Algebraic & Geometric Topology 16 (2016) 1279–1308

A lower bound on tunnel number degeneration

TRENTON SCHIRMER

We prove a theorem that bounds the Heegaard genus from below under special kinds of toroidal amalgamations of 3–manifolds. As a consequence, we conclude that $t(K_1 \# K_2) \geq \max\{t(K_1), t(K_2)\}$ for any pair of knots $K_1, K_2 \subset S^3$, where $t(K)$ denotes the tunnel number of K.

57M25, 57N10

The tunnel number $t(K)$ of a knot $K \subset S^3$ can be defined by the equation $t(K)+1 = g(S^3 - \eta(K))$, where $g(\cdot)$ denotes Heegaard genus and $\eta(K)$ is an open regular neighborhood of K. In more intuitive terms, the tunnel number of a knot is the minimal number of "tunnels" that must be drilled through $S^3 - \eta(K)$ in order to make the resulting manifold a handlebody.

The behavior of $t(K)$ under the operation of connected sum has been studied extensively. It is not difficult to see that $t(K_1 \# K_2) \leq t(K_1) + t(K_2) + 1$, although it takes some work to find examples where equality is achieved in this bound; on this see Kobayashi and Rieck [5], Moriah and Rubinstein [8], and Morimoto, Sakuma and Yokota [10]. Morimoto [9] was the first to find pairs of knots in S^3 which satisfy $t(K_1 \# K_2) < t(K_1) + t(K_2)$. Soon after, Kobayashi [3] constructed an infinite family of examples for which the degeneration $t(K_1) + t(K_2) - t(K_1 \# K_2)$ can be arbitrarily large.

Perhaps most difficult is the task of finding lower bounds on $t(K_1 \# K_2)$. Norwood [13] employed a group-theoretic argument to show that $t(K_1 \# K_2) \geq 2$ for any pair of nontrivial knots in S^3, and Scharlemann and Schultens [15] subsequently used topological arguments to show $t(K_1 \# \cdots \# K_n) \geq n$. In the case that K_1 and K_2 are small, Morimoto and Schultens [11] proved that $t(K_1 \# K_2) \geq t(K_1) + t(K_2)$, and Kobayashi and Rieck [4] subsequently proved that this inequality holds even under the assumption that K_1 and K_2 are meridionally small. Scharlemann and Schultens [16] also proved that the lower bound $t(K_1 \# K_2)/(t(K_1) + t(K_2)) \geq \frac{2}{5}$ holds for any pair of nontrivial knots in S^3 (in fact, they derive a more general analogue involving iterated connected sums there).

In this paper, we prove that

$$t(K_1 \# K_2) \geq \max\{t(K_1), t(K_2)\}$$

Published: 1 July 2016 DOI: 10.2140/agt.2016.16.1279

for any pair of knots $K_1, K_2 \subset S^3$. This bound was previously unknown, although there are many examples that show it to be best possible, including those of Morimoto [9], Nogueira [12], and Li and Qui [7]. Moreover, this lower bound gives a negative answer to [7, Question 1.5].

A rough outline of the strategy of our proof is as follows. Suppose without loss of generality that $\max\{t(K_1), t(K_2)\} = t(K_2)$ and let $K_1 \# K_2$ be realized via the satellite construction with K_1 as the companion and K_2 as the pattern. This means that $K_1 \# K_2$ lies in $V = \overline{\eta(K_1)}$, and if $h\colon V \to S^3$ is the standard unknotted embedding of the solid torus V, then $h(K_1 \# K_2) = K_2$. If $\mathcal{G}$ is a thin generalized Heegaard surface of $S^3 - \eta(K_1 \# K_2)$, then $\mathcal{G}$ can be isotoped to intersect $\overline{S^3 - V}$ in a particularly nice way. Taking into account certain information contained in the intersection $\mathcal{G} \cap (\overline{S^3 - V})$, we can then construct a so-called *doppelgänger surface* $\mathcal{Q}$ inside of a solid torus $W = \overline{S^3 - h(V)}$ which, in certain important respects, imitates the placement of the surface $\mathcal{G} \cap (\overline{S^3 - V})$ in $(\overline{S^3 - V})$. As a result, $\mathcal{Q} \cup h(\mathcal{G} \cap V)$ forms a generalized Heegaard surface of $\overline{S^3 - h(V)} = S^3 - \eta(K_2)$ which amalgamates to a surface of lower genus than the amalgamation of $\mathcal{G}$. This yields the desired lower bound.

In Section 1, we introduce *generalized compression bodies*, which form the basic pieces of $\overline{S^3 - V} - \mathcal{G}$ and $W - \mathcal{Q}$, and we prove a series of essential cutting and pasting lemmas about them. Section 2 then describes and works out the basic topology of so-called *spoke graphs* and *spoke surfaces*, which form the building blocks of the doppelgänger surface $\mathcal{Q}$. Section 3 then constructs $\mathcal{Q}$ in detail and proves that it has the desired properties, culminating in the main technical result of the paper, Theorem 3.21. In Section 4, the bound

$$t(K_1 \# K_2) \geq \max\{t(K_1), t(K_2)\}$$

is proved (Theorem 4.1), and some topics related to it are briefly discussed.

Throughout this paper, $N(Y, X)$ denotes a *closed* regular neighborhood of Y in X, $E(Y, X) = \overline{X - N(Y, X)}$, and $\mathrm{Fr}(Y, X) = N(Y, X) \cap E(Y, X)$, or equivalently, $\mathrm{Fr}(Y, X) = \overline{\partial N(Y) - \partial X}$. We assume throughout that $N(Y, X)$ behaves well with respect to intersection, so that $N(Y_1, X) \cap N(Y_2, X) = N(Y_1 \cap Y_2, X)$. If $\mathcal{X}$ is a topological space, $|\mathcal{X}|$ denotes the number of components of $\mathcal{X}$. An embedding of manifolds $f\colon Y \to X$ is said to be *proper* if f is transverse to ∂X and $f(\partial Y) \subset \partial X$. A *proper isotopy* is a homotopy through proper embeddings (note that this does *not* imply that the boundary remains fixed). As an informal aid to the reader, topological spaces that are allowed to have multiple connected components will usually be denoted in calligraphic font, eg $\mathcal{A}$, $\mathcal{X}$, and $\mathcal{Y}$, whereas connected topological spaces will usually be denoted in standard font, eg A, X, and Y.

1 Generalized compression bodies

Definition 1.1 Let $\mathcal{F}$ be a compact orientable surface, and let $\mathcal{V} = (\mathcal{F} \times I) \cup (2\text{–}handles) \cup (3\text{–}handles)$, where the 2–handles are attached along essential, non-boundary parallel curves in $\mathcal{F} \times \{0\}$, and 3–handles are attached along all spherical components of $\mathcal{F} \times I \cup (2\text{–}handles)$ that are disjoint from $\mathcal{F} \times \{1\}$. Then $\mathcal{V}$ is called a *generalized compression body over* $\mathcal{F}$, or simply a *generalized compression body*. Let $\partial_+\mathcal{V} = \mathcal{F} \times \{1\}$, $\partial_v\mathcal{V} = (\partial\mathcal{F}) \times I$, and $\partial_-\mathcal{V} = \overline{\partial\mathcal{V} - (\partial_+\mathcal{V} \cup \partial_v\mathcal{V})}$. If $\mathcal{V}$ is connected and $\partial_v\mathcal{V} = \varnothing$, $\mathcal{V}$ is a *compression body*. If $\mathcal{V}$ is connected and $\partial_-\mathcal{V} = \varnothing$, $\mathcal{V}$ is a *handlebody*.

Observation 1.2 Suppose $\mathcal{V}$ is a generalized compression body, A_1 and A_2 are disjoint components of $\partial_v\mathcal{V}$, and $h\colon A_1 \to A_2$ is an orientation reversing homeomorphism which preserves $\partial_+\mathcal{V}$. Then $\mathcal{V}/h$ is a generalized compression body over $(\partial_+\mathcal{V})/h$.

Observation 1.3 If $\mathcal{V}$ is a generalized compression body and $\mathcal{W}$ is obtained by compressing $\mathcal{V}$ along a properly embedded disk D such that $\partial D \subset \partial_+\mathcal{V}$, then $\mathcal{W}$ is again a generalized compression body. Going the other way, if $\mathcal{W}$ is obtained from $\mathcal{V}$ by attaching an oriented 1–handle along $\partial_+\mathcal{V}$, then $\mathcal{W}$ is a generalized compression body.

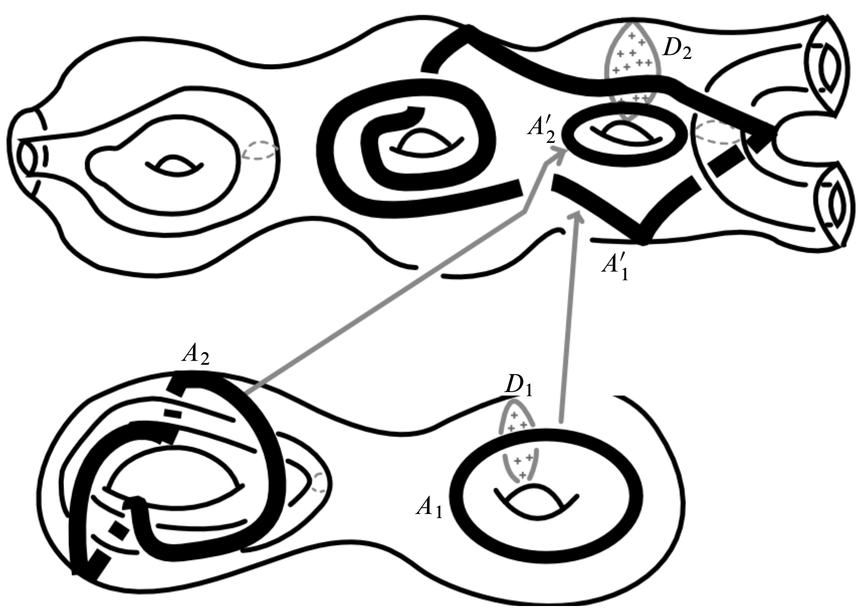

Figure 1: Primitive disk set for a paired union of annuli.

Definition 1.4 Let $\mathcal{V}$ be a generalized compression body and let $\mathcal{A} = A_1 \cup \cdots \cup A_n$ and $\mathcal{A}' = A'_1 \cup \cdots \cup A'_n$ be disjoint unions of annuli embedded in $\partial_+\mathcal{V}$ satisfying $\mathcal{A} \cap \mathcal{A}' = \varnothing$. Let $\mathcal{D} = D_1 \cup \cdots \cup D_n$ be a disjoint union of compressing disks for $\mathcal{V}$ such that $\partial\mathcal{D} \subset \partial_+\mathcal{V}$. If $D_i \cap (A_i \cup A'_i)$ consists of a single spanning arc in one of A_i or A'_i for all $1 \le i \le n$, and $D_i \cap (A_j \cup A'_j) = \varnothing$ if $j > i$, then $\mathcal{D}$ is said to be a *primitive disk set* for $\mathcal{A} \cup \mathcal{A}'$, and the component of $A_i \cup A'_i$ which meets D_i is said to be *dual* to D_i. The above orderings of the components of $\mathcal{A} \cup \mathcal{A}'$ and $\mathcal{D}$ will be called the *primitive ordering* associated with $\mathcal{D}$. See Figure 1.

Remark 1.5 The choice of primitive ordering is essential to Definition 1.4, and some fixed choice is always assumed to be present when we are dealing with a primitive disk set $\mathcal{D}$ for a paired union $\mathcal{A} \cup \mathcal{A}'$ of annuli. For the most part, however, the primitive ordering will only be specified explicitly when necessary.

Proposition 1.6 *Let $\mathcal{V}$ be a generalized compression body and let $\mathcal{A} = A_1 \cup \cdots \cup A_n$ and $\mathcal{A}' = A'_1 \cup \cdots \cup A'_n$ be disjoint unions of annuli embedded in $\partial_+\mathcal{V}$ satisfying $\mathcal{A} \cap \mathcal{A}' = \varnothing$. Let $h\colon \mathcal{A} \to \mathcal{A}'$ be an orientation reversing homeomorphism such that $h(A_i) = A'_i$ for all $1 \le i \le n$, and suppose that $\mathcal{A} \cup \mathcal{A}'$ admits a primitive disk set. Then $\mathcal{V}/h$ is a generalized compression body over $(\partial_+\mathcal{V})/h$.*

Proof We proceed by induction on n. In the base case $n = 0$, there is nothing to prove. If $n > 0$, suppose without loss of generality that A_n is dual to D_n (the argument is the same if A'_n is dual D_n). By Observation 1.3, $E(D_n, \mathcal{V})$ is again a generalized compression body, and the result of reattaching $N(D_n, \mathcal{V})$ to $E(D_n, \mathcal{V})$ via the map $h|_{N(D_n \cap A_n, A_n)}$ again results in a generalized compression body $\mathcal{V}'$ since this amounts to trivially attaching a ball to $\partial_+ E(D_n, \mathcal{V})$ along a disk on its boundary. But observe that $\mathcal{V}/h|_{A_n}$ is obtained from $\mathcal{V}'$ by identifying a pair of disks in $\partial_+\mathcal{V}'$, which is the same as a 1–handle attachment, so Observation 1.3 tells us that $\mathcal{V}/h|_{A_n}$ is a generalized compression body. Since $D_1 \cup \cdots \cup D_{n-1}$ was disjoint from $A_n \cup A'_n$ and D_n, it remains a primitive disk set for $A_1 \cup \cdots \cup A_{n-1}$ and $A'_1 \cup \cdots \cup A'_{n-1}$ in $\mathcal{V}/h|_{A_n}$, and the desired conclusion follows by induction. □

Proposition 1.7 *Suppose $\mathcal{V}$ is a generalized compression body, let $\mathcal{A} = A_1 \cup \cdots \cup A_n$ be a disjoint union of annuli embedded in $\partial_+\mathcal{V}$, and let $\mathcal{D} = D_1 \cup \cdots \cup D_n$ be a disjoint union of disks properly embedded in $\mathcal{V}$ such that $\partial\mathcal{D} \subset \partial_+\mathcal{V}$. If $D_i \cap A_i$ consists of a single spanning arc in A_i for all $1 \le i \le n$, and $D_i \cap A_j = \varnothing$ whenever $i < j$, then manifold $\mathcal{W}$ obtained by attaching 2–handles along $\mathcal{A}$ is again a generalized compression body.*

Proof The proposition is well known in the case that $\mathcal{V}$ is a compression body. The proof in the general case here is essentially the same as that of Proposition 1.6. □

Definition 1.8 An annulus A properly embedded in a generalized compression body V is said to be *spanning* if one component of ∂A lies on $\partial_+ V$, and the other lies on $\partial_- V$. A is said to be *horizontal* if $\partial A \subset \partial_+ V$.

Proposition 1.9 *If F is a compact surface and $\mathcal{A}$ is a disjoint union of incompressible spanning annuli embedded in $F \times I$, then $F \times I$ can be reparameterized so that $\mathcal{A} = \mathcal{C} \times I$ for some disjoint union of essential simple closed curves $\mathcal{C} \subset F$.*

Proof This is a well known fact which often appears in the literature, so the following proof is merely a sketch. If $\mathcal{C} = \mathcal{A} \cap (F \times \{0\})$, then $\mathcal{C} \times I$ is another union of spanning annuli $\mathcal{A}'$ such that $\mathcal{A}' \cap (F \times \{1\})$ is isotopic to $\mathcal{A} \cap (F \times \{1\})$ in $F \times \{1\}$ (this follows from the π_1–injectivity of $\mathcal{A}$ and $\mathcal{A}'$). This allows $\mathcal{A}'$ to be properly isotoped so that $\partial\mathcal{A}'$ and $\partial\mathcal{A}$ are parallel and disjoint in $\mathcal{F} \times \{0, 1\}$. Since $\mathcal{A}$ and $\mathcal{A}'$ are both incompressible, and $F \times I$ is irreducible, any simple closed curves in $\mathcal{A}' \cap \mathcal{A}$ that are trivial in either of $\mathcal{A}'$ or $\mathcal{A}$ can be eliminated via further isotopy of $\mathcal{A}'$ using standard inner-most disk arguments. Again, since each component of $\mathcal{A} \cup \mathcal{A}'$ is π_1–injective, any remaining components of $\mathcal{A} \cap \mathcal{A}'$ must come from pairs of annuli A, A' with isotopic boundaries on $\mathcal{F} \times \{0, 1\}$. Thus these components of $\mathcal{A} \cap \mathcal{A}'$ can also be removed, uppermost ones first, using the fact that any incompressible horizontal annulus in $F \times I$ with parallel boundary components will cobound a solid torus with an annulus on $F \times \{1\}$. Once $\mathcal{A}'$ has been made disjoint from $\mathcal{A}$, the components of $\mathcal{A} \cup \mathcal{A}'$ will cobound solid tori with annuli in $F \times \{0, 1\}$, so that $\mathcal{A}'$ can finally be isotoped onto $\mathcal{A}$. Extending this proper isotopy of $\mathcal{A}'$ to an ambient isotopy of $\mathcal{F} \times I$ yields the desired reparameterization. □

Proposition 1.10 *Let $\mathcal{V}$ be a compression body and let $\mathcal{A} = \mathcal{A}_s \cup \mathcal{A}_h$ be a disjoint union of incompressible annuli properly embedded in $\mathcal{V}$, so that every component of $\mathcal{A}_s$ is spanning and every component of $\mathcal{A}_h$ is horizontal. Then $E(\mathcal{A}, \mathcal{V})$ is a generalized compression body $\mathcal{V}'$ such that $\partial_v(\mathcal{V}') = \mathrm{Fr}(\mathcal{A}_s, \mathcal{V})$ and $\mathrm{Fr}(\mathcal{A}_h, \mathcal{V}) \subset \partial_+\mathcal{V}'$. Moreover, for an appropriate ordering of the components of $\mathcal{A}_h = A_1 \cup \cdots \cup A_n$, if we set $\mathrm{Fr}(A_i, \mathcal{V}) = A_i' \cup A_i''$, $\mathcal{A}' = A_1' \cup \cdots \cup A_n'$, and $\mathcal{A}'' = A_1'' \cup \cdots \cup A_n''$, then the collection $\mathcal{A}' \cup \mathcal{A}''$ admits a primitive disk set in $\mathcal{V}'$.*

Proof From Schultens [18, Lemma 2], we have that $\mathcal{W} = E(\mathcal{A}_h, \mathcal{V})$ is a union of compression bodies. The annuli $\mathcal{A}_s$ remain spanning in $\mathcal{W}$, and thus are disjoint from some union of disks $\mathcal{E}$ properly embedded in $\mathcal{W}$ such that $E(\mathcal{E}, \mathcal{W}) \cong \partial_-\mathcal{W} \times I$;

see, eg Saito, Scharlemann and Schultens [14, Lemma 3.1.5]. By Proposition 1.9, we may assume that $E(\mathcal{E},\mathcal{W})$ has been parameterized so that $\mathcal{A}_s$ has the form $\mathcal{C}\times I$ in $\partial_-\mathcal{W}\times I$, where $\mathcal{C}\subset\partial_-\mathcal{W}$ is a disjoint union of simple closed curves. Thus $\mathcal{V}'=E(\mathcal{A}_s,\mathcal{W})=E(\mathcal{A},\mathcal{V})$ is a generalized compression body, since it is just $\partial_- W$ cut along $\mathcal{C}\times I\cup\{2\text{–}handles\}\cup\{3\text{–}handles\}$, and it satisfies $\partial_v\mathcal{V}'=\mathrm{Fr}(\mathcal{A}_s,\mathcal{V})$ as claimed.

We prove the second part by induction on n. There is nothing to prove if $n=0$. If $n>0$, then from Bonahon and Otal [2, Lemma 9], we have that $\mathcal{A}_h$ is boundary-compressible in $\mathcal{V}$ via some disk D_1. We may choose D_1 so that it is disjoint from $\mathcal{A}_s$, and we assign the label A_1 to the component of $\mathcal{A}_h$ which has been boundary-compressed by D_1. By the first part of this lemma, $E(\mathcal{A}_s\cup A_1,\mathcal{V})$ is a generalized compression body, and so by induction, $\mathrm{Fr}(\mathcal{A}_h-A_1,\mathcal{V})$ admits a primitive disk set $\mathcal{D}'$ in $\mathcal{V}'=E(\mathcal{A},\mathcal{V})$ with respect to an appropriate choice of numbering for $\mathcal{A}_h-A_1$, with A_2 as the lowest indexed annulus. We may also choose (appropriately indexed) $\mathcal{D}'$ so that $\mathcal{D}'\cap D_1=\varnothing$, and since $D_1\cap(\mathcal{A}_h-A_1)=\varnothing$, it follows that $\mathcal{D}=\mathcal{D}'\cup(D_1\cap\mathcal{V}')$ is a primitive disk set for $\mathrm{Fr}(\mathcal{A}_h,\mathcal{V})$ in $\mathcal{V}'$ as required. □

Definition 1.11 A graph $\mathcal{X}$ embedded in a 3–manifold M is said to be *properly embedded* if $\mathcal{X}$ is transverse to ∂M and $\mathcal{X}\cap\partial M$ is a union of elements from the set of univalent vertices of $\mathcal{X}$. Also, $\mathcal{X}$ is said to be *unknotted* if it can be isotoped into ∂M via an isotopy $\Phi\colon\mathcal{X}\times I\to M$ such that the function $\Phi|_{\{t\}\times X}$ is a proper embedding for all $0\le t<1$.

Observation 1.12 If X is an unknotted tree properly embedded in the 3–ball B, and F is the surface which results from removing an open collar of $\partial(\overline{\partial B-N(X)})$ from $\overline{\partial B-N(X)}$, then there is homeomorphism $h\colon F\times I\cong E(X,B)$ such that $h(F\times\{1\})=F$ and $h(F\times\{0\})=\mathrm{Fr}(X,B)$.

Definition 1.13 Let X be a graph embedded in a handlebody V such that $E(X,V)\cong\partial V\times I$. Then X is called a *spine* of V.

Proposition 1.14 *Suppose X is a graph embedded in the interior of a handlebody V, and that there is a disjoint union of compressing disks $\mathcal{D}$ properly embedded in V so that the following are true:*

(1) *$E(\mathcal{D},V)$ is a union of balls.*
(2) *$X\cap B$ is a properly embedded, unknotted tree for each component B of $E(\mathcal{D},V)$.*
(3) *For each component D of $\mathcal{D}$, $D\cap X$ is a single point on an edge of X.*

Then X is a spine of V.

Proof For each component B of $E(\mathcal{D}, V)$, let $\mathcal{E} \subset \partial B$ be the set of "scars" left behind by cutting along $\mathcal{D}$, ie $\mathcal{E} = N(\mathcal{D}, V) \cap B$, and let $X' = X \cap B$. Then hypotheses 2 and 3 give us a homeomorphism $h\colon F \times I \to E(X', B)$ as per Observation 1.12, where we may take $F = \overline{\partial B - \mathcal{E}}$. These homeomorphisms can then be pasted together to form a homeomorphism between $E(X, V)$ and $\partial V \times I$. □

Proposition 1.15 *Suppose V is a handlebody with spine X, and that Y is a subgraph of X without simply connected components. Then $E(Y, V)$ is a compression body.*

Proof $E(Y, V)$ is obtained from $E(X, V) \cong \partial V \times I$ by attaching 2–handles and 3–handles along $\partial V \times \{0\}$. □

Definition 1.16 A *Heegaard splitting* (V, W, G) of a compact, connected, orientable 3–manifold M is a decomposition $M = V \cup W$, where each of V and W is a compression body, and $G = \partial_+ V = \partial_+ W = V \cap W$. Also, G is called a *Heegaard surface*. The *Heegaard genus* $g(M)$ of a manifold is the minimal genus of a Heegaard surface for M.

Definition 1.17 A *generalized Heegaard splitting* $((V_1, W_1, G_1), \dots, (V_n, W_n, G_n))$ of a compact, orientable, connected 3–manifold M is a decomposition $M = M_1 \cup \dots \cup M_n$ such that the following conditions hold:

- (V_i, W_i, G_i) forms a Heegaard splitting for the submanifold M_i.
- For all $i < j$, M_i meets M_j only along components of $\partial_- W_i$ that are identified with components of $\partial_- V_j$.

We include the case $n = 1$ corresponding to standard Heegaard splittings. Let $\mathcal{G}_+ = G_1 \cup \dots \cup G_n$, and let $\mathcal{G}_-$ denote the union of surfaces $M_i \cap M_j$ for $i < j$. Then $\mathcal{G} = \mathcal{G}_+ \cup \mathcal{G}_-$ is called a *generalized Heegaard surface* with *thick surfaces* $\mathcal{G}_+$ and *thin surfaces* $\mathcal{G}_-$.

Definition 1.18 A Heegaard splitting (V, W, G) is said to be:

- *Stabilized* if there exist compressing disks $D \subset V$ and $D' \subset W$ such that $D \cap D'$ is a single point.
- *Reducible* if there exist compressing disks $D \subset V$ and $D' \subset W$ such that $\partial D = \partial D'$.
- *Weakly reducible* if it is not stabilized or reducible, and there are compressing disks $D \subset V$ and $D' \subset W$ such that $D \cap D' = \varnothing$.
- *Strongly irreducible* if it is not stabilized and, for all compressing disks $D \subset V$ and $D' \subset W$, $D \cap D' \neq \varnothing$.

Remark 1.19 There is a process of *untelescoping* a weakly reducible Heegaard splitting (V, W, G) whereby it is changed into a generalized Heegaard splitting of the form $((V_1, W_1, G_1), \dots, (V_n, W_n, G_n))$ satisfying $g(G) = \Sigma g(G_i) - \Sigma g(F_i)$. Conversely, given a generalized Heegaard splitting $((V_1, W_1, G_1), \dots, (V_n, W_n, G_n))$ for M, one can always use the process of *amalgamation* to change it into a standard Heegaard splitting (V, W, G) of M satisfying the same equation. The interested reader is referred to Saito, Scharlemann and Schultens [14] for the details of these processes and a proof of the following lemma.

Proposition 1.20 (Scharlemann and Thompson [17]) *Suppose that we have a weakly reducible Heegaard splitting (V, W, G) of a compact, orientable, connected, irreducible 3–manifold M. Then (V, W, G) can be untelescoped to a generalized Heegaard splitting $((V_1, W_1, G_1), \dots, (V_n, W_n, G_n))$ such that (V_i, W_i, G_i) is a strongly irreducible splitting of M_i for each $1 \le i \le n$, and the thin surfaces F_i are incompressible in M for each $1 \le i < n$. In this case, the generalized splitting is said to be* fully untelescoped. *A standard Heegaard splitting that is strongly irreducible will also be considered* fully untelescoped.

Proposition 1.21 (Scharlemann and Schultens [15]) *If $\mathcal{G}$ is the union of the thick and thin surfaces of a fully untelescoped Heegaard splitting of M, and T is an incompressible surface properly embedded in M, then $\mathcal{G}$ can be isotoped so that it meets T only in simple closed curves that are nontrivial in both $\mathcal{G}$ and T.*

2 Spoke surfaces in the solid torus

Convention 2.1 Throughout this section, we set $W = S^1 \times D^2$ and parameterize it using polar coordinates (ϕ, r, θ) with $0 \le \phi, \theta \le 2\pi$ and $0 \le r \le 1$.

Definition 2.2 Let $X \subset W$ be an embedded graph with one central vertex x_0 at $(\phi_0, 0, 0)$, a finite number of outer vertices $\{x_1, \dots, x_n\} \subset \{\phi_0\} \times \partial D^2$, one radial edge connecting each outer vertex x_i to the central vertex x_0, and one longitudinal edge $l_i = S^1 \times \{x_i\}$ connecting x_i to itself for $1 \le i \le n$. Then X is said to be a *connected spoke graph* in W. A finite, disjoint union of connected spoke graphs $\mathcal{X}$ in W is simply called a *spoke graph*, $\mathrm{Fr}(\mathcal{X}, W)$ is called a *spoke surface*, and $E(\mathcal{X}, W)$ is a *spoke chamber*.

Definition 2.3 Suppose X is a connected spoke graph with central vertex at $(\phi_0, 0, 0)$. Let $\mathcal{D}$ be a disjoint union of disks embedded in $\mathring{W}$ such that for each component D of $\mathcal{D}$,

- $D \cap X$ is a connected subarc of a radial edge of X, and
- $D \cap (\{\phi_0\} \times D^2) = D \cap X$.

Then $X \cup \partial\mathcal{D}$ is called a *stabilized* spoke graph with stabilizing disk set $\mathcal{D}$, $\mathrm{Fr}(X \cup \partial\mathcal{D})$ is a stabilized spoke surface, and $E(X \cup \partial D, W)$ is a stabilized spoke chamber. Moreover, $\overline{\partial\mathcal{D} - X}$ is called the set of *stabilizing arcs* of X.

Definition 2.4 Let X be a connected stabilized spoke graph with stabilizing disk set $\mathcal{D}$ and central vertex at $(\phi_0, 0, 0)$. Let $\mathcal{A} = E(X \cap \partial W, \partial W)$, and let $\mathcal{E} = (\{\phi_0\} \times D^2) \cap V$, where $V = E(X, W)$. Then the *standard disk set* of X is the union of disks $\mathcal{D}_X = \mathcal{E} \cup (\mathcal{D} \cap V) \cup \mathrm{Fr}(\mathcal{E} \cup \mathcal{A}, V)$.

Observation 2.5 The standard disk set of X cuts $E(X, W)$ into a union of balls. Thus $E(X, W)$ is a handlebody.

Definition 2.6 Let X be a connected, stabilized spoke graph with stabilizing disk set $\mathcal{D}$, and whose central vertex has ϕ–coordinate ϕ_1. Let X' be another connected, stabilized spoke graph disjoint from X with stabilizing disk set $\mathcal{D}'$, whose central vertex has ϕ–coordinate $\phi_2 \neq \phi_1$, and suppose the following properties are also satisfied:

- The set of longitudinal edges of X' is precisely the set of core curves of the annuli $E(X \cap \partial W, \partial W)$.
- $\mathcal{D} \cap (\{\phi_2\} \times D^2) = \varnothing = \mathcal{D}' \cap (\{\phi_1\} \times D^2)$.
- Every component D of $\mathcal{D}$ meets precisely one component D' of $\mathcal{D}'$ in a single arc that has one endpoint on $\partial D - X$ and the other on $\partial D' - X'$, and conversely, each component of $\mathcal{D}'$ meets precisely one component of $\mathcal{D}$ in this way.
- Let h be the projection $S^1 \times D^2 \to S^1$. Then $h|_{\mathring{\mathcal{D}}}$ is a circle-valued Morse function without singularities, and for every stabilizing arc α of X, $h|_\alpha$ is Morse with only one critical point occurring at $\alpha \cap \mathcal{D}'$. The corresponding condition holds for the stabilizing disks and arcs of X'.

Then X and X' are said to be *dual* to one another. See Figure 2.

Remark 2.7 One way to obtain a dual graph X' is to rotate a copy of X slightly in the ϕ and θ directions, and then isotope the stabilizing arcs of X' slightly so that they clasp those of X in one to one fashion. However, we use the Morse condition on the stabilizing arcs because it allows us complete flexibility in the choice of radial edges of X' at which to base its stabilizing arcs, while still avoiding knottedness.

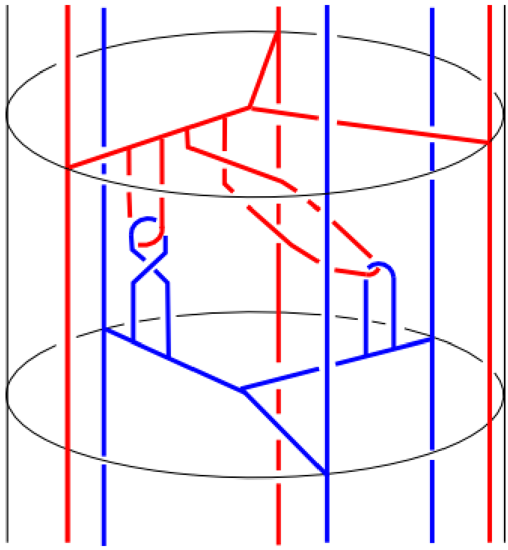

Figure 2: Dual stabilized spoke graphs.

Lemma 2.8 *Let X and X' be a dual pair of connected spoke surfaces in W, and let B be the component of $E(X,W) - \mathring{N}(\mathcal{D}_X)$ that contains the central vertex of X', where $\mathcal{D}_X$ is the standard disk set of X from Definition 2.4. Then B is a ball, and $X' \cap B$ is an unknotted tree properly embedded in B.*

Proof It is clear that B is a ball; we show that $X' \cap B$ is unknotted. We retain the notation of Definition 2.6 throughout. If $\mathcal{C}$ is the set of stabilizing arcs of X', then the Morse condition on $\mathcal{C}$ ensures that $h|_{\mathcal{C}\cap B}$ is Morse without singularities, and since $h|_{\mathring{\mathcal{D}}}$ is also Morse without singularities, we can slide the endpoints of $\mathcal{C} \cap B$ off of $N(\mathcal{D}) \cap \partial B$ without introducing any further singularities. The components of $N(\mathcal{D}, W) \cap \partial B$ can then be "pushed in" so that $\partial B - \partial W$ is level with respect to h, and the arcs of $\mathcal{C} \cap B$ can then be properly isotoped horizontally with respect to h until they are vertical. After these isotopies, it is clear that $X' \cap B$ is unknotted. □

Lemma 2.9 *If X and X' are connected, dual, stabilized spoke graphs, then $E(X \cup X', W)$ is homeomorphic to $\mathrm{Fr}(X,W) \times I$ via a map sending $\mathrm{Fr}(X',W)$ to $\mathrm{Fr}(X,W) \times \{0\}$ and $\mathrm{Fr}(X,W)$ to $\mathrm{Fr}(X,W) \times \{1\}$.*

Proof Take the double of W to obtain $S^1 \times S^2$, let X_d be the double of X, and let X'_d be the double of X'. Then $E(X_d, S^1 \times S^2)$ is a handlebody since the double $\mathcal{E}$ of the standard disk set $\mathcal{D}_X$ cuts $E(X_d, S^1 \times S^2)$ into balls. Moreover, $E(X_d, S^1 \times S^2)$, $\mathcal{E}$, and X'_d satisfy the hypotheses of Proposition 1.14 (Lemma 2.8 handles the only subtle aspect of this). Thus X_d' is a spine of $E(X_d, S^1 \times S^2)$, and we obtain a parameterization $E(X_d \cup X_d', S^1 \times S^2) \cong \partial N(X_d) \times I$. By Proposition 1.9, the spanning annuli $(\partial W) \cap E(X_d \cup X_d', S^1 \times S^2)$ can be assumed to be vertical with respect to this parameterization, and the result follows. □

Remark 2.10 Besides being a steppingstone to Proposition 2.18 below, the significance of Lemma 2.9 is that it allows us to isotope a connected, stabilized spoke surface $S \subset W$ back and forth between small neighborhoods of dual spoke graphs lying on opposite sides of S in W. This kind of isotopy will play an essential role in the final doppelgänger construction.

Definition 2.11 Let $\mathcal{X}$ be a disjoint union of stabilized spoke graphs embedded in W. Then two components X_1 and X_2 of $\mathcal{X}$ are said to be:

- θ*–adjacent* if the closure A of a component of $\overline{\partial W - N(\mathcal{X})}$ meets $N(X_1, W)$ in one boundary component and $N(X_2, W)$ in the other (in this case, A is said to be a *spanning annulus of* θ*–adjacency*).
- ϕ*–adjacent* if there is a subarc $\beta \subset S^1$ such that the endpoints of $\beta \times \{0\} \subset W$ are the central vertices of X_1 and X_2, and $\beta \times \{0\}$ meets $\mathcal{X}$ only in these endpoints.

In all cases, an arc of the form $\beta \times \{0\}$ that connects the central vertices of X_1 and X_2 shall be called a *spanning arc* of X_1 and X_2 (there are only two such arcs).

Definition 2.12 Let $\mathcal{X}$ be a disjoint union of stabilized spoke graphs whose components are ordered $X_0, \dots X_n$ so that X_i is ϕ–adjacent to X_{i+1} for all $0 \le i < n$, and let α be the spanning arc of X_0 and X_n that meets all components of $\mathcal{X}$. Then α is said to be the *binding arc* of $\mathcal{X}$ with respect to the given ordering of its components.

Suppose further that, for all $1 < i \le n$, X_i is θ–adjacent to X_{j_i} for some $j_i < i$, and let A_i be a spanning annulus of θ–adjacency connecting X_i to X_{j_i}. Then $\mathcal{A} = A_1 \cup \dots \cup A_n$ is said to form an *adjacency chain* for $\mathcal{X}$ with respect to the given ordering of its components.

Definition 2.13 Let $\mathcal{X}$ be a disjoint union of stabilized spoke graphs (possibly with detached longitudes in the sense of Definition 2.17 below), and let X be a connected stabilized spoke graph (possibly with detached longitudes) obtained from $\mathcal{X}$ by rotating each component of $\mathcal{X}$ in the ϕ–direction so that all of their central vertices coincide at a single vertex x_0. Then $\mathcal{X}$ is said to be a *decomposition* of X. If $d(x_0, v) < \epsilon$ for every central vertex v occurring in a component of $\mathcal{X}$, where $d\colon W \times W \to \mathbb{R}$ is the flat metric, then $\mathcal{X}$ is said to be an ϵ*–small* decomposition of X.

Definition 2.14 Suppose $\mathcal{X}$ is an ϵ–small decomposition of X with $\epsilon < \pi/2$. Then $\mathcal{X}$ is said to be a *good* decomposition of X if the components $X_0, \dots, X_n$ of $\mathcal{X}$ can be ordered so that the following conditions are satisfied:

- X_i is ϕ–adjacent to X_{i+1} for all $1 \le i < n$.
- The binding arc of $\mathcal{X}$ with respect to this ordering has length less than 2ϵ.
- $\mathcal{X}$ admits an adjacency chain with respect to this ordering.

Observation 2.15 Let $\mathcal{X}$ be an ϵ–small decomposition of X, and let α be the binding arc of $\mathcal{X}$. Then $\mathrm{Fr}(\mathcal{X} \cup \alpha, W)$ is isotopic to $\mathrm{Fr}(X, W)$, and $E(\mathcal{X}, W)$ is obtained from $E(X, W) \cong E(\mathcal{X} \cup \alpha, W)$ via 2–handle attachments to $\mathrm{Fr}(\mathcal{X} \cup \alpha, W)$ along meridians of α.

Definition 2.16 Let X_1 and X_2 be a pair of components in $\mathcal{X}$ which are θ–adjacent, with spanning annulus of θ–adjacency A. Let D be disk embedded in W such that $\partial D = \alpha \cup e_1 \cup \beta \cup e_2$, where α is a spanning arc for X_1 and X_2, e_i is the radial edge of X_i nearest to A for $i = 1, 2$, $\beta \subset \partial W$ is an arc joining e_1 to e_2 that spans A, and $\mathring{D} \cap \mathcal{X} = \varnothing$. Then D is said to be a *spanning disk* of X_1 and X_2 that *cuts* A and is based at α. See Figure 3.

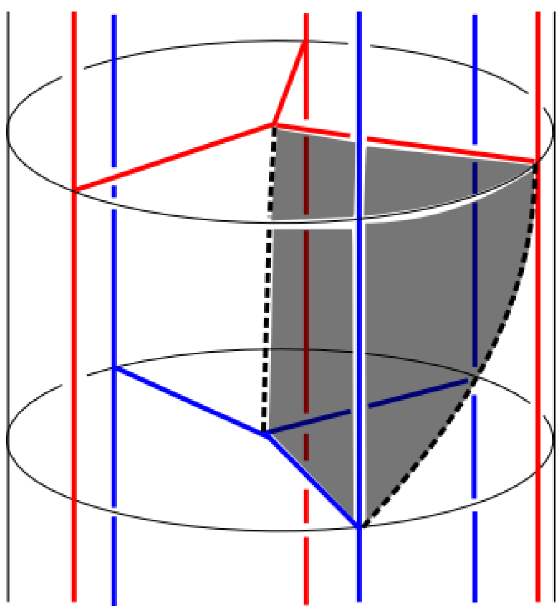

Figure 3: A spanning disk.

Definition 2.17 Let $h_\epsilon\colon W \to W$ be the dilation $(\phi, r, \theta) \mapsto (\phi, (1-\epsilon)r, \theta)$. Let X be a stabilized spoke graph, let $\mathcal{L}$ be a union of longitudinal edges of X, let $\mathcal{E}$ be the union of those radial edges of X which meet $\mathcal{L}$, and let $\mathcal{S}$ be the union of stabilizing arcs attached to $\mathcal{E}$. Then the graph X' obtained by removing $\mathcal{E} \cup \mathcal{L} \cup \mathcal{S}$ from X and attaching $h(\mathcal{E} \cup \mathcal{L} \cup \mathcal{S})$ is said to be a *spoke graph obtained by ϵ–small detachments of the longitudes $h(\mathcal{L})$*.

In the following proposition, it is important to remember the conventions made in the final paragraph before Section 1 regarding regular neighborhoods. In particular, since $\mathcal{X}$ and $\mathcal{Y}$ are disjoint subsets of W, we assume their regular neighborhoods are chosen small enough so that $N(\mathcal{Y}, W) \cap N(\mathcal{X}, W) = \varnothing$.

Proposition 2.18 *Suppose that X and X' are dual stabilized spoke graphs embedded in W, that $\epsilon < \pi/2$, and that $d(X, X') > \epsilon$ (as usual, d is the flat metric). Suppose $\mathcal{X}$ is an $\epsilon/8$–small, good decomposition of X, and that $\mathcal{A} = A_1 \cup \cdots \cup A_k$ is an adjacency chain of annuli for $\mathcal{X} = X_0 \cup \cdots \cup X_k$. Let Y be a spoke subgraph of X' which does not meet $\mathcal{A}$, and suppose $\mathcal{Y}$ is obtained from $\epsilon/8$–small longitudinal detachments of Y, followed by an $\epsilon/8$–small decomposition (which need not be "good"). Let $\mathcal{A}'$ be the subset of $E(\mathcal{X} \cup \mathcal{Y}, \partial W)$ consisting of those annuli which meet* $\mathrm{Fr}(\mathcal{X}, W)$ *on one boundary component, and* $\mathrm{Fr}(\mathcal{Y}, W)$ *on the other. Then $E(\mathcal{X} \cup \mathcal{Y}, W)$ is a generalized compression body V such that $\partial_v V = \mathcal{A}'$ and $\partial_- V = \mathrm{Fr}(\mathcal{Y}, W)$.*

Proof We have chosen the various stabilizations and detachments small enough to ensure that they can all be carried simultaneously without creating new intersections. As in the proof of Lemma 2.9, we double W to obtain $S^1 \times S^2$, let X_d denote the double of X, and let X_d' be the double of X'. Then if Y' is the graph obtained by detaching some of the longitudes of Y, it remains isotopic to a subgraph of X_d' that is a spine of $E(X_d, S^1 \times S^2)$. Thus $E(X_d \cup Y', S^1 \times S^2)$ is a compression body with negative boundary $\partial N(Y')$.

If $\mathcal{Y}$ is any $\epsilon/8$–small decomposition of Y', then $E(X_d \cup \mathcal{Y}, S^1 \times S^2)$ is also a compression body because Observation 2.15 tells us that it is obtained from $E(X_d \cup Y', S^1 \times S^2)$ via 2–handle attachments along $\partial N(Y', W) = \partial_- E(X_d \cup Y', S^1 \times S^2)$. The annuli $\mathcal{A}'$ are the spanning annuli in the collection $(\partial W) \cap E(X_d \cup \mathcal{Y}, S^1 \times S^2)$ of incompressible annuli properly embedded in $E(X_d \cup \mathcal{Y}, S^1 \times S^2)$, and thus form the vertical boundary of the generalized compression body $E(X \cup \mathcal{Y}, W)$ cut off by $\partial W \cap E(X_d \cup \mathcal{Y}, S^1 \times S^2)$.

Define X_{i_j} so that each spanning annulus A_i in our adjacency chain $\mathcal{A}$ meets $N(X_i, W)$ and $N(X_{j_i}, W)$ as in Definition 2.12. The hypothesis that Y does not meet $\mathcal{A}$ implies the existence of a spanning disk D_i, disjoint from $\mathcal{Y}$, for each pair X_i, X_{j_i}. If α is the binding arc of $\mathcal{X}$, Observation 2.15 tells us $V' = E(\mathcal{X} \cup \alpha \cup \mathcal{Y}, W)$ is homeomorphic to $E(X \cup \mathcal{Y}, W)$, and is thus a compression body. Let α_i be the spanning arc of X_i and X_{j_i} that is a subset of the binding arc α for $\mathcal{X}$, and let C_i be the annulus $\mathrm{Fr}(\alpha_i, E(\mathcal{X} \cup \mathcal{Y}, W)) \subset \partial_+ V'$ (whose core curve is just a meridian of α_i). Then the disks $\mathcal{D} = (D_1 \cap V') \cup \cdots \cup (D_k \cap V')$ and annuli $\mathcal{C} = C_1 \cup \cdots \cup C_k$ satisfy the hypotheses of Proposition 1.7 in V', and hence (again remembering Observation 2.15) we conclude that $E(\mathcal{X} \cup \mathcal{Y}, W)$ is a generalized compression body. □

3 The doppelgänger

Convention 3.1 Throughout this section, M is a compact, connected, orientable, irreducible 3–manifold, $\mathcal{G}$ is a generalized Heegaard surface of M, T is a separating essential torus properly embedded in M, and $E(T, M) = M_1 \cup M_2$.

Definition 3.2 $\mathcal{G}$ and T are said to be *well-configured* with respect to M_1 if the following conditions hold:

(1) $\mathcal{G} \cap T$ consists only of simple closed curves which are essential in T and $\mathcal{G}$.

(2) Each component of $\mathcal{G} \cap M_1$ separates M_1.

(3) For each component V of $E(\mathcal{G}, M)$, $T \cap V$ consists only of annuli which are spanning or horizontal.

Convention 3.3 For the remainder of the section, we assume that $\mathcal{G}$ and T are well-configured with respect to M_1.

Observation 3.4 Let $\mathcal{H} = E(\mathcal{G}, M)$, which is a disjoint union of compression bodies. Let $\mathcal{A} = \mathrm{Fr}(T \cap \mathcal{H}, \mathcal{H})$, $\mathcal{A}^1 = \mathcal{A} \cap M_1$, and $\mathcal{A}^2 = \mathcal{A} \cap M_2$. Then conditions 1 and 3 of Definition 3.2, together with Proposition 1.10, imply that $\mathcal{V} = E(T \cap \mathcal{H}, \mathcal{H})$ is a generalized compression body satisfying $\partial_v \mathcal{V} = \mathcal{A}_s$, where $\mathcal{A}_s$ is the union of spanning annuli in $\mathcal{A}$. Moreover, if $\mathcal{A}_h = \mathcal{A} - \mathcal{A}_s$ is the subset of horizontal annuli in $\mathcal{A}$, and $\mathcal{A}_h^i = \mathcal{A}^i \cap \mathcal{A}_h$ for $i = 1, 2$, then there is a primitive disk set $\mathcal{D}$ for $\mathcal{V}$ with respect to some ordering of $\mathcal{A}_h^1 \cup \mathcal{A}_h^2$.

Convention 3.5 The notation of Observation 3.4 is fixed for the remainder of the section. Moreover, we fix a choice of a primitive disk set $\mathcal{D} = D_1 \cup \cdots \cup D_n$, which imposes the primitive orderings $\mathcal{A}_h^i = A_1^i \cup \cdots \cup A_n^i$ for $i = 1, 2$. Here it is understood that, for all $1 \le j \le n$, $A_j^1 \cup A_j^2$ is the frontier of a single component of $T \cap \mathcal{H}$.

Definition 3.6 Let V be a component of $\mathcal{V} = E(T \cap \mathcal{H}, \mathcal{H})$. Let $\mathcal{D}^V = \mathcal{D} \cap V$, $\mathcal{A}^V = \mathcal{A} \cap V$, $\mathcal{A}_s^V = \mathcal{A}_s \cap V = \partial_v V$, $\mathcal{A}_h^V = \mathcal{A}_h \cap V$, and let $\mathcal{A}_p^V$ consist of those components of $\mathcal{A}_h^V$ that are dual to some component of $\mathcal{D}^V$.

Lemma 3.7 *For every component V of $\mathcal{V}$, $\partial_+ V - \mathcal{A}_p^V$ is connected.*

Proof In the case $|\mathcal{A}_p^V| = 0$ there is nothing to prove, so assume $|\mathcal{A}_p^V| > 0$. Order the components of $\mathcal{D}^V = D_1 \cup \cdots \cup D_k$ so that $i < j$ if D_i has lower index than D_j with respect to the primitive ordering of $\mathcal{D}$. Similarly, order $\mathcal{A}_p^V = A_1 \cup \cdots \cup A_k$ so that $i < j$ if A_i has lower index than A_j with respect to the primitive ordering of $\mathcal{A}_h$ (so D_i is dual to A_i for all $1 \le i \le k$). The fact that $\mathcal{D}$ is primitive implies that ∂D_1 meets $\mathcal{A}_p^V$ only in a single spanning arc of A_1, so that A_1 meets a single component of $\overline{\partial_+ V - \mathcal{A}_p^V}$. Likewise, ∂D_2 is disjoint from $\mathcal{A}_p^V - A_1$ and meets A_2 only in a single spanning arc. Since ∂D_2 does not change the component of $\partial_+ V - \mathcal{A}_p^V$ on which it lies when it passes through A_1, it follows that A_2 also meets the same component of $\overline{\partial_+ V - A_p^V}$ on each side. Continuing in this way for the remaining components of $\mathcal{A}_p^V$, we see that every component of $\mathcal{A}_p^V$ meets a single component of $\overline{\partial_+ V - \mathcal{A}_p^V}$. Since $\partial_+ V$ is connected, this implies that $\partial_+ V - \mathcal{A}_p^V$ is also connected. □

Definition 3.8 Let V be a component of $\mathcal{V}$, and index the annuli $\mathcal{A}_h^V = A_1 \cup \cdots \cup A_m$ so that $i < j$ implies that A_i has lower index than A_j with respect to the primitive ordering on $\mathcal{A}_h$. A set $\boldsymbol{A} = \{A_{i_1}, \dots, A_{i_k}\}$ of components of $\mathcal{A}_h^V - \mathcal{A}_p^V$ is said to be *connective* if $(\partial_+ V - \mathcal{A}_h^V) \cup A_{i_1} \cup \cdots \cup A_{i_k}$ is connected. Moreover, $\boldsymbol{A}$ is *minimal* if:

1. $|\boldsymbol{A}|$ is minimal among all connective sets of V.
2. For every other connective set of annuli $\boldsymbol{A}' = \{A_{j_1}, \dots, A_{j_k}\}$ satisfying $|\boldsymbol{A}| = |\boldsymbol{A}'|$, $i_l \le j_l$ for all $1 \le l \le k$.

Note that conditions (1) and (2) define a unique minimal connective set with respect to any given ordering.

Convention 3.9 For the remainder of the section, let $W \cong S^1 \times D^2$ be a solid torus parameterized as in Convention 2.1. Furthermore, let $T_i = \mathrm{Fr}(T, M) \cap M_i$ for $i = 1, 2$ and let $h\colon T_1 \to \partial W$ be a homeomorphism such that each component of $h(\mathcal{G} \cap T_1)$ is a longitude of the form $S^1 \times \{x\}$. Let $\pi\colon T_2 \to T_1$ be the projection which collapses T_2 onto T_1 along the I–fibers of $N(T, M)$ (we assume that $\pi(T_2 \cap \mathcal{G}) = T_1 \cap \mathcal{G}$). We let $M' = W \cup_{h \circ \pi} M_2$ for the remainder of the section.

Observation 3.10 If a component V of $\mathcal{V} \cap M_1$ meets T_1 at all, then $\partial_+ V$ must meet T_1 since the annuli of $\mathcal{A}^V = T_1 \cap V$ are all either horizontal or spanning. However, it is possible that $\mathcal{A}^V$ consists entirely of horizontal annuli, so that $\partial_- V$ does not meet T_1, and this is a case which requires special treatment at certain points in our construction.

Definition 3.11 For each component V of $\mathcal{V} \cap M_1$ which meets T_1, let $\mathcal{B}^V$ denote $\overline{\partial W - h(\mathcal{A}^V)}$, which is a union of annuli.

Lemma 3.12 *For each component V of $\mathcal{V} \cap M_1$ which meets T_1, and each component B of $\mathcal{B}^V$, both components of $h^{-1}(\partial B)$ lie on the same component of $\overline{\partial V - \mathcal{A}^V}$.*

Proof If the curves of $h^{-1}(\partial B)$ lie on distinct components F_1 and F_2 of $\partial V - \mathcal{A}^V$, then we can construct an embedded curve in M_1 that is the union of a spanning arc α of $h^{-1}(B)$ and an arc β properly embedded in V with $\partial\alpha = \partial\beta$. This curve would then intersect the surface F_1 in a single point, which contradicts the assumption that each component of $\mathcal{G} \cap M_1$ (and hence each component of $\overline{\partial V - \mathcal{A}^V}$) is separating in M_1. □

Definition 3.13 Let V be a component of $\mathcal{V} \cap M_1$, and let F be a component of $\overline{\partial V - \mathcal{A}^V}$ that meets T_1. Let $\mathcal{B}_F$ be the union of those components B of $\mathcal{B}^V$ such that $h^{-1}(\partial B) \subset F$. If X is a connected, stabilized spoke graph, possibly with detached longitudes, whose nondetached longitudinal edges are the core curves of $\mathcal{B}_F$, then X is said to be a *doppelgänger spoke graph* for F.

Definition 3.14 Let V be a component of $\mathcal{V} \cap M_1$ that meets T_1, and let $\mathcal{X} \cup \mathcal{Y}$ be a spoke graph constructed as follows:

(1) Let $\mathcal{B}^V_+$ be the union of those annuli in $\mathcal{B}^V$ that appear in $\mathcal{B}_F$ for some component F of $\overline{\partial_+ V - \mathcal{A}^V}$, let X be a connected stabilized spoke graph whose longitudinal edges are the core curves of $\mathcal{B}^V_+$, and let X' be its dual. Suppose $d(X, X') = \epsilon < \pi/2$, where d is the flat metric on W as in Proposition 2.18.

(2) Suppose $A = \{A_{i_1}, \dots A_{i_k}\}$ is the minimal connective set of components of $\overline{\mathcal{A}^V_h - \mathcal{A}^V_p}$ defined in Definition 3.8. Label the components of $\overline{\partial_+ V - \mathcal{A}^V}$ that meet A_{i_1} as F_0 and F_1, and inductively label the remaining components of $\overline{\partial_+ V - \mathcal{A}^V}$ by setting F_j equal to the component of $\overline{\partial_+ V - \mathcal{A}^V}$ that meets A_{i_j} and has not been labeled yet. At each stage of the induction, such a component of $\overline{\partial_+ V - \mathcal{A}^V}$ will always exist because A was chosen to be minimal (otherwise, we could remove an element of A and still have a connective set). Moreover, since A is connective, $\overline{\partial_+ V - \mathcal{A}^V} = F_0 \cup \dots \cup F_k$. Thus there exists a *good* $\epsilon/8$–small decomposition $\mathcal{X}$ of X such that $\mathcal{X} = X_0 \cup \dots \cup X_k$, where X_j is a doppelgänger spoke graph of F_j. Moreover, $\mathcal{X}$ can be chosen so that $C_{i_1} \cup \dots \cup C_{i_k}$ forms an adjacency chain for $\mathcal{X}$, where here C_{i_j} denotes the component of $\partial W - \mathring{N}(X, W)$ that contains $h(A_{i_j})$; see Definition 2.12.

(3) Let $\mathcal{B}_-^V = \mathcal{B}^V - \mathcal{B}_+^V$, which is the union of those annuli in $\mathcal{B}^V$ which appear in $\mathcal{B}_F$ for some component F of $\partial_- V$ that meets T_1. Every component of $\partial W - X$ contains at most one component of $\mathcal{B}_-^V$, for if two components of $\mathcal{B}_-^V$ both lie in the same component of $\partial W - X$, this would imply the existence of a component of $\mathcal{A}^V$ whose boundary components both lie in $\partial_- V$, contrary to part 3 of Definition 3.2. Thus we may assume that the core curves of $\mathcal{B}_-^V$ form a subset of the longitudinal edges of the dual X' of X.

(4) Define the *prohibited* longitudinal edges of X' to be those that lie in the same component of $\partial W - X$ as a component of $h(\mathcal{A}_p^V) \cup h(A_{i_1}) \cup \cdots \cup h(A_{i_k})$ (there is at most one component of this set lying inside each component of $\partial W - X$). A subgraph Y' of X' is said to be *admissible* if it possesses every core curve of $\mathcal{B}_-^V$ as a longitudinal edge, but no prohibited longitudinal edges.

(5) It is possible that $\mathcal{B}_-^V = \varnothing$, in which case we set $\mathcal{Y} = \varnothing$. Otherwise, let $F_1', \ldots, F_l'$ be the components of $\partial_- V$ that meet T_1. Let Y' be an admissible subgraph of X', and let Y be obtained from Y' via an $\epsilon/8$–small detachment of those longitudes of Y' that are not core curves of $\mathcal{B}_-^V$. Finally, let $\mathcal{Y} = Y_1 \cup \cdots \cup Y_l$ be an $\epsilon/8$–small decomposition of Y, where Y_j is a doppelgänger spoke graph for F_j' with $1 \le j \le l$.

Then $\mathcal{X} \cup \mathcal{Y} = X_0 \cup \cdots \cup X_k \cup Y_1 \cup \cdots \cup Y_l$ is said to be a *doppelgänger spoke graph* of V, and it is said to be *perfect* if $\mathrm{Fr}(X_j, W) \cong F_j$ and $\mathrm{Fr}(Y_r, W) \cong F_r'$ for all $0 \le j \le k$ and $1 \le r \le l$.

Observation 3.15 The construction of Definition 3.14 was tailored to the hypotheses of Proposition 2.18. It implies that if $\mathcal{X} \cup \mathcal{Y}$ is a doppelgänger spoke graph associated with V, then $U = E(\mathcal{X} \cup \mathcal{Y}, W)$ is a generalized compression body satisfying $\partial_- U = \mathrm{Fr}(\mathcal{Y}, W)$. Then $\partial_+ U$ is the union of $\mathrm{Fr}(\mathcal{X}, W)$ with those components of $\overline{\partial W - N(\mathcal{X} \cup \mathcal{Y}, W)}$ whose boundary components both lie in $\mathrm{Fr}(\mathcal{X}, W)$.

If V does not meet T_1 as our hypotheses require, then U simply does not exist. Also, if V does meet T_1 but some component F' of $\partial_- V$ does not, then $\partial_- U$ will not contain a component corresponding to F'.

Parts (2) and (4) of Definition 3.14 allow us to deduce the existence of a primitive disk set $\mathcal{E}^U$ in U which will serve as a substitute for the disk set $\mathcal{D}^V$ of V, as the following lemma shows.

Lemma 3.16 *Let V be a component of $\mathcal{V} \cap M_1$ that meets T_1, let $\mathcal{X} \cup \mathcal{Y}$ be a doppelgänger spoke graph of V, and let $U = E(\mathcal{X} \cup \mathcal{Y}, W)$. Suppose $\mathcal{A}_h^V = A_1 \cup \cdots \cup A_m$ is ordered as in Definition 3.8, and that $\mathcal{A}_p^V = A_{p_1} \cup \cdots \cup A_{p_q}$. Then there is* *an ordered, disjoint collection of disks* $\mathcal{E}^U = E_{p_1} \cup \cdots \cup E_{p_q}$ *properly embedded in* U *with the following properties:*

(1) $\mathcal{E}^U \cap \partial_v U = \varnothing$.

(2) $E_{p_j} \cap h(A_{p_j})$ *is a single spanning arc for all* $1 \leq j \leq q$.

(3) $E_{p_j} \cap h(A_l) = \varnothing$ *for all* $1 \leq j \leq q$ *and* $p_j < l \leq m$.

Proof For each component A_l of $\mathcal{A}_h^V$, let C_l denote the closure of the component of $\partial W - \mathring{N}(\mathcal{X} \cup \mathcal{Y})$ that contains $h(A_l)$. Let $\boldsymbol{A} = \{A_{i_1}, \dots, A_{i_k}\}$ be the minimal connective set for $\mathcal{A}_h^V$. Order the components of $\mathcal{X} = X_0 \cup \cdots \cup X_k$ as in part 2 of Definition 3.14 so that, for all $1 \leq j \leq k$, one component of ∂C_{i_j} lies on $\mathrm{Fr}(X_j, W)$ and the other on some component $\mathrm{Fr}(X_r, W)$ of lower index $r < j$. The following technical claim is of central importance.

Claim A *If* C_{p_l} *meets* $\mathrm{Fr}(X_s, W)$ *and* $\mathrm{Fr}(X_r, W)$, *where* $s < r$, *then* $i_r < p_l$.

In particular, Claim A implies that no component of A_p^V lies in the minimal set $\boldsymbol{A}$, so C_{p_l} and C_{i_j} will be distinct for all $1 \leq l \leq q$ and $1 \leq j \leq k$.

Proof of Claim A Notice that both C_{p_l} and C_{i_r} connect $\mathrm{Fr}(X_r, W)$ to a component of $\mathrm{Fr}(\mathcal{X}, W)$ of lower index. It follows that if $p_l < i_j$, then the annulus A_{i_j} could be replaced by A_{p_l} in $\boldsymbol{A}$ and still yield a connective set, which violates condition 2 of Definition 3.8. If $p_l = i_r$, then $A_{p_l} = A_{i_r}$ is both a component of A_p^V and a member of $\boldsymbol{A}$. As a member of the minimal set $\boldsymbol{A}$, we know that the components of ∂A_{p_l} lie on distinct components of $\partial_+ V - (\mathring{A}_{p_l} \cup \mathring{A}_{p_l+1} \cup \cdots \cup \mathring{A}_m)$. But this is incompatible with the fact that, as a component of A_p^V, there must be a disk properly embedded in V whose boundary lies in $\partial_+ V - (\mathring{A}_{p_l+1} \cup \cdots \cup \mathring{A}_m)$ and meets A_{p_l} in a single spanning arc. □

The remainder of the proof will mostly be devoted to the construction of a collection $K_{i_1} \cup \cdots \cup K_{i_k}$ of auxiliary "flap" disks associated with the adjacency chain $C_{i_1} \cup \cdots \cup C_{i_k}$. These flap disks will then become pieces of the so-called "disks with flaps" that will form the components of $\mathcal{E}^U$. See Figure 4 for an example of a disk with very simple flaps.

Recalling Definitions 2.11 and 2.12, let α be the binding arc of $\mathcal{X}$, let α'_j be the spanning arc of X_{j-1} and X_j that lies in α, and let α_j denote the slightly shorter arc $U \cap \alpha'_j$. Recalling Definition 2.16, let K'_{i_1} be the spanning disk of X_0 and X_1 that cuts C_{i_1} and is based at α'_1. Condition 4 of Definition 3.14 ensures that K'_{i_1} is disjoint

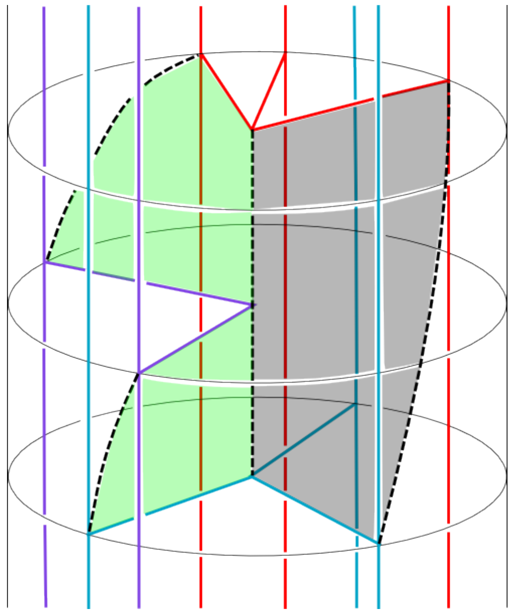

Figure 4: A disk with flaps.

from $\mathcal{Y}$. Set $K_{i_1} = K'_{i_1} \cap U$ and observe that $\partial K_{i_1} = \alpha_1 \cup \beta_1$, where β_1 is an arc lying on $\partial_+ U$ that meets $h(A_p^V)$ only in the annulus $h(A_{i_1}) \subset C_{i_1}$.

The remaining flap disks K_{i_j} are defined recursively. Suppose that the disks K_{i_s} for $s < j$ have already been constructed in U and satisfy the following properties:

(1) $\partial K_{i_s} = \alpha_s \cup \beta_s$, where β_s is an arc on $\partial_+ U$ that meets $h(A^V)$ only in a subset of $h(A_{i_1} \cup \cdots \cup A_{i_s}) \subset C_{i_1} \cup \cdots \cup C_{i_s}$.

(2) K_{i_s} meets $\alpha \cap U$ only in α_s.

(3) $K_{i_s} \cap K_{i_r} = \varnothing$ for all $r < s < j$.

We will construct K_{i_j} so that it also satisfies these properties. To begin with, we know that C_{i_j} meets $\mathrm{Fr}(X_j, W)$ and $\mathrm{Fr}(X_r, W)$ for some $r < j$. Let D be the spanning disk for X_r and X_j that cuts C_{i_j} and is based at $\alpha'_{r+1} \cup \cdots \cup \alpha'_j$. If we construct the disks K_{i_s} correctly in the "obvious" way (or rather, the way that is most directly imaginable after some careful thought), then $K = K_{i_{r+1}} \cup \cdots \cup K_{i_{j-1}} \cup (D \cap U)$ will already be an embedded disk in U satisfying property 1, and after a small isotopy it could be made to satisfy 2 and 3 as well (in the case that $r = j - 1$ we just set $K = D \cap U$). A formal description of these so-called "obvious" disks would be rather tedious and obscure, so we will take a more abstract route in the description of K_{i_j}. However, it is a helpful exercise to visualize what K_{i_2} could look like.

As it stands, our hypotheses on the disks K_{i_s} allow the possibility that $D \cap U$ meets $\mathcal{K} = K_{i_{r+1}} \cup \cdots \cup K_{i_{j-1}}$ not only along $\alpha_{r+1} \cup \cdots \cup \alpha_{j-1}$, which is desirable, but also (transversely, as always) along other "undesirable" arcs and simple closed curves disjoint from $\alpha_{r+1} \cup \cdots \cup \alpha_j$. We eliminate the undesirable intersections as follows. First, the simple closed curves of intersection can be eliminated via isotopy of $D \cap U$ in U. If $D \cap U$ meets $\mathcal{K}$ along any undesirable arcs, then it must meet some component K_{i_s} of $\mathcal{K}$ along an undesirable arc γ which bounds a disk $\delta \subset K_{i_s}$ with the following properties:

- $\partial\delta \subset \gamma \cup (\partial K_{i_s} - \alpha_s)$.
- $\mathring{\delta} \cap D = \varnothing$.

In other words, δ is an outermost disk of intersection whose boundary does not contain the arc α_s. If D'' is the component of $(D \cap U) - \gamma$ that contains α_j in its boundary, $D' = D'' \cup \delta$ is a disk which can be isotoped slightly away from γ so that it intersects $\mathcal{K}$ in strictly fewer arcs than $D \cap U$. It is also possible that D' contains fewer of the arcs α_s for $r < s < j$, so let $\mathcal{K}'$ denote the union of those K_{i_s} that meet D'' along α.

Certainly D' meets $\mathcal{K}'$ in fewer undesirable arcs than D met $\mathcal{K}$. Thus, by repeating this process, we eventually will obtain a disk (call it D' as well) that still contains α_j in its boundary, as well as a corresponding subset of components of $\mathcal{K}$ (call it $\mathcal{K}'$ as well) such that D' meets $\mathcal{K}'$ only along desirable arcs. Then $D' \cup \mathcal{K}'$ will be a disk embedded in U that satisfies property 1 above and, after a small isotopy, property 2 as well. If necessary, another round of outermost disk surgeries that preserve the piece of D' that contains α_j can eliminate any intersections of $D' \cup \mathcal{K}'$ with $K_{i_1} \cup \cdots \cup K_{i_{j-1}}$. The result is the desired disk K_{i_j} that satisfies property 3 as well.

We are now ready to construct the disks E_{p_j}. The first kind we construct are those with index p_j such that both components of ∂C_{p_j} lie on a single component $\mathrm{Fr}(X_r, W)$ of $\mathrm{Fr}(\mathcal{X}, W)$, where X_r has central vertex $v = (\phi_0, 0, 0)$. In the notation of part 4 of Definition 3.14, Y' must be disjoint from C_{p_j} since the longitude of X' that lies in C_{p_j} is prohibited. Hence $\mathcal{Y}$ is disjoint from the component of $(\{\phi_0\} \times D^2) \cap E(\mathcal{X}, W)$ that meets C_{p_j}. We define E_{p_j} to be this component, whose boundary lies in $\partial_+ U$ and meets $h(\mathcal{A}^V)$ only in a single spanning arc of A_{p_j}. Hence it satisfies conclusions (1), (2), and (3) of our lemma.

The remaining components of $\mathcal{E}^U$ will be "disks with flaps"; see Figure 4 for a simple example. Such disks have indexes p_j such that the components of ∂C_{p_j} lie on distinct components $\mathrm{Fr}(X_s, W)$ and $\mathrm{Fr}(X_r, W)$ of $\mathrm{Fr}(\mathcal{X}, W)$ with $s < r$. Let D be the spanning disk of X_s and X_r that cuts C_{p_j} and is based on the spanning arc $\alpha'_{s+1} \cup \cdots \cup \alpha'_r$ of X_s and X_r. Then $D \cap U$ will meet $\mathcal{K} = K_{i_{s+1}} \cup \cdots \cup K_{i_r}$ along "desirable" arcs $\alpha_{s+1} \cup \cdots \cup \alpha_r$, and possibly also along undesirable simple closed curves and arcs, which form the self intersections of the immersed disk $E' = (D \cap U) \cup \mathcal{K}$ whose boundary lies entirely on $\partial_+ U$. Moreover, $D \cap U$ meets $h(\mathcal{A}^V)$ only in A_{p_j}, and Claim A and property 1 of the flap disks ensure $\mathcal{K}$ can only meet $h(\mathcal{A}^V)$ in components of the form $h(A_l)$ with $l < p_j$. Hence the boundary of E' already has the desired properties. As with the construction of the flap disks, we can now eliminate the self-intersections of E' by performing isotopies and a sequence of outermost disk surgeries on $D \cap U$ that preserve its intersection with $h(A_{p_j})$. The result is our desired "disk with flaps" E_{p_j} that does satisfy conclusions (1)–(3) of our lemma.

One technicality to note is that, as constructed, it is possible for some of the components of $\mathcal{E}^U$ to intersect, but again a sequence of appropriate innermost/outermost disk surgeries can be employed to turn it into a disjoint union of disks with the desired properties if necessary. □

Proposition 3.17 *For every component V of $\mathcal{V} \cap M_1$ that meets T_1, there is a perfect doppelgänger spoke graph embedded in W.*

Proof For any doppelgänger spoke graph X_F of a component F of $\overline{\partial V - \mathcal{A}^V}$, the surface $\mathrm{Fr}(X_F, W)$ will have the same number of boundary components as F. Moreover, the genus of $\mathrm{Fr}(X_F, W)$ is the same as the total number of stabilizing arcs and detached longitudes that occur in X_F. In particular, if X_F is unstabilized and has no detached longitudes, then $\mathrm{Fr}(X_F, W)$ will be planar. Thus we can always find a doppelgänger spoke graph X_F that satisfies $\mathrm{Fr}(X_F, W) \cong F$ after attaching a sufficient number of stabilizing arcs and/or detached longitudes.

In part 1 of Definition 3.14, we have the flexibility to stabilize X as often as we need, with stabilizing arcs based on radial edges of our choosing. This allows us to choose the number of stabilizing arcs that will eventually occur in the components of the spoke graph $\mathcal{X}$ defined in part 2 of Definition 3.14. It follows from this and the previous paragraph that we may choose $\mathcal{X}$ so that each of its components X_j satisfies $\mathrm{Fr}(X_j, W) \cong F_j$.

As noted in Remark 2.7, the Morse condition of Definition 2.6 grants us enough flexibility to choose the radial edges on which the stabilizing arcs of X' (the dual of X) will be based, and this in turn allows us to control the component of $\mathcal{Y}$ on which they will eventually occur in part 5 of Definition 3.14. Likewise, we can choose the components of $\mathcal{Y}$ on which the detached longitudes of Y shall occur after the decomposition described in part 5 of Definition 3.14.

So, as in the case with $\mathcal{X}$, we may distribute detached longitudes and stabilizing arcs among the components of $\mathcal{Y}$ however we please. But there is an important difference:

the total number of stabilizing arcs and detached longitudes that can occur in $\mathcal{Y}$ is bounded above by $s+a$, where s denotes the number of stabilizing arcs that occur in X, and a denotes the maximal number of detached longitudes that can occur on an admissible subgraph of X' (as defined in Definition 3.14(4)). Therefore, to complete the proof, we must show that a total of $s+a$ stabilizing arcs and detached longitudes is always sufficient to create a spoke graph $\mathcal{Y}$ that satisfies the equation $\mathrm{Fr}(Y_j, W) \cong F'_j$ for each of its components Y_j.

As in Definition 3.14(5), let $\mathcal{F}' = F'_1 \cup \cdots \cup F'_l$ be the union of those components of $\partial_- V$ that meet T_1. By the first paragraph of this proof, the total number of stabilizing arcs and detached longitudes necessary to ensure that $\mathrm{Fr}(Y_j, W) \cong F'_j$ for all $1 \le j \le l$ is equal to $\sum g(F'_j)$, where $g(F'_j)$ is the genus of F'_j. Since $g(F) = 1 - \frac{1}{2}\chi(F) + |\partial F|$ for any connected, compact surface F, we obtain

$$\sum g(F'_j) = |\mathcal{F}'| - \frac{\chi(\mathcal{F}') + |\mathcal{A}^V_s|}{2}. \tag{1}$$

The fact that this quantity is less than $s+a$ is ultimately derived from the inequality

$$\chi(\partial_+ V) - \chi(\partial_- V) \le -2(|\mathcal{A}^V_p| + |\partial_- V| - 1). \tag{2}$$

The truth of (2) can be seen as follows: $V' = E(\mathcal{D}^V, V)$ is a generalized compression body with the same negative boundary as V. Furthermore, $\partial_+ V'$ is connected since $\partial_+ V - \partial \mathcal{D}^V$ is connected, as can be seen using essentially the same proof as that of Lemma 3.7. Now $\partial_- V'$ is obtained from $\partial_+ V'$ via surgeries along disks, and there must be at least $|\partial_- V'| - 1$ such surgeries since $\partial_+ V'$ is connected. Thus $\chi(\partial_+ V') - \chi(\partial_- V') \le -2(|\partial_- V'| - 1)$. Inequality (2) now follows from the fact that $\partial_- V' = \partial_- V$, and the fact that $\chi(\partial_+ V') = \chi(\partial_+ V) + 2|\mathcal{D}^V| = \chi(\partial_+ V) + 2|\mathcal{A}^V_p|$.

Since no component of $\partial_- V$ is a disk or sphere, and since $\mathcal{F}' \subset \partial_- V$, we have $\chi(\partial_- V) \le \chi(\mathcal{F}')$. Thus, from (2), we easily obtain the analogue $\chi(\partial_+ V) - \chi(\mathcal{F}') \le -2(|\mathcal{A}^V_p| + |\mathcal{F}'| - 1)$. In conjunction with equation (1), we obtain

$$\sum g(F'_j) \le 1 - |\mathcal{A}^V_p| - \frac{\chi(\partial_+ V) + |\mathcal{A}^V_s|}{2}. \tag{3}$$

Our choice of $\mathcal{X}$ has ensured that $\chi(\partial_+ V) = \chi(\mathrm{Fr}(\mathcal{X}, W))$. We then compute $\chi(\mathrm{Fr}(\mathcal{X}, W)) = -|\partial \,\mathrm{Fr}(\mathcal{X}, W)| - 2s + 2|\,\mathrm{Fr}(\mathcal{X}, W)| = -2|\mathcal{A}^V_h| - |\mathcal{A}^V_s| - 2s + 2|\mathcal{X}|$, and so deduce

$$\sum g(F'_j) \le 1 - |\mathcal{A}^V_p| + |\mathcal{A}^V_h| + s - |\mathcal{X}|. \tag{4}$$

The detached longitudes of Y, as described in part 5 of Definition 3.14, all come from core curves of the annuli of $\partial W - X$ that contain a component of $h(\mathcal{A}^V_h)$. On the

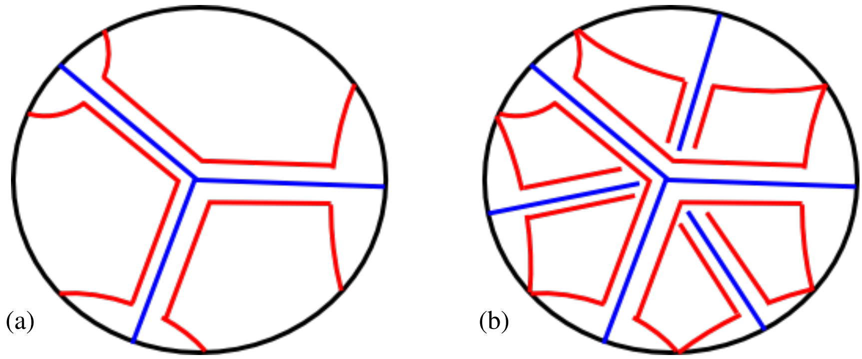

Figure 5: On the left, an image of a component X of $\mathcal{X} \cup \mathcal{Y}$ of Definition 3.18 in blue, projected onto the meridian disk of W. We have chosen $\mathrm{Fr}(X, W)$, in red, to lie close to X in the usual way near the radial edges, but it spreads out near ∂W to meet $h(\partial \mathcal{G}^V)$. On the right, a spoke graph Y and its dual Y' (both in blue), together with $\mathrm{Fr}(Y, W)$ and $\mathrm{Fr}(Y', W)$, chosen as in Lemma 3.20.

other hand, the number of prohibited annuli (part 4 of Definition 3.14) is equal to $|\mathcal{A}_p^V| + |\boldsymbol{A}|$, where $\boldsymbol{A}$ denotes the set of annuli defined in part 2 of Definition 3.14. Hence there will be at most $a = |\mathcal{A}_h^V| - |A_p^V| - |\boldsymbol{A}|$ detached longitudes which may occur in $\mathcal{Y}$. However, $|\boldsymbol{A}| = |\mathcal{X}| - 1$ by the minimality of $|\boldsymbol{A}|$. Plugging this into inequality (4) yields

$$\sum g(F_j') \leq s + a. \tag{5}$$

The proposition now follows. □

Definition 3.18 Let V be a component of $\mathcal{V} \cap M_1$ that meets T_1, let $\mathcal{G}^V$ denote the union of those components G of $\mathcal{G} \cap M_1$ such that $N(G, M_1) \cap V \neq \varnothing$ and $G \cap T_1 \neq \varnothing$. Let $\mathcal{X} \cup \mathcal{Y}$ be a perfect doppelgänger spoke graph for V, and let $\mathcal{Q}^V = \mathrm{Fr}(\mathcal{X} \cup \mathcal{Y}, W)$ where we have chosen $\mathrm{Fr}(\mathcal{X} \cup \mathcal{Y}, W)$ so that $\partial\, \mathrm{Fr}(\mathcal{X} \cup \mathcal{Y}, W) = h(\partial \mathcal{G}^V)$ as in Figure 5(a). Then $\mathcal{Q}^V$ is called the *doppelgänger surface* of V, and the closure of the component of $W - \mathcal{Q}^V$ that does *not* contain $\mathcal{X} \cup \mathcal{Y}$ is the *doppelgänger chamber* of V.

Remark 3.19 Suppose now that two components V and V' of $\mathcal{V} \cap M_1$ that both meet T_1 are adjacent in the obvious sense, which means that $N(G, M_1)$ meets both V and V' for some component G of $\mathcal{G} \cap M_1$ that meets T_1. If Q is the component of $\mathcal{Q}^V$ corresponding to G, and Q' is the component of $\mathcal{Q}^{V'}$ corresponding to G, then $\partial Q = \partial Q'$. Identifying Q with $\mathrm{Fr}(Y, W)$ and Q' with $\mathrm{Fr}(Y', W)$ in Lemma 3.20 below, we see that in fact Q is isotopic to Q' via an isotopy which fixes ∂Q. Moreover, the proof will show that this isotopy, performed ambiently, will push the doppelgänger chamber U of V into the side of Q' that does not contain the doppelgänger chamber U' of V'. We shall call such an isotopy a *flipping isotopy*.

Lemma 3.20 *Let X be a connected, unstabilized spoke graph and let X' be its dual. Suppose Y is obtained from X by attaching a total of n stabilizing arcs and detached longitudes, and that Y' is obtained from X' by attaching a total of n stabilizing arcs and detached longitudes, in any fashion. Then if we choose* $\mathrm{Fr}(Y, W)$ *and* $\mathrm{Fr}(Y', W)$ *so that their boundaries coincide as in Figure 5(b),* $\mathrm{Fr}(Y, W)$ *will be isotopic to* $\mathrm{Fr}(Y', W)$ *via an isotopy which remains fixed on ∂W.*

Proof If Y has detached longitudes, then there is a sequence of edge slides that change Y into a connected stabilized spoke graph $\widetilde{Y}$ without detached longitudes (see Figure 6), and these correspond to an isotopy of $\mathrm{Fr}(Y, W)$ to $\mathrm{Fr}(\widetilde{Y}, W)$, one which we can choose to be supported outside of a small open collar of ∂W.

If $\widetilde{Y}'$ is the dual of $\widetilde{Y}$, and we choose $\mathrm{Fr}(\widetilde{Y}', W)$ so that its boundary coincides with the boundary of $\mathrm{Fr}(\widetilde{Y}, W)$ (as in Figure 5(b)), then Lemma 2.9 (see also Remark 2.10) tells us that the region of W trapped between $\mathrm{Fr}(\widetilde{Y}', W)$ and $\mathrm{Fr}(\widetilde{Y}, W)$ can be parameterized as a “pinched” thickened surface whose vertical boundary has been collapsed. Hence we can isotope $\mathrm{Fr}(\widetilde{Y}, W)$ onto $\mathrm{Fr}(\widetilde{Y}', W)$ via an isotopy that is fixed on ∂W.

Finally, since $\widetilde{Y}'$ is obtained from X' by attaching n stabilizing arcs, using (a reversed version of) the same kind of isotopy described in the first paragraph, we may slide the stabilizing arcs of $\widetilde{Y}'$ along X' so that they coincide with the stabilizing arcs and detached longitudes of Y'. This corresponds to an isotopy of $\mathrm{Fr}(\widetilde{Y}', W)$ onto $\mathrm{Fr}(Y', W)$ that is fixed on $\partial\,\mathrm{Fr}(\widetilde{Y}', W)$. □

Theorem 3.21 *If the generalized Heegaard surface $\mathcal{G} \subset M$ amalgamates to a minimal genus Heegaard surface of M, then $g(M') \le g(M)$ (here M' is the manifold obtained by gluing W to M_2 as described in Convention 3.9).*

Proof Our assumption that T and $\mathcal{G}$ are well-configured does not eliminate the possibility that $T \cap \mathcal{G} = \varnothing$. But in this case, T will be parallel to a component of $\mathcal{G}_-$, the thin part of $\mathcal{G}$, which implies the stronger conclusion $g(M_1) + g(M_2) = g(M)$. So we assume $T \cap \mathcal{G} \neq \varnothing$.

Suppose that $V_1 \cup \cdots \cup V_n$ is the union of all components of $\mathcal{V} \cap M_1$ that meet T_1, and let $\{U_1, \ldots, U_n\}$ be the corresponding set of doppelgänger chambers in W. If the U_i are embedded exactly as in Definition 3.18, then they will intersect one another badly. However, using Lemma 3.20 (and Remark 3.19 preceding it) we can show:

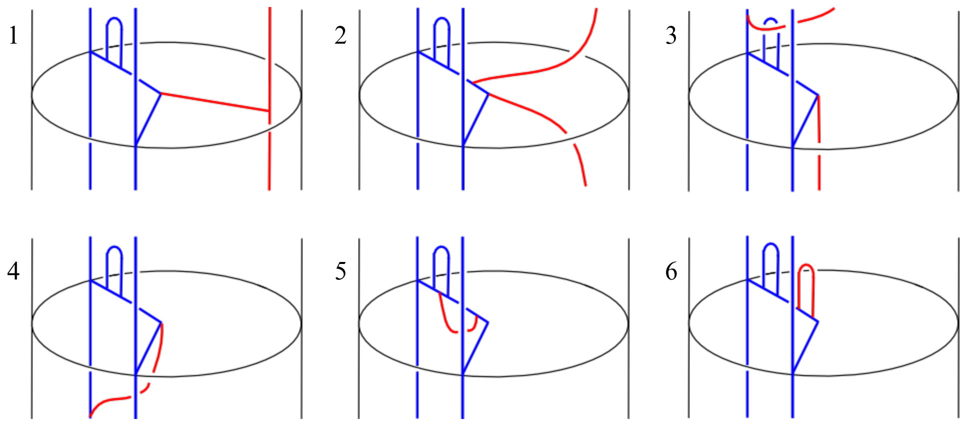

Figure 6: Edge slides that turn a detached longitude into a stabilized arc.

Claim B *The U_i can be isotoped in W, via isotopies which leave ∂W fixed, so that:*

(1) $U_1 \cup \cdots \cup U_n = W$.

(2) *For all $i \neq j$, $U_i \cap U_j = \varnothing$ unless V_i is adjacent to V_j as in Remark 3.19.*

(3) *If V_i is adjacent to V_j, then $U_i \cap U_j = (\partial_+ U_i \cap \partial_+ U_j) \cup (\partial_- U_i \cap \partial_- U_j)$, and it is a single component of $\partial_\pm U_i$.*

We shall prove Claim B at the end. Assuming it is established, let $\mathcal{Q}$ denote the union of all surfaces of the form $U_i \cap U_j$ with $i \neq j$, which is the same as the union of all surfaces $\partial_+ U_i \cup \partial_- U_i$. Moreover, for all $1 \leq i \leq n$, let $\widetilde{U}_i$ be the component of $E(\mathcal{Q}, W)$ corresponding to U_i. Thus $\widetilde{U}_i$ is just the result of carving out a collar of $\partial_+ U_i \cup \partial_- U_i$ from U_i; we need it here so that it will match up snugly with the components of $\mathcal{V} \cap M_2$, but morally it should just be thought of as U_i. The following facts are now easily verified for all $1 \leq i \leq n$:

- $h(\partial_v V_i) = \partial_v \widetilde{U}_i$.
- $h(\partial_+ V_i \cap T_1) = \partial_+ \widetilde{U}_i$.

Hence by Observation 1.2, the result of attaching the vertical annuli of $(\mathcal{V} \cap M_2)$ to the vertical annuli of $\widetilde{U}_1 \cup \cdots \cup \widetilde{U}_n$ along $h \circ \pi$ will be a generalized compression body, and it will have no vertical boundary components (they have all been glued together).

Moreover, if $\mathcal{E}$ is the union of all the (deformed versions of) the disk sets $\mathcal{E}^{U_i}$ defined in Lemma 3.16, then Lemma 3.16 implies that $\mathcal{E} \cup (\mathcal{D} \cap M_2)$ admits an ordering which makes it a primitive disk set for the ordered union of annuli $h(\mathcal{A}_h^1) \cup \mathcal{A}_h^2$ (see Definition 1.4 and Convention 3.5) with respect to the map $h \circ \pi$ (see Convention 3.9).

By Proposition 1.6, the result of gluing $E(\mathcal{Q}, W) = \widetilde{U}_1 \cup \cdots \cup \widetilde{U}_n$ to $E(\mathcal{G}, M_2) = \mathcal{V} \cap M_2$ along $h \circ \pi$ is a generalized compression body. But this is just $E(\mathcal{Q} \cup_{h\circ\pi} (\mathcal{G} \cap M_2), M')$. Hence $\mathcal{G}' = \mathcal{Q} \cup_{h\circ\pi} (\mathcal{G} \cap M_2)$ is a generalized Heegaard surface in M'. The surface $\mathcal{G}'$ may have fewer components than $\mathcal{G}$, but in any event it will amalgamate to a surface of genus less than or equal to the genus of the surface that $\mathcal{G}$ amalgamates to in M. Thus we will be finished once we prove Claim B.

Proof of Claim B We will describe an algorithm which uses flipping isotopies to embed the U_i in the desired fashion. Start by embedding U_1 as described in Definition 3.18. Suppose now that V_1 is adjacent to V_i along some component G of $\mathcal{G} \cap M_1$ (such a component G will be unique by Convention 3.3 above), which means that $N(G, M_1)$ meets V_1 and V_i in components F_1 of $\partial_\pm V_1$ and F_i of $\partial_\pm V_i$. If R_1 and R_i are the corresponding components of $\partial_\pm U_1$ and $\partial_\pm U_i$, then Remark 3.19 and Lemma 3.20 tell us that R_1 can be isotoped onto R_i via a flipping isotopy which is fixed on ∂W. Moreover, this flipping isotopy, if performed ambiently, will push U_1 into $\overline{W - U_i}$, where U_i is now embedded as described in Definition 3.18 (and U_1 is now distorted).

Now suppose V_1 is adjacent to another component V_j of $\mathcal{V} \cap M_1$ along a component G' of $\mathcal{G} \cap M_1$, and suppose R_1' and R_j' are the components of $\partial_\pm U_1$ and $\partial_\pm U_j$, respectively, that correspond to G'. Reverse the flipping isotopy of the previous paragraph, so that U_1 will again be embedded in the standard way, and U_i will be distorted. We can then perform a further flipping isotopy taking R_1' to R_j', which will push U_1 and the distorted version of U_i into $\overline{W - U_j}$, where we assume U_j is now embedded in the standard way.

Repeating the process of the previous two paragraphs, we can eventually embed (distorted versions of) the doppelgänger chambers of every other component of $\mathcal{V} \cap M_1$ that is adjacent to V_1 in a way that satisfies part 2 of Claim B. Once this is done, we can then embed the components of $\mathcal{V} \cap M_1$ that are adjacent to those components of $\mathcal{V} \cap M_1$ adjacent to V_1, and so on. Eventually this process will terminate, and we will have embedded every doppelgänger U_i of every component of $\mathcal{V} \cap M_1$ that meets T_1, and it is easy to verify that part 1 of Claim B will then be satisfied as well. The claim and theorem now follow. □

Corollary 3.22 *If $\mathcal{G}$ amalgamates to a minimal genus Heegaard surface of M, then $g(M_2) \le g(M) + 1$.*

Proof Since the core c of W can be embedded in the surface $\mathcal{G}'$ constructed in the proof of Theorem 3.21, we can stabilize $\mathcal{G}'$ once (if necessary) to obtain a generalized Heegaard splitting of $E(c, M') \cong M_2$ of genus at most $g(M) + 1$. □

4 The main result

Theorem 4.1 *If K_1 and K_2 are knots in S^3, then $t(K_1 \# K_2) \geq \max\{t(K_1), t(K_2)\}$.*

Proof The proposition is trivial if one of K_1 or K_2 is the unknot, so suppose that both are nontrivial knots. Assume also that $\max\{t(K_1), t(K_2)\} = t(K_2)$. Let T be the "swallow-follow" torus in $E(K_1 \# K_2)$ that swallows the K_2 summand and follows the K_1 summand; see Figure 7. We apply Theorem 3.21 by setting $M = E(K_1 \# K_2)$, noting that one component of $E(T, M)$ is homeomorphic to $E(K_1)$, which will correspond to M_1. What needs to be shown is that the untelescoped minimal splitting $\mathcal{G}$ can be isotoped so that it meets T only in essential simple closed curves, and such that each component of $\mathcal{G} \cap E(K_1)$ is separating.

By its definition as a swallow-follow torus, T is isotopic to $A \cup B$, where A is the decomposing annulus of the connected sum in $M = E(K_1 \# K_2)$, and B is the subannulus of $\partial M - A$ that lies in the component of $M - A$ corresponding to $E(K_1)$. By Proposition 1.21, $\mathcal{G}$ can be isotoped to intersect A only in essential simple closed curves, and since each boundary component of $\mathcal{G} \cap E(K_1)$ is then a standard meridional curve of $\partial E(K_1)$, every component of $\mathcal{G} \cap E(K_1)$ is separating in $E(K_1)$ (otherwise we could obtain a nonseparating surface in S^3). The hypotheses of Theorem 3.21 (which assume Conventions 3.1, 3.3, and 3.9) can then be satisfied by isotoping T sufficiently close to $A \cup B$.

Now M_2 is the component of $E(T, M)$ which is *not* homeomorphic to $E(K_1)$, but is instead homeomorphic to $E(L)$, where L is the link in S^3 that has K_2 as one component and a meridian μ of K_2 as its other component, and $T = \partial N(\mu)$ under this correspondence; see Figure 7. Furthermore, the slope in which $\mathcal{G}$ has been made to intersect $T = \partial N(\mu)$ is the standard longitudinal slope determined by the meridian disk $\delta \subset S^3$ with $\partial\delta = \mu$ and $|\delta \cap K_2| = 1$. Thus the slope of the trivial Dehn filling of $\partial N(\mu) = T$ that yields $E(K_2)$ meets each component of $T \cap \mathcal{G}$ exactly once, and Theorem 3.21 applies to $M' = E(K_2)$, yielding

$$t(K_1 \# K_2) = g(E(K_1 \# K_2)) - 1 \geq g(E(K_2)) - 1 = t(K_2). \qquad \square$$

This proof also works if the knots K_1 and K_2 are embedded in homology spheres (or any pair of compact 3–manifolds in which every closed embedded surface is separating). In general, however, it is important to keep in mind the delicacy of Theorem 3.21 (and Corollary 3.22), whose assumptions are encoded in Conventions 3.1, 3.3 and 3.9. In particular, the assumption of Convention 3.3 that T and $\mathcal{G}$ are well-configured cannot always be satisfied as can be shown using straightforward examples in $S^1 \times F$, where F is a closed genus $g > 1$ surface. Thus Corollary 3.22 cannot be applied to prove the following plausible conjecture in the case $g = 1$ in any obvious way.

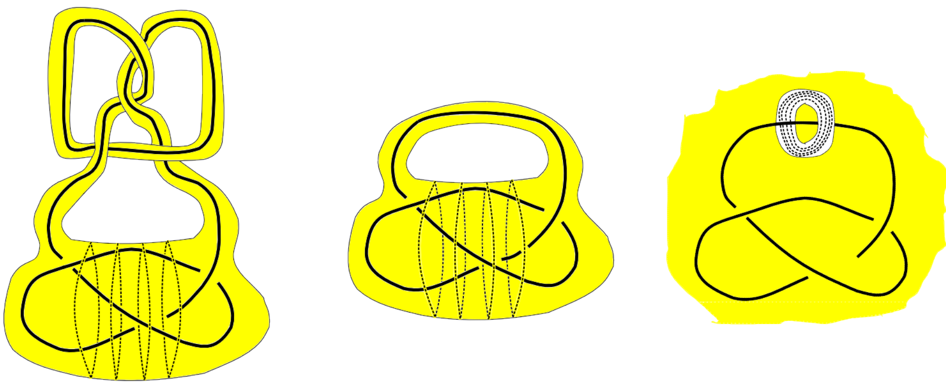

Figure 7: In the first diagram, M_2 (in yellow) is seen as situated in $E(K_1 \# K_2, S^3)$, and $\mathcal{G} \cap T$ is indicated with dashed lines lying on the "swallow-follow" torus T. In the second diagram, M_2 is re-embedded in $E(K_2, S^3)$, and in the final diagram, we see M_2 and $\mathcal{G} \cap T$ as they look after inverting the image of T under this re-embedding.

Conjecture 4.2 *Suppose M is a compact 3–manifold and T is a separating, incompressible, orientable, genus g surface properly embedded in M. If M_1 and M_2 are the components of $E(T, M)$, then $g(M) \geq \max\{g(M_1), g(M_2)\} - g$.*

Similarly, the need for T and $\mathcal{G}$ to be well-configured is what keeps us from applying Theorem 3.21 and Corollary 3.22 to prove the analogue of Theorem 4.1 for satellite knots.

Theorem 4.1 has some relation to the "rank-genus conjecture" for knot complements in S^3. If we define $r(K)$ to be the minimal number of generators for $\pi_1(S^3 - K)$, then the rank-genus conjecture states:

Conjecture 4.3 *For all knots $K \subset S^3$, $r(K) = g(E(K, S^3)) = t(K) + 1$.*

Since a genus g Heegaard splitting of a knot complement induces a g–generator presentation of $\pi_1(S^3 - K)$, it is clear that $r(K) \leq t(K) + 1$, but it remains unknown whether it is possible for this inequality to be strict. Boileau and Zieschang [1] described closed Seifert fibered 3–manifolds M that satisfy $g(M) > r(M)$, where $r(M)$ is the rank of $\pi_1(M)$. More recently, Li [6] constructed closed hyperbolic 3–manifolds satisfying the same inequality.

Hence it seems likely that the rank-genus conjecture fails for knot complements, although it remains unknown. A pair of knots in S^3 whose tunnel number degenerated enough to violate Theorem 4.1 would have given a counterexample, since the following analogue of Theorem 4.1 for rank is trivial (thanks to Richard Weidmann for pointing out the simple line of proof below).

Proposition 4.4 *For any knots* $K_1, K_2 \subset S^3$, $r(K_1 \# K_2) \geq \max\{r(K_1), r(K_2)\}$.

Proof We have that $\pi_1(E(K_1 \# K_2))$ is an amalgamated free product $\pi_1(E(K_1)) *_{\mathbb{Z}} \pi_1(E(K_2))$ that retracts onto each of its factors. □

The fact that Theorem 4.1 is true indicates that the class of knot pairs that experience high tunnel number degeneration is not a good place to look for counterexamples to the rank-genus conjecture after all. In any event, the simplicity of the proof of Proposition 4.4 makes a striking contrast to our proof of Theorem 4.1.

Acknowledgments I would like to thank Maggy Tomova and Charlie Frohman for their support and guidance throughout my career. I would like to thank Jesse Johnson for the same, as well as for many helpful conversations and observations about this paper in particular.

References

[1] **M Boileau**, **H Zieschang**, *Heegaard genus of closed orientable Seifert* 3*–manifolds*, Invent. Math. 76 (1984) 455–468 MR746538

[2] **F Bonahon**, **J-P Otal**, *Scindements de Heegaard des espaces lenticulaires*, Ann. Sci. École Norm. Sup. 16 (1983) 451–466 MR740078

[3] **T Kobayashi**, *A construction of arbitrarily high degeneration of tunnel numbers of knots under connected sum*, J. Knot Theory Ramifications 3 (1994) 179–186 MR1279920

[4] **T Kobayashi**, **Y Rieck**, *Heegaard genus of the connected sum of* m*–small knots*, Comm. Anal. Geom. 14 (2006) 1037–1077 MR2287154

[5] **T Kobayashi**, **Y Rieck**, *Knot exteriors with additive Heegaard genus and Morimoto's conjecture*, Algebr. Geom. Topol. 8 (2008) 953–969 MR2443104

[6] **T Li**, *Rank and genus of* 3*–manifolds*, J. Amer. Math. Soc. 26 (2013) 777–829 MR3037787

[7] **T Li**, **R Qiu**, *On the degeneration of tunnel numbers under a connected sum*, Trans. Amer. Math. Soc. 368 (2016) 2793–2807

[8] **Y Moriah**, **H Rubinstein**, *Heegaard structures of negatively curved* 3*–manifolds*, Comm. Anal. Geom. 5 (1997) 375–412 MR1487722

[9] **K Morimoto**, *There are knots whose tunnel numbers go down under connected sum*, Proc. Amer. Math. Soc. 123 (1995) 3527–3532 MR1317043

[10] **K Morimoto**, **M Sakuma**, **Y Yokota**, *Examples of tunnel number one knots which have the property "$1+1=3$"*, Math. Proc. Cambridge Philos. Soc. 119 (1996) 113–118 MR1356163

[11] **K Morimoto**, **J Schultens**, *Tunnel numbers of small knots do not go down under connected sum*, Proc. Amer. Math. Soc. 128 (2000) 269–278 MR1641065

[12] **J M Nogueira**, *Tunnel number degeneration under the connected sum of prime knots*, Topology Appl. 160 (2013) 1017–1044 MR3049251

[13] **F H Norwood**, *Every two-generator knot is prime*, Proc. Amer. Math. Soc. 86 (1982) 143–147 MR663884

[14] **T Saito**, **M Scharlemann**, **J Schultens**, *Lecture notes on generalized Heegaard splittings*, preprint (2005) `arXiv:math/0504167`

[15] **M Scharlemann**, **J Schultens**, *The tunnel number of the sum of n knots is at least n*, Topology 38 (1999) 265–270 MR1660345

[16] **M Scharlemann**, **J Schultens**, *Annuli in generalized Heegaard splittings and degeneration of tunnel number*, Math. Ann. 317 (2000) 783–820 MR1777119

[17] **M Scharlemann**, **A Thompson**, *Thin position for 3–manifolds*, from: "Geometric topology", (C Gordon, Y Moriah, B Wajnryb, editors), Contemp. Math. 164, Amer. Math. Soc. (1994) 231–238 MR1282766

[18] **J Schultens**, *Additivity of tunnel number for small knots*, Comment. Math. Helv. 75 (2000) 353–367 MR1793793

Department of Mathematics, Oklahoma State University
Stillwater, OK 74078, USA

`trentschirmer@gmail.com`

`http://www.trentschirmer.com`

Received: 3 December 2012 Revised: 3 August 2015