C-essential surfaces in (3-manifold, graph) pairs

SCOTT A. TAYLOR AND MAGGY TOMOVA

Let T be a graph in a compact, orientable 3-manifold M and let Γ be a subgraph. T can be placed in bridge position with respect to a Heegaard surface H. We show that if H is what we call (T, Γ)-c-weakly reducible in the complement of T then either a "degenerate" situation occurs or H can be untelescoped and consolidated into a collection of "thick surfaces" and "thin surfaces". The thin surfaces are c-essential (c-incompressible and essential) in the graph exterior and each thick surface is a strongly irreducible bridge surface in the complement of the thin surfaces. This strengthens and extends previous results of Hayashi-Shimokawa and Tomova to graphs in 3-manifolds that may have non-empty boundary.

1. Introduction

Thin position has been an important tool in knot theory and 3-manifold topology. Gabai [5] defined thin position for a knot in S^3 in his solution of the Poenaru Conjecture; it was also put to good use by Gordon and Luecke [7] in their solution of the knot complement problem. Thin position was extended to graphs in S^3 by Scharlemann and Thompson [16] who used it to give a new proof of Waldhausen's classification of Heegaard splittings of S^3. Goda et al. [6] also used thin position for graphs in S^3 to prove that an unknotting tunnel for a tunnel number one knot in minimal bridge position can be slid and isotoped to lie in the bridge surface. They used this fact to give a new proof of the classification of unknotting tunnels for 2-bridge knots. Goda et al.'s result is also foundational to Cho and McCullough's [3] recent important work on tunnel number one knots.

In tandem with their development of thin position for knots and graphs in S^3, Scharlemann and Thompson [16] described a very different thin position for 3–manifolds (based on earlier work of Casson and Gordon [2]). In this type of thin position, a 3-manifold is decomposed along incompressible surfaces (called "thin surfaces") into codimension 0 submanifolds each of which contains a strongly irreducible Heegaard surface (called a "thick surface").

A locally thin position for a 3-manifold can be obtained by "untelescoping" a Heegaard surface for the 3-manifold into thin and thick surfaces. In practice, a strongly irreducible Heegaard surface functions very much like an incompressible surface. Part of the power of thin position for a 3-manifold comes from the fact that an irreducible Heegaard surface in a 3-manifold is either strongly irreducible, or can be untelescoped into a collection of incompressible surfaces and strongly irreducible surfaces. In either case, we have surfaces that function like incompressible surfaces. This aspect of thin position for 3-manifolds has been used, for example, in work on the virtual Haken conjecture [12, 13]. The virtual Haken conjecture was recently proven by Agol et al. [1].

Hayashi and Shimokawa [9] found a way of uniting (at least in spirit) the notions of thin position for a knot and thin position for a 3-manifold. Their version of thin position involves placing a properly embedded 1-manifold in a 3-manifold into bridge position with respect to a Heegaard surface for the 3-manifold and then untelescoping the Heegaard surface using compressing disks and bridge disks. Coward [4] used Hayashi and Shimokawa's thin position to find an algorithm for determining the bridge number of a hyperbolic knot in S^3.

Tomova [22] strengthened Hayashi and Shimokawa's untelescoping operation by weakening the hypotheses on the sort of disks used in the untelescoping operation: she used so-called "c-disks". Tomova [21] used this new version of thin position in her work on the relationship between multiple bridge surfaces for a given knot, producing a version of the "alternate Heegaard genus bounds distance" theorem of Scharlemann and Tomova [18]. Scharlemann and Tomova [17] also use this version of thin position to show that 2-bridge knots have essentially unique bridge surfaces. Tomova's version of thin position, however, requires that the 3-manifold be boundary-less.

In the present work, we show that (except in a few degenerate situations) a bridge surface for a graph in a 3-manifold, possibly with boundary, can be transformed using untelescoping-like operations into a type of thin position for the graph where the thin surfaces are essential in the graph complement and the thick surfaces are strongly irreducible bridge surfaces in the complement of the thin surfaces. The union of thin and thick surfaces is called a "multiple Heegaard splitting" for the graph in the 3-manifold.

In the thin position of Hayashi-Shimokawa and Tomova the proof that thin surfaces are not parallel to a component of the link relies on the difficult classification of certain bridge surfaces [10, 11]. So in our situation, proving

that thin surfaces are essential relies on the classification of bridge surfaces for certain graphs in compressionbodies [20]. (See Theorem 3.4.)

Since in the past it has been helpful to be able to untelescope a bridge surface for a knot along c-disks (and not just compressing disks), we prove a relative version of the theorem. The set up is as follows: We have a compact orientable 3-manifold M, a Heegaard surface $H \subset M$ and a graph $T \subset M$ in bridge position with respect to H. Inside T we specify a subgraph Γ that contains no vertices of T, and when we untelescope, we will use compressing disks for $H - T$ and so-called "cut disks" that intersect Γ. If M is a closed 3-manifold and T is a link then if Γ is chosen to equal T, we obtain Tomova's thin position, although our conclusion that the thin surfaces are c-essential is stronger than her conclusion. If M is a 3-manifold and T is a properly embedded 1-manifold, then choosing $\Gamma = \varnothing$ gives Hayashi and Shimokawa's thin position. We note that, even in this case, if $\partial M \neq \varnothing$ our conclusion that the thin surfaces are essential (i.e., incompressible and not boundary parallel) in the exterior of the graph is stronger than Hayashi and Shimokawa's conclusion that the thin surfaces are incompressible in the complement of the graph and not "T-parallel" to a component of ∂M.

We prove

Corollary 9.4 (Rephrased). *Let M be a compact, orientable 3-manifold containing a properly embedded graph T and let Γ be a subgraph of T disjoint from the vertices of $T - \partial T$. Furthermore, assume that $M - T$ is irreducible and that no sphere in M intersects T exactly once. Suppose K is a bridge surface for (M, T, Γ) such that no circle or edge component of Γ is isotopically core with respect to K. Then one of the following holds:*

- *there is a multiple Γ-Heegaard splitting $\mathcal{H}$ for (M, T, Γ) so that each thin surface is T-essential and each thick surface is (T, Γ)-c-strongly irreducible in the component of $M - \mathcal{H}^-$ containing it. Furthermore, $\mathcal{H}$ is obtained by applying "untelescoping-like" operations to K.*

- *K contains a generalized stabilization.*

- *K is perturbed.*

- *K has a removable path.*

All of the terms are defined in the next section.

In a forthcoming paper, we will demonstrate how this result can be used to show that certain edges in a Heegaard spine can be made level with respect to a bridge surface.
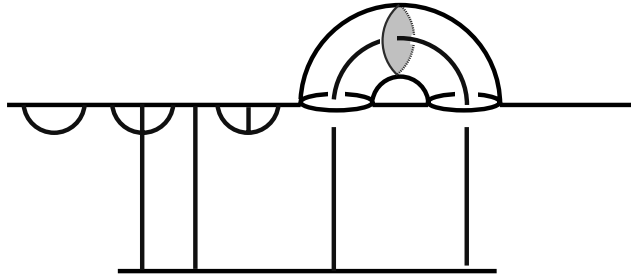
Figure 1: The five types of components in a Γ-compressionbody. From left to right we have: bridge arc, vertical pod, vertical arc, pod, ghost arc. The red disc is a cut disc for the ghost arc.

2. Definitions

2.1. Surfaces in (M, T)

Let T be a finite graph. Unless otherwise specified we assume that T has no valence 2 vertices as such vertices can generally be deleted and their adjacent edges amalgamated. T may contain closed loops without vertices. We say that T is *properly embedded* in a 3-manifold M if $T \cap \partial M$ is the set of all valence 1 vertices of T. We will denote the pair (M, T).

Suppose that $F \subset M$ is a surface such that $\partial F \subset (\partial M \cup T)$. Then F is *T–compressible* if there exists a compressing disk for $F - T$ in $M - T$. If F is not T–compressible, it is *T–incompressible*. F is *T–∂–compressible* if there exists a disk $D \subset M - T$ with interior disjoint from F such that ∂D is the endpoint union of an arc γ in F and an arc δ in ∂M. We require that γ not be parallel in $F - T$ to an arc of $\partial F - T$. If F is not T–∂– compressible, it is *T–∂–incompressible*. Finally suppose Γ is some subgraph of T. We will say that F is *Γ-cut-compressible*, if there exists a disk D^c that only intersects F it its boundary, the boundary is essential in $F - T$, so that $|D^c \cap T| = 1$ and that point of intersection is contained in Γ. We also require that ∂D^c is not parallel in $F - T$ to a puncture $T \cap F$. We call D^c a *Γ-cut-disk*. A *Γ-c-disk* will be either a T-compressing disk or a Γ-cut-disk. If Γ is understood, we refer to it as simply a c-disk. Figure 1 depicts, among other things, a cut disk. A surface F in M is called *T-parallel* if F is boundary parallel in $M - \mathring{\eta}(T)$ and *T-essential* if it is T-incompressible and not T-parallel. The surface F will be called *T-c-essential* if it is essential and cut-incompressible.

2.2. Trivially embedded graphs in compressionbodies

Let C be a compressionbody and T be a properly embedded graph in C. A connected component τ of T is *trivial* in C if it is one of four types (see figure 1):

(1) *Bridge arc:* a single edge with both endpoints in $\partial_+ C$ which is parallel to an arc in $\partial_+ C$. The disk of parallelism is called a *bridge disk*.

(2) *Vertical edge:* a single edge with one endpoint in $\partial_+ C$ and one endpoint in $\partial_- C$ that is isotopic to $\{\text{point}\} \times I$.

(3) *Pod:* a graph with a single vertex in the interior of C and with all valence 1 vertices lying in $\partial_+ C$ so that there is a disk D with $\partial D \subset \partial_+ C$ inessential in $\partial_+ C$ and so that $\tau \subset D$. The disk D is called a *pod disk*. Each of the components of $D - \tau$ will be called a *bridge disk* as these components play the same role as bridge disks for bridge arcs.

(4) *Vertical pod:* a graph with a single vertex in the interior of C and with one valence 1 vertex lying in $\partial_- C$ and all other valence 1 vertices in $\partial_+ C$ so that if the edge adjacent to $\partial_- C$ is removed the resulting graph is a pod and if instead all but one of the edges adjacent to $\partial_+ C$ are removed, the result is a vertical edge with a valence 2 vertex in its interior. The edges that have one endpoint in $\partial_+ C$ are called *pod legs* and the other edge is called a *pod handle*.

If all components of T are simultaneously trivially embedded, then we say that T is *trivially embedded* in C.

2.3. Trivially embedded graphs in Γ-compressionbodies

Definition 2.1. Let C be a compressionbody containing a properly embedded graph T and let Γ be a subgraph of T. Suppose that there is a collection of (T, Γ)-cut-disks, $\mathcal{D}^c$, such that:

(1) for each edge of Γ there is at most one disk in $\mathcal{D}^c$ intersecting it,

(2) each edge of Γ intersected by $\mathcal{D}^c$ has both endpoints on $\partial_- C$,

(3) boundary reducing C along all cut-disks $\mathcal{D}^c$ produces a union of compressionbodies $C_1, \ldots, C_n$,

(4) for each i the graph $C_i \cap T$ is trivially embedded in C_i.

Then we call the triple of the compressionbody and the graphs a Γ-*compressionbody containing a trivially embedded graph* $T - \Gamma$ and denote it (C, T, Γ).

Note that if $\Gamma = \emptyset$ then $(C, T, \Gamma) = (C, T)$ is a compressionbody containing a trivially embedded graph.

If (C, T, Γ) is a Γ-compressionbody and γ is an arc of Γ in (C, T, Γ) that has both of its endpoint in $\partial_- C$ and contains no vertices, then γ is called a *ghost arc*. Each ghost arc intersects a Γ-cut disk in C.

2.4. Heegaard surfaces and Γ-Heegaard surfaces

Let (M, T) be a (3-manifold, graph) pair. A *Heegaard splitting* for (M, T) is a decomposition of M into two compressionbodies, C_1 and C_2, such that $T_i = T \cap C_i$ is trivially embedded in C_i for $i \in \{1, 2\}$. The surface $H = \partial_+ C_1 = \partial_+ C_2$ is called a *Heegaard surface* for (M, T). We will also say that T is in *bridge position* with respect to H and that H is a *bridge splitting* of (M, T).

Suppose now that Γ is a subgraph of T and H is a surface in (M, T) transverse to T so that H splits M into compressionbodies C_1 and C_2 such that (C_i, T_i, Γ_i) is a Γ_i-compressionbody containing a trivially embedded graph $T_i - \Gamma_i$ for $i \in \{1, 2\}$. In this case we say H is Γ-*Heegaard surface* or a Γ-*bridge surface* for (M, T, Γ). If $\Gamma = \emptyset$ then H is simply a Heegaard surface.

Suppose H is a Γ-Heegaard surface for (M, T, Γ) splitting it into triples (C_1, T_1, Γ_1) and (C_2, T_2, Γ_2). We will say that H is *T-reducible* if there exists a sphere S disjoint from T such that $S \cap H$ is a single curve essential in $H - T$, otherwise H is *T-irreducible*. We will say that H is *T-weakly reducible* if H has T-compressing disks on opposite sides with disjoint boundaries. Otherwise H is said to be *T-strongly irreducible*. We will say that H is (T, Γ)-*c-weakly reducible* if H has Γ-c-disks on opposite sides with disjoint boundaries. Otherwise H is said to be (T, Γ)-*c-strongly irreducible*.

2.5. Multiple bridge splittings

To prove our main theorem we will extend to graphs the definition of multiple bridge splittings for the pair (3-manifold, 1-manifold) introduced by Hayashi and Shimokawa in [11] and generalized by Tomova in [22].

Definition 2.2. Suppose M is a 3-manifold containing a properly embedded graph T and let Γ be a subgraph of T. A disjoint union of surfaces $\mathcal{H}$ is a multiple Γ-Heegaard splitting for (M, T, Γ) if:

- the closure of each component of $(M, T) - \mathcal{H}$ is a Γ-compressionbody C_i containing a trivially embedded graph $T_i = (T - \Gamma) \cap C_i$,

- for each i, $\partial_+ C_i$ is attached to some $\partial_+ C_j$ and each component of $\partial_- C_i$ is either contained in ∂M or is attached to a component of some $\partial_- C_k$ with $i \neq k$.

Let $\mathcal{H}^+ = \cup \partial_+ C_i$ and $\mathcal{H}^- = \cup \partial_- C_i$. The components of $\mathcal{H}^+$ and $\mathcal{H}^-$ are called *thick* and *thin* surfaces respectively.

Note that in our definition each thick surface is separating in the complement of the thin surfaces.

2.6. Generalized stabilizations, perturbations and removable paths

Several geometric operations can be used to produce new Γ-bridge surfaces from old ones. These are generalizations of stabilizations for Heegaard splittings of manifolds and usually we work with bridge surfaces that are not obtained from others via these operations. A more detailed discussion of these operations can be found in [10, 17, 20]. A Γ-bridge surface H for (M, T, Γ) will be called *stabilized* if there is a pair of T-compressing disks on opposite sides of H that intersect in a single point. The Γ-bridge surface is *meridionally stabilized* if there is a (T, Γ)-cut-disk and a T-compressing disk on opposite sides of H that intersect in a single point.

As we are considering manifolds with boundary there are two other geometric operations that can be used to obtain a new bridge surface from an old one. Suppose H is a Γ-Heegaard splitting for (M, T, Γ) decomposing M into compressionbodies C_1 and C_2. Let F be a component of $\partial_- C_1 \subset \partial M$ and let T' be a collection of vertical edges in $F \times [-1, 0]$ so that $T' \cap (F \times \{0\}) = T \cap F$. Let H' be a minimal genus Heegaard surface for $(F \times [-1, 0], T')$ which does not separate $F \times \{-1\}$ and $F \times \{0\}$ and which intersects each edge in T' exactly twice. H' can be formed by tubing two parallel copies of F along a vertical arc not in T'. We can form a Γ-Heegaard surface H'' for $M \cup (F \times [-1, 0])$ by *amalgamating* H and H'. This is simply the usual notion of amalgamation of Heegaard splittings (see [19]). In fact, H'' is a Γ-Heegaard surface for $(M \cup (F \times [-1, 0]), T \cup T')$.

Since $(M \cup (F \times [-1, 0]), T \cup T')$ is homeomorphic to (M, T), we may consider H'' to be a Γ-Heegaard surface for (M, T, Γ). H'' is called a *boundary stabilization* of H. A similar construction can be used to obtain a new Γ-Heegaard splitting of (M, T, Γ) by tubing two parallel copies of F along a vertical arc that does lie in $T' \cap \Gamma$. In this case H'' will be called *meridionally boundary stabilized*.

If a Γ-bridge surface is stabilized, boundary stabilized, meridionally stabilized, or meridionally boundary stabilized we will say that it contains a *generalized stabilization*.

A Γ-bridge surface is called *cancellable* if there is a pair of bridge disks D_i on opposite sides of H such that $\emptyset \neq (\partial D_1 \cap \partial D_2) \subset (H \cap T)$. If $|\partial D_1 \cap \partial D_2| = 1$ we will call the bridge surface *perturbed*. Unlike the case when T is a 1-manifold, a perturbed bridge surface cannot necessarily be unperturbed by an isotopy as that may result in edges that are no longer trivially embedded. The pair $\{D_1, D_2\}$ is called the *associated cancelling pair of disks*.

Suppose that $\zeta \subset T$ is a 1-manifold which is the the union of edges in T (possibly a closed loop with zero or more vertices of T). We say that ζ is a *removable path* if the following hold:

(1) Either the endpoints of ζ lie in ∂M or ζ is a cycle in T.

(2) ζ intersects H exactly twice,

(3) If ζ is a cycle, there exists a cancelling pair of disks $\{D_1, D_2\}$ for ζ with $D_j \subset C_j$. Furthermore, there exists a compressing disk E for H such that $|E \cap T| = 1$ and if $E \subset C_j$ then $|\partial E \cap \partial D_{j+1}| = 1$ (indices run mod 2) and E is otherwise disjoint from a complete collection of bridge disks for $T - H$ containing $D_1 \cup D_2$.

(4) If the endpoints of ζ lie in ∂M, there exists a bridge disk D for the bridge arc component of $\zeta - H$ such that $D - T$ is disjoint from a complete collection of bridge disks Δ for $T - H$. Furthermore, there exists a compressing disk E for H on the opposite side of H from D such that $|E \cap D| = 1$ and E is disjoint from Δ.

See figure 2 for an example of a removable path.

If T has a removable path ζ, ζ can be isotoped to lie in a spine for one of the compressionbodies C_1 or C_2. A condition weaker than being removable is the following: A circle or edge component γ of T is *isotopically core* if it can be isotoped in $M - (T - \gamma)$ to lie in the 1-dimensional portion of a spine of one of the compressionbodies C_1 or C_2. (Recall that the spine of a compressionbody is the union of the negative boundary of the compressionbody with a 1–complex such that the exterior of the spine is a product region.)
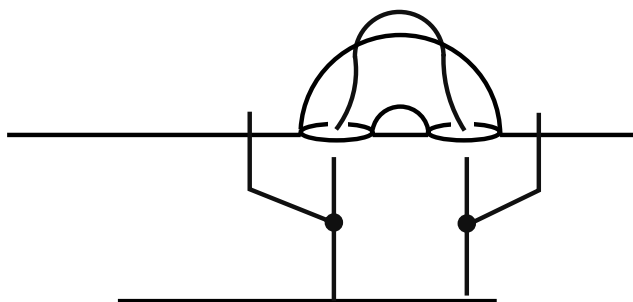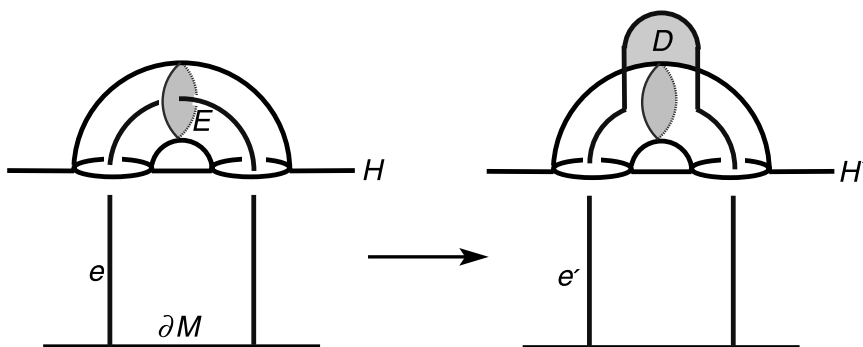
Figure 2: An example of a removable path.

Figure 3: The conversion of a ghost edge into a removable edge.

2.7. Γ-c-Heegaard surfaces and removable edges

This section presents a technical result that sometimes allows a Γ-c-Heegaard surface to be converted into a Heegaard surface with removable edges.

Suppose that H is a Γ-Heegaard surface for (M, T, Γ) and that $e \subset \Gamma$ is an edge disjoint from H with both endpoints in ∂M. Let E be a cut disk intersecting e whose boundary is in H. By isotoping e so that $e \cap E$ moves through ∂E to the other side, we convert e into a removable path e'. Let T' be the new graph. Let H' be the new Γ-Heegaard surface for (M, T'). Let D be the bridge disk for the bridge arc component of $e' - H'$. See figure 3.

Lemma 2.3. *If H' is stabilized or meridionally stabilized then so is H. If H' is boundary-stabilized or meridionally boundary stabilized, so is H and the stabilization is along the same component of ∂M. If H' is perturbed then so is H. If T' contains a removable path other than e' then either that path is removable in T or H is meridionally stabilized.*

Proof. **Case 1:** Suppose that H' is stabilized or meridionally stabilized by disks D_1 and D_2 on opposite sides of H' which intersect once. Out of all such pairs of stabilizing or meridionally stabilizing pairs, choose D_1 and D_2 so that $|(D_1 \cup D_2) \cap (D \cup E)|$ is minimal. The next claim shows that (D_1, D_2) also stabilizes or meridionally stabilizes H.

    Claim: $|(D_1 \cup D_2) \cap e| \leq 1$.

    Since (D_1, D_2) (meridionally) stabilize H', $|(D_1 \cup D_2) \cap e'| \leq 1$. Without loss of generality, choose the labelling so that D_1 is on the same side of H as D. Then $D_1 \cap e = \varnothing$, and so $(D_1 \cup D_2) \cap e = D_2 \cap e$. If D_2 is disjoint from and not parallel to E, then clearly $|e \cap D_2| \leq 1$. If D_2 is parallel to E, then $|D_2 \cap e| = |E \cap e| = 1$. If D_2 is not disjoint from E, then by the minimality of $|(D_1 \cup D_2) \cap (E \cup D)|$ it follows that $D_2 \cap E$ is a collection of arcs. After possibly a small isotopy, we may assume that $D \cap E$ is disjoint from $D_2 \cap E$. Then using D to isotope e' back to e guarantees that e is disjoint from D_2. $\square$(Claim)

    Case 2: Suppose that H' is boundary stabilized or meridionally boundary stabilized. In this case there exists a Γ-c-disk D' for H such that compressing H' along D' produces surfaces H_1 and H_2 and the surface H_2 bounds a product region with ∂M containing only vertical arcs of T while the surface H_1 is a Γ-Heegaard surface for (M, T, Γ).

    If D' is disjoint from or parallel to E then it is clear that H is boundary stabilized or meridionally boundary stabilized. Suppose, therefore that D' intersects E. We may assume that D' was chosen so as to minimize $D' \cap E$. This implies that $D' \cap E$ is a non-empty collection of arcs. Then, as in Case 1, isotoping e' back to e shows that $e \cap D' = \varnothing$. Hence, H is boundary stabilized or meridionally boundary stabilized and the stabilization is along the same component as that of H'.

    Case 3: Suppose that (D_1, D_2) are a perturbing pair of disks for H' such that D_1 is on the same side of H as D. Notice that $\partial D_1 \cup \partial D_2$ is disjoint from e' since two components of $e' - H'$ are vertical arcs in the compressionbody containing them. Thus, unless $e \cap D_2 \neq \varnothing$, (D_1, D_2) is a perturbing pair for H. If $D_2 \cap e \neq \varnothing$, then D_2 intersects the neighbourhood of E used to push e to e'. An argument similar to that of Cases 1 and 2 shows that ∂D can be assumed to be disjoint from ∂D_2, and so e is disjoint from D_2, as desired.

    Case 4: Suppose that T' contains a removable path $\zeta \neq e$. An argument similar to the previous cases shows that ζ is a removable path in T, unless ζ is not a cycle and the compressing disk E from condition (4) of the definition

of removable path is equal to the present disk E. Suppose, therefore, that this is the case. By the definition of removable path, there is a bridge disk D' for $e' - H'$ which is disjoint from E. A small isotopy of $D \cup D'$ creates a compressing disk E' for $H' - T'$ intersected once by E. The pair (E, E') therefore, shows that H' is stabilized. Since $\partial E'$ is non-separating on H' and therefore on H, E' is a compressing disk for H disjoint from T. The boundary of E' intersects E exactly once and $E \cap T = E \cap e$ and so H is meridionally stabilized. $\square$

3. Properties of compressionbodies containing properly embedded graphs

In this section, we will generalize many of the well-known results for compressionbodies to the case when the compressionbody contains a graph embedded in a specific way.

Lemma 3.1. *Suppose C is a Γ-compressionbody containing a trivially embedded graph $T - \Gamma$. Compressing or cut-compressing C results in a union of Γ-compressionbodies each containing a trivially embedded graph.*

Proof. Given any Γ-c-disk D^c for (C, T, Γ) we can always find a collection of pairwise disjoint Γ-c-disks $\mathcal{D}^c$ containing D^c so that (C, T, Γ) c-compressed along $\mathcal{D}^c$ is a collection 3-balls and components homeomorphic to $G \times I$ where G is a component of $\partial_- C$. Both types of components may contain trivially embedded graphs. The result follows. $\square$

The next lemma shows that the negative boundary of a Γ-compressionbody is Γ-c-incompressible.

Lemma 3.2. *Let (C, T, Γ) be a Γ-compressionbody containing a trivially embedded graph $T - \Gamma$. Assume that no component of $\partial_- C$ is a 2–sphere intersecting T one or fewer times. Then $C - T$ is irreducible, no sphere in C intersects T exactly once, and $\partial_- C$ is Γ-c-incompressible.*

Proof. Suppose that $S \subset C$ is either a sphere intersecting T transversally one or fewer times and which does not bound a ball in $C - T$ or a Γ-compressing disk for $\partial_- C$. An innermost disk argument shows that we may choose S so that it is disjoint from all the Γ-c-disks for $\partial_+ C$. Let Δ be a maximal collection of Γ-c-disks for $\partial_+ C$, chosen so that each disk of Δ intersecting T intersects a unique ghost arc of T exactly once and so that each ghost arc of T intersects Δ exactly once. Boundary-reducing C along Δ creates a

Γ-compressionbody (C', T', Γ') that is possibly disconnected. C' is the union of trivial compressionbodies and $S \subset C'$. No component of T' is a ghost arc since every ghost arc intersects a disk in Δ. In a trivial compressionbody, there are no essential spheres or disks. Hence, if S is a sphere, it either bounds a 3–ball in C' or is boundary parallel in C'.

If S is a sphere bounding a 3-ball in C', then since every component of T' is adjacent to $\partial_+ C'$, S bounds a 3-ball in $C' - T'$. Hence if S bounds a 3-ball in C' it also bounds one in $C - T$, a contradiction. Thus, if S is a sphere, then it is in a component C'' of C' homeomorphic to $S^2 \times [0,1]$ and is boundary-parallel. At least two components of $T \cap C''$ are adjacent to $\partial_- C \cap C''$ and these components are also adjacent to $\partial C'' - \partial_- C$. Hence these components also intersect S and so S intersects T twice, a contradiction.

If S is a disk, then ∂S bounds a disk $E \subset \partial_- C = \partial_- C'$. If S is a compressing disk, then E is adjacent to at least one component of T' and if S is a cut disk, then E is adjacent to at least two components of T'. The union $S \cup E$ is a sphere in C' which either bounds a 3–ball in C' or is boundary-parallel. If the sphere bounds a 3-ball, then each intersection point of $T' \cap E$ is joined in T' to an intersection point of $T' \cap S$, implying that S intersects T' at least once if it is a compressing disk and at least twice if it is a cut disk. This contradicts the definition of compressing disk and cut disk, respectively. Suppose the sphere is parallel to a boundary component P of C'. If $P \subset \partial_+ C'$, then E is a subset of a sphere component P' of $\partial_- C'$. The disk $E' = P - E$ must intersect T' once or twice (corresponding to S being a compressing or cut disk). Since the positive boundary of a connected compressionbody is connected, the sphere $S \cup E'$ bounds a 3-ball in C'. This gives rise to a contradiction as before. Hence, $P \subset \partial_- C'$. Since no component of T' is adjacent to two components of $\partial_- C'$, each intersection point of $T' \cap E$ corresponds to an intersection point of $T' \cap S$. Hence, D can be neither a compressing or a cut disk, a contradiction. $\square$

The next lemma does not concern bridge surfaces *per se*, but will be useful in the proof of the main theorem when we try to show that the thin surfaces are not T-parallel. We use it to guarantee that if there is a thin surface that is not T-parallel, then we can find an innermost such surface.

Lemma 3.3. *Suppose that C is a compressionbody with $F = \partial_+ C$. Let $T \subset C$ be a properly embedded graph such that $F - \mathring{\eta}(T)$ is parallel to $\partial(C - \mathring{\eta}(T)) - F$. Let $G \subset C$ be a T-incompressible surface disjoint from F and transverse to T. Then each component of $G - \mathring{\eta}(T)$ is parallel to $\partial(C - \mathring{\eta}(T)) - F$.*

Proof. Let $N = C - \mathring{\eta}(T)$. Then N is homeomorphic to $(F - \mathring{\eta}(T)) \times I$ with the homeomorphism taking $F - \mathring{\eta}(T)$ to $(F - \mathring{\eta}(T)) \times \{0\}$. The boundary of $G - \mathring{\eta}(T)$ is contained in $(F - \mathring{\eta}(T)) \times \{1\}$. Since G is T-incompressible, by [24, Corollary 3.2], each component of $G - \mathring{\eta}(T)$ is parallel to $(F - \mathring{\eta}(T)) \times \{1\}$. Hence, each component of $G - \mathring{\eta}(T)$ is parallel to $\partial(C - \mathring{\eta}(T)) - F$. $\square$

Finally we recall the classification of Heegaard splittings of pairs (M, T) where M is a compressionbody and $\partial_+ M - \mathring{\eta}(T)$ is parallel to $\partial(M - \mathring{\eta}(T)) - \partial_+ M$, [20]. This result is key to proving our main theorem.

Theorem 3.4. [Theorem 3.1, 20] *Let M be a compressionbody and T be a properly embedded graph so that $\partial_+ M - \mathring{\eta}(T)$ is parallel to $\partial(M - \mathring{\eta}(T)) - \partial_+ M$. Let H be a Heegaard surface for (M, T). Assume that T contains at least one edge. Then one of the following occurs:*

(1) H is stabilized;

(2) H is boundary stabilized;

(3) H is perturbed;

(4) T has a removable path disjoint from $\partial_+ M$;

(5) M is a 3–ball, T is a tree with a single interior vertex (possibly of valence 2), and $H - \mathring{\eta}(T)$ is parallel to $\partial M - \mathring{\eta}(T)$ in $M - \mathring{\eta}(T)$;

(6) $M = \partial_- M \times I$, H is isotopic in $M - \mathring{\eta}(T)$ to $\partial_+ M - \mathring{\eta}(T)$.

For our purposes we need to strengthen the second conclusion of the above theorem:

Theorem 3.5. *Let M be a compressionbody and T be a properly embedded graph so that $\partial_+ M - \mathring{\eta}(T)$ is parallel to $\partial(M - \mathring{\eta}(T)) - \partial_+ M$. Let H be a Heegaard surface for (M, T). Assume that T contains at least one edge. Then one of the following occurs:*

(1) H is stabilized;

(2) H is boundary stabilized along $\partial_- M$;

(3) H is perturbed;

(4) T has a removable path disjoint from $\partial_+ M$;

(5) M is a 3–ball, T is a tree with a single interior vertex (possibly of valence 2), and $H - \mathring{\eta}(T)$ is parallel to $\partial M - \mathring{\eta}(T)$ in $M - \mathring{\eta}(T)$;

(6) $M = \partial_- M \times I$, H is isotopic in $M - \mathring{\eta}(T)$ to $\partial_+ M - \mathring{\eta}(T)$.

Proof. Suppose H is boundary stabilized along $\partial_+ M$. Then H is obtained by amalgamating a minimal genus Heegaard surface for $\partial_+ M \times [-1, 0]$ which does not separate $\partial_+ M \times \{-1\}$ and $\partial_+ M \times \{0\}$ and which intersects each edge in $T \cap (\partial_+ M \times [-1, 0])$ exactly twice, together with a Heegaard surface $\tilde{H}$ for M. Without loss of generality we will assume that H is obtained from $\tilde{H}$ after a single boundary stabilization along $\partial_+ M$.

By Theorem 3.4, $\tilde{H}$ satisfies one of six possible conclusions. If $\tilde{H}$ is stabilized, perturbed or if T has a removable path disjoint from $\partial_+ M$ then the same is true for H as boundary stabilizations preserve all of these properties. If $\tilde{H}$ is boundary stabilized, the stabilization must be along $\partial_- M$, and so the same is true for H.

Suppose that M is a 3–ball, T is a tree with a single interior vertex, and $H - \mathring{\eta}(T)$ is parallel to $\partial M - \mathring{\eta}(T)$ in $M - \mathring{\eta}(T)$. Let A and B be the components of $M - \tilde{H}$ so that B is a ball and let κ be one of the edges of $T \cap M$. Then H can be recovered from $\tilde{H}$ by tubing H to the boundary of a collar of ∂M along a vertical tube τ in A. We can choose τ to be arbitrarily close to $\kappa \cap A$; in particular, we may assume that the disk of parallelism between τ and κ intersects some bridge disk that contains $\kappa \cap B$ in its boundary only in the point $\kappa \cap \tilde{H}$. We conclude that H is perturbed.

Suppose then that $M = \partial_- M \times I$ and H is isotopic in $M - \mathring{\eta}(T)$ to $\partial_+ M - \mathring{\eta}(T)$. Let A and B be the components of $M - \tilde{H}$ so that A contains $\partial_+ M$. The argument in this case is identical to the one above as long as there is at least one bridge disk in B. If $T \cap B$ is a product, then $T \cap M$ is a collection of vertical arcs and thus H can be obtained from the Heegaard surface $\tilde{H}$ by stabilizing along $\partial_- M$. $\qquad\square$

4. Haken's lemma

Let H be a Γ-bridge surface for (M, T, Γ) and suppose D is a T-compressing disk for some component G of ∂M. It is a classic result of Haken [8] that in the case $T = \emptyset$, there is a compressing disk D' for G so that D' intersects H in a unique essential curve. This result was extended to the case where T is a 1-manifold and $\Gamma = \emptyset$ in [11] and to the case where T is a 1-manifold and $\Gamma = T$ in [22]. This result implies that, in most cases, if ∂M has a compressing disk, then H is weakly reducible.

Here we will use a generalization of this result to cut-disks. We are grateful to Jesse Johnson for showing us a proof of the following theorem.

Theorem 4.1. *Suppose M is a compact orientable manifold, T is a properly embedded graph in M and Γ is a subgraph of T. Assume that $M - T$ is irreducible and that no sphere in M intersects T exactly once transversally. Suppose H is a Γ-Heegaard splitting for (M, T, Γ) such that there is a Γ-c-disk for H on each side. If there exists a Γ-c-disk for some component of ∂M then H is Γ-c-weakly reducible.*

Proof. Let D^* be a Γ-c disk for ∂M.

By Lemma 3.2, the negative boundary of a Γ-c-compressionbody is (T, Γ)-c-incomressible so D^* must intersect H. Let D_a and D_b be c-disks for H above and below it. Let H_a be isotopic to H but with the c-disk D_a shrunk to a small disk neighbourhood of a point. In particular in H_a we may assume that D_a is disjoint from D^*. Furthermore we may assume that H_a is isotoped so that no curve of $D^* \cap H_a$ bounds a disk in $H_a - T$. As H_a and D^* must intersect, an innermost in D^* circle of intersection bounds a Γ-c-disk Δ_a for H_a. The disk Δ_a is disjoint from D_a and therefore either we are done, or Δ_a is on the same side of H_a as D_a. Similarly there is a surface H_b isotopic to H so that an innermost curve of intersection between H_b and D^* bounds a c-disk for H_b on the same side as D_b.

Now consider the isotopy ρ between H_a and H_b. The surface $H_a = \rho^{-1}(-1)$, intersects D^* in a simple closed curve which gives a Γ c-disk for H on one side and the surface $H_b = \rho^{-1}(1)$, intersects D^* in a simple closed curve which gives a Γ c-disk for H on the opposite side. As $\rho^{-1}(t)$ always intersects D^* in at least one essential curve there are two cases to consider:

(1) There is some regular value r where $\rho^{-1}(r)$ intersects D^* in two disjoint essential curves σ_a and σ_b so that the interiors of the disks these curves bound on D^* intersect $\rho^{-1}(r)$ only in inessential curves and near their boundaries these disks lie on opposite sides of H.

(2) There is a critical point c so that $\rho^{-1}(c - \epsilon) \cap D^*$ contains an essential curve with the properties of σ_a and $\rho^{-1}(c + \epsilon) \cap D^*$ contains an essential curve with the properties of σ_b.

As the curves of intersection in $\rho^{-1}(c - \epsilon) \cap D^*$ are disjoint from the curves $\rho^{-1}(c + \epsilon) \cap D^*$ in either case after an isotopy of D^* to remove any inessential circles of intersection we can obtain c-disks D_a and D_b with disjoint boundaries σ_a and σ_b that lie on opposite sides of H as desired. $\square$

5. Multiple Γ-bridge splittings of (M, T, Γ)

5.1. Complexity

We define a complexity on the thick surfaces $\mathcal{H}^+$ of a multiple Γ-Heegaard splitting of (M, T, Γ).

Definition 5.1. Let X be a set with an order $\leq$. Let Y and Z be two finite multisets of elements of X. Write $Y = (y_1, y_2, \ldots, y_n)$ and $Z = (z_1, z_2, \ldots, z_m)$ so that for all i, $y_i \geq y_{i+1}$ and $z_i \geq z_{i+1}$. We say that $Y < Z$ if and only if one of the following occurs:

- There exists $j \leq \min(n, m)$ so that for all $i < j$, $y_i = z_i$ and $y_j < z_j$.

- $n < m$ and for all $i \leq n$, $y_i = z_i$.

Definition 5.2. Let S be a closed connected surface embedded in M transverse to a properly embedded graph $T \subset M$. The complexity of S is $c(S) = 4g(S) + |S \cap T|$.

If $\mathcal{H}$ is a Γ-multiple bridge splitting for (M, T, Γ), let the complexity of $\mathcal{H}$, denoted $c(\mathcal{H})$, be the multiset $\{c(S) | S \in \mathcal{H}^+\}$. If $\mathcal{H}$ and $\mathcal{H}'$ are two multiple Γ-Heegaard splittings for (M, T, Γ), their complexities will be compared as in Definition 5.1.

The next lemma is immediate from the definition of complexity.

Lemma 5.3. *Suppose that $\mathcal{H}$ is a multiple Γ-Heegaard splitting of (M, T, Γ) and suppose that $\mathcal{J}$ is a multiple Γ-Heegaard splitting of (M, T, Γ) such that $\mathcal{J}^+$ is a proper subset of $\mathcal{H}^+$. Then the complexity of $\mathcal{J}$ is strictly less than the complexity of $\mathcal{H}$.*

Next, we will show that c-compressing a surface always decreases its complexity

Lemma 5.4. *Suppose S is meridional surface in (M, T) of non-positive euler characteristic. If S' is a component of the surface obtained from S by compressing along a c-disk, then $c(S) > c(S')$.*

Proof. Suppose S' is obtained from S after a single c-compression along a c-disk D^*. If D^* is non-separating then $g(S') = g(S) - 1$ and $|S \cap T| \geq |S' \cap T| - 2$ so $c(S) \geq c(S') - 2$ and thus the complexity has strictly decreased. If D^* is separating then S' has two components, S'_1 and S'_2, and

$g(S') = g(S'_1) + g(S'_1)$ and $|S \cap T| \geq |S'_1 \cap T| + |S'_2 \cap T| - 2$. As D^* is a c-disk each of S'_1 and S'_2 must either have positive genus or at least 4 punctures. Therefore each of S'_1 and S'_2 either has genus strictly less than $g(S)$ or has at least 2 fewer punctures than S. In either case we can conclude that $c(S) > c(S'_1)$ and $c(S) > c(S'_2)$ as desired.   □

Remark 5.5. In place of $c(S)$, we can use any complexity function that strictly decreases when a surface is c-compressed or when the number of intersections between the graph and the surface is decreased. For example, Tomova's complexity in [22] can also be used.

6. Untelescoping

We are interested in obtaining a multiple Γ-Heegaard splitting of (M, T, Γ) with the property that every thin surface is Γ-c-incompressible. The following lemma will be useful and follows directly from Theorem 4.1.

Lemma 6.1. *Suppose that $\mathcal{H}$ is a multiple Γ-Heegaard splitting of (M, T, Γ). If some component of $\mathcal{H}^-$ is Γ-c-compressible in M then one of the following occurs:*

- *Some component of $\mathcal{H}^+$ is a (T, Γ)-c-weakly reducible Γ-Heegaard splitting of the component of $M - \mathcal{H}^-$ containing it; or*

- *Some component of $M - \mathcal{H}$ is a product compressionbody with one boundary component in $\mathcal{H}^-$ and the other in $\mathcal{H}^+$. The graph T intersects this component in vertical arcs and vertical pods with pod handles in $T - \Gamma$.*

In the first case one can reduce the complexity of the splitting by c-compressing some thick surface along a pair of disjoint c-disks on opposite sides. This operation is called *untelescoping* and was first introduced by Scharlemann and Thompson in [16] in the context of 3-manifolds. The concept was generalized to weakly reducible bridge surfaces for a manifold containing a properly embedded tangle in [11] and then to (τ, τ)-c-weakly reducible τ-bridge surfaces for a manifold containing a properly embedded tangle τ in [22]. In the following definition, we extend the construction to (T, Γ)-c-weakly reducible Γ-bridge surfaces for (M, T, Γ) where T is a properly embedded graph.

Let H be a (T, Γ)-c-weakly reducible Γ-bridge splitting of (M, T, Γ) and let $\mathcal{D}_1$ and $\mathcal{D}_2$ be collections of pairwise disjoint Γ-c-disks above and below

H such that $\mathcal{D}_1 \cap \mathcal{D}_2 = \emptyset$. Then we can obtain a multiple Γ-bridge splitting for (M, T) with one thin surface obtained from H by (T, Γ)-c-compressing it along $\mathcal{D}_1 \cup \mathcal{D}_2$ and two thick surfaces, obtained by (T, Γ)-c-compressing H along $\mathcal{D}_1$ and $\mathcal{D}_2$ respectively. Although this operation is similar to the one described in [22] we include the details here with the modifications required.

Let $(A, A \cap T)$ and $(B, B \cap T)$ be the two Γ-compressionbodies into which H decomposes (M, T, Γ) and let $\mathcal{D}_A \subset A$ and $\mathcal{D}_B \subset B$ be collections of pairwise disjoint Γ-c-disks such that $\mathcal{D}_A \cap \mathcal{D}_B = \emptyset$. Let $A' = A - \mathring{\eta}(\mathcal{D}_A)$ and $B' = B - \mathring{\eta}(\mathcal{D}_B)$. Then by Lemma 3.1 A' and B' are each the disjoint union of Γ-compressionbodies containing trivial graphs $A' \cap T$ and $B' \cap T$, respectively.

Take small collars $\eta(\partial_+ A')$ of $\partial_+ A'$ and $\eta(\partial_+ B')$ of $\partial_+ B'$. Let $C^1 = cl(A' - \eta(\partial_+ A'))$, $C^2 = \eta(\partial_+ A') \cup \eta(\mathcal{D}_B)$, $C^3 = \eta(\partial_+ B') \cup \eta(\mathcal{D}_A)$ and $C^4 = cl(B' - \eta(\partial_+ B'))$. Note that C_1 and C_4 are Γ-compressionbodies containing trivial graphs because they are homeomorphic to A' and B' respectively. C_2 and C_3 are obtained by taking surface $\times I$ containing vertical arcs and attaching 2-handles, some of which may contain segments of Γ as their cores. Therefore C_2 and C_3 are also Γ-compressionbodies containing trivial graphs. We conclude that we have obtained a multiple Γ-bridge splitting $\mathcal{H}$ of (M, T, Γ) with thick surfaces $\partial_+ C_1$ and $\partial_+ C_2$ that can be obtained from H by Γ-c-compressing along $\mathcal{D}_A$ and $\mathcal{D}_B$ respectively and a thin surface $\partial_- C_2 = \partial_- C_3$ obtained from H by Γ-c-compressing along both sets of c-disks. We say that $\mathcal{H}$ is obtained by untelescoping H using Γ-c-disks. The next remark follows directly from Lemma 5.4.

Remark 6.2. Suppose $\mathcal{H}'$ is a multiple Γ-bridge splitting of (M, T, Γ) obtained from another multiple Γ-bridge splitting $\mathcal{H}$ of (M, T, Γ) via untelescoping. Then $c(\mathcal{H}') < c(\mathcal{H})$.

Suppose that H is a multiple Γ-Heegaard splitting for (M, T, Γ) with a thick surface that can be untelescoped to become a multiple Γ-Heegaard splitting $\mathcal{H}$ for (M, T, Γ). It is natural to ask: If $\mathcal{H}$ contains a generalized stabilization must H contain a generalized stabilization? If $\mathcal{H}$ contains a perturbed thick surface must H also be perturbed? If T has a path which is removable with respect to $\mathcal{H}$, is that path removable with respect to H? The answer is positive in all cases.

Lemma 6.3. *Let H be a (multiple) Γ-Heegaard splitting for the triple (M, T, Γ) obtained by amalgamating the multiple Γ-Heegaard splitting $\mathcal{H}$ for (M, T, Γ). (We assume that the amalgamation actually produces a (multiple)*

Γ-*Heegaard splitting which is not necessarily the case). Then the follow-ing hold*

(1) *Suppose the Γ-Heegaard splitting induced by $\mathcal{H}^+$ on some component of $(M, T, \Gamma) - \mathcal{H}^-$ contains a generalized stabilization. Furthermore sup-pose that if the generalized stabilization is a (meridional) boundary sta-bilization, then it is along a component of ∂M. Then H contains a gen-eralized stabilization of the same type and if the generalized stabilization is a (meridional) boundary stabilization then it must also be along a component of ∂M.*

(2) *If the Γ-Heegaard splitting induced by $\mathcal{H}^+$ on some component of $(M, T, \Gamma) - \mathcal{H}^-$ is perturbed, then so is H.*

(3) *If the Γ-Heegaard splitting induced by $\mathcal{H}^+$ on some component of $(M, T, \Gamma) - \mathcal{H}^-$ contains a removable path so that the path is either a cycle or both of its endpoints are contained in ∂M, then so does H.*

(4) *If T contains a circle or edge component γ that is isotopic into the spine of one of the components of $M - \mathcal{H}$, then γ is isotopic into the spine of one of the compressionbodies of $M - H$.*

Proof. Without the loss of much generality, we may assume that H is obtained from $\mathcal{H}$ by a single amalgamation (i.e., we can obtain $\mathcal{H}$ from H by untelescoping in one step). The last observation is the easiest to deal with: as H is obtained from $\mathcal{H}$ by an amalgamation, the spine of a compressionbody after amalgamation can be obtained by gluing together the 1–complexes forming the spines of two compressionbodies before amal-gamation. Thus, a loop in one of those spines remains in a spine after amalgamation.

We now proceed to prove the other observations. Suppose that H decom-poses (M, T) into compressionbodies A and B. Let N_i for $i = 1, 2$ be the closure of the components of $M - \mathcal{H}^-$ on either side of $\mathcal{H}^-$. Let H_i be the Γ-Heegaard surface of N_i induced by $\mathcal{H}^+$. Assume that H_1 contains a gen-eralized stabilization, or is perturbed, or that $T \cap N_1$ contains a removable path satisfying the hypotheses of (3).

Consider a collar $\mathcal{H}^- \times [-1, 1]$ of $\mathcal{H}^-$ where $\mathcal{H}^- = \mathcal{H}^- \times \{0\}$. Recall that H_1 can be obtained from $\mathcal{H}^- \times [0, 1]$ by attaching handles (possi-bly with cores running along the knot) to $\mathcal{H}^- \times \{1\}$. Similarly, H_2 can be obtained from $\mathcal{H}^- \times [-1, 0]$ by attaching handles to $\mathcal{H}^- \times \{-1\}$. The Heegaard surface H can be obtained by extending the handles of H_1 through $\mathcal{H}^- \times [-1, 1]$ and attaching them to H_2.
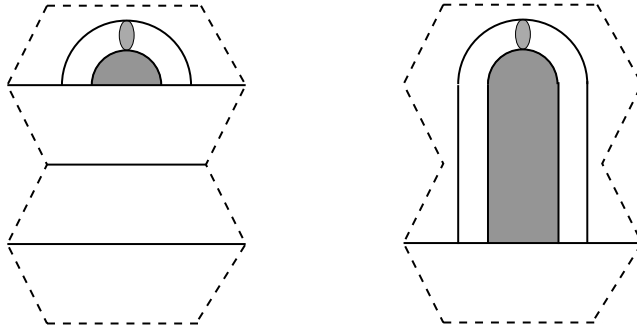
Figure 4: If H_1 is stabilized, so is H.

Case 1a: H_1 is stabilized or meridionally stabilized. Let D and E be the (T, Γ)-c-disks with boundary on H_1 defining the (meridional) stabilization. We may assume that the handles attached to $\mathcal{H}^- \times \{1\}$ include one that has ∂D as a core. The intersection of ∂E with $\mathcal{H}^- \times \{1\}$ then consists of a single arc. The possibly once punctured disk $E' = E \cup (\partial E \times [-1, 1])$ is then a c-disk with boundary on H. The c-disks E' and D intersect exactly once and define a (meridional) stabilization of H. See figure 4.

Case 1b: H_1 is boundary stabilized or meridionally boundary stabilized. Let C_1^+ and C_2^+ be the two Γ-compressionbodies into which H_1 divides N_1. Without loss of generality in this case C_1^+ contains some component $G \subset \partial N$ and a c-disk D so that removing $\eta(D)$ from C_1^+ decomposes it into a compressionbody $C_1'^+$ and a component R homeomorphic to $G \times I$ and adding R to C_2^+ results in a compressionbody $C_2'^+$. Let $\tilde{H}_1 = \partial_+ C_1'^+ = \partial_+ C_2'^+$. Amalgamating the multiple bridge splitting with thick surfaces $\tilde{H}_1$ and H_1 gives a bridge splitting for N with bridge surface $\tilde{H}$ which can be obtained from H by c-compressing along the disk D. Therefore H is boundary stabilized or meridionally boundary stabilized.

Case 2: H_1 is perturbed. Let D and E be the two bridge disks for H_1 that intersect in one or two points and these points lie in T. Both disks are completely disjoint from $\mathcal{H}^- \times \{1\}$ so in particular extending the 1-handles from H_1 across $\mathcal{H}^- \times [-1, 1]$ has no effect on these disks. Therefore H is also perturbed.

Case 3: Suppose ζ is a removable path in N_1 and suppose first that ζ is a cycle. Let C_1 and C_2 be the two Γ-compressionbodies $\mathrm{cl}(N_1 - H_1)$. As ζ is a removable cycle there exists a cancelling pair of disks $\{D_1, D_2\}$ for ζ with $D_j \subset C_j$. These disks are both completely disjoint from $\mathcal{H}^- \times \{1\}$ so in
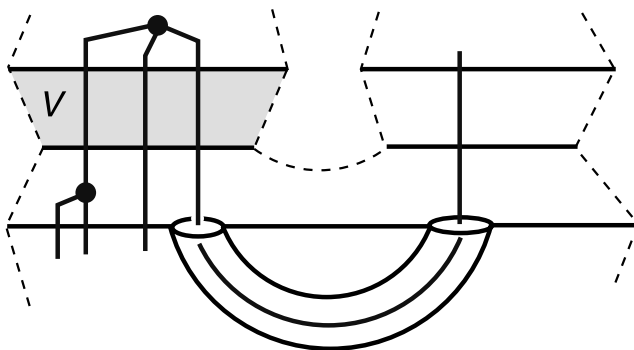
Figure 5: Removing the parallel thick and thin surfaces that form the boundary of V destroys bridge position. Note the presence of a pod handle and a ghost arc.

particular extending the 1-handles from H_1 across $\mathcal{H}^- \times [-1, 1]$ has no effect on these disks. As ζ is a removable cycle there exists a compressing disk E in C_1, say, so that $|E \cap T| = 1$ and $|\partial E \cap \partial D_2| = 1$ and E is otherwise disjoint from a complete collection of bridge disks for $(T \cap N_1) - H_1$ containing $D_1 \cup D_2$. We may assume that the handles attached to $\mathcal{H}^- \times \{1\}$ include one that has ∂E as a core. Therefore E satisfies all the desired properties as a compressing disk for H.

Suppose then that ζ is a path. Again we may assume that the handles attached to $\mathcal{H}^- \times \{1\}$ include one that has ∂E as a core. The bridge disk for the component of $\zeta - H_1$ is also a bridge disk for $\zeta - H$ satisfying all of the required properties. $\qquad\square$

7. Consolidation

We now return to the second conclusion of Lemma 6.1. Scharlemann and Thompson in [16] encounter a similar situation and are able to simply remove the two parallel surfaces from the collection of surfaces comprising the multiple Heegaard splitting. In the context of a 3-manifold without an embedded graph this results in a new multiple Heegaard splitting. When there is a graph present though, it is possible that after the removal of the two surfaces the collection of the remaining surfaces is no longer a multiple Γ-Heegaard surface; see figure 5 for an example. Thus, we will need to analyse this situation very carefully. We begin by introducing some additional definitions.

7.1. Parallelism and consolidation

Suppose that $\mathcal{H}$ is a multiple Γ-Heegaard splitting and that V is the closure of a component of $M - \mathcal{H}$ such that $(V, T \cap V)$ is homeomorphic to $(F \times I, \text{points} \times I)$ for a component $F \subset \partial V$. We allow the possibility that $T \cap V = \varnothing$. We say that V is a *product region* of $\mathcal{H}$. A product region is a *parallelism* of $\mathcal{H}$ if it is disjoint from ∂M. If $\mathcal{H}$ is a multiple Γ-Heegaard splitting with a parallelism V, we say that $\mathcal{H}' = \mathcal{H} - (\mathcal{H} \cap V)$ is obtained from $\mathcal{H}$ by *consolidating* V. We also say that $\mathcal{H}'$ is obtained from $\mathcal{H}$ by *consolidation*. In what follows we develop criteria to guarantee that $\mathcal{H}'$ is also a multiple Γ-Heegaard splitting.

If V is a parallelism, let $V_\pm$ be the closure of the components of $M - \mathcal{H}$ not in V that are adjacent to $\mathcal{H}^\pm \cap V$ respectively. Notice that each of $V_\pm$ has exactly one component. Let $T_\pm = T \cap V_\pm$.

We say that V satisfies the *consolidation criteria* if all of the following hold:

(C1) No circle component of T is contained entirely in $V_- \cup V \cup V_+$.

(C2) No edge of T with endpoints at vertices of T is contained entirely in $V_- \cup V \cup V_+$.

(C3) If $v \in T_+$ is a vertex adjacent to a handle, then there does not exist a ghost arc τ in T_- connected in $V_- \cup V \cup V_+$ to v.

(C4) If v is a vertex of T_+ then there is at most one edge of $T \cap (V_- \cup V \cup V_+)$ with one endpoint at v and one at $\partial_- V_+$.

Lemma 7.1. *Suppose that $\mathcal{H}$ is a multiple Γ-Heegaard splitting of (M, T, Γ) and that V is a parallelism of $\mathcal{H}$ satisfying the consolidation criteria. Let $\mathcal{H}'$ be obtained by consolidating V. Then $\mathcal{H}'$ is a multiple Γ-Heegaard splitting of (M, T).*

Proof. The multiple Γ-Heegaard splitting $\mathcal{H}$ defines a cobordism from $\partial_+ V_-$ to $\partial_- V_+$, with only 2-handles and 3-handles. We can attach all the 2-handles before all the 3-handles and so $\mathcal{H}'$ is a multiple Heegaard splitting of M.

Let $W = V_- \cup V \cup V_+$. Suppose that D is a disk in V_+ with ∂D the union of two arcs, one running along T and the other running along $\partial_+ V_+$. Using the product structure of V and the fact that T_- is trivially embedded in V_-, $\partial D \cap \partial_+ V_+$ can be extended through V and V_- to be a disk whose boundary is the union of an arc on T and an arc on $\partial_+ V_-$. Thus, by (C1) and (C2), $T \cap W$ is trivially embedded if and only if each vertex of $T \cap W$

is joined by at most a single edge to $\partial_- V_+$ and if any such edge is disjoint from a complete collection of Γ-c-disks in W for $\partial_+ V_-$. Conditions (C3) and (C4) ensure that this occurs. $\qquad \square$

Lemma 7.2. *Let V be a parallelism of $\mathcal{H}$ satisfying the consolidation criteria. Let $\mathcal{H}'$ be the multiple Γ-Heegaard splitting obtained by consolidating V. Then $W = V_- \cup V \cup V_+$ is a product region of $\mathcal{H}'$ if and only if V_- and V_+ are product regions of $\mathcal{H}$.*

Proof. The "if" direction is obvious, so we prove only the "only if" direction. Suppose that W is a product region of $\mathcal{H}'$ and let $W_\pm$ be the closure of the component of $M - \mathcal{H}'$ adjacent to $\mathcal{H}'^\pm \cap V$. Notice that W_- and W_+ are also closures of components of $M - \mathcal{H}$.

Let e be a component of $T \cap W$. Since W is a product region, e contains no vertices and joins $\partial_- V_+$ to $\partial_+ V_-$. Since those two surfaces are the boundary of W, they are parallel. Hence, there are no ghost arcs in V_- or V_+. Thus, e intersects each component of $\mathcal{H}^- \cap W$ exactly once. Since V is a product region, e intersects each component of $\mathcal{H} \cap W$ exactly once. Hence, e is the union of vertical arcs in V_-, V_+, and V, and so V_- and V_+ are product regions. $\qquad \square$

Finally, we note the following:

Lemma 7.3. *Suppose that $\mathcal{H}$ is a multiple Γ-c-Heegaard surface for (M, T, Γ). Suppose that no pod handle of T is adjacent to $\mathcal{H}^-$ and that $\mathcal{K}$ is obtained by consolidating a parallelism V of $\mathcal{H}$. If $T - \mathcal{K}$ has an edge joining a vertex of T to $\mathcal{K}^-$ or an edge joining two vertices of T, then $T \cap V_+$ has a vertex joined in T to a ghost arc of $T \cap V_-$.*

Proof. Suppose that e is an edge of $T - \mathcal{K}$ that either joins two vertices of T or that joins a vertex to $\mathcal{K}^-$. Since no pod handle of T is adjacent to $\mathcal{H}^-$, e must intersect V. It does so in vertical arcs since V is a parallelism. Each component of $T \cap V_-$ adjacent to $\partial_- V$ is either a vertical arc, a ghost arc, or a vertical pod. In fact, it must be a ghost arc or vertical arc since no pod handle of T is adjacent to a thin surface of $\mathcal{H}$. Each vertical arc in V_- has an endpoint on the thick surface $\mathcal{K}^+ \cap V_-$, so $e \cap V_-$ must contain a ghost arc and does not contain a vertex of T. Hence $T \cap V_+$ has a vertex joined in T to a ghost arc of $T \cap V_-$. $\qquad \square$

Corollary 7.4. *Suppose that $\mathcal{H}$ is a multiple Γ-c-Heegaard splitting of (M, T, Γ) and that V is a parallelism of $\mathcal{H}$ satisfying the following:*

- *No pod handle of T is adjacent to $\mathcal{H}^-$,*

- *No vertex of $T \cap V_+$ is adjacent in T to a ghost arc of $T \cap V_-$.*

Then V satisfies consolidation criteria (C2), (C3), and (C4).

Proof. (C3) is satisfied by hypothesis. Let $\mathcal{K}$ be the result of consolidating V. If (C2) were not satisfied, then we would have an edge of $T - \mathcal{K}$ joining two vertices of T, contradicting Lemma 7.3. If (C4) were not satisfied, then we would have an edge of $T - \mathcal{K}$ joining a vertex to a component of $\mathcal{K}^-$, also contradicting Lemma 7.3.   □

8. Combining untelescoping and consolidation

We will usually use consolidation in conjunction with untelescoping. When the untelescoping operation is performed along a single pair of Γ-c-disks both thin and thick surfaces are produced. There are two possibilities for a parallelism to arise: either one of the new thick surfaces is parallel to one of the new thin surfaces or a new thick surface is parallel to a thin surfaces that existed before the untelescoping. In order to consider these possibilities in more detail we introduce the following terminology:

Definition 8.1. Suppose that H is a Γ-c-weakly reducible Γ-c-Heegaard splitting of (M, T, Γ). A pair of disjoint Γ-c-disks D_1, D_2 on opposite sides of H are called *maximally separated* if whenever $\{i, j\} = \{1, 2\}$ the following is true:

   If E is a Γ-c-disk for H on the same side of H as D_i such that ∂D_i separates ∂E from ∂D_j, then ∂E is parallel in $H - T$ to ∂D_i.

   Notice that if D_1 and D_2 are both non-separating then they are automatically maximally separated, because the condition is vacuously true. It is clear that if H is Γ-c-weakly reducible then there exist a maximally separated pair of Γ-c-disks on opposite sides of H.

Lemma 8.2. *Let M be connected. Suppose that D_1 and D_2 are a pair of maximally separated Γ-c-disks for a Γ-c-Heegaard surface H for (M, T, Γ).*

Assume that Γ is disjoint from the vertices of $T - \partial T$. Let $\mathcal{H}$ be the result of untelescoping H using $D_1 \cup D_2$. Then the following hold:

(1) *$\mathcal{H}$ has at most two parallelisms and each satisfies the consolidation criteria.*

(2) *After consolidating those parallelisms, the resulting multiple Γ-Heegaard surface has no new pod handles.*

(3) *Suppose that V_+ is a component of $M - \mathcal{H}$ adjacent to a parallelism V. Then V_+ is one of the following:*
   - *A 3-ball containing a single pod and the adjacent disk of $D_1 \cup D_2$ is not a cut disk.*
   - *A 3-ball containing a single arc and the adjacent disk of $D_1 \cup D_2$ is not a cut disk.*
   - *A product compressionbody that is either disjoint from T or which intersects T in vertical arcs and vertical pods.*

Proof. Untelescoping does not produce pod handles, so $\mathcal{H}$ has no new pod handles. Since Γ is disjoint from the vertices of $T - \partial T$, if either D_1 or D_2 is a cut disk, it does not intersect a pod or vertical pod of T. If D_1 and D_2 are non-separating then after the untelescoping there cannot be any regions of parallelism as the resulting thin surface has complexity strictly less than the complexity of either of the thick surfaces. Therefore we only need to consider the case where at least one of the disks is separating.

Let H_i be the union of the components of $\mathcal{H}^+$ obtained by compressing H along D_i. Since D_i is a single c-disk, each H_i contains at most two components. The thin surface $\mathcal{H}^-$ is obtained by compressing H along $D_1 \cup D_2$ and so contains at most three components. Let $C_i^{\pm}$ be the closures of the components of $M - \mathcal{H}$ adjacent to H_i with C_i^- adjacent to $\mathcal{H}^-$. Notice that $T \cap C_i^-$ consists of vertical arcs and at most one ghost arc. Since Γ is disjoint from the vertices of $T - \partial T$, the ghost arc does not join up to a vertex of $T - \partial T$. See figure 6.

One component of H_i is compressible in C_i^- and so C_i^- contains at most one product region and if it contains such a product region then it is not connected, i.e., the c-disk D_i is separating. Thus $\mathcal{H}^-$ contains at most two parallelisms and such a parallelism is adjacent to exactly one of D_1 or D_2. The two parallelisms are shaded in figure 7.

Since $T \cap C_i^-$ consists of vertical arcs and at most one ghost arc and since that ghost arc, if it exists, joins distinct components of $\mathcal{H}^-$, it is straightforward to verify that each parallelism of $\mathcal{H}$ satisfies the consolidation criteria.
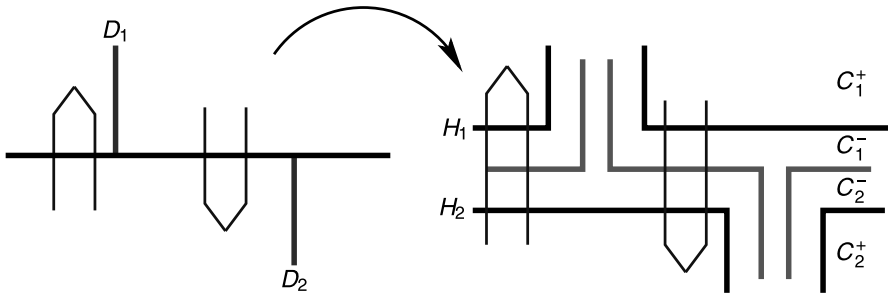
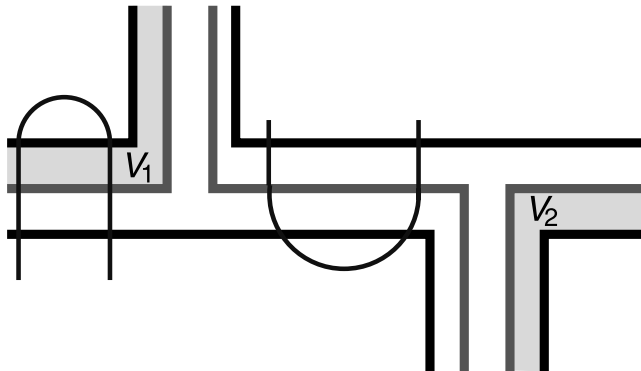Figure 6: Untelescoping a Heegaard splitting using two disks.

Figure 7: V_1 and V_2 are the two possible parallelisms.

Furthermore, consolidating one of the parallelisms does not alter the other region or the components of $M - \mathcal{H}$ on either side of it and so it will continue to satisfy the consolidation criteria and therefore can also be consolidated.

Let V be a parallelism of $\mathcal{H}$. Without loss of generality, assume that $V \subset C_1^-$. Note that this implies that D_1 is separating and that ∂D_2 is on the side of ∂D_1 not adjacent to V. Let $H_1' = \mathcal{H}^+ \cap V_+$. If H_1' had a Γ-c-disk then its boundary could be extended through V to lie on H. The boundary of this extended disk is on the inside of D_1 with respect to D_2 and does not have boundary parallel to ∂D_1. This contradicts the choice of D_1.

Notice that V_+ cannot be a 3-ball disjoint from T, since D_1 and D_2 were c-disks. Since H_1' does not have any Γ-c-disk in V_+, if there exists a bridge arc of $T \cap V_+$, then V_+ is a 3-ball and $\tau = T \cap V_+$ is connected. Furthermore, if D_1 intersected T (i.e., was a cut disk), then ∂D_1 would bound a once-punctured disk in H, since ∂V_+ is a 2-sphere. This contradicts the definition of cut disk, so D_1 is not a cut disk.

Similarly, if $T \cap V_+$ contains a bridge pod, V_+ must be a 3-ball and $T \cap V_+$ is connected. D_1 is not a cut disk since Γ is disjoint from the vertices of $T - \partial T$. Thus, we may assume that $T \cap V_+$ consists of vertical arcs and vertical pods. Since H_1' is incompressible in V_+, V_+ must be a product compressionbody. □

Lemma 8.3. *Suppose that (M, T) is a (3-manifold, graph) pair. Let $\mathcal{K}$ be a multiple Γ-Heegaard splitting for (M, T, Γ) such that the following hold:*

(1) Γ is disjoint from the vertices of $T - \partial T$ and no pod handle of $T - \mathcal{K}$ is adjacent to a component of $\mathcal{K}^-$.

(2) $\mathcal{K}$ does not have any parallelisms.

(3) No circle component of Γ is isotopically core for $\mathcal{K}$.

If some component H of $\mathcal{K}^+$ is Γ-c-weakly reducible in $M - \mathcal{K}^-$ then there is a non-empty sequence of untelescopings and consolidations of $\mathcal{K}$ converting it into a multiple Γ-Heegaard splitting $\mathcal{J}$ for (M, T) satisfying (1), (2), and (3).

Proof. Let $\mathcal{H}$ be obtained by untelescoping $H \subset \mathcal{K}$ using a pair of maximally separated Γ-c-disks D_1 and D_2. By Lemma 8.2, no pod handle of $T - \mathcal{H}$ is adjacent to a component of $\mathcal{H}^-$. Since Γ is disjoint from the vertices of $T - \partial T$, no vertex of $\mathcal{H}$ is connected in T to a ghost arc of $T - \mathcal{H}$. Consequently, all parallelisms of $\mathcal{H}$ satisfy consolidation criteria (C2), (C3), and (C4). Furthermore, we are guaranteed that consolidating a parallelism satisfying (C1)–(C4) produces a multiple Γ-bridge surface for (M, T, Γ) where all parallelisms continue to satisfy (C2)–(C4).

Let N be the component of $M - \mathcal{K}^-$ containing H. Any parallelism of $\mathcal{H}$ is contained in N and by Lemma 7.2, consolidating any such parallelisms does not create additional parallelisms. Thus, it suffices to show that each parallelism of $\mathcal{H} \cap N$ satisfies consolidation criterion (C1) and if V and W are two such parallelisms then consolidating one of them does not cause the other to stop satisfying (C1).

We observe, first, that the component F of $\mathcal{H}^- \cap N$ adjacent to both D_1 and D_2 is not parallel to either component of $\mathcal{H}^+ \cap N$ since the complexity of F differs from the complexity of each of those surfaces. Thus, the component F will never lie in the boundary of a parallelism of $\mathcal{H}$. Consequently, it suffices to consider only the situation in which all parallelisms of $\mathcal{H}$ lie on the same side of F in N. There can be at most two such parallelisms. Let
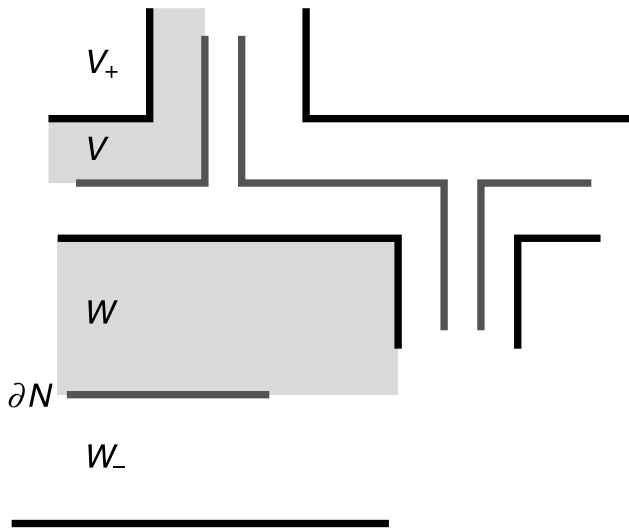
Figure 8: The shaded regions denote product regions of $\mathcal{H}$. V is a parallelism and W is adjacent to ∂N.

V be the parallelism of $\mathcal{H} \cap N$ adjacent to $\mathcal{H}^- - \partial N$, if such exists. Let W be the parallelism of $\mathcal{H} \cap N$ adjacent to ∂N, if such exists. See figure 8.

Since each compressionbody of $N - \mathcal{H}$ adjacent to $\mathcal{H}^-$ intersects T only in vertical arcs and at most one ghost arc, the parallelism V satisfies (C1). If $V = \varnothing$, then W_+ is a compressionbody of $N - \mathcal{H}$ adjacent to $\mathcal{H}^-$. The ghost arc in such a compressionbody, if it exists, is disjoint from ∂W. Hence, if $V = \varnothing$ the parallelism W satisfies the consolidation criteria. Suppose, therefore, that $V \neq \varnothing$ and $W \neq \varnothing$. Let $\mathcal{H}'$ be the result of consolidating V in $\mathcal{H}$.

Assume that W does not satisfy (C1) and let $\tau \subset (W_- \cup W \cup W_+)$ be a circle component of T. Since $\tau \cap W$ consists of vertical arcs, $\tau \cap W_+$ must have a bridge arc. By Lemma 8.2, such a bridge arc intersects V_+. By that lemma, V_+ is a 3–ball, the component of $D_1 \cup D_2$ adjacent to it is a not a cut-disk and $T \cap V_+$ consists only of a single bridge arc. Thus, $\tau \cap W \cup W_+$ consists of a single arc. Hence, $\tau \cap W_-$ consists of a single arc with both endpoints on $\partial N \cap \mathcal{K}^-$. Such an arc is a ghost arc in W_- and so must belong to Γ. Hence $\tau \subset \Gamma$. See figure 9. Since τ is the union of a ghost arc in τ_- with an arc in $W \cup W_+$ that is isotopic into $\partial_- W_-$. Thus, τ can be isotoped to lie in the spine of the compressionbody W_-. Hence, by Lemma 6.3, τ is isotopically core with respect to $\mathcal{K}$, a contradiction.   □
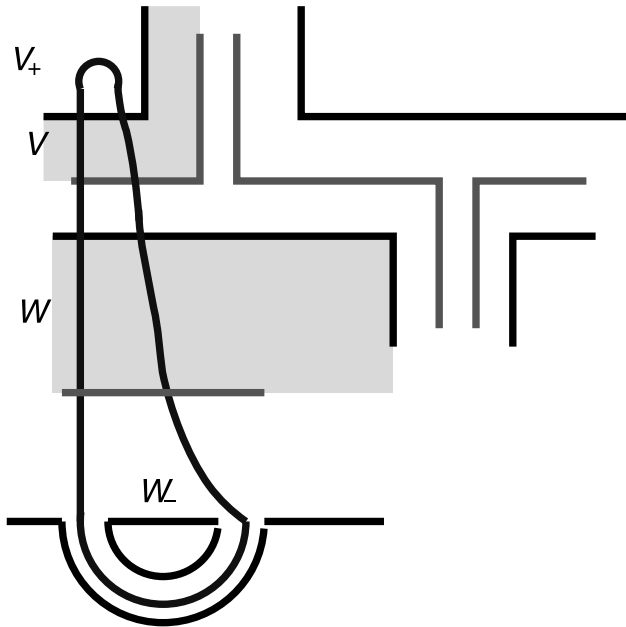
Figure 9: The circle τ.

9. Untelescoping and essential surfaces

In this section, we put together the previous results on untelescoping and consolidation and prove the main theorem. We begin by naming the main properties we seek.

Definition 9.1. Suppose that $\mathcal{K}$ and $\mathcal{H}$ are multiple Γ-c-Heegaard splitting for (M, T, Γ). If $\mathcal{H}$ is obtained from $\mathcal{K}$ by a sequence of untelescopings and consolidations, then we say that $\mathcal{H}$ is obtained by *thinning* $\mathcal{K}$. If $\mathcal{H}$ has the property that each component of $\mathcal{H}^+$ is Γ-c-irreducible in $M - \mathcal{H}^-$, each component of $\mathcal{H}^-$ is Γ-c-incompressible, and no component of $\mathcal{H}^-$ bounds a parallelism with a component of $\mathcal{H}^+$, then we say that $\mathcal{H}$ is *slim*. If $\mathcal{H}$ is slim and if each component of $\mathcal{H}^-$ is T-essential, then $\mathcal{H}$ is *locally thin*.

We begin with Theorem 9.2 that shows that we can obtain slim multiple Heegaard splittings. Theorem 9.3 will strengthen this to produce locally thin multiple Heegaard splittings.

Theorem 9.2. *Let M be a compact, orientable 3-manifold containing a properly embedded graph T and let Γ be a subgraph of T disjoint from the vertices of $T - \partial T$. Furthermore assume that $M - T$ is irreducible and that no sphere in M intersects T exactly once. Suppose $\mathcal{K}$ is a Γ-bridge surface for (M, T, Γ) such that the following hold:*

- *$\mathcal{K}$ has no parallelisms,*

- *no pod handle of $T - \mathcal{K}$ is adjacent to $\mathcal{K}^-$,*

- *no component of Γ is isotopically core with respect to $\mathcal{K}$.*

Then there is a slim multiple Γ-bridge splitting $\mathcal{H}$ for (M, T, Γ) obtained by thinning $\mathcal{K}$.

Proof. We define a sequence $(\mathcal{H}_i)$ of multiple Γ-c-Heegaard splittings such that each $\mathcal{H}_i$ satisfies the hypotheses of the theorem and so that for all i, $\mathcal{H}_{i+1}$ is obtained from $\mathcal{H}_i$ by a sequence of untelescopings using Γ-c-disks and consolidations. Assume that such an $\mathcal{H}_i$ does not satisfy the conclusions of the theorem. Since $\mathcal{H}_i$ has no parallelisms, if a component of $\mathcal{H}_i^-$ is Γ-c-compressible, then by Lemma 6.1, there is a component of $\mathcal{H}_i^+$ that is Γ-c-weakly reducible in $M - \mathcal{H}_i^-$. Hence, if $\mathcal{H}_i$ does not satisfy the conclusions of the theorem, some component of $\mathcal{H}_i^+$ is Γ-c-weakly reducible in $M - \mathcal{H}_i^-$.

We define $(\mathcal{H}_i)$ inductively as follows. Let $\mathcal{H}_1 = \mathcal{K}$. Assume that $\mathcal{H}_i$ has been defined. If $\mathcal{H}_i$ satisfies the conclusions of the theorem, let $\mathcal{H}_{i+1} = \mathcal{H}_i$. If $\mathcal{H}_i$ does not satisfy the conclusions of the theorem, let $\mathcal{H}_{i+1}$ be the multiple Γ-c-Heegaard splitting of (M, T, Γ) obtained by applying Lemma 8.3 to $\mathcal{H}_i$. By that lemma, $\mathcal{H}_{i+1}$ is obtained from $\mathcal{H}_i$ by a sequence of untelescopings using Γ-c-disks and consolidations and it satisfies the hypotheses of the theorem. Hence, we obtain a sequence $(\mathcal{H}_i)$ as desired.

Since both untelescoping and consolidation strictly reduce complexity, this sequence is eventually constant at a Γ-multiple Heegaard splitting $\mathcal{H}$ of (M, T, Γ) satisfying the conclusions of the theorem.   □

We can now prove our main theorem.

Theorem 9.3. *Let M be a compact, orientable 3-manifold containing a properly embedded graph T and let Γ be a subgraph of T disjoint from the vertices of $T - \partial T$. Furthermore assume that $M - T$ is irreducible and that no sphere in M intersects T exactly once. Suppose $\mathcal{K}$ is a multiple Γ-bridge surface for (M, T, Γ) such*

- *$\mathcal{K}$ has no parallelisms,*

- *no pod handle of $T - \mathcal{K}$ is adjacent to $\mathcal{K}^-$,*

- *no circle or edge component of Γ is isotopically core with respect to $\mathcal{K}$.*

Then there is a multiple Γ-Heegaard splitting $\mathcal{H}$ for (M, T, Γ) obtained by thinning $\mathcal{K}$ such that one of the following holds:

- *$\mathcal{H}$ is locally thin, or*

- *some component of $\mathcal{H}^+$ contains a generalized stabilization in $M - \mathcal{H}^-$, or*

- *some component of $\mathcal{H}^+$ is perturbed in $M - \mathcal{H}^-$, or*

- *T has a removable path.*

Proof. Let $\mathcal{H}$ be the multiple Γ-Heegaard splitting of (M, T, Γ) obtained by applying Theorem 9.2 to $\mathcal{K}$. By that theorem, each component of $\mathcal{H}^+$ is Γ-c-weakly reducible in $M - \mathcal{H}^-$, each component of $\mathcal{H}^-$ is Γ-c-incompressible, and no component of $\mathcal{H}^-$ bounds a parallelisms with a component of $\mathcal{H}^+$.

Suppose that a component F of $\mathcal{H}^-$ is T-parallel. Let C_F be the compressionbody so that $F = \partial_+ C_F$ is parallel to the boundary of a regular neighbourhood of some components of ∂M together with some subset of T. By Lemma 3.3 we may assume that F is innermost, i.e., C_F does not contain any other thin surfaces. Let H_{C_F} be the Γ-Heegaard splitting for C_F given by the unique thick surface of $\mathcal{H}$ contained in C_F.

Case 1: $T \cap C_F = \varnothing$.

Recall that F is not parallel to H_{C_F}. If $\partial_- C_F = \varnothing$, then by [23], H is stabilized. If $\partial_- C_F \neq \varnothing$, by [15] H_{C_F} is stabilized or boundary stabilized along $\partial_- C_F$. See also [14].

Case 2: H_{C_F} is a Heegaard splitting.

Since no thin surface is parallel to a thick surface, by Theorem 3.5 one of the following occurs:

- H_{C_F} has a generalized stabilization and if this is a boundary stabilization, it is along a component of ∂M;

- H_{C_F} is perturbed; or

- $T \cap C_F$ has a removable path disjoint from F.

All of these conclusions are possible conclusions in the theorem at hand.

Case 3: H_{C_F} is a Γ-Heegaard splitting but not a Heegaard splitting.

Let A and B be the two compressionbodies into which H_{C_F} divides C_F. Since H_{C_F} is a Γ-Heegaard splitting but not a Heegaard splitting, there exists an edge $e \subset \Gamma$ in either A or B which is disjoint from H_{C_F} and which has both endpoints on ∂C_F.

Case 3a: $\partial e \subset F$.

Since F is parallel to $\partial M \cup T$ and since $e \subset T$ is an edge with both endpoints on F, $T = e$ and $F = S^2$. Then by [10, Lemma 2.1] and [9, Theorem 1.1], either H_{C_F} is stabilized, meridionally stabilized, or perturbed. See Case 2 of the proof of [22, Lemma 5.2] for details.

Case 3b: $\partial e \subset \partial M$

In this case e is disjoint not only from H_{C_F} but also from $\mathcal{H}$. Then e is an edge of Γ with both endpoints on ∂M which is disjoint from $\mathcal{H}$. Hence e is isotopically core with respect to $\mathcal{K}$, a contradiction.

Case 3c: One endpoint of e is on F and one endpoint of e is on ∂M.

We may assume that the hypotheses of cases (3a) and (3b) do not apply. Let $e_1, \ldots, e_n$ be the union of edges of $T \cap C_F$ with one endpoint on F, one endpoint on ∂M, and which are disjoint from H_{C_F}. Perform a slight isotopy of each of them to convert T into a graph T' and each edge e_i into e_i' so that each e_i' is a removable edge of T' as in Lemma 2.3. Let H_{C_F}' be the new Γ-Heegaard surface and notice that H_{C_F}' is, in fact, a Heegaard surface for C_F. Since F is not parallel to H_{C_F}, by Theorem 3.5 one of the following occurs:

- H_{C_F}' has a generalized stabilization and if this is a boundary stabilization, it is along a component of ∂M;

- H_{C_F}' is perturbed; or

- $T \cap C_F$ has a removable path with both endpoints in ∂M.

Note that if the last option occurs the removable path is not equal to any of the e_i' since each of those edges has one endpoint on F. By Lemma 2.3 one of the following occurs:

- H_{C_F} has a generalized stabilization and if this is a boundary stabilization, it is along a component of ∂M;

- H_{C_F} is perturbed; or

- $T \cap C_F$ has a removable path disjoint from F.

All of these are conclusions for the theorem at hand. □

Combining the above result with Lemma 6.3 gives us an almost immediate corollary for bridge surfaces:

Corollary 9.4. *Let M be a compact, orientable 3-manifold containing a properly embedded graph T and let Γ be a subgraph of T disjoint from the vertices of $T - \partial T$. Furthermore assume that $M - T$ is irreducible and that no sphere in M intersects T exactly once. Suppose K is a bridge surface for (M, T, Γ) such no circle or edge component of Γ is isotopically core with respect to K. Then one of the following holds:*

- *there is a slim multiple Γ-Heegaard splitting $\mathcal{H}$ for (M, T, Γ) obtained by thinning K. Furthermore, if K is not thin, then $\mathcal{H}^- \neq \varnothing$.*

- *K contains a generalized stabilization,*

- *K is perturbed, or*

- *K has a removable path.*

Proof. By Theorem 9.3, either T has a removable path or there is a multiple Γ-Heegaard splitting $\mathcal{H}$ for (M, T, Γ) obtained by thinning K such that one of the following holds:

- $\mathcal{H}$ is locally thin, or

- some component of $\mathcal{H}^+$ contains a generalized stabilization in $M - \mathcal{H}^-$, or

- some component of $\mathcal{H}^+$ is perturbed in $M - \mathcal{H}^-$.

Note that K can be recovered from $\mathcal{H}$ by a sequence of amalgamations so by Lemma 6.3 if $\mathcal{H}^+$ contains a generalized stabilization or perturbation in $M - \mathcal{H}^-$ so does K. Thus the only statement we need to verify is that if K is not thin, then $\mathcal{H}^- \neq \varnothing$. To see this, notice that when a bridge surface is untelescoped using a pair of maximally separated disks, the thin surface adjacent to both disks is never adjacent to one of the parallelisms that is possibly created by the untelescoping. Thus it cannot be consolidated until it is adjacent to a thick surface that is untelescoped to create a thick surface parallel to it. Thus, throughout the thinning process there is always a thin surface, although the thin surface may change as the sequence progresses. □

Acknowledgment

References

[1] I. Agol, D. Groves and J. Manning, *The virtual Haken conjecture*, `arXiv:1204.2810`, 2012.

[2] A. Casson and C. Gordon, *Reducing Heegaard splittings*, Topol. Appl. **27**(3) (1987), 275–283.

[3] S. Cho and D. McCullough, *The tree of knot tunnels*, Geom. Topol. **13**(2) (2009), 769–815.

[4] A. Coward, *Algorithmically detecting the bridge number of hyperbolic knots*, `arXiv:0710.1262`.

[5] D. Gabai, *Foliations and the topology of* 3-*manifolds. III*, J. Differ. Geom. **26**(3) (1987), 479–536.

[6] H. Goda, M. Scharlemann and A. Thompson, *Levelling an unknotting tunnel*, Geom. Topol. **4** (2000), 243–275 (electronic).

[7] C. McA. Gordon, C. McA. and J. Luecke, *Knots are determined by their complements*, J. Amer. Math. Soc. **2**(2) (1989), 371–415.

[8] W. Haken, *Some results on surfaces in* 3-*manifolds*, Stud. Mod. Topol. (1968), 39–98.

[9] C. Hayashi and K. Shimokawa, *Heegaard splittings of the pair of the solid torus and the core loop*, Rev. Mat. Complut. **14**(2) (2001), 479–501.

[10] C. Hayashi and K. Shimokawa, *Heegaard splittings of trivial arcs in compressionbodies*, J. Knot Theory Ramifications **10**(1) (2001), 71–87.

[11] C. Hayashi and K. Shimokawa, *Thin position of a pair (3-manifold, 1-submanifold)*, Pacific J. Math. **197**(2) (2001), 301–324.

[12] M. Lackenby, *Heegaard splittings, the virtually Haken conjecture and property* (τ), Invent. Math. **164**(2) (2006), 317–359.

[13] J. Maher, *Heegaard gradient and virtual fibers*, Geom. Topol. **9** (2005), 2227–2259 (electronic).

[14] M. Yoav and S. Eric, *Closed essential surfaces and weakly reducible Heegaard splittings in manifolds with boundary*, J. Knot Theory Ramifications **13**(6) (2004), 829–843.

[15] M. Scharlemann and A. Thompson, *Heegaard splittings of* (surface) $\times I$ *are standard*, Math. Ann. **295**(3) (1993) 549–564.

[16] M. Scharlemann and A. Thompson, Thin position for 3-manifolds. In *Geometric topology (Haifa, 1992)*, volume 164 of Contemporary Mathematics, 231–238. Amer. Math. Soc., Providence, RI, 1994.

[17] M. Scharlemann and M. Tomova, *Uniqueness of bridge surfaces for 2-bridge knots*. Mathematical Proceedings, Cambridge Philosophical Society,

[18] M. Scharlemann and M. Tomova, *Alternate Heegaard genus bounds distance*, Geom. Topol. **10** (2006), 593–617 (electronic).

[19] J. Schultens, *The classification of Heegaard splittings for (compact orientable surface)* $\times S^1$, Proc. London Math. Soc. (3), **67**(2) (1993), 425–448.

[20] S. Taylor and M. Tomova, Heegaard surfaces for certain graphs in compressionbodies, To be published by *Revista Matemática Complutense*. Doi: 10.1007/s13163-011-0075-6 `arXiv:0910.3019`

[21] M. Tomova, *Multiple bridge surfaces restrict knot distance*, Algebr. Geom. Topol. **7** (2007), 957–1006.

[22] M. Tomova, *Thin position for knots in a 3-manifold*, J. London Math. Soc. (2) **80** (2009), 85–98.

[23] F. Waldhausen, *Heegaard-Zerlegungen der 3-Sphäre*, Topology, **7** (1968), 195–203.

[24] F. Waldhausen, On irreducible 3-manifolds which are sufficiently large, Ann. Math. (2) **87** (1968), 56–88.

Colby College
Department of Mathematics
5832 Mayflower Hill Waterville
Maine 04901-8858
E-mail address: sataylor@colby.edu

Mathematics Department
University of Iowa MacLean Hall
Iowa City
IA 52242
E-mail address: maggy-tomova@uiowa.edu