOI: 10.2140/agt.2013.13.1125

an n–strand tangle product $K_1 *_S K_2$ isotopic to the unknot. Hence, we will have to restrict our hypothesis to achieve a meaningful lower bound for $\beta(K_1 *_S K_2)$ in terms of $\beta(K_1)$ and $\beta(K_2)$. Similarly, if U is the unknot, there exist tangle products $U *_S U$ of arbitrarily high bridge number. This observation implies, in the absence of additional information, that there does not exist an upper bound for $\beta(K_1 *_S K_2)$ in terms of $\beta(K_1)$ and $\beta(K_2)$. The following is the main theorem.

Theorem 1.1 *Given an n–strand tangle product $K_1 *_S K_2$ such that there exists a minimal bridge sphere for $K_1 *_S K_2$ of distance at least three and the product sphere S is c–incompressible, then $\beta(K_1 *_S K_2) \geq \beta(K_1) + \beta(K_2) - n(10n-6)$.*

In [2], the key additional hypothesis needed to produce a lower bound on the bridge number of a Conway product in terms of the bridge number of the two factor links was that bridge position and thin position coincide for the Conway product. In contrast, the result presented here is heavily dependent on the hypothesis that the tangle product has a minimal bridge sphere of distance at least three. This hypothesis allows for a much stronger structure theorem then that found in [2] and, thus, a more restrictive lower bound on the bridge number of a tangle product. It is important to note that having the property that bridge position is thin position and having the property that a minimal bridge sphere is distance at least three are believed to be independent conditions.

Acknowledgements Research partially supported by NSF and JSPS grants.

2 Preliminaries

In this paper we study smooth links in S^3. Our central tool will be the standard height function on S^3, $h\colon S^3 \to [-1,1]$. The level surfaces of h foliate S^3 into concentric 2–spheres and two exceptional points. The *bridge number* of a link K is the minimal number of maxima of $h|_K$ over all isotopic morse embeddings of K.

A *tangle* is an ordered pair (B,T) where B is a 3–ball and $T \subset B$ is a properly embedded collection of arcs and loops. An *untangle* is a tangle (B,T) such that T is a collection of boundary parallel arcs. We say an embedded surface in S^3 is k–punctured if it meets K transversely in k points. A *bridge sphere* for a link K is an k–punctured sphere decomposing (S^3,K) into two untangles (H_1,T_1) and (H_2,T_2).

Definition 2.1 A bridge sphere Σ is *h–level* if there exists a regular value r such that Σ is isotopic to $h^{-1}(r)$.

Given an embedded morse surface S in S^3, let F_S be the singular foliation on S induced by $h|_S$. A *saddle* is any leaf of this foliation homeomorphic to the wedge of two circles. By standard position, we can assume that all saddles of F_S are disjoint from K.

Definition 2.2 Let K be a link embedded in S^3 and S be a surface embedded in S^3 which meets K transversely. We say that the pair (K, S) is in bridge position with respect to the standard height function on S^3 if h is a morse function when restricted to both K and S and there exist $a, b \in [-1, 1]$ such that

(1) all maxima of K and all maxima of S lie in $h^{-1}((b, 1))$;
(2) all minima of K and all minima of S lie in $h^{-1}((-1, a))$;
(3) all saddles of S and all intersection points $S \cap K$ lie in $h^{-1}((a, b))$.

Lemma 2.3 *For any embeddings of K and S in S^3 there is an isotopy of first S and subsequently K such that the resulting pair (K, S) is in bridge position. Moreover, these isotopies fix S and K outside of a neighborhood of their maxima and minima, preserve the number of maxima of $h|K$, and the number of saddles of F_S.*

Proof After a small isotopy, we can assume that $h|_K$ and $h|_S$ are morse functions and S intersects K transversely. Additionally, by general position, we can assume that both K and S are disjoint from both $h^{-1}(1)$ and $h^{-1}(-1)$. Let b_o be the largest value among the heights of all saddles of F_S and all points of $S \cap K$. Let a_o be the smallest value among the heights of all saddles of F_S and all points of $S \cap K$. Let $a = a_o - \frac{1}{2}(1 + a_o)$ and $b = b_o + \frac{1}{2}(1 - b_o)$. Let M_1 be the highest maximum of S that lies below $h^{-1}(b)$. Let α be a monotone arc connecting M_1 to any point in S^3 above $h^{-1}(b)$. Since M_1 is the highest maximum below $h^{-1}(b)$, we can choose α so that the interior of α is disjoint from both K and S. The portion of the boundary of a regular neighborhood of α lying above M_1 is a monotone disk D such that D is disjoint from K and S except in its boundary. Let D_M be a regular neighborhood of M_1 in S. The monotone disk D together with D_M cobound a 3–ball whose intersection with S is D_M. Isotope D_M to D across this 3–ball. This isotopy fixes K, is supported in a neighborhood of M_1 in S, and raises one maximum of S above $h^{-1}(b)$. Repeat this process until all maxima of S lie above $h^{-1}(b)$. Now that there are no maxima of S between $h^{-1}(a)$ and $h^{-1}(b)$ we can similarly isotope all the maxima of K above $h^{-1}(b)$ via an isotopy that fixes S and is supported on a neighborhood of the maxima of h_K. Hence, we have achieved (1) in the definition of (K, S) bridge position while preserving the number of maxima of K and the number of saddles of S. By a symmetric argument, we can also achieve (2) in the definition of (K, S) bridge position while preserving the number of minima of K and the number of saddles of S. After these isotopies, our choice of a and b guarantee that (3) in the definition of (K, S) bridge position is satisfied. □

Definition 2.4 A punctured surface S is *taut* with respect to an h–level bridge sphere Σ if $\mathcal{F}_S$ contains the fewest number of saddles subject to (K, S) being in bridge position.

This notion of taut is different then that found in [2]. Specifically, a taut surface in [2] is one that has a minimal number of saddles subject to h_K having a minimal number of maxima. In this paper, a taut surface has a minimal number of saddles subject to a specified bridge sphere appearing as a level sphere of the height function.

3 The saddle structure of n–punctured spheres

In this section we use the notion of taut introduced in the previous section to develop constraints on $\mathcal{F}_S$ when S is an embedded essential punctured sphere. For a more detailed discussion of the following definitions and their applications see [2].

Any given saddle $\sigma = s_1^\sigma \vee s_2^\sigma$, lies in a level sphere $S_\sigma = h^{-1}(h(\sigma))$. Let D_1^σ be the closure of the component of $S_\sigma - s_1^\sigma$ that is disjoint from s_2^σ and D_2^σ be the closure of the component of $S_\sigma - s_2^\sigma$ that is disjoint from s_1^σ.

A subdisk D in $\mathcal{F}_S$ is monotone if its boundary is entirely contained in a leaf of $\mathcal{F}_S$ and the interior of D is disjoint from every saddle in $\mathcal{F}_S$. In practice, we will use the term subdisk in a slightly broader sense, allowing ∂D to be immersed in S (ie ∂D is a saddle). We say a monotone disk is *outermost* if its boundary is s_i^σ for some saddle σ and label the disk D_σ. Similarly, if s_i^σ bounds an outermost disk D_σ, we say σ is an outermost saddle. It is usually the case that only one of s_1^σ and s_2^σ is the boundary of an outermost disk, so, our convention is to relabel so that $\partial D_\sigma = s_1^\sigma$. We say σ is an *inessential saddle* if σ is an outermost saddle and D_σ is disjoint from K.

Suppose σ is an outermost saddle. The sphere S_σ cuts S^3 into two 3–balls. The one that contains D_σ is again cut by D_σ into two 3–balls B_σ and B'_σ. We chose the labeling of B_σ and B'_σ so that $\partial B_\sigma = D_1^\sigma \cup D_\sigma$. We say a saddle σ is *standard* if there is a monotone disk E_σ in S such that $\partial E_\sigma = \sigma$ and E_σ is disjoint from K.

By general position arguments, we can assume every saddle σ in $\mathcal{F}_S$ has a bicollared neighborhood in S that is disjoint from K and all other singular leaves of $\mathcal{F}_S$. The boundary of this bicollared neighborhood consists of three circles c_1^σ, c_2^σ and c_3^σ where

c_1^σ and c_2^σ are parallel to s_1^σ and s_2^σ respectively. We can assume c_1^σ, c_2^σ and c_3^σ are level with respect to h and that c_1^σ and c_2^σ lie in the same level surface.

Figure 1 illustrates all of the terminology outlined above.

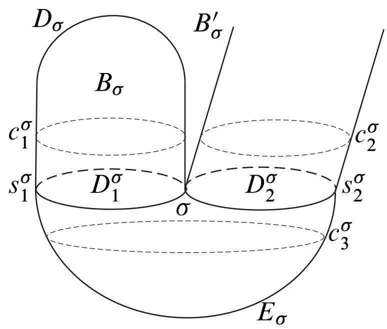

Figure 1: σ is a standard, outermost saddle

For the remainder of the paper S will always denote a punctured sphere. We will denote the point $h^{-1}(1)$ as $+\infty$ and the point $h^{-1}(-1)$ as $-\infty$.

Lemma 3.1 *Let σ be an outermost saddle in $\mathcal{F}_S$. There is an ambient isotopy of S that fixes K, lowers minima of S, raises maxima of S, and fixes S outside of a neighborhood of the maxima and minima of S such that, after this isotopy, B_σ does not contain $+\infty$ or $-\infty$.*

Proof See [5, Lemma 1]. □

The following lemma is an extension of [5, Lemma 2] to our alternative notion of taut.

Lemma 3.2 *Suppose Σ is an h–level bridge sphere for a link K. If $\mathcal{F}_S$ contains an inessential saddle, then S is not taut with respect to Σ.*

Proof Suppose (K, S) is in bridge position with respect to Σ, an h–level bridge sphere. Let σ be an inessential saddle in $\mathcal{F}_S$. We can assume D_σ contains a unique maximum and, by Lemma 3.1, B_σ does not contain $+\infty$. Let (a, b) be the interval in the definition of (K, S) bridge position. There exists an open interval (p, q) such that $h(\sigma) \in (p, q) \subset (a, b)$ and $h^{-1}(p, q)$ contains no saddles other than σ, no maxima or minima of $\mathcal{F}_S$, no maxima or minima of K and no points of $K \cap S$. Let $s = h(\sigma)$. As described in the following paragraph, *horizontally shrink and vertically lower* B_σ so that the result of the isotopy, call it B_σ^*, is contained in $h^{-1}([s, q))$.

An isotopy is *level-preserving* if for every $r \in [-1,1]$ the isotopy fixes $h^{-1}(r)$ setwise. By definition, a level-preserving isotopy preserves the saddle structure of $\mathcal{F}_S$ and the number and height of critical values of h_K. Since $\mathcal{F}_{B_\sigma}$ is a collection of disks, then we can preform a level-preserving isotopy resulting in each leaf of $\mathcal{F}_{B_\sigma}$ being a standard metric disk of radius ρ_r for each r. After a second level-preserving isotopy, we can assume each leaf of $\mathcal{F}_{B_\sigma}$ is centered over a common point in $h^{-1}(s)$. After a third level-preserving isotopy that contracts leaves of $\mathcal{F}_{B_\sigma}$ when $\rho_r > \rho_s$ and expands leaves of $\mathcal{F}_{B_\sigma}$ when $\rho_r < \rho_s$, we can assume that B_σ is a right cylinder. Horizontally shrink B_σ until it is contained in $h^{-1}([s,q))$. Finally, undo the three level-preserving isotopies. The result is an isotopy supported in a neighborhood of B_σ that lowers B_σ into a neighborhood of $h^{-1}(s)$ and preserves the saddle structure of $\mathcal{F}_S$ and the number of critical values of h_K.

Let S^* be the image of S under this isotopy. Similarly, let D^*_σ be the image of D_σ under this isotopy and let m be the unique maximum of D^*_σ. Let M be the level surface containing m. Hence, $M \cap S^*$ consists of the point m and a collection of circles.

Prior to the isotopy, we can assume both c_1^σ and c_2^σ are contained in M. Since σ is the unique critical point of h_S in $h^{-1}(p,q)$, then c_1^σ and c_2^σ are incident to a common component of $M - S$. Thus, after the isotopy m and c_2^σ are incident to a common component of $M - S^*$. We can choose a point n in c_2^σ and an arc α in M that is disjoint from S^* except at its boundary $\{m,n\}$. Additionally, let β be an arc in S^* that does not meet K, has boundary $\{m,n\}$ and is transverse to $\mathcal{F}_C$ everywhere except where it passes through $s_1^\sigma \cap s_2^\sigma$ so that α and β cobound a vertical disk F that is disjoint from S except along β. Additionally, after possibly rechoosing α, we can assume F is disjoint from K since h_K has no critical values in (p,q). Isotope S^* along F to effectively cancel a saddle with a maximum; see Figure 2. Let S^{**} be the image of S^* under this second isotopy. Notice that the ambient isotopies described fix all of S^3 below $h^{-1}(s)$, hence, $h^{-1}(s)$ remains a bridge sphere for K isotopic to Σ. The only way for (K^{**}, S^{**}) to fail to be in bridge position is for maxima of S^{**} or K to lie below a saddle of S^{**} or a point of $K \cap S$. In this case use the isotopy from Lemma 2.3 to raise the maxima of S^{**} and subsequently the maxima of K. Thus, we have produced an ambient isotopy of S and K which reduces the number of saddles of S but preserves both Σ as an h–level bridge sphere and (K,S) in bridge position. Hence S is not taut with respect to Σ. $\square$

Definition 3.3 We say σ is a *removable saddle* if σ is an outermost saddle where D_σ has a unique maximum(minimum) and $h|_{K \cap B_\sigma}$ has a local endpoint maximum (minimum) at every point of $K \cap D_\sigma$; see Figure 3. Otherwise, we say σ is nonremovable.

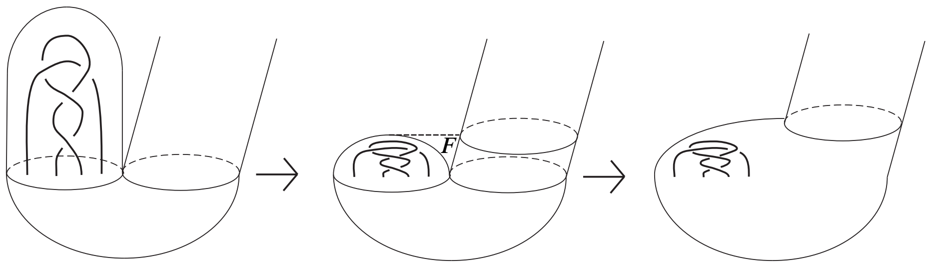

Figure 2

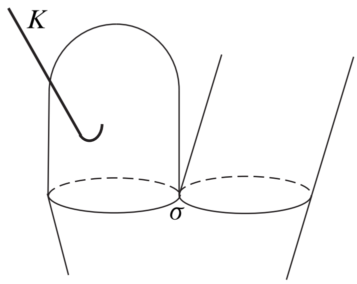

Figure 3

The following is an adaptation of [2, Lemma 3] to the notion of taut presented here.

Lemma 3.4 *Suppose Σ is an h–level bridge sphere for a link K. If F_S contains a removable saddle, then S is not taut with respect to Σ.*

Proof Suppose (K,S) is in bridge position with respect to Σ, an h–level bridge sphere. Let σ be an removable saddle in F_S. We can assume D_σ contains a unique maximum and, by Lemma 3.1, B_σ does not contain $+\infty$. Applying the isotopy presented in Lemma 3.2, we see that each point $x_i \in K \cap D_\sigma$ together with the image of x_i in B_σ^* bound monotone subarcs of K^*; see Figure 4. Since σ is removable neither the x_i nor the x_i^* is a maximum or minimum. Thus, we have eliminated a saddle of S while preserving (K,S) bridge position, a contradiction to tautness of S. □

The punctured sphere S decomposes S^3 into two 3–balls B_1 and B_2. Let σ be a saddle in F_S and L be the level sphere containing c_1^σ and c_2^σ. Then $L-(c_1^\sigma \cup c_2^\sigma)$ is composed of two disks and an annulus A. If a collar of ∂A in A is contained in B_1, then we say σ is *unnested* with respect to B_1. If not, we say σ is nested with respect to B_1. We define nested and unnested with respect to B_2 similarly. Note that nested

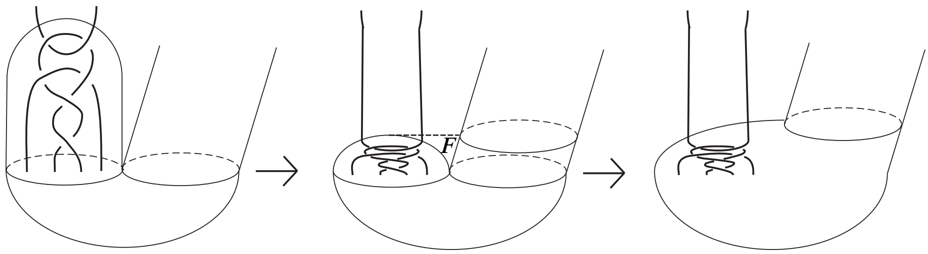

Figure 4

with respect to B_1 is the same as unnested with respect to B_2 and nested with respect to B_2 is unnested with respect to B_1.

Two saddles $\sigma = s_1^\sigma \vee s_2^\sigma$ and $\tau = s_1^\tau \vee s_2^\tau$ in F_S are *adjacent* if, up to labeling, s_i^σ and s_j^τ cobound a monotone annulus in S that is disjoint from K.

The following lemma is an extension of [5, Lemma 3] to our alternative notion of taut.

Lemma 3.5 *Suppose Σ is an h–level bridge sphere for a link K. If F_S contains adjacent saddles σ and τ where σ is a standard saddle and σ and τ are nested with respect to different 3–balls, then S is not taut with respect to Σ.*

Proof Due to the symmetry of the argument we can assume that E_σ has a unique maximum. Since σ and τ are adjacent, then, up to relabeling, s_1^σ and s_1^τ cobound a monotone annulus A in S that is disjoint from K. After a small isotopy of K and S, we can assume that $K \cup S$ meets D_2^σ in a collection of points and simple closed curves. After a small tilt of D_2^σ, $A \cup E_\sigma \cup D_2^\sigma$ is a monotone disk. Eliminate the saddle τ by applying the isotopy from the proof of Lemma 3.2 to the 3–ball cobounded by the monotone disk $A \cup E_\sigma \cup D_2^\sigma$ and the level disk D_1^τ that is disjoint from D_2^τ; see Figure 5. As in the proof of Lemma 3.4, we see that each point of $K \cap D_2^\sigma$ together with its image under the isotopy cobound monotone subarcs of K. Similarly, each curve $S \cap D_2^\sigma$ together with its image under the isotopy cobounds monotone annulus in S. Thus, we have produced an ambient isotopy of S and K which reduces the number of saddles of S but preserves both Σ as an h–level bridge sphere and (K, S) in bridge position. Hence, S is not taut with respect to Σ. □

The previous lemmas in this section are independent of bridge sphere distance. Below we define the distance of a bridge sphere and obtain additional constraints on taut punctured spheres.

Let Σ be a $2n$–punctured bridge sphere separating (S^3, K) into two n–strand untangles (H_1, T_1) and (H_2, T_2). Let $\mathcal{C}_n$ be the curve complex for Σ. Let $\mathcal{V}_1$ be the set of

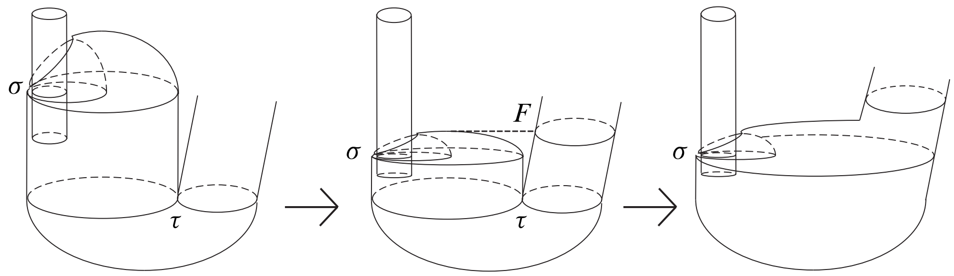

Figure 5

all isotopy classes of essential simple closed curves in $\partial(H_1) - T_1$ that bound disks in $H_1 - T_1$. Define $\mathcal{V}_2$ analogously. The *distance* of Σ, denoted $d(\Sigma)$, is the distance between $\mathcal{V}_1$ and $\mathcal{V}_2$ in $\mathcal{C}_n$ where the metric structure of $\mathcal{C}_n$ arises from assigning a length of one to each edge.

Lemma 3.6 *Suppose K is a link with a bridge sphere Σ of distance three or greater, then K is not split and K is prime.*

Proof This follows immediately from Bachman and Schleimer [1, Theorem 5.1]. □

Lemma 3.7 *Suppose there exists a bridge sphere Σ for K of distance three or greater. If S is a c–incompressible punctured sphere that is taut with respect to an h–level embedding of Σ, then there do not exist standard saddles σ, ρ and τ in $\mathcal{F}_S$ such that σ is adjacent to ρ and ρ is adjacent to τ.*

Proof By Lemma 3.6, we know K is nonsplit and prime. We will proceed by proving the contrapositive of the above statement. Suppose that three such saddles σ, ρ and τ do exist. If $\mathcal{F}_S$ contains inessential or removable saddles, then S is not taut by Lemmas 3.2 or 3.4. If σ and ρ are nested with respect to different three balls, then S is not taut by Lemma 3.5. Similarly, if ρ and τ are nested with respect to different three balls then S is not taut by Lemma 3.5. Hence, we can assume that all three saddles σ, ρ and τ have a common nesting and $\mathcal{F}_S$ contains no inessential and no removable saddles.

Assume that E_ρ has a unique maximum and $h(\sigma) > h(\tau)$. If E_ρ has a unique minimum or $h(\sigma) < h(\tau)$ the proof follows similarly. By Lemma 2.3, we can assume there is a partitioning of the critical values of $h|_K$ and $h|_S$ as in Definition 2.2. By standard Morse theory arguments, we can assume that the collection $\mathcal{C}$ of all critical points of $h|_K$, all critical points of $h|_S$, and all points of $S \cap K$ occur at distinct heights.

Choose c so that c is strictly between $h(\sigma)$ and the height of the next higher element of $\mathcal{C}$. Hence, $h^{-1}(c)$ is a bridge sphere for K that is isotopic to Σ. Although an abuse of notation, we will refer to $h^{-1}(c)$ as Σ. Recall the definition of c_1^σ and c_2^σ from Section 3 and Figure 1. We can assume that c_1^σ and c_2^σ lie on Σ.

Let A_τ be the monotone annulus with boundary $s_2^\rho \cup s_1^\tau$. Since σ is assumed to be higher than τ then A_τ intersects Σ in a single simple closed curve c_τ. In particular, c_τ is isotopic to both s_2^ρ and s_1^τ; see Figure 6.

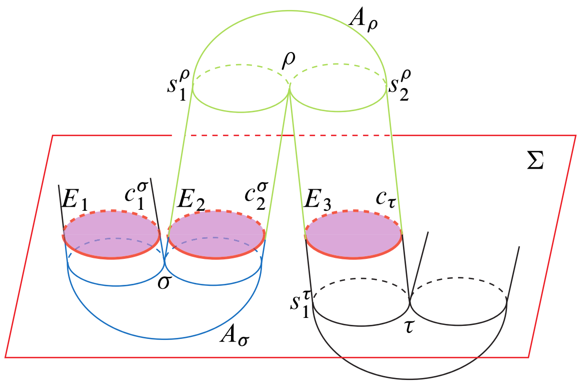

Figure 6

Claim If any of c_1^σ, c_2^σ or c_τ bounds a zero or once punctured disk in Σ, then S is not taut.

Proof of claim Since all three saddles σ, ρ and τ have a common nesting, each of c_1^σ, c_2^σ and c_τ bound pairwise disjoint disks E_1, E_2 and E_3 respectively in Σ. To prove the claim it suffices to show that each of the disks E_1, E_2 and E_3 meet K in at least two points.

Suppose, to form a contradiction, that E_1 is disjoint from K. $S \cap E_1$ is a collection of disjoint simple closed curves. An innermost such curve γ bounds a disk D in Σ that is disjoint from K and meets S only in its boundary. By c–incompressibility of S, γ also bounds a disk D' in S that is disjoint from K. Since K is nonsplit, D and D' cobound a 3–ball disjoint from K. Hence, we can eliminate γ as a curve of intersection by isotoping D' across this 3–ball and just past D. This isotopy leaves K fixed and can only decrease the number of saddles in F_C. By repeating this process, we can eliminate all curves of intersection of S with $\mathrm{int}(E_1)$. Again, by c–incompressibility of S, ∂E_1 bounds a disk D' in C which is disjoint from K. Since K is nonsplit D' and E_1 cobound a 3–ball. Isotope D' across this 3–ball to E_1 while fixing K. If D' contains σ, then this isotopy eliminates σ, ρ and τ while creating no new saddles. In this case S was not taut. If D' does not contain σ then, after the isotopy, σ is an inessential saddle. Hence, S is not taut, by Lemma 3.2.

Suppose, to form a contradiction, that E_1 meets K in exactly one point. Then $S \cap E_1$ is a collection of disjoint simple closed curves. An innermost such curve γ bounds a disk D in Σ that meets S only in its boundary. If D is disjoint from K then apply the argument in the preceding paragraph to eliminate γ. Hence we can assume D meets K exactly once. By c–incompressibility of S, γ also bounds a disk D' in S that meets K exactly once. Since K is prime, D and D' cobound a 3–ball containing an unknotted arc of K. Hence, we can eliminate γ as a curve of intersection isotopying D'_K across this 3–ball and just past D_K. This isotopy leaves K fixed and can only decrease the number of saddles in $\mathcal{F}_C$. By repeating this process, we can eliminate all curves of intersection of S with $\mathrm{int}(E_1)$. Again, by c–incompressibility of S, ∂E_1 bounds a disk D' in C which meets K exactly once and $E_1 \cup D$ is the boundary of a three ball containing an unknotted arc of K. Isotope D' across this 3–ball to E_1 while fixing K. If D' contains σ, then this isotopy eliminates σ, ρ and τ while creating no new saddles. In this case S was not taut. If D' does not contain σ then, after the isotopy, σ is a removable saddle. Hence, S is not taut, by Lemma 3.4.

By applying nearly identical arguments, we can show that both E_2 and E_3 meets K in at least two points. The claim then follows. □

Let M be the three ball above Σ and N be the three ball below Σ. By construction c_1^σ and c_2^σ cobound an annulus A_σ properly embedded in N and disjoint from K. Similarly, c_2^σ and c_τ cobound and annulus A_ρ properly embedded in M and disjoint from K. Since $(M, K \cap M)$ and $(N, K \cap N)$ are both untangles then A_σ and A_ρ are boundary compressible in $M - K$ and $N - K$. Let H_σ be the disk in M gotten by boundary compressing A_σ and H_ρ be the disk in M gotten by boundary compressing A_ρ. Both ∂H_σ and ∂H_ρ are disjoint from c_2^σ. By the above claim, ∂H_ρ and ∂H_σ are essential in Σ. Since $\partial H_\sigma \in \mathcal{A}$, $\partial H_\rho \in \mathcal{B}$, and both ∂H_σ and ∂H_ρ are disjoint from c_2^σ, $d(B) = d(\mathcal{A}, \mathcal{B}) < 3$. □

We summarize the results of the previous lemmas using the following definition.

Definition 3.8 A singular foliation $\mathcal{F}_S$ for a closed surface S with k–marked points is *admissible* if it is induced by the standard height function on S^3 via some Morse embedding of S into S^3 such that the following hold:

(1) there do not exist standard saddles σ, ρ and τ in $\mathcal{F}_S$ such that σ is adjacent to ρ and ρ is adjacent to τ;

(2) every outermost disk of $\mathcal{F}_S$ contains at least one marked point.

The following structure lemma is of independent interest. Bachman and Schleimer showed that twice the genus plus the number of boundary components of an essential surface serves as an upper bound for the distance of any bridge sphere for a knot [1]. In other words, a high distance bridge sphere forces a high *intrinsic* complexity for any essential embedded surface in a knot complement. In contrast, Theorem 3.9 implies that if K has a bridge sphere of distance 3 or greater, then any essential punctured sphere can be isotoped so that the number of saddles in the induced foliation is bounded with respect to the number of punctures. Colloquially, if there exists a bridge sphere that is not low distance, then every essential punctured sphere has low *extrinsic* complexity.

Theorem 3.9 *Suppose there exists a bridge sphere Σ for K of distance three or greater and S is taut with respect to an h–level embedding of Σ, then the following hold:*

(1) *there do not exist standard saddles σ, ρ and τ for S such that σ is adjacent to ρ and ρ is adjacent to τ;*
(2) *every outermost disk of the foliation of S induced by the standard height function contains at least one point of $K \cap S$.*

Proof By Lemmas 3.2 and 3.7, $\mathcal{F}_S$ is admissible with marked points $S \cap K$ when S is taut. □

For a fixed surface type and number of puncture points, the number of saddles in any admissible singular foliation is bounded.

Lemma 3.10 *If S is topologically a 2–sphere with k–marked points and $\mathcal{F}_S$ is admissible, then the number of saddles in $\mathcal{F}_S$ is at most $5k-8$.*

Proof Let $\mathcal{F}_S$ be an admissible singular foliation for S that has a maximal number of saddles.

Claim All marked points are contained in the collection of outermost disks of $\mathcal{F}_S$.

Proof of claim Suppose not. Hence, there is a puncture point x not on an outermost disk of $\mathcal{F}_S$. There is a nonsingular leaf α in $\mathcal{F}_S$ containing x. Isotope S in a neighborhood of α as in Figure 7. The resulting singular foliation remains admissible but has one more saddle then $\mathcal{F}_S$, contradicting the maximality of the number of saddles of $\mathcal{F}_S$. □

Claim All outermost disks of $\mathcal{F}_S$ contain exactly one puncture point.

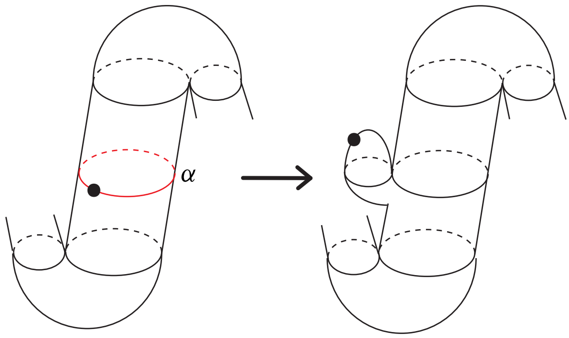

Figure 7

Proof of claim Every outermost disk of F_S meets K in at least one point, by the definition of admissible. Suppose there is an outermost disk D_σ that contains two puncture points x and y. Let α_x and α_y be the closed curves in F_S containing x and y respectively. Up to relabeling we can assume that α_y bounds a monotone disk in F_S that contains x. Alter F_S in a neighborhood of α_y as in Figure 8. The resulting singular foliation remains admissible but has one more saddle than F_S, contradicting the maximality of the number of saddles of F_S. □

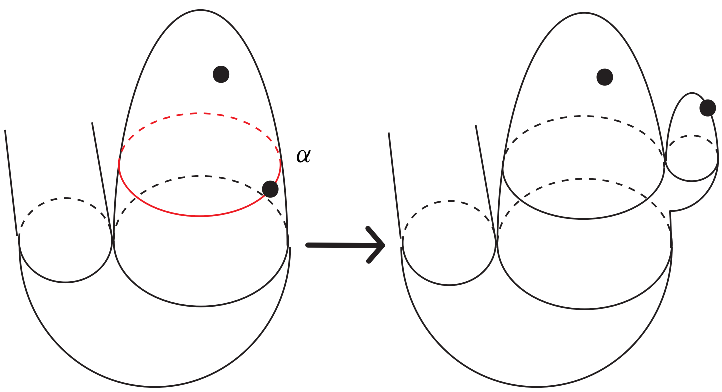

Figure 8

We proceed by induction on $k \geq 2$.

Suppose $k = 2$, then the maximal number of saddles is two, as depicted in Figure 9.

Assume F_S has at most $5(l-1)-8$ saddles for $k = l-1$. We will show F_S has at most $5l-8$ saddles for $k = l$.

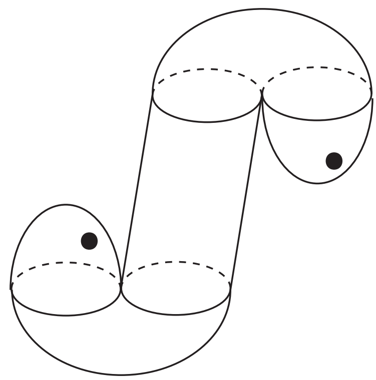

Figure 9

Choose F_S to have the maximal number of saddles of any admissible singular foliation. Since all puncture points are contained in outermost disks of F_S and all outermost disks of F_S contain exactly one puncture point, then S can be decomposed into a collection of 1–punctured disks, annuli and pairs of pants glued together along a collection of circles so that the 1–punctured disks are in one-to-one correspondence with outermost disks of F_S, the annuli are in one-to-one correspondence with the standard saddles of F_S and the pairs of pants are in one-to-one correspondence with the nonstandard saddles. Since 1–punctured disks and annuli contribute zero to Euler characteristic, pairs of pants contribute -1, and $\chi(S) = 2 - l$, then F_S contains exactly $l - 2$ nonstandard saddles. Since we can assume $l > 2$ there exists a nonstandard saddle for F_S. Label this saddle σ. By passing to an outermost nonstandard saddle and possibly relabeling c_1^σ, c_2^σ and c_3^σ we can assume that both c_1^σ and c_2^σ bound disks F_1^σ and F_2^σ respectively in F_S such that both F_1^σ and F_2^σ are disjoint from all nonstandard saddles. Each of F_1^σ and F_2^σ meets a unique outermost disk of F_S. By condition (2) of the definition of admissible and the previous claim, each of these disks contains a unique marked point. By condition (1) of the definition of admissible, each of F_1^σ and F_2^σ contains at most two saddles, all of which are standard. If one of F_1^σ or F_2^σ contains fewer than two saddles, then we would contradict the fact that F_S was chosen to contain a maximal number of saddles. Hence, we can assume that each of F_1^σ and F_2^σ contains exactly two saddles. Replace a neighborhood of the union of F_1^σ and F_2^σ in F_S with a single monotone disk M containing a unique marked point. Call the resulting singular foliation F_S^* and notice that F_S^* is a foliation for a sphere with $l - 1$ marked points. Notice that F_S^* has five fewer saddles than F_S. If we can show that F_S^* is admissible, then, by the induction hypothesis, F_S has at most $5(l-1) - 8 + 5 = 5k - 8$ saddles and we have proven the theorem.

Beginning with an embedding of a sphere with l marked points in S^3 realizing F_S, use the isotopy from Lemma 3.2 to eliminate the outermost saddles in F_1^σ and F_2^σ and iterate this process until σ is outermost and can be eliminated similarly. Since the isotopy in Lemma 3.2 preserves the singular foliation induced by the height function outside a neighborhood of $F_1^\sigma \cup F_2^\sigma$, then this process produces an embedding of a sphere with $l-1$ marked points into S^3 with induced singular foliation F_S^*.

If F_S^* fails Definition 3.8(1), then, by inclusion of the complement of M in F_S^* into F_S, then F_S also fails criteria (1), a contradiction.

Let D_τ be an outermost disk of F_S^*. If D_τ is again an outermost disk in F_S via inclusion, then D_τ contains one marked point by the admissibility of F_S. If D_τ is not an outermost disk for F_S, then D_τ contains M and, thus, contains a marked point.

Hence, F_S^* is admissible. □

4 Tangle products

In this section we define tangle product and use the combinatorial result of the previous section to show that, under suitable hypothesis, the bridge number of a tangle product is superadditive up to constant error.

Definition 4.1 A graph G is an n–star graph if G has n edges and $n+1$ vertices such that n of the vertices are valence one and one of the vertices is valence n. Denote by ∂G the set of valence one vertices.

Definition 4.2 Let K_1 and K_2 be links embedded in distinct copies of S^3, S_1^3 and S_2^3. Let G_1 and G_2 be n–star graphs embedded in S_1^3 and S_2^3 respectively such that $G_i \cap K_i = \partial G_i$. Let $\mu(G_i)$ be a small, closed, regular neighborhood of G_i in S_i^3 such that $(\mu(G_i), K_i \cap \mu(G_i))$ is an untangle tangle. Let $B_i = S_i^3 - \mathrm{int}(\mu(G_i))$. A link in S^3 obtained by gluing ∂B_1 to ∂B_2 via a homeomorphism such that points in $\partial(B_1) \cap K_1$ are mapped to points in $\partial(B_2) \cap K_2$ is called an *n–strand tangle product* of K_1 and K_2 and is denoted by $K_1 *_S K_2$. The image of ∂B_1 and ∂B_2 under this identification is called the product sphere and is denoted S.

Proof of Theorem 1.1 Let Σ be a minimal bridge sphere for $K_1 *_S K_2$ of distance at least three. We can assume that Σ is h–level and that S is taut with respect to Σ. By Theorem 3.9, F_S is admissible. By Lemma 3.10, F_S contains at most $10n-8$ saddles, since $k = 2n$.

If the set of saddles F_S is nonempty, then it contains at least two outermost disks. Since $K \cap S$ contains $2n$ points and F_S contains at least two outermost disks, there exists an outermost saddle σ such that D_σ meets $K_1 *_S K_2$ in at most n points. We can eliminate σ via the ambient isotopy that horizontally shrinks and vertically lowers B_σ as in the proof of Lemma 3.2. This isotopy produces at most n new maxima for $h_{K_1 *_S K_2}$; see Figure 10. By repeating this process, we can eliminate all saddles of F_S at the cost of creating at most n new maxima per saddle. Thus, we can assume we have an embedding of $K_1 *_S K_2$ with at most $\beta(K_1 *_S K_2) + n(10n-8)$ maxima such that F_S contains no saddles. Denote this embedding of $K_1 *_S K_2$ by K^*.

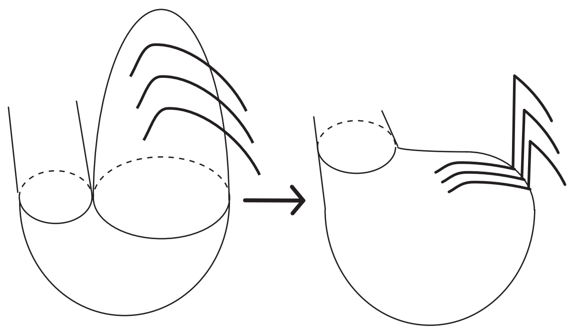

Figure 10

If F_S contains no saddles, there is a level-preserving isotopy of S^3 taking S to a standard round 2–sphere. Such an isotopy preserves the number of maxima of h_{K^*}. Recall that S decomposes S^3 into two 3–balls B_1 and B_2. The link K_1 can be recovered from the tangle $(B_1, K^* \cap B_1)$ by gluing an untangle (B^3, R) to $(B_1, K^* \cap B_1)$ along their common $2n$–punctured sphere boundary; see Figure 11. The number of maxima of the resulting embedding of K_1 is at most n more than the number of maxima of $h|_{K^*}$ in B_1. By a similar argument, we can produce an embedding of K_2 with at most n more maxima than the number of maxima of $h|_{K^*}$ in B_2. Hence, $\beta(K_1) - n + \beta(K_2) - n \leq \beta(K_1 *_S K_2) + n(10n-8)$, or $\beta(K_1 *_S K_2) \geq \beta(K_1) + \beta(K_2) - n(10n-6)$.

This ends the proof. □

Remark 4.3 With more detailed analysis the constant $-n(10n-6)$ that appears in the statement of Theorem 1.1 can be improved. However, the author believes it can not be improved beyond a quadratic expression in n using the techniques presented in this paper.

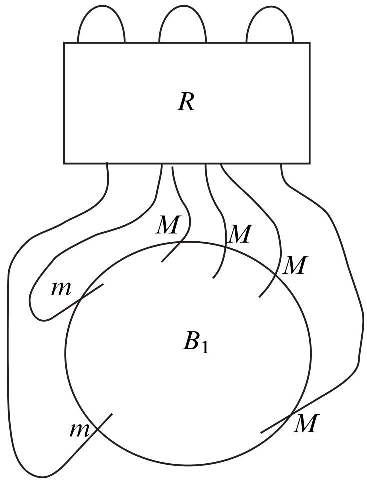

Figure 11

References

[1] **D Bachman**, **S Schleimer**, *Distance and bridge position*, Pacific J. Math. 219 (2005) 221–235 MR2175113

[2] **R C Blair**, *Bridge number and Conway products*, Algebr. Geom. Topol. 10 (2010) 789–823 MR2629764

[3] **M Scharlemann**, **M Tomova**, *Conway products and links with multiple bridge surfaces*, Michigan Math. J. 56 (2008) 113–144 MR2433660

[4] **H Schubert**, *Über eine numerische Knoteninvariante*, Math. Z. 61 (1954) 245–288 MR0072483

[5] **J Schultens**, *Additivity of bridge numbers of knots*, Math. Proc. Cambridge Philos. Soc. 135 (2003) 539–544 MR2018265

Department of Mathematics, David Rittenhouse Lab
209 South 33rd Street, Philadelphia PA 19104, USA

ryblair@math.upenn.edu

http://www.math.upenn.edu/~ryblair

Received: 21 January 2012 Revised: 6 June 2012