Algebraic & Geometric Topology 18 (2018) 1361–1409

Thin position for knots, links, and graphs in 3–manifolds

SCOTT A TAYLOR
MAGGY TOMOVA

We define a new notion of thin position for a graph in a 3–manifold which combines the ideas of thin position for manifolds first originated by Scharlemann and Thompson with the ideas of thin position for knots first originated by Gabai. This thin position has the property that connect-summing annuli and pairs-of-pants show up as thin levels. In a forthcoming paper, this new thin position allows us to define two new families of invariants of knots, links, and graphs in 3–manifolds. The invariants in one family are similar to bridge number, and the invariants in the other family are similar to Gabai's width for knots in the 3–sphere. The invariants in both families detect the unknot and are additive under connected sum and trivalent vertex sum.

57M25, 57M27; 57M50

1 Introduction

In [3], Gabai introduced *width* as an extremely useful knot invariant. Width is a certain function (whose exact definition isn't needed for our purposes) from the set of height functions on a knot in S^3 to the natural numbers. It becomes an invariant after minimizing over all possible height functions. A particular height function of a knot is *thin* if it realizes the minimum width. Thin embeddings produce very useful topological information about the knot (see, for example, Gabai's proof [3] of Property R, Gordon and Luecke's solution [5] to the knot complement problem, and Thompson's proof [21] that small knots have thin position equal to bridge position). Scharlemann and Thompson [15] extended Gabai's width for knots to a width for graphs in S^3 and gave a new proof of Waldhausen's classification [22] of Heegaard splittings of S^3. In [16], they also applied a similar idea to handle structures of 3–manifolds, producing an invariant of 3–manifolds also called width. A handle decomposition which attains the width is said to be *thin*. Thin handle decompositions for 3–manifolds have been very useful for understanding the structure of Heegaard splittings of 3–manifolds.

There have been a number of attempts (by, for example, Bachman [1], Howards, Rieck and Schultens [7], Hayashi and Shimokawa [6], Johnson [8] and the authors [19]) to

Published: 9 April 2018 DOI: 10.2140/agt.2018.18.1361

define width of knots (and later for tangles and graphs) in a 3–manifold by using various generalizations of the Scharlemann–Thompson constructions. These definitions have been used in various ways, however they have never been as useful as Scharlemann and Thompson's thin position for 3–manifolds. For instance, although Scharlemann and Thompson's thin position has the property that all thin surfaces in a thin handle decomposition for a closed 3–manifold are essential surfaces, the same is not true for the thin positions applied to knot and graph complements. (The papers [1] and [19] are exceptions. The former, however, applies only to links in closed 3–manifolds, and in the latter there are a number of technical requirements which limit its utility.)

We define an *oriented multiple vp-bridge surface* as a certain type of surface $\mathcal{H}$ in a 3–manifold M transverse to a graph $T \subset M$. The components of $\mathcal{H}$ are partitioned into *thick surfaces* $\mathcal{H}^+$ and *thin surfaces* $\mathcal{H}^-$. We exhibit a collection of thinning moves which give rise to a partial order, denoted by $\to$, on the set of *oriented multiple vp-bridge surfaces* (terms to be defined later) for a (3–manifold, graph) pair (M, T). These thinning moves include the usual kinds of destabilization and untelescoping moves known to experts, but we also include several new ones, corresponding to the situation when portions of the graph T are cores of compressionbodies in a generalized Heegaard splitting of M (in the sense of [16]). More significantly, we also allow untelescoping using various generalizations of compressing discs. Throughout the paper, we show how these generalized compressing discs arise naturally when considering bridge surfaces for (3–manifold, graph) pairs. If $\mathcal{H}$ and $\mathcal{K}$ are oriented bridge surfaces, we say that $\mathcal{H} \to \mathcal{K}$ if certain kinds of carefully constructed sequences of thinning moves produce $\mathcal{K}$ from $\mathcal{H}$. If no such sequence can be applied to $\mathcal{H}$ then we say that $\mathcal{H}$ is *locally thin*. If the reader allows us to defer some more definitions until later, we can state our results as:

Main Theorem *Let M be a compact, orientable 3–manifold and $T \subset M$ a properly embedded graph such that no vertex has valence 2 and no component of ∂M is a sphere intersecting T two or fewer times. Assume also that no sphere in M intersects T exactly once transversely. Then $\to$ is a partial order on $\overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$. Furthermore, if $\mathcal{H} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$, then there is a locally thin $\mathcal{K} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$ such that $\mathcal{H} \to \mathcal{K}$. Additionally, if $\mathcal{H}$ is locally thin then the following hold:*

(1) *Each component of $\mathcal{H}^+$ is sc-strongly irreducible in the complement of the thin surfaces.*

(2) *No component of $(M, T) \setminus \mathcal{H}$ is a trivial product compressionbody between a thin surface and a thick surface.*

(3) *Every component of $\mathcal{H}^-$ is c-essential in the exterior of T.*

(4) *If there is a 2–sphere in M which intersects T three or fewer times and which is essential in the exterior of T, then some component of $\mathcal{H}^-$ is such a sphere.*

The properties of locally thin surfaces are proved in part by using sweepout arguments; see Theorems 7.6 and 8.2. The existence of a locally thin $\mathcal{K}$ with $\mathcal{H} \to \mathcal{K}$ is proved using a new complexity which decreases under thinning sequences; see Theorem 6.17. Although this complexity behaves much as Gabai's or Scharlemann–Thompson's widths do, we view it as being more like the complexities used to guarantee that hierarchies of 3–manifolds terminate. In the sequel [20] we will show how powerful these locally thin positions for (3–manifold, graph) pairs are. In that paper, we construct two families of nonnegative half-integer invariants of (3–manifold, graph) pairs. The invariants of one family are similar to the bridge number and tunnel number of a knot. The invariants of the other family are very similar to Gabai's width for knots in S^3. We prove that these invariants (under minor hypotheses) are additive for both connect sum and trivalent vertex sum and detect the unknot.

In Sections 2 and 3 we establish our notation and important definitions including the definition of a multiple vp-bridge surface. We describe our simplifying moves in Sections 4 and 5. In Section 6, we define a complexity for oriented multiple vp-bridge surfaces and show it decreases under our simplifying moves. Section 6 also uses the simplifying moves to define a partial order $\to$ on the set $\overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$ of oriented multiple vp-bridge surfaces for (M,T). The main theorem, Theorem 6.17, shows that given $\mathcal{H} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$ there is a least element $\mathcal{K} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$ with respect to the partial order $\to$ such that $\mathcal{H} \to \mathcal{K}$. The least elements are called *locally thin*. In Section 7, we study the important properties of locally thin multiple vp-bridge surfaces. Theorem 7.6 lists a number of these properties, one of which is that each component of $\mathcal{H}^-$ is essential in the exterior of T. Section 8 sets us up for working with connected sums in [20] by showing that if there is a sphere in M, transversely intersecting T in three or fewer points, and which is essential in the exterior of T, then there is such a sphere that is a thin level for any locally thin multiple vp-bridge surface.

Acknowledgements Some of this paper is similar in spirit to [19], but here we operate under much weaker hypotheses and obtain much stronger results. We have been heavily influenced by the work of Gabai [3], Scharlemann and Thompson [16], and Hayashi and Shimokawa [6]. Throughout we assume some familiarity with the theory of Heegaard splittings, as in Scharlemann's [12]. We thank Ryan Blair, Marion Campisi, Jesse Johnson, Alex Zupan, and the attendees at the 2014 Thin Manifold conference for helpful conversations. Thanks also to the referee for many helpful comments and, in particular, finding a subtle but serious error in the original version of Section 6. The resolution of this error led to stronger results and simplified proofs. Tomova is supported by an NSF CAREER grant and Taylor by grants from the Colby College Division of Natural Sciences.

2 Definitions and notation

We let $I = [-1, 1] \subset \mathbb{R}$, D^2 be the closed unit disc in $\mathbb{R}^2$, and B^3 be the closed unit ball in $\mathbb{R}^3$. For a topological space X, we let $|X|$ denote the number of components of X. All surfaces and 3–manifolds we consider will be orientable, smooth or PL, and (most of the time) compact. If S is a surface, then $\chi(S)$ is its Euler characteristic.

A (3*–manifold, graph*) *pair* (M, T) (or simply a *pair*) consists of a compact, orientable 3–manifold M (possibly with boundary) and a properly embedded graph $T \subset M$. We do not require T to have vertices, so T can be empty or a knot or link. Since T is properly embedded in M, all valence-1 vertices lie on ∂M. We call the valence 1–vertices of T the *boundary vertices* or *leaves* of T and all other vertices the *interior vertices* of T. We require that no vertex of T have valence 0 or 2, but we allow a graph to be empty.

For any subset X of M, let $\eta(X)$ be an open regular neighborhood of X in M and $\overline{\eta(X)}$ its closure. If S is a (orientable, by convention) surface properly embedded in M and transverse to T, we write $S \subset (M, T)$. If $S \subset (M, T)$, we abuse notation slightly and write

$$(M, T) \setminus S = (M \setminus S, T \setminus S) = (M \setminus \eta(S), T \setminus \eta(S)).$$

We also write $S \setminus T$ for $S \setminus \eta(T)$. Observe that $\partial(M \setminus T)$ is the union of $(\partial M) \setminus T$ with $\partial\overline{\eta(T)}$. A surface $S \subset (M, T)$ is *∂–parallel* if $S \setminus T$ is isotopic relative to its boundary into $\partial(M \setminus T)$. We say that $S \subset (M, T)$ is *essential* if $S \setminus T$ is incompressible in $M \setminus T$, not ∂–parallel, and not a 2–sphere bounding a 3–ball in $M \setminus T$. We say that the graph $T \subset M$ is *irreducible* if whenever $S \subset (M, T)$ is a 2–sphere we have $|S \cap T| \neq 1$. The pair (M, T) is *irreducible* if T is irreducible and if the 3–manifold $M \setminus T$ is irreducible (ie does not contain an essential sphere).

We will need notation for a few especially simple (3–manifold, graph) pairs. The pair (B^3, arc) will refer to any pair homeomorphic to the pair (B^3, T) with T an arc properly isotopic into ∂B^3. The pair $(S^1 \times D^2, \text{core loop})$ will refer to any pair homeomorphic to the pair $(S^1 \times D^2, T)$ where T is the product of S^1 with the center of D^2.

Finally, we will often convert vertices of T into boundary components of M and vice versa. More precisely, if V is the union of all the interior vertices of T, we say that $(\mathring{M}, \mathring{T}) = (M \setminus \eta(V), T \setminus \eta(V))$ is obtained by *drilling out the vertices of* T. Similarly, we will sometimes refer to *drilling out* certain edges of T, ie removing an open regular neighborhood of those edges and incident vertices from both M and T.

2.1 Compressing discs of various kinds

We will be concerned with several types of discs which generalize the classical definition of a compressing disc for a surface in a 3–manifold.

Definition 2.1 Suppose that $S \subset (M, T)$ is a surface. Suppose that D is an embedded disc in M such that the following hold:

(1) $\partial D \subset (S \setminus T)$, the interior of D is disjoint from S, and D is transverse to T.

(2) $|D \cap T| \leq 1$.

(3) D is not properly isotopic into $S \setminus T$ in $M \setminus T$ via an isotopy which keeps the interior of D disjoint from S until the final moment. Equivalently, there is no disc $E \subset S$ such that $\partial E = \partial D$ and $E \cup D$ bounds either a 3–ball in M disjoint from T or a 3–ball in M whose intersection with T consists entirely of a single unknotted arc with one endpoint in E and one endpoint in D.

Then D is an *sc-disc*. More specifically, if $|D \cap T| = 0$ and ∂D does not bound a disc in $S \setminus T$, then D is a *compressing disc*. If $|D \cap T| = 0$ and ∂D does bound a disc in $S \setminus T$, then D is a *semicompressing disc*. If $|D \cap T| = 1$ and ∂D does not bound an unpunctured disc or a once-punctured disc in $S \setminus T$, then D is a *cut disc*. If $|D \cap T| = 1$ and ∂D does bound an unpunctured disc or a once-punctured disc in $S \setminus T$, then D is a *semicut disc*. A *c-disc* is a compressing disc or cut disc. The surface $S \subset (M, T)$ is *c-incompressible* if S does not have a c-disc; it is *c-essential* if it is essential and c-incompressible.

Remark 2.2 Semicut discs arise naturally when T has an edge containing a local knot, as in Figure 1. Semicompressing discs occur in part because even though a 3–manifold M may be irreducible, there is no guarantee that a given 3–dimensional submanifold is also irreducible.

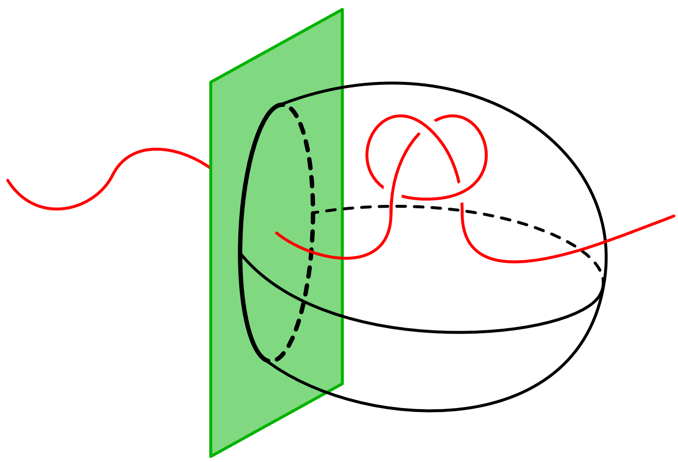

Figure 1: Except in very particular situations, the black disc is a semicut disc for the green surface.

3 Compressionbodies and multiple vp-bridge surfaces

3.1 Compressionbodies

In this section we generalize the idea of a compressionbody to our context.

Definition 3.1 Suppose that H is a closed, connected, orientable surface. We say that $(H \times I, T)$ is a *trivial product compressionbody* or a *product region* if T is isotopic to the union of vertical arcs, and we let $\partial_{\pm}(H \times I) = H \times \{\pm 1\}$. If B is a 3–ball and if $T \subset B$ is a (possibly empty) connected, properly embedded, ∂–parallel tree having at most one interior vertex, then we say that (B, T) is a *trivial ball compressionbody*. We let $\partial_+ B = \partial B$ and $\partial_- B = \varnothing$. A *trivial compressionbody* is either a trivial product compressionbody or a trivial ball compressionbody. Figure 2 shows both types of trivial compressionbodies.

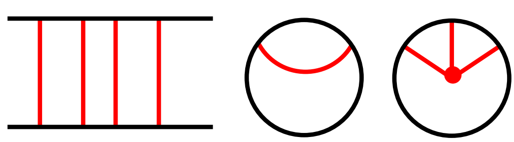

Figure 2: On the left is a trivial product compressionbody; in the center is a trivial ball compressionbody with T an arc; on the right is a trivial ball compressionbody with T a tree having a single interior vertex.

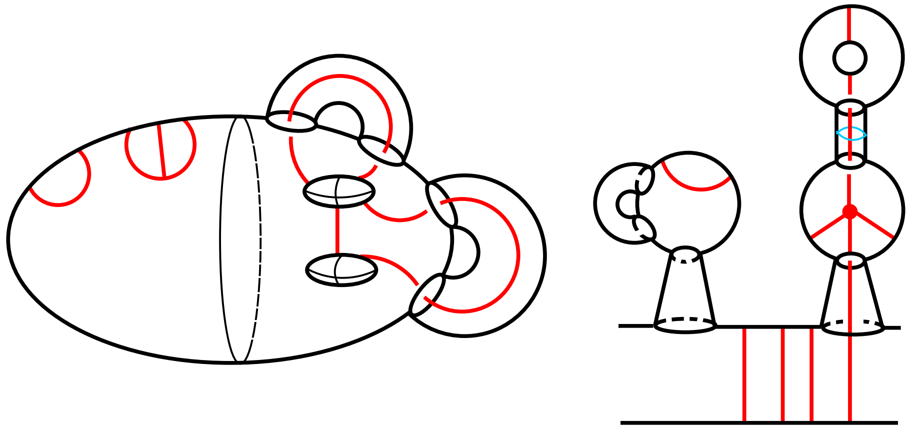

Figure 3: On the left is an example of a vp-compressionbody (C, T) with $\partial_- C$ the union of spheres. On the right is an example of a vp-compressionbody (C, T) with $\partial_- C$ the union of two connected surfaces, one of which is a sphere twice-punctured by T.

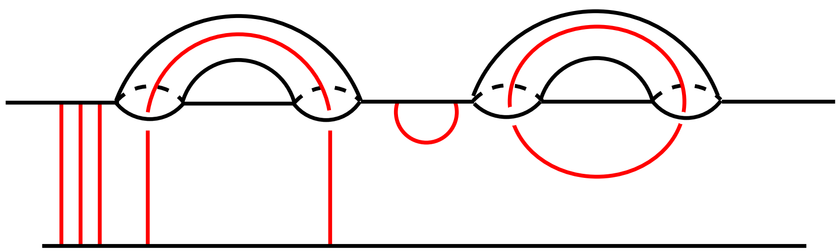

Figure 4: A vp-compressionbody (C, T). From left to right we have three vertical arcs, one ghost arc, one bridge arc, and one core loop in T.

A pair (C, T) is a *vp-compressionbody* if there is some component denoted by $\partial_+ C$ of ∂C and a collection of pairwise disjoint sc-discs $\mathcal{D} \subset (C, T)$ for $\partial_+ C$ such that the result of ∂–reducing (C, T) using $\mathcal{D}$ is a union of trivial compressionbodies. Observe that trivial compressionbodies are vp-compressionbodies as we may take $\mathcal{D} = \varnothing$. Figure 3 shows two different vp-compressionbodies. We will usually represent vp-compressionbodies more schematically as in Figure 4.

The set $\partial C \setminus \partial_+ C$ is denoted by $\partial_- C$. If no two discs of $\mathcal{D}$ are parallel in $C \setminus T$ then $\mathcal{D}$ is a *complete collection of discs* for (C, T). An edge of T which is disjoint from $\partial_+ C$ (and so has endpoints on $\partial_- C$ and the vertices of T) is a *ghost arc*. An edge of T with one endpoint in $\partial_+ C$ and one endpoint in $\partial_- C$ is a *vertical arc*. A component of T which is an arc having both endpoints on $\partial_+ C$ is a *bridge arc*. A component of T which is homeomorphic to a circle and is disjoint from ∂C is called a *core loop*. C is a *compressionbody* if $(C, \varnothing)$ is a vp-compressionbody. A compressionbody C is a *handlebody* if $\partial_- C = \varnothing$. A *bridge disc* for $\partial_+ C$ in C is an embedded disc in C with boundary the union of two arcs α and β such that $\alpha \subset \partial_+ C$ joins distinct points of $\partial_+ C \cap T$ and β is a bridge arc of T.

Remark 3.2 Suppose that (C, T) is a vp-compressionbody and that $(\mathring{C}, \mathring{T})$ is the result of drilling out the vertices of T. Considering the components of $\partial \mathring{C} \setminus \partial C$ as components of $\partial_- C$, we see that $(\mathring{C}, \mathring{T})$ is also a vp-compressionbody having the same complete collection of discs as (C, T). Furthermore, every component of $\mathring{T}$ is a vertical arc, ghost arc, bridge arc, or core loop. The "vp" stands for "vertex-punctured" as this notion of compressionbody is a generalization of the compressionbodies used in [19]: the vp-compressionbody $(\mathring{C}, \mathring{T})$ satisfies [19, Definition 2.1] with $\Gamma = \mathring{T}$ (in the notation of that paper). The notation vpwill also be helpful as a reminder that the first step in calculating many of the various quantities we consider is to drill out the vertices of T and treat them as boundary components of M.

The next two lemmas establish basic properties of vp-compressionbodies.

Lemma 3.3 *Suppose that (C, T) is a vp-compressionbody such that no spherical component of $\partial_- C$ intersects T exactly once. Suppose $P \subset (C, T)$ is a closed surface transverse to T. If D is an sc-disc for P, then if $|D \cap T| = 1$ either ∂D is essential on $P \setminus T$ or ∂D bounds a disc in P intersecting T exactly once. Furthermore, if P is a sphere, then after some sc-compressions it becomes the union of spheres, each either bounding a trivial ball compressionbody in (C, T) or parallel to a component of $\partial_- C$.*

Proof Suppose first that there is an sc-disc D for P such that ∂D bounds an unpunctured disc $E \subset P$ but $|D \cap T| = 1$. Then $E \cup D$ is a sphere in (C, T) intersecting T exactly once. Every sphere in C must separate C. Let $W \subset C \setminus (E \cup D)$ be the component disjoint from $\partial_+ C$. The fundamental group of every component of $\partial_- C$ injects into the fundamental group of C and every curve on every component of $\partial_- C$ is homotopic into $\partial_+ C$. Thus, any component of $\partial_- C$ contained in W must be a sphere. Drilling out the vertices of T along with all edges of T disjoint from $\partial_+ C$ creates a new vp-compressionbody (C', T'). As before, any essential curve in $\partial_- C'$ is not null-homotopic in C' and is homotopic into $\partial_+ C = \partial_+ C'$. Thus, $\partial_- C' \cap W$ can contain no essential curves. There is an edge $e \subset (T \cap W)$ with an endpoint in D. Beginning with e, traverse a path across edges of $T \cap W$ and components of $\partial_- C \cap W$ (each necessarily a sphere) so that no edge of $T \cap W$ is traversed twice. The path terminates when it reaches a component of $\partial_- C \cap W$ which is a once-punctured sphere, contrary to our hypotheses. Thus, no such sc-disc D can exist.

Suppose now that P is a sphere. Use c-discs to compress P as much as possible. By the previous paragraph, we end up with the union P' of spheres, each intersecting T no more times than does P. Let Δ be a complete collection of discs for (C, T) chosen so as to minimize $|\Delta \cap P'|$ up to isotopy of Δ. If some component P_0 of P' is disjoint from Δ, then it is contained in the union of trivial vp-compressionbody obtained by ∂–reducing (C, T) using Δ. Standard results from 3–manifold topology show that P_0 is either ∂–parallel to a component of $\partial_- C$ or is the boundary of a trivial ball compressionbody in (C, T).

If $\Delta \cap P' \neq \varnothing$, then it consists of circles. Since we have minimized $|\Delta \cap P'|$ up to isotopy, each circle of $\Delta \cap P'$ which is innermost on Δ bounds a semicompressing or semicut disc $D \subset \Delta$ for P'. By the previous paragraph, compressing P' using D creates an additional component of P' and preserves the property that each component of P' intersects T no more times than does P. The result follows by repeatedly performing such compressions until P' becomes disjoint from Δ. □

Lemma 3.4 *Suppose that (C, T) is a (3–manifold, graph) pair such that no component of $\partial_- C$ is a sphere intersecting T exactly once. Then the following hold:*

(1) *If $P \subset \partial_- C$ is an unpunctured sphere or a twice-punctured sphere, the result $(\widehat{C}, \widehat{T})$ of capping off P with a trivial ball compressionbody is still a vp-compressionbody.*

(2) *If p is a point in the interior of C (possibly in T), then the result of removing an open regular neighborhood of p from C (and T if $p \in T$) is a vp-compressionbody.*

Proof First, suppose that $P \subset \partial_- C$ is a unpunctured or twice-punctured sphere. Let Δ be a complete collection of sc-discs for (C, T). Let (C', T') be the result of ∂–reducing (C, T) using Δ. Then (C', T') is the union of vp-compressionbodies one of which is a product vp-compressionbody containing P. Capping off P with a trivial ball compressionbody converts this product vp-compressionbody into a trivial ball compressionbody. Thus, the result of ∂–reducing $(\widehat{C}, \widehat{T})$ using Δ is the union of trivial vp-compressionbodies. Thus, $(\widehat{C}, \widehat{T})$ is a vp-compressionbody.

Now suppose that p is a point in the interior of C. Let Δ be a complete collection of sc-discs for (C,T). By general position, we may isotope Δ to be disjoint from T. Let (C',T') be the result of ∂–reducing (C,T) using Δ. Each component of (C',T') is a trivial vp-compressionbody, one of which (W,T_W) contains p. If (W,T_W) is a trivial ball compressionbody either with T_W an arc containing p or with p an interior vertex of T_W, then the result of removing $\eta(p)$ from (W,T_W) is again a trivial compressionbody, and the result follows. Suppose, therefore, that if $p\in T_W$, then either $(W,T_W)\neq(B^3,\text{arc})$ or p is not a vertex of T_W.

If $p\in T$, there is a subarc of an edge of T_W joining ∂_+W to p. Let E be the frontier of a regular neighborhood of that edge. Then E is a semicut disc for ∂_+W cutting off a vp-compressionbody from (W,T_W) which is $(S^2\times I,\text{two vertical arcs})$. (The fact that E is a semicut disc follows from the considerations of the previous paragraph.) We may isotope E so that ∂E is disjoint from the remnants of Δ in ∂_+W. The disc E is then a cut disc or semicut disc for (C,T) such that $\Delta\cup E$ is a collection of s.c.-discs such that ∂–reducing $(C\setminus\eta(p),T\setminus\eta(p))$ is the union of trivial compressionbodies. Hence, $(C\setminus\eta(p),T\setminus\eta(p))$ is a vp-compressionbody.

If $p\notin T$, the proof is similar except we can pick any (tame) arc joining ∂_+C to p which is disjoint from T. □

Lemma 3.5 *Suppose that (C,T) is a vp-compressionbody such that no component of ∂_-C is a 2–sphere intersecting T exactly once. The following are true:*

1. *(C,T) is a trivial compressionbody if and only if there are no sc-discs for ∂_+C.*
2. *There are no c-discs for ∂_-C.*
3. *If D is an sc-disc for ∂_+C, then reducing (C,T) using D is the union of vp-compressionbodies. Furthermore, there is a complete collection of discs for (C,T) containing D.*

Proof (1) From the definition of vp-compressionbody, if there is no sc-disc for ∂_+C, then (C,T) is a trivial compressionbody. The converse requires a little more work, but follows easily from standard results in 3–dimensional topology.

(2) This is similar to the proof of Lemma 3.3. Suppose that ∂_-C has a c-disc P. As in the previous lemma, there is no sc-disc D for P such that ∂D bounds an unpunctured disc on P but $|D\cap T|=1$. Consequently, compressing P using any sc-disc D creates a new disc P' intersecting T no more often than did P and with $\partial P'=\partial P$. Since ∂P is essential on $\partial_-C\setminus T$, the disc P' is a c-disc for ∂_-C. Thus, we may assume

that P is disjoint from a complete collection of discs for (C,T). It follows easily that P is ∂–parallel in (C,T) and so is not a c-disc for $\partial_- C$, contrary to our assumption.

(3) Let (C,T) be a vp-compressionbody and suppose that D is an sc-disc for $\partial_+ C$. Let $(\widehat{C},\widehat{T})$ be the result of capping off all zero and twice-punctured sphere components of $\partial_- C$ with trivial ball compressionbodies. By Lemma 3.4, $(\widehat{C},\widehat{T})$ is a vp-compressionbody.

If D is not an sc-disc for $(\widehat{C},\widehat{T})$, it is ∂–parallel. Boundary-reducing $(\widehat{C},\widehat{T})$ with D results in two vp-compressionbodies: one a trivial ball compressionbody and the other equivalent to $(\widehat{C},\widehat{T})$. Removing regular neighborhoods of certain points in the interior of $\widehat{C}$, converts $(\widehat{C},\widehat{T})$ back into (C,T). By Lemma 3.4, ∂–reducing (C,T) using D results in vp-compressionbodies. A collection of sc-discs for those compressionbodies, together with D, gives a collection Δ of sc-discs such that ∂–reducing (C,T) using Δ results in trivial vp-compressionbodies. Thus, the lemma holds if D is ∂–parallel in $(\widehat{C},\widehat{T})$.

Now suppose that D is not ∂–parallel in $(\widehat{C},\widehat{T})$. Choose a complete collection Δ of sc-discs such that ∂–reducing $(\widehat{C},\widehat{T})$ using Δ results in the union (C',T') of trivial vp-compressionbodies. Out of all possible choices, choose Δ to intersect D minimally. We prove the lemma by induction on $|D \cap \Delta|$.

If $|D \cap \Delta| = 0$, then D is ∂–parallel in (C',T'). In this case, the result follows easily. Suppose, therefore, that $|D \cap \Delta| \geq 1$. The intersection $D \cap \Delta$ is the union of circles and arcs.

Suppose, first, that there is a circle of intersection. Let $\zeta \subset D \cap \Delta$ be innermost on Δ. Compressing D using the innermost disc $E \subset \Delta$ results in a disc D' and a sphere P. By Lemma 3.3, if $|E \cap T| = 1$, then $|D' \cap T| = 1$ and $|P \cap T| = 2$. On the other hand, if $|E \cap T| = 0$, then both D' and P are disjoint from T. By Lemma 3.3, there is a sequence of sc-compressions of P which result in zero and twice-punctured spheres. (If $|E \cap T| = 0$, there are no twice-punctured spheres.) These spheres either bound trivial ball compressionbodies in $(\widehat{C},\widehat{T})$ or are parallel to components of $\partial_- \widehat{C}$. But since $\partial_- \widehat{C}$ contains no unpunctured or twice-punctured sphere components, the latter situation is impossible. The sphere P is thus obtained by tubing together inessential spheres. It follows that P is also inessential: it bounds a 3–ball either disjoint from T or intersecting T in an unknotted arc. Since D is a unpunctured or once-punctured disc, this ball gives an isotopy of Δ reducing the intersection with D, a contradiction. Thus, there are no circles of intersection in $D \cap \Delta$.

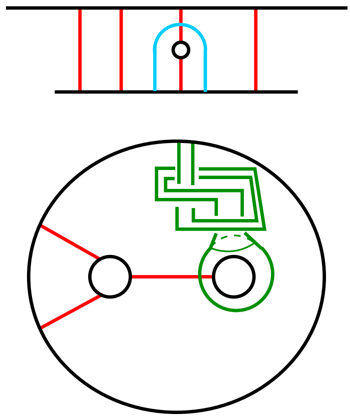

Figure 5: Above is an example of a vp-compressionbody (C, T) where $\partial_- C$ has a semicut disc (shown in blue). Below is an example of an sc-disc with the property that boundary-reducing the vp-compressionbody along that disc does not result in the union of vp-compressionbodies.

Let ζ be an arc of intersection in $D \cap \Delta$ which is outermost in Δ. Let $E \subset \Delta$ be the outermost disc. We may choose ζ and E so that E is disjoint from T. Boundary-reducing D using E results in two discs D' and D'', at most one of which is once-punctured. Observe that a small isotopy makes both disjoint from D. By our inductive hypotheses applied to D and D', the result of ∂–reducing (C, T) along both of them (either individually or together) is the union (W, T_W) of vp-compressionbodies. Remove the regular neighborhoods of points corresponding to the zero and twice-punctured sphere components of $\partial_- C$. By Lemma 3.4, we still have vp-compressionbodies, which we continue to call (W, T_W). Let $\mathcal{D}$ be the union of sc-discs for $\partial_+ W$, including D' and D'', such that the result of ∂–reducing (W, T_W) using $\mathcal{D}$ is the union of trivial compressionbodies. The disc D is contained in one of these trivial compressionbodies and, therefore, must be ∂–parallel. The result of ∂–reducing (C, T) along $D \cup \mathcal{D}$ is then the union of trivial compressionbodies, as desired. □

Remark 3.6 It is not necessarily the case that if (C, T) is an irreducible vp-compressionbody then there is no sc-disc for $\partial_- C$. To see this, let (C, T) be the result of removing the interior of a regular neighborhood of a point on a vertical arc in an irreducible vp-compressionbody $(\widetilde{C}, \widetilde{T})$. Then there is an sc-disc for $\partial_- C$ which is boundary-parallel in $\widetilde{C} \setminus \widetilde{T}$ and which cuts off from (C, T) a compressionbody which is $S^2 \times I$ intersecting T in two vertical arcs. See the top diagram in Figure 5. Similarly, if $\partial_- C$ contains 2–spheres disjoint from T, there will be a semicompressing disc for any component of $\partial_- C$ which is not a 2–sphere disjoint from T.

We also cannot drop the hypothesis that no component of $\partial_- C$ is a sphere intersecting T exactly once. The bottom diagram in Figure 5 shows an sc-disc with the property that boundary-reducing the vp-compressionbody along that disc does not result in the union of vp-compressionbodies.

In what follows, we will often use Lemma 3.5 without comment.

3.2 Multiple vp-bridge surfaces

The definition of a multiple vp-bridge surface for the pair (M,T) which we are about to present is a version of Scharlemann and Thompson's "generalized Heegaard splittings" [16] in the style of [6], but using vp-compressionbodies. We will also make use of orientations in a similar way to what shows up in Gabai's definition of thin position [3] and the definition of Johnson's "complex of surfaces" [8].

Definition 3.7 A connected closed surface $H \subset (M,T)$ is a *vp-bridge surface* for the pair (M,T) if $(M,T) \setminus H$ is the union of two distinct vp-compressionbodies $(H_\uparrow, T_\uparrow)$ and $(H_\downarrow, T_\downarrow)$ with $H = \partial_+ H_\uparrow = \partial_+ H_\downarrow$. If $T = \varnothing$, then we also call H a *Heegaard surface* for M.

A *multiple vp-bridge surface* for (M,T) is a closed (possibly disconnected) surface $\mathcal{H} \subset (M,T)$ such that

- $\mathcal{H}$ is the disjoint union of $\mathcal{H}^-$ and $\mathcal{H}^+$, each of which is the union of components of $\mathcal{H}$;
- $(M,T) \setminus \mathcal{H}$ is the union of embedded vp-compressionbodies (C_i, T_i) with $\mathcal{H}^- \cup \partial M = \bigcup \partial_- C_i$ and $\mathcal{H}^+ = \bigcup \partial_+ C_i$;
- each component of $\mathcal{H}$ is adjacent to two distinct compressionbodies.

If $T = \varnothing$, then $\mathcal{H}$ is also called a *multiple Heegaard surface* for M. The components of $\mathcal{H}^-$ are called *thin surfaces* and the components of $\mathcal{H}^+$ are called *thick surfaces*. We denote the set of multiple vp-bridge surfaces for (M,T) by $\mathrm{vp}\mathbb{H}(M,T)$.

Note that each component of $\mathcal{H}^+$ is a vp-bridge surface for the component of $(M,T) \setminus \mathcal{H}^-$ containing it. In particular, if $H \in \mathrm{vp}\mathbb{H}(M,T)$ is connected, then it is a vp-bridge surface and $H^- = \varnothing$. Also, observe that the components of ∂M are not considered to be thin surfaces; the surfaces ∂M and $\mathcal{H}^-$ play different roles in what follows. We now introduce orientations and flow lines.

Definition 3.8 Suppose that $\mathcal{H}$ is a multiple vp-bridge surface for (M,T). Suppose that each component of $\mathcal{H}$ is given a transverse orientation so that the following hold:

- If (C,T_C) is a component of $(M,T)\setminus\mathcal{H}$, then the transverse orientations of the components of $\partial_- C\cap\mathcal{H}^-$ either all point into or all point out of C.
- If (C,T_C) is a component of $(M,T)\setminus\mathcal{H}$ and if the transverse orientation of $\partial_+ C$ points into (respectively, out of) C, then the transverse orientations of the components of $\partial_- C\cap\mathcal{H}^-$ point out of (respectively, into) C.

A *flow line* is a nonconstant oriented path in M transverse to $\mathcal{H}$, not disjoint from $\mathcal{H}$, and always intersecting $\mathcal{H}$ in the direction of the transverse orientation. If S_1 and S_2 are components of $\mathcal{H}$, then a flow line from S_1 to S_2 is a flow line which starts at S_1 and ends at S_2. The multiple vp-bridge surface $\mathcal{H}$ is an *oriented* multiple vp-bridge surface if each component of $\mathcal{H}$ has a transverse orientation as above and there are no closed flow lines.

If there is a flow line from a thick surface $H\subset\mathcal{H}^+$ to a thick surface $J\subset\mathcal{H}^+$, then we may consider J to be *above* H and H to be *below* J. Reversing the transverse orientation on $\mathcal{H}$ interchanges the notions of above and below.

The set of oriented multiple vp-bridge surfaces for (M,T) is denoted by $\overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$. Note that there is a forgetful map from $\overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$ to $\mathrm{vp}\mathbb{H}(M,T)$. Any of our results for $\mathrm{vp}\mathbb{H}(M,T)$ can be turned into results for $\overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$, though the converse isn't true.

Given thick surfaces H and J, it is not necessarily the case that H is above J or vice versa, even if they are in the same component of M. See Figure 6 for a depiction of an oriented multiple vp-bridge surface. Not all multiple vp-bridge surfaces can be oriented. For example, consider circular thin position (defined in [9]); although we can define "above" and "below", the set of thick surfaces below a given thick surface H will equal the set of thick surfaces above H. Notice, however, that every connected multiple vp-bridge surface, once it is given a transverse orientation is an oriented multiple vp-bridge surface since it separates M.

Finally, for this section, we observe that cutting open along a thin surface induces oriented multiple bridge surfaces of the components.

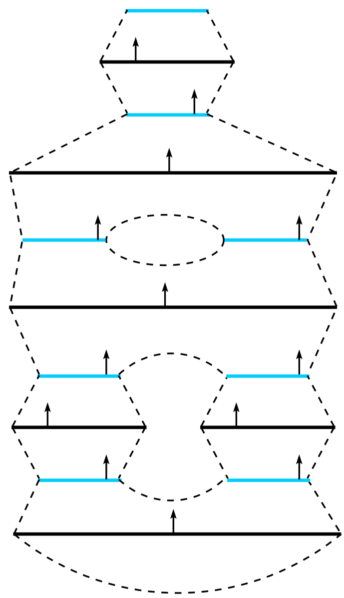

Figure 6: An example of an oriented multiple vp-bridge surface. Blue horizontal lines represent thin surfaces or boundary components. Black horizontal lines represent thick surfaces.

Lemma 3.9 *Suppose that* $\mathcal{H} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$ *and that* $F \subset \mathcal{H}^-$ *is a component. Let* (M',T') *be a component of* $(M,T) \setminus F$ *and let* $\mathcal{K} = (\mathcal{H} \setminus F) \cap M'$*. Then* $\mathcal{K} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M',T')$*.*

The proof of Lemma 3.9 follows immediately from the definitions, as the orientation on $\mathcal{H}$ restricts to an orientation on $\mathcal{K}$ and the flow lines for $\mathcal{K}$ form a subset of the flow lines for $\mathcal{H}$.

4 Simplifying bridge surfaces

This section presents a host of ways of replacing certain types of multiple vp-bridge surfaces by new ones that are closely related but are "simpler" (we will make this concept precise in Section 6). These simplifications are similar to the notion of destabilization and weak reduction for Heegaard splittings. Versions of many of these have appeared in other papers (eg [6; 17; 18; 19; 10]). The operations are: (generalized) destabilization, unperturbing, undoing a removable arc, untelescoping, and consolidation.

4.1 Destabilizing

Given a Heegaard splitting one can always obtain a Heegaard splitting of higher genus by adding a canceling pair of a one-handle and a two-handle, or (if the manifold has boundary) by tubing the Heegaard surface to the frontier of a collar neighborhood of a component of the boundary of the manifold. In the case where the manifold contains a graph, the core of the 1–handle, the cocore of the 2–handle, or the core of the tube might be part of the graph. (Though in this paper, we do not need to consider the case when *both* the 1–handle and the 2–handle contain portions of the graph.) In the realm of Heegaard splittings, the higher-genus Heegaard splitting is said to be either a stabilization or a boundary-stabilization of the lower-genus one. Observe that drilling out edges of T disjoint from $\mathcal{H} \in \mathrm{vp}\mathbb{H}(M,T)$ preserves the fact that $\mathcal{H}$ is a multiple vp-bridge surface. This suggests we also need to consider boundary-stabilization along portions of the graph T. Without further ado, here are our versions of destabilization:

Definition 4.1 Suppose that $\mathcal{H} \in \mathrm{vp}\mathbb{H}(M,T)$ and let H be a component of $\mathcal{H}^+$. There are six situations in which we can replace H by a new thick surface H' that is obtained from H by compressing along an sc-disc D. If H satisfies any of these conditions we say that H and $\mathcal{H}$ contain a *generalized stabilization*. See Figure 7 for examples.

- There is a pair of compressing discs for H which intersect transversely in a single point and are contained on opposite sides of H and in the complement of all other surfaces of $\mathcal{H}$. In this case we say that H and $\mathcal{H}$ are *stabilized*. The pair of compressing discs is called a *stabilizing pair*. The surface H' is obtained from H by compressing along either of the discs.

- There is a pair of a compressing disc and a cut disc for H which intersect transversely in a single point and are contained on opposite sides of H and in the complement of all other surfaces of $\mathcal{H}$. In this case we say that H and $\mathcal{H}$ are *meridionally stabilized*. The pair of compressing disc and cut disc is called a *meridional stabilizing pair*. The surface H' is obtained by compressing H along the cut disc.

- There is a separating compressing disc D for H contained in the complement of all other surfaces of $\mathcal{H}$ such that the following hold. Let W be the component of $M \setminus \mathcal{H}^-$ containing H. Compressing H along D produces two connected surfaces, H' and H'', where H' is a vp-bridge surface for W and H'' bounds a trivial product compressionbody disjoint from H' with a component S of ∂M. In this case we say that H and $\mathcal{H}$ are *boundary-stabilized* along S.

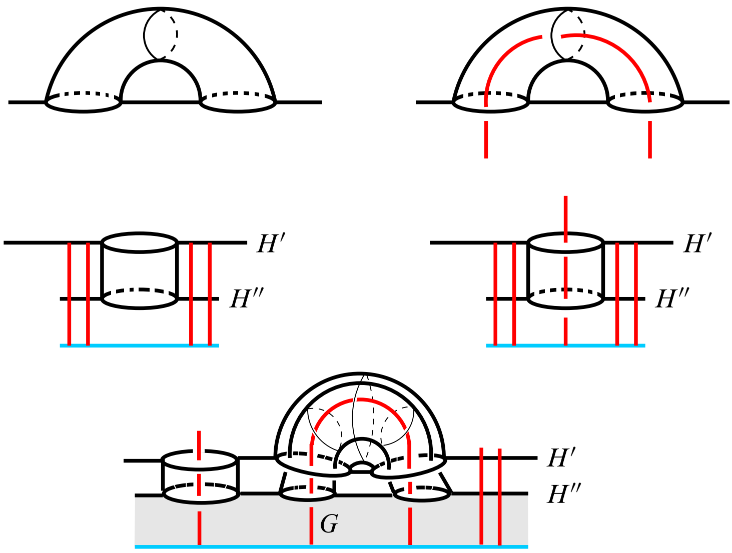

Figure 7: Depictions of stabilization, meridional stabilization, ∂–stabilization, meridional ∂–stabilization, and meridional ghost ∂–stabilization. The disc D in the final picture corresponds to the core of the left-most tube. In the last three cases, portions of the surfaces H' and H'', which appear after compressing along the disc D, have been labeled. In the final case, we have shaded the product region between H'' and S.

- There is a separating cut disc D for H contained in the complement of all other surfaces of $\mathcal{H}$ such that the following hold. Let W be the component of $M \setminus \mathcal{H}^-$ containing H. Compressing H along D produces two connected surfaces, H' and H'', where H' is a vp-bridge surface for W and H'' bounds a trivial product compressionbody disjoint from H' with a component S of ∂M. In this case we say that H and $\mathcal{H}$ are *meridionally boundary-stabilized* along S.

- Let G be a nonempty collection of vertices and edges of T disjoint from $\mathcal{H}$. Let $\widetilde{M} = M \setminus G$. If H and $\mathcal{H}$ as a multiple vp-bridge surface of $\widetilde{M}$ are (meridionally) boundary stabilized along a component of $\partial \widetilde{M}$ which is not a component of ∂M, then H and $\mathcal{H}$ are (*meridionally*) *ghost boundary-stabilized* along G.

Remark 4.2 In the definitions of (meridional) (ghost) ∂–stabilization, it's important to note that the statement that H' is a vp-bridge surface for the component of $M \setminus \mathcal{H}^-$ containing it is a precondition of being able to destabilize. Not every sc-compression of a thick surface resulting in a ∂–parallel surface is a destabilization. Performing a (meridional) (ghost) ∂–destabilization moves one or more components of ∂M from one side of H to the other side of H'. This is the reason we don't place transverse orientations on the components of ∂M.

Remark 4.3 Suppose that $H \subset \mathcal{H}^+$ has a generalized stabilization and let H' be the surface obtained from H by sc-compressing as in the definition above. It is easy to check (as in the classical settings) that $\mathcal{K} = (\mathcal{H} \setminus H) \cup H'$ is a multiple vp-bridge surface for (M, T). If $\mathcal{H} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$, the transverse orientation on $\mathcal{H}$ induces a transverse orientation on $\mathcal{K}$. Clearly, no new nonconstant closed flow lines are created. In particular, if $\mathcal{H} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$, there is a natural way of thinking of $\mathcal{K}$ as an element of $\overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$. We say that the (oriented) multiple vp-bridge surface $\mathcal{K}$ is obtained by *destabilizing* $\mathcal{H}$ (and that the thick surface H' is obtained by *destabilizing* the thick surface H).

4.2 Perturbed and removable bridge surfaces

We can sometimes push a bridge surface across a bridge disc and obtain another bridge surface. This operation is called unperturbing.

Definition 4.4 Let $\mathcal{H} \in \mathrm{vp}\mathbb{H}(M, T)$ and let $H \subset \mathcal{H}^+$ be a component. Suppose that there are bridge discs D_1 and D_2 for H in $M \setminus \mathcal{H}^-$, on opposite sides, disjoint from the vertices of T, and which have the property that the arcs $\alpha_1 = \partial D_1 \cap H$ and $\alpha_2 = \partial D_2 \cap H$ share exactly one endpoint and have disjoint interiors. Then H and $\mathcal{H}$ have a *perturbation*. The discs D_1 and D_2 are called a *perturbing pair* of discs for H and $\mathcal{H}$.

Remark 4.5 The type of perturbation we have defined here might better be called an "arc-arc"-perturbation. There are also perturbations where the bridge discs are allowed to contain vertices of T, but we will not need them in this paper.

Lemma 4.6 *Let $\mathcal{H}$ be an (oriented) multiple vp-bridge surface for (M, T). Suppose that $H \subset \mathcal{H}^+$ is a perturbed component with perturbing discs D_1 and D_2. Let E be the frontier of the neighborhood of D_1. Then compressing H along E and discarding the resulting twice-punctured sphere component results in a new surface H' such that $\mathcal{K} = (\mathcal{H} - H) \cup H'$ is an (oriented) multiple vp-bridge surface for (M, T).*

Proof This is nearly identical to Lemma 3.1 of [17]. We can alternatively think of H' as obtained from H by an isotopy along D_1. On the side of H containing D_1,

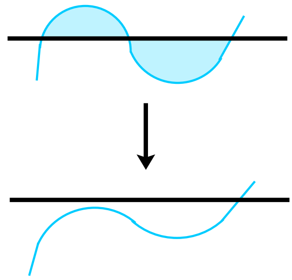

Figure 8: Unperturbing H

this isotopy removed a bridge arc and so H is still the positive boundary of a vp-compressionbody to that side. Let (C, T_C) be the vp-compressionbody containing D_2. Let D be the frontier of a regular neighborhood of D_2 in C, so that D cuts off a (B^3, arc) containing D_2 from (C, T_C). Note that D is an sc-disc for (C, T_C). Let Δ be a complete set of sc-discs for (C, T_C) containing D and chosen so as to minimize $|\partial\Delta \cap \partial D_1|$. Observe that no component of $\Delta \setminus D$ is inside the (B^3, arc) cut off by D. Suppose $E \subset \Delta \setminus D$ is a disc with boundary intersecting ∂D_1, and which contains the intersection point of $\partial\Delta \cap \partial D_1$ closest to the point $\partial D \cap \partial D_1$. Let E' be the disc obtained by tubing E to a parallel copy of D, along a subarc of ∂D_1. It is not difficult to confirm that $(\Delta \setminus E) \cup E'$ is still a complete collection of sc-discs for (C, T_C). However, it intersects ∂D_1 fewer times than Δ, a contradiction. Thus, ∂D_1 is disjoint from $\Delta \setminus D$.

Boundary-reduce (C, T_C) using $\Delta \setminus D$. We arrive at the union of vp-compressionbodies, one of which contains ∂D_1. We can now see that the isotopy of H across D_1, either combines two bridge arcs into another bridge arc or combines a vertical arc and a bridge arc into a bridge arc. Thus, the result of unperturbing is still a multiple vp-bridge surface.

If $\mathcal{H}$ is oriented we make $\mathcal{K}$ oriented by using the transverse orientations induced from $\mathcal{H}$. Clearly, no new closed flow lines are created. □

We say that the (oriented) multiple vp-bridge surface $\mathcal{K}$ constructed in the proof is obtained by *unperturbing* $\mathcal{H}$. See Figure 8 for a schematic depiction of the unperturbing operation.

4.3 Removable pairs

Suppose that $\mathcal{H}$ is an (oriented) multiple vp-bridge surface for (M, T) such that no component of $\mathcal{H}^- \cup \partial M$ is a sphere intersecting T exactly once. Let H be a component

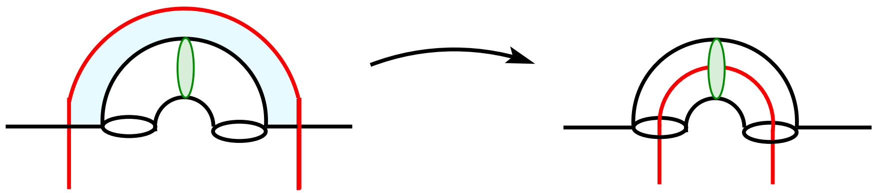

Figure 9: Undoing a removable arc

of $\mathcal{H}^+$, with $\mathcal{D}_\uparrow$ and $\mathcal{D}_\downarrow$ complete sets of discs for $(H_\uparrow, T_\uparrow)$ and $(H_\downarrow, T_\downarrow)$ respectively. Suppose that there exists a bridge disc D for H in $H_\uparrow$ (or $H_\downarrow$) with the following properties:

- It is disjoint from the vertices of T.
- It is disjoint from $\mathcal{D}_\uparrow$ (resp. $\mathcal{D}_\downarrow$).
- The arc $\partial D \cap H$ intersects a single component D^* of $\mathcal{D}_\downarrow$ (resp. $\mathcal{D}_\uparrow$), D^* is a disc and $|D \cap D^*| = 1$.

Then $\mathcal{H}$ and H are *removable*. The discs D and D^* are called a *removing pair*. See the left side of Figure 9.

Example 4.7 Suppose that $H \in \mathrm{vp}\mathbb{H}(M, T)$ is connected and that M' is obtained from M by attaching a 2–handle to ∂M or Dehn-filling a torus component of ∂M. Let α be either a cocore of the 2–handle or a core of the filling torus. Using an unknotted path in $M - H$, isotope α so that it intersects H exactly twice. Then $H \in \mathrm{vp}\mathbb{H}(M, T \cup \alpha)$ is removable. The component α is called the *removable component of $T \cup \alpha$*.

Lemma 4.8 *Suppose that $\mathcal{H} \in \mathrm{vp}\mathbb{H}(M, T)$ is removable. Then there is an isotopy of $\mathcal{H}$ in M to $\mathcal{K} \in \mathrm{vp}\mathbb{H}(M, T)$ supported in the neighborhood of the removing pair such that $\mathcal{K}$ intersects T two fewer times than $\mathcal{H}$ does. Furthermore, if $\mathcal{H}$ is oriented, so is $\mathcal{K}$.*

Proof Let H be the thick surface which is removable. We will construct an isotopy from H to a surface H' supported in a regular neighborhood of the removing pair and let $\mathcal{K} = (\mathcal{H} - H) \cup H'$. We will show that $\mathcal{K}$ is a multiple vp-bridge surface. Assuming it is, if $\mathcal{H} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$, we give H' the normal orientation induced by H. It is then easy to show that $\mathcal{K} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$.

Without the loss of much generality, we may assume that $\mathcal{H} = H$ is connected.

Let $D \subset H_\uparrow$ and $D^* \subset H_\downarrow$ be the removing pair and let $\mathcal{D}_\uparrow$ and $\mathcal{D}_\downarrow$ be the corresponding complete set of discs from the definition of "removable". Isotope T across D so that $T \cap D$ lies in $H_\downarrow$. Let T' be the resulting graph and let D^*_c be the cut disc that D^* gets converted into. Equivalently, we may isotope H across D and let H' be the resulting surface. See Figure 9.

The graph $T'_\uparrow = T' \cap H_\uparrow$ is obtained from $T_\uparrow = T \cap H_\uparrow$ by removing a component of $T_\uparrow$. After creating T' from T, the collection $\mathcal{D}_\uparrow$ remains a set of discs that decompose $(H_\uparrow, T'_\uparrow)$ into trivial compressionbodies, although there may now be discs in $\mathcal{D}_\uparrow$ which are parallel or which are boundary-parallel in $H_\uparrow \setminus T'_\uparrow$. Thus, $(H_\uparrow, T'_\uparrow)$ is a vp-compressionbody.

To show that $(H_\downarrow, T'_\downarrow)$ is a vp-compressionbody, note that cut-compressing $(H_\downarrow, T')$ along D^*_c results in the same collection of compressionbodies as compressing $(H_\downarrow, T)$ along D^*. Therefore $\mathcal{D}_\downarrow$ with D^* replaced by the induced cut disc D^*_c is a complete collection of sc-discs for $(H_\downarrow, T'_\downarrow)$ and so $(H_\downarrow, T'_\downarrow)$ is a vp-compressionbody. We conclude that $\mathcal{K}$ is an (oriented) multiple vp-bridge surface. □

The surface $\mathcal{K}$ in the preceding lemma is said to be obtained by *undoing a removable arc* of $\mathcal{H}$.

5 Untelescoping and consolidation

If we let T be empty in everything discussed so far and if we ignore the transverse orientations, then we are in Scharlemann–Thompson's set-up for thin position. We need a way to recognize when the multiple bridge surface can be "thinned" and a way to show that this thinning process eventually terminates. Scharlemann and Thompson thin by switching the order in which some 1–handle and 2–handle are added and they use Casson–Gordon's criterion [2] to recognize that this is possible by finding disjoint compressing discs on opposite sides of a thick surface. In this section, we use compressions along sc-weak reducing pairs of discs in place of handle exchanges.

5.1 Untelescoping

Suppose that $\mathcal{H} \in \mathrm{vp}\mathbb{H}(M, T)$. If $\mathcal{H}$ has the property that there is a component $H \subset \mathcal{H}^+$, and disjoint sc-discs D_- and D_+ for H on opposite sides such that D_- and D_+ are disjoint from $\mathcal{H}^-$, we say that $\mathcal{H}$ is *sc-weakly reducible*, that H is the *sc-weakly reducible component* and that $\{D_-, D_+\}$ is a *sc-weakly reducing pair*. If $\mathcal{H}$

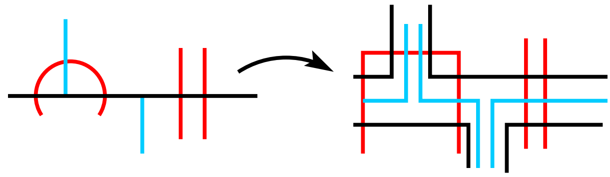

Figure 10: Untelescoping H. The red curves are portions of T. The blue lines on the left are sc-discs for H. Note that if a semicut or cut disc is used then a ghost arc is created.

is not sc-weakly reducible, we say it is *sc-strongly irreducible*. If D_- and D_+ are c-discs, we also say that $\mathcal{H}$ is *c-weakly reducible*, etc. Suppose that no component of $\mathcal{H}^- \cup \partial M$ is a sphere intersecting T exactly once. Then, given an sc-weakly reducible $\mathcal{H} \in \mathrm{vp}\mathbb{H}(M,T)$, we can create a new $\mathcal{K} \in \mathrm{vp}\mathbb{H}(M,T)$ by *untelescoping* $\mathcal{H}$ as follows:

Definition 5.1 Let $\{D_-, D_+\}$ be an sc-weakly reducing pair for an sc-weakly reducible component H of $\mathcal{H}^+$. Let N be the component of $M \setminus \mathcal{H}^-$ containing H. Let F be the result of compressing H using both D_- and D_+. Let $H_\pm$ be the result of compressing H using only $D_\pm$ and isotope each of $H_\pm$ slightly into the compressionbody containing $D_\pm$, respectively. Let $\mathcal{K}^- = \mathcal{H}^- \cup F$ and $\mathcal{K}^+ = (\mathcal{H}^+ \setminus H) \cup (H_- \cup H_+)$. See Figure 10 for a schematic picture. The component of F adjacent to copies of both D_- and D_+ is called the *doubly spotted* component. (The terminology is taken from [12].)

Lemma 5.2 *If $\mathcal{H} \in \mathrm{vp}\mathbb{H}(M,T)$ and if $\mathcal{K}$ is obtained by untelescoping $\mathcal{H}$, then $\mathcal{K} \in \mathrm{vp}\mathbb{H}(M,T)$.*

Proof Let $H \subset \mathcal{H}^+$ be the component which is untelescoped using discs $\{D_-, D_+\}$. Let W_+ and W_- be the two compressionbody components of $M \setminus \mathcal{H}$ that have copies of H as their positive boundaries. Let $\mathcal{D}_\pm$ be a complete collections of discs for the compressionbodies $W_\pm$ containing $D_\pm$. The discs $\mathcal{D}_\pm \setminus D_\pm$ after an isotopy are a complete collection of discs for the components of $(M,T) \setminus \mathcal{K}$ adjacent to $H_\pm$ and not adjacent to F. An isotopy of the disc $D_\pm$ makes it into an sc-disc for $H_\mp$. Boundary-reducing the submanifold bounded by $H_\pm$ and F using $D_\mp$ creates the union of product compressionbodies. Thus, $\mathcal{K} \in \mathrm{vp}\mathbb{H}(M,T)$. □

To extend this operation to oriented multiple vp-bridge surfaces we simply give H_- and H_+ the transverse orientations induced from H. We defer until Corollary 5.9 the proof that if $\mathcal{H}$ is oriented, then so is $\mathcal{K}$.

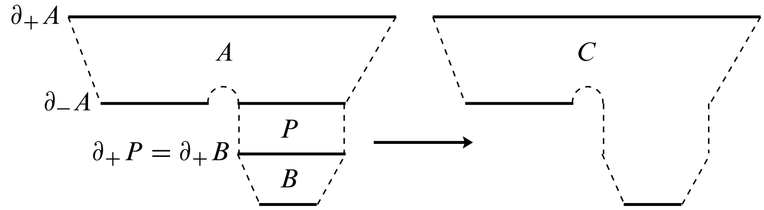

Figure 11: The vp-compressionbodies A, B, P, and C in Lemma 5.4

5.2 Consolidation

Untelescoping usually creates product compressionbodies which need to be removed as in Scharlemann–Thompson thin position. In our situation, though, this process is complicated by the presence of the graph T. We call the operation *consolidation*.

Definition 5.3 Suppose that $\mathcal{H}$ is an (oriented) multiple vp-bridge surface for (M, T) and that (P, T_P) is a product compressionbody component of $(M, T) \setminus \mathcal{H}$ which is adjacent to a component of $\mathcal{H}^-$ (and, therefore, not adjacent to a component of ∂M). Let $\mathcal{K} = \mathcal{H} \setminus (\partial_- P \cup \partial_+ P)$. If $\mathcal{H}$ is oriented, give each component of $\mathcal{K}$ the induced orientation from $\mathcal{H}$. We say that $\mathcal{K}$ is obtained from $\mathcal{H}$ by *consolidation* or by *consolidating* (P, T_P). (These terms were introduced in [19].)

The next two lemmas verify that consolidation is a valid operation in $\mathrm{vp}\mathbb{H}(M, T)$. See Figure 11 for a schematic depiction of the vp-compressionbodies in the following lemma.

Lemma 5.4 *Suppose that (P, T_P) is a trivial product compression body and that (A, T_A) and (B, T_B) are vp-compressionbodies with interiors disjoint from each other and from the interior of P. Assume also that $\partial_- P \subset \partial_- A$ and $\partial_+ B = \partial_+ P$. Let $(C, T) = (A, T_A) \cup (P, T_P) \cup (B, T_B)$ and assume that T is properly embedded in C. Then (C, T) is a vp-compressionbody.*

Proof We can dually define a vp-compressionbody to be a 3–manifold containing a properly embedded 1–manifold obtained by taking a collection of trivial vp-compressionbodies and adding to their positive boundary some 1–handles and some 1–handles containing a single piece of tangle as their core. With this dual definition, the lemma is obvious. □

Lemma 5.5 *Suppose that $\mathcal{H}$ is an (oriented) multiple vp-bridge surface for (M, T) and that $\mathcal{K}$ is obtained by consolidating a product region (P, T_P) of $\mathcal{H}$. Then $\mathcal{K}$ is an (oriented) multiple vp-bridge surface for (M, T).*

Proof This follows immediately from Lemma 5.4 and the observation that any closed flow line for $\mathcal{K}$ could be isotoped to be a closed flow line for $\mathcal{H}$. □

5.3 Elementary thinning sequences

As mentioned before, untelescoping often produces product regions. These product regions, in general, are of two types: they can be between a thin and thick surface neither of which existed before the untelescoping or they can be between a newly created thick surface and a thin surface (or a boundary component) that existed before the untelescoping operation. In fact, consolidating product regions of the first type can create additional product regions of the second type. The next definition specifies the order in which we will consolidate, before untelescoping further.

Definition 5.6 Suppose that $\mathcal{H}$ is an sc-weakly reducible oriented multiple vp-bridge surface for (M,T). Let $\mathcal{H}_1$ be obtained by untelescoping $\mathcal{H}$ using an sc-weak reducing pair. Let $\mathcal{H}_2$ be obtained by consolidating all trivial product compressionbodies of $\mathcal{H}_1 \setminus \mathcal{H}$. There may now be trivial product compressionbodies in $M \setminus \mathcal{H}_2$. Let $\mathcal{H}_3$ be obtained by consolidating all those products. We say that $\mathcal{H}_3$ is obtained from $\mathcal{H}$ by an *elementary thinning sequence*.

See Figure 12 for a depiction of the creation of $\mathcal{H}_2$ from $\mathcal{H}$.

To understand the effect of an elementary thinning sequence, we examine the untelescoping operation a little more carefully.

Lemma 5.7 *Suppose that H is a connected (oriented) vp-bridge surface and that $D_\uparrow$ and $D_\downarrow$ are an sc-weak reducing pair. Let $H_- \subset H_\downarrow$ and $H_+ \subset H_\uparrow$ be the new thick surfaces created by untelescoping H. Let F be the union of the new thin surfaces. Then the following are equivalent for a component Φ of F:*

(1) *Φ is not doubly spotted and is adjacent to a remnant of $D_\uparrow$ (resp. $D_\downarrow$).*
(2) *The disc $D_\uparrow$ (resp. $D_\downarrow$) is separating and Φ bounds a product region in $H_\uparrow$ (resp. $H_\downarrow$) with a component of H_+ (resp. H_-).*

Proof Suppose Φ is only adjacent to $D_\uparrow$. In this case $D_\uparrow$ must be separating as otherwise Φ would have two spots from $D_\uparrow$, and as H is connected, it would also have to have a spot coming from $D_\downarrow$. Compressing H along $D_\uparrow$ then results in two

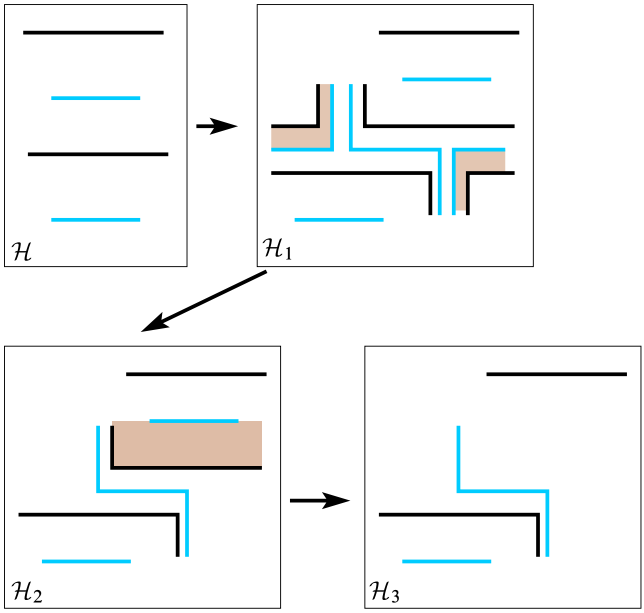

Figure 12: The surface $\mathcal{H}_2$ is created by untelescoping and consolidation. One or both of the compressionbodies $M \setminus \mathcal{H}_2$ shown in the figure may be product regions adjacent to $\mathcal{H}^-$. We consolidate those product regions to obtain $\mathcal{H}_3$.

components. Let H' be the component that doesn't contain $\partial D_\downarrow$. Then H' is not affected by compressing along $D_\downarrow$ to obtain F. Thus H' is parallel to Φ.

Conversely if Φ is parallel to some component H' of H_+ say, then Φ must be disjoint from the compressing disc $D_\downarrow$ and is therefore not double spotted. $\square$

Using the notation from Definition 5.6, we have:

Lemma 5.8 *Suppose that $\mathcal{H} \in \mathrm{vp}\mathbb{H}(M,T)$ and that $(M,T)\setminus\mathcal{H}$ has no trivial product compressionbodies adjacent to $\mathcal{H}^-$. Let $\mathcal{H}_1$, $\mathcal{H}_2$, and $\mathcal{H}_3$ be the surfaces in an elementary thinning sequence beginning with the untelescoping of a component $H \subset \mathcal{H}^+$. Then the doubly spotted component of $\mathcal{H}_1$ persists into $\mathcal{H}_3$ and no component of $(M,T)\setminus\mathcal{H}_3$ is a trivial product compressionbody adjacent to $\mathcal{H}_3^-$.*

Proof Let H_- and H_+ be the thick surfaces resulting from untelescoping the thick surface $H \subset \mathcal{H}^+$ and let F be the thin surface, with F_0 the doubly spotted component. Since F is obtained by compressing using an sc-disc, F is not parallel to either of H_- or H_+. In creating $\mathcal{H}_2$ we remove all components of F which are not doubly spotted (Lemma 5.7). The doubly spotted surface is not parallel to the remaining components of H_- or H_+ since we can obtain it by an sc-compression of each of them. Thus, the doubly spotted component persists into $\mathcal{H}_2$. Let H'_- and H'_+ be the components of H_- and H_+ remaining in $\mathcal{H}_2$. If either of H_- or H_+ bounds a trivial product compressionbody with $\mathcal{H}^-$, we create $\mathcal{H}_3$ by consolidating those trivial product compressionbodies.

Suppose that a component $(W, T_W) \subset (M, T) \setminus \mathcal{H}_3$ contains $F \subset \partial_- W$. Since H_- and H_+ each had an sc-compression producing the doubly spotted components of F, (W, T_W) must contain an sc-disc for $\partial_+ W$. Consequently, (W, T_W) is not a trivial product compressionbody. The result then follows from the assumption that no component of $(M, T) \setminus \mathcal{H}$ was a trivial product compressionbody adjacent to $\mathcal{H}^-$. □

Corollary 5.9 *Suppose that $\mathcal{H}, \mathcal{K}$ are multiple vp-bridge surfaces for (M, T) such that $M \setminus \mathcal{H}$ has no trivial product compressionbodies adjacent to $\mathcal{H}^-$. Assume that $\mathcal{K}$ is obtained from $\mathcal{H}$ using an elementary thinning sequence. Then the following are true:*

(1) $\mathcal{K}^- \neq \varnothing$.

(2) *$\mathcal{K}$ has no trivial product compressionbodies disjoint from ∂M.*

(3) *If $\mathcal{H}$ is oriented, so is $\mathcal{K}$.*

Proof By Lemma 5.8, the doubly spotted component of $\mathcal{K}^- \setminus \mathcal{H}^-$ does not get consolidated during the elementary thinning sequence and $\mathcal{K}$ has no trivial product compressionbodies adjacent to $\mathcal{K}^-$.

Suppose that $\mathcal{H}$ is oriented. We wish to show that $\mathcal{K}$ is oriented. We have described how to give transverse orientations to the components of $\mathcal{H}_1$ and these induce transverse orientations on $\mathcal{H}_2$, and $\mathcal{K}$. It follows immediately from the construction that the transverse orientations are coherent on the vp-compressionbodies. We need only show that we cannot create closed flow lines. Since consolidation does not create closed flow lines, it suffices to show that $\mathcal{H}_1$ does not have any closed flow lines.

Suppose that α is a closed flow line for $\mathcal{H}_1$. It must intersect $H_\pm$. As we have noted before, the (possibly disconnected) surface $H_\pm$ is obtained from H by compressing along an sc-disc $D_\pm$. We can recover H from $H_\pm$ by tubing (possibly along an arc component of $T \setminus \mathcal{H}_1$). We can isotope α to be disjoint from the tube, at which point it becomes a closed flow line for $\mathcal{H}$, a contradiction. Thus, $\mathcal{H}_1$, $\mathcal{H}_2$, and $\mathcal{H}_3$ are all oriented. □

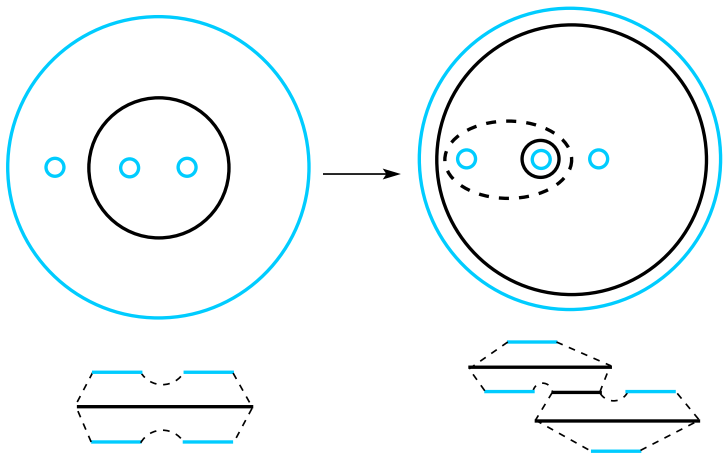

Figure 13: On the left in black is the original Heegaard surface and on the right is a multiple Heegaard surface which "should", by all rights, be thinner. The thick surfaces are in solid black and the thin surface is in dashed black. Below each figure is a schematic representation with the boundary components in blue, the thick surfaces in long black lines, and the thin surface in a short black line.

6 Complexity

The theory of 3–manifolds is rife with various complexity functions on surfaces which guarantee certain processes terminate. In [16], Scharlemann and Thompson used a version of Euler characteristic as their measure of complexity to ensure that untelescoping (and consolidation) of Heegaard surfaces will eventually terminate. Since that foundational paper, similar complexities have been used by many authors, eg [6; 8]. The requirement for a complexity is that it decreases under all possible types of compressions and any other moves that "should" simplify the decomposition. In our context, we need a complexity that decreases under destabilizing a generalized stabilization, unperturbing, undoing a removable arc, and applying an elementary thinning sequence. The next example demonstrates some of the difficulties that arise in our context.

6.1 An example

Traditionally thin position in the style of Scharlemann–Thompson [16] is done only for irreducible 3–manifolds. However, the following example (see Figure 13) shows that, at an informal level, it should be possible to define a thin position for reducible 3–manifolds.

Let P be the result of removing a regular neighborhood of two points from a 3–ball. Choose one component of ∂P as $\partial_+ P$ and the other two as $\partial_- P$. Let M be the result of gluing two copies of P along $H = \partial_+ P$. Then there is a certain sense in which the splitting of M can be untelescoped to a simpler splitting, but the new thick surface appear more complicated. Figure 13 shows the original Heegaard surface and another, ostensibly thinner, multiple Heegaard surface. The surface on the right can be obtained from the one on the left by thinning using semicompressing discs.

Although this example concerns a reducible manifold, we will run into similar problems when we have thin surfaces which are spheres twice-punctured by the graph T. If, in the example on the left, we add in a single ghost arc on each side of the Heegaard surface and four vertical arcs, one adjacent to each boundary component, we obtain an irreducible pair (M, T) with a connected vp-bridge surface that can be thinned to the surface on the right using semicut discs. Observe that neither of the vp-compressionbodies in the example on the left contains a compressing disc or a cut disc and that neither is a trivial vp-compressionbody.

6.2 Index of vp-compressionbodies

We introduce the index of a vp-compressionbody (see below) as a first step is developing a useful complexity for oriented multiple vp-bridge surfaces. This index is a proxy for counting handles. The index of a compressionbody without an embedded graph was first defined by Scharlemann and Schultens [13].

Definition 6.1 For a vp-compressionbody (C, T_C) such that T_C does not have interior vertices, define

$$\mu(C, T_C) = 3(-\chi(\partial_+ C) + \chi(\partial_- C)) + 2(|\partial_+ C \cap T| - |\partial_- C \cap T|) + 6.$$

If T_C does have interior vertices, drill them out and then calculate μ. For convenience, define $\mu(\varnothing) = 0$.

Remark 6.2 The +6 isn't strictly needed, but it allows us to work with nonnegative integers.

Observe that $\mu(B^3, \varnothing) = 0$, $\mu(B^3, \text{arc}) = 4$, and the index of any other trivial vp-compressionbody is 6. Since the Euler characteristic of a closed surface is even, the index is always even. The next lemma is proved by considering the effect of a ∂–reduction on μ.

Lemma 6.3 *Suppose that (C, T_C) is a vp-compressionbody such that no component of $\partial_- C$ is a sphere intersecting T_C exactly once. If $D \subset (C, T_C)$ is an sc-disc for $\partial_+ C$ and if (C_1, T_1) and (C_2, T_2) are the result of ∂–reducing (C, T_C) using D (we allow (C_2, T_2) to be empty), then*

$$\mu(C_1, T_1) + \mu(C_2, T_2) = \mu(C, T_C) - 6 + 4|D \cap T_C| + 6\delta, \tag{1}$$

where $\delta = 1$ if D is separating and 0 otherwise. Consequently, $\mu(C_1, T_1) < \mu(C, T)$. Furthermore, for any vp-compressionbody, $\mu(C, T_C) \geq 0$ with

$$\begin{aligned} \mu(C, T_C) &= 0 \quad \textit{if } (C, T_C) = (B^3, \varnothing), \\ \mu(C, T_C) &= 4 \quad \textit{if } (C, T_C) = (B^3, \text{arc}), \\ \mu(C, T_C) &\geq 6 \quad \textit{otherwise.} \end{aligned}$$

Proof If T_C has interior vertices, drill them out. Suppose first that $D \subset (C, T_C)$ is an sc-disc for $\partial_+ C$. Recall that $|D \cap T_C| \in \{0, 1\}$. Let Δ be a complete collection of sc-discs containing D. Let $(C', T') = (C_1, T_1) \cup (C_2, T_2)$. We prove the lemma by induction on $|\Delta|$.

Let (C', T') be the result of ∂–reducing (C, T_C) using D. Considering the effect of ∂–reduction on Euler characteristic and the number of punctures produces (1). Notice that $\Delta \setminus D$ is a complete collection of sc-discs for (C', T').

If $|\Delta| = 1$, then (C', T') is the union of trivial vp-compressionbodies. If $\delta = 0$, then (C, T_C) is either $(S^1 \times D^2, \varnothing)$ or $(S^1 \times D^2, \text{ core loop})$. In either case, $\mu(C, T_C) = 6$ and $\mu(C_1, T_1) = \mu(C', T')$ is either 0 or 4. Suppose $\delta = 1$. Since D is an sc-disc, neither (C_1, T_1) nor (C_2, T_2) is $(B^3, \varnothing)$. Similarly, if $|D \cap T_C| = 1$, then neither can be (B^3, arc). In particular,

$$\mu(C_1, T_1) = \mu(C, T_C) + 4|D \cap T_C| - \mu(C_2, T_2) < \mu(C, T_C).$$

The proof of the inductive step is similar; we apply the inductive hypothesis to (C_2, T_2) to conclude that $\mu(C_2, T_2) \geq 6$ when it is nontrivial. □

The next lemma considers the effect of consolidation on the index. See Figure 11 for a diagram.

Lemma 6.4 *Suppose that $\mathcal{H} \in \mathrm{vp}\mathbb{H}(M, T)$ is a multiple vp-compressionbody. Suppose (A, T_A), (P, T_P) and (B, T_B) are vp-compressionbodies with (P, T_P) a product vp-compressionbody such that $\partial_- P \subset \partial_- A$ and $\partial_+ B = \partial_+ P$. Let $C = A \cup P \cup B$ and $T = T_A \cup T_P \cup T_B$. Then $\mu(C, T_C) = \mu(A, T_A) + \mu(B, T_B) - 6$.*

Proof By the definition of product vp-compressionbody, $-\chi(\partial_+ P) = -\chi(\partial_- P)$ and $|\partial_+ P \cap T| = |\partial_- P \cap T|$. Let $\alpha = \partial_- A \setminus \partial_- P$. Recall that $\partial_+ C = \partial_+ A$ and $\partial_- C = \alpha \cup \partial_- B$. We have

$$\begin{aligned}\mu(A, T_A) + \mu(B, T_B) &= 3(-\chi(\partial_+ A) + \chi(\alpha)) + 2(|\partial_+ A \cap T| - |\alpha \cap T|) + 6 \\ &\quad + 3\chi(\partial_- B) - 2|\partial_- B \cap T| + 6 \\ &\quad + 3\chi(\partial_- P) - 2|\partial_- P \cap T| - 3\chi(\partial_+ B) + 2|\partial_+ B \cap T| \\ &= \mu(C, T) + 6.\end{aligned}$$

□

For a thick surface $H \subset \mathcal{H}^+$, let $\mu_\downarrow(H) = \mu(H_\downarrow)$ and $\mu_\uparrow(H) = \mu(H_\uparrow)$. We now define the oriented indices $I_\uparrow(\mathcal{H})$ and $I_\downarrow(\mathcal{H})$. These will contribute to a complexity which decreases under all relevant moves. Informally, for each thick surface we calculate the sum of the number of "handles" which are immediately above some thick surface which is either H or above H and the number of handles which are immediately below some thick surface which is either equal to H or below H. We place these numbers into a nonincreasing sequence and compare the results lexicographically. Instead of working with handles, however, we use the indices of vp-compressionbodies.

Definition 6.5 Let $\mathcal{H} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$. Each vp-compressionbody $(C,T_C) \subset (M,T) \backslash \mathcal{H}$ is adjacent to a single thick surface $H \subset \mathcal{H}^+$. The transverse orientation on H either points into or out of C. If it points into C, then $(C, T_C) = H_\uparrow$, and if it points out of C, then $(C, T_C) = H_\downarrow$. In the former case we say that (C, T_C) is an *upper* vp-compressionbody for $\mathcal{H}$, and we say it is a *lower* vp-compressionbody in the latter case.

Consider the set of flow lines beginning at H. A vp-compressionbody component (other than $H_\downarrow$) of $(M, T) \backslash \mathcal{H}$ intersecting one of these flow lines is said to be *above* H. We say that a vp-compressionbody is *below* H if reversing the transverse orientation of $\mathcal{H}$ makes it above H. Define $\mathcal{H}^H_\uparrow$ to be the set of all upper compression bodies $J_\uparrow$ above H. Define $\mathcal{H}^H_\downarrow$ to be the set of all lower compression bodies $J_\downarrow$ below H. Since there are no closed flow lines, the sets $\mathcal{H}^H_\uparrow$ and $\mathcal{H}^H_\downarrow$ are disjoint.

Define the *upper index* and *lower index* of H to be (respectively)

$$I_\uparrow(H) = 6 - 6|\mathcal{H}^H_\uparrow| + \sum_{J_\uparrow \in \mathcal{H}^H_\uparrow} \mu_\uparrow(J) \quad \text{and} \quad I_\downarrow(H) = 6 - 6|\mathcal{H}^H_\downarrow| + \sum_{K_\downarrow \in \mathcal{H}^H_\downarrow} \mu_\downarrow(K).$$

In Lemma 6.11 below, we verify that both $I_\uparrow(H)$ and $I_\downarrow(H)$ are nonnegative.

To package the indices for thick surfaces into an invariant for $\mathcal{H}$, let $\vec{c}(\mathcal{H})$, the *oriented complexity* of $\mathcal{H}$, be the *nonincreasing* sequence whose terms are the quantities $I(H) = I_\uparrow(H) + I_\downarrow(H)$ for each thick surface $H \subset \mathcal{H}^+$.

6.2.1 Oriented complexity decreases under generalized destabilization, unperturbing, and undoing a removable arc

Lemma 6.6 *Assume that no component of ∂M is a sphere intersecting T in two or fewer points. Suppose that $\mathcal{K}$ is obtained from $\mathcal{H}$ by a generalized destabilization, unperturbing, or undoing a removable arc. Then $\vec{c}(\mathcal{K}) < \vec{c}(\mathcal{H})$.*

Proof Let $H \subset \mathcal{H}^+$ be the thick surface to which we apply the generalized destabilization, unperturbing, or undoing a removable arc. Let H' be the new thick surface, so that $\mathcal{K} = (\mathcal{H} \setminus H) \cup H'$.

Suppose first that we are performing a destabilization or meridional destabilization. In this case, $-\chi(H') = -\chi(H) - 2$ and $|H' \cap T|$ is either $|H \cap T|$ or $|H \cap T| + 2$. Thus, $\mu_\downarrow(H') < \mu_\downarrow(H)$ and $\mu_\uparrow(H') < \mu_\uparrow(H)$. It follows that $I(H') < I(H)$ and for every thick surface $J \subset \mathcal{H}^+ \setminus H$, the index $I(J)$ does not increase under the destabilization or meridional destabilization.

The cases when we unperturb or undo a removable arc are very similar: we simply use the fact that $|H' \cap T| = |H \cap T| - 2$.

Now suppose that we perform a ∂–destabilization, meridional ∂–destabilization, ghost ∂–destabilization, or meridional ghost ∂–destabilization. Since indices are calculated by drilling out the interior vertices of T, we may assume that there are none. In all these cases, there is a (possibly disconnected) closed subsurface $S \subset \partial M$ and a (possibly empty) subset Γ of edges of T each disjoint from $\mathcal{H}$. These are such that H' is the result of compressing H along a separating sc-disc D and then discarding a component H'' which is the frontier of the regular neighborhood of $S \cup \Gamma$. In particular, the Euler characteristic of the discarded component is $\chi(S) - 2|\Gamma|$.

We claim that $\mu_\downarrow(H') < \mu_\downarrow(H)$ and $\mu_\uparrow(H') < \mu_\uparrow(H)$. Let $p = |D \cap T|$. Observe that $|H' \cap T| = |H \cap T| - |S \cap T| + p$. (If $p = 1$, this follows from the fact that D is separating.) Also note that $-\chi(H') = -\chi(H) + \chi(S) - 2|\Gamma| - 2$. Hence

$$-3\chi(H') + 2|H' \cap T| = -3\chi(H) + 2|H \cap T| + 3\chi(S) - 6|\Gamma| - 2|S \cap T| + 2p - 6.$$

Without loss of generality, we may assume that $S \cup \Gamma \subset H_\uparrow$. The (ghost) (meridional) ∂–stabilization then moves $S \cup \Gamma$ to the lower compressionbody $H'_\downarrow$. That is, $S \cup \Gamma \subset H'_\downarrow$. In particular, $\partial_- H'_\uparrow = \partial_- H_\uparrow \setminus S$ and $\partial_- H'_\downarrow = \partial_- H_\downarrow \cup S$. Thus, we have

$$\mu_\uparrow(H') = \mu_\uparrow(H) + (3\chi(S) - 6|\Gamma| - 2|S \cap T| + 2p - 6) + (-3\chi(S) + 2|S \cap T|),$$
$$\mu_\downarrow(H') = \mu_\downarrow(H) + (3\chi(S) - 6|\Gamma| - 2|S \cap T| + 2p - 6) + (3\chi(S) - 2|S \cap T|).$$

The first term in parentheses in each equation comes from the change of H to H' and the second term comes from the movement of $S \cup \Gamma$ from $\partial_- H_\uparrow$ to $H'_\downarrow$.

Simplifying, and using the fact that $2p \in \{0, 2\}$, we obtain

$$\mu_\uparrow(H') \le \mu_\uparrow(H) - 6|\Gamma| - 4,$$
$$\mu_\downarrow(H') \le \mu_\downarrow(H) - 6|\Gamma| + 6\chi(S) - 4|S \cap T| - 4.$$

In particular, $\mu_\uparrow(H') < \mu_\uparrow(H)$. The situation for $\mu_\downarrow$ requires more analysis. Let $S_0 \subset S$ be the subset which is the union of all spherical components of S and let $S_1 = S \setminus S_0$. We have

$$\mu_\downarrow(H') \le \mu_\downarrow(H) - 6|\Gamma| + 12|S_0| - 4|S_0 \cap T| - 4.$$

By assumption, each component of S_0 intersects T at least three times, so $4|S_0 \cap T| \ge 12|S_0|$. Thus, $\mu_\downarrow(H') < \mu_\downarrow(H)$.

Since $S \subset \partial M$ and does not belong to $\mathcal{H}'$, we can conclude that for each thick surface $J \subset \mathcal{H} \setminus H$, the indices $I_\uparrow(J)$ and $I_\downarrow(J)$ do not increase under the (meridional) (ghost) ∂–destabilization. Furthermore, since $I_\uparrow(H') + I_\downarrow(H') < I_\uparrow(H) + I_\downarrow(H)$, we have

$$\vec{c}(\mathcal{K}) < \vec{c}(\mathcal{H})$$

as desired. □

6.2.2 Oriented complexity decreases under consolidation

Lemma 6.7 *Suppose that $\mathcal{K} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$ is obtained from $\mathcal{H} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$ by consolidating a thick surface $H \subset \mathcal{H}^+$ with a thin surface $Q \subset \mathcal{H}^-$. Then $\vec{c}(\mathcal{K}) < \vec{c}(\mathcal{H})$.*

Proof Without loss of generality, we may suppose that $Q \subset \partial_-(H_\downarrow)$. (If not, reverse orientations so that above and below are interchanged.) That is, $H_\downarrow$ is the product compressionbody bounded by H and Q. Let $C \ne H_\downarrow$ be the other vp-compressionbody such that $Q \subset \partial_- C$. Let $J \subset \mathcal{H}^+ \setminus H$ be another thick surface. We will show that $I_\uparrow(J)$ and $I_\downarrow(J)$ are unchanged by the consolidation.

The vp-compressionbodies of $(M,T)\setminus\mathcal{K}$ are obtained from those of $(M,T)\setminus\mathcal{H}$ by replacing C, $H_\downarrow$, and $H_\uparrow$ with their union. By Lemma 6.4, we have

$$\mu(C\cup H_\downarrow\cup H_\uparrow)=\mu(C)+\mu(H_\uparrow)-6. \tag{2}$$

The consolidation does not affect flow lines, and so if there is no flow line from J to H or from H to J, then $I_\uparrow(J)$ and $I_\downarrow(J)$ are clearly unaffected.

If there is a flow line from J to H, then (2) implies that $I_\uparrow(J)$ decreases by 6. But we also have $|\mathcal{K}^J_\uparrow|=|\mathcal{H}^J_\uparrow|-1$, and so $I_\uparrow(J)$ also increases by 6. Thus, $I_\uparrow(J)$ is unchanged by the consolidation. Clearly, $I_\downarrow(J)$ is also unchanged by the consolidation, because the consolidation happens above J.

If there is a flow line from H to J, then clearly $I_\uparrow(J)$ is unchanged by the consolidation. On the other hand, $I_\downarrow(J)$ decreases by 6 because we have removed $\mu(H_\downarrow)$ from the sum. However, $I_\downarrow(J)$ also increases by 6 since $|\mathcal{K}^J_\downarrow|=|\mathcal{H}^J_\downarrow|-1$. Thus, $I_\downarrow(J)$ is also unchanged by the consolidation.

Thus, $\vec{c}(\mathcal{K})$ is simply obtained from $\vec{c}(\mathcal{H})$ by removing the term $I_\uparrow(H)+I_\downarrow(H)$. Since the sequence was nonincreasing, we have $\vec{c}(\mathcal{K})<\vec{c}(\mathcal{H})$. □

6.2.3 Oriented complexity decreases under an elementary thinning sequence

Suppose that $\mathcal{H}$, $\mathcal{H}_1$, $\mathcal{H}_2$, $\mathcal{H}_3=\mathcal{K}$ are the multiple vp-bridge surfaces in an elementary thinning sequence obtained by untelescoping a thick surface $H\subset\mathcal{H}^+$ using sc-discs D_- and D_+. As we've done before, let H_- and H_+ be the new thick surfaces and F the new thin surface. We will generally work with $\mathcal{H}_2$ and $\mathcal{H}_3$ (rather than $\mathcal{H}_1$), so $H_\pm$ is obtained from H by compressing along $D_\pm$ and discarding a component if $\delta_\pm=1$. The surface F is then obtained from $H_\pm$ by compressing along $D_\mp$ and possibly discarding a component.

Lemma 6.8 *The following hold for $H_-,H_+\subset\mathcal{H}_2^+$:*

1. $\mu_\downarrow(H_-)<\mu_\downarrow(H)$.
2. $\mu_\uparrow(H_+)<\mu_\uparrow(H)$.
3. $\mu_\downarrow(H_-)+\mu_\downarrow(H_+)=\mu_\downarrow(H)+6$.
4. $\mu_\uparrow(H_-)+\mu_\uparrow(H_+)=\mu_\uparrow(H)+6$.

Proof Claims (1) and (2) follow immediately from Lemma 6.3. Claim (4) can be obtained from the proof of claim (3) by interchanging $+$ and $-$ and $\uparrow$ and $\downarrow$. We prove claim (3).

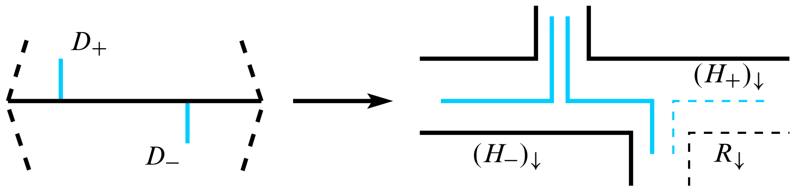

Figure 14: The region between the thin surface and thick surface, both indicated with dashed lines, is consolidated in the passage from $\mathcal{H}_1$ to $\mathcal{H}_2$ when ∂D_- separates H. The vp-compressionbody $R_\downarrow$ is then a subset of $(H_+)_\downarrow$.

Let $p_- = |D_- \cap T|$. Let $\delta_- = 1$ if ∂D_- separates H and 0 otherwise. If $\delta_- = 1$, let $R_\downarrow$ be the vp-compressionbody such that $(H_-)_\downarrow \cup R_\downarrow$ is the result of ∂–reducing $H_\downarrow$ using D_-. See Figure 14. If $\delta_- = 0$, then let $R_\uparrow = \varnothing$ and recall that $\mu(R_\uparrow) = 0$ by convention.

From the definition of index (see Lemma 6.3) we have

$$\mu_\downarrow(H_-) + \mu(R_\downarrow) = \mu_\downarrow(H) - 6 + 4p_- + 6\delta_-.$$

If $\delta_- = 0$, notice that $\mu_\downarrow(H_+) = 12 - 4p_-$, since a single compression creates F from H_+. If $\delta_- = 1$, then before the consolidation that creates $\mathcal{H}_2$ from $\mathcal{H}_1$, the index of $(H_+)_\downarrow$ is again $12 - 4p_-$. The consolidation removes a surface parallel to $\partial_+ R_\uparrow$ from the negative boundary of the vp-compressionbody and replaces it with $\partial_- R_\uparrow$. Recalling the additional (+6) term in the definition of μ, we have in either case

$$\mu_\downarrow(H_+) = 12 - 4p_- + \mu(R_\downarrow) - 6\delta_-.$$

Thus,

$$\mu_\downarrow(H_-) + \mu_\downarrow(H_+) = \mu_\downarrow(H) + 6. \qquad \square$$

Corollary 6.9 *The following hold for $H_-, H_+ \subset \mathcal{H}_2^+$:*

(1) $I_\downarrow(H_-) < I_\downarrow(H)$.

(2) $I_\uparrow(H_-) = I_\uparrow(H)$.

(3) $I_\downarrow(H_+) = I_\downarrow(H)$.

(4) $I_\uparrow(H_+) < I_\uparrow(H)$.

Proof We prove (2) and (4); claims (1) and (3) follow by reversing the orientation on $\mathcal{H}$.

Observe that each flow line beginning at H extends to a flow line beginning at H_- and can be restricted to a flow line beginning at H_+. Thus, the set $(\mathcal{H}_2)_\uparrow^{H_-}$ is obtained from the set $\mathcal{H}^H_\uparrow$ by removing the vp-compressionbody $H_\uparrow$ and replacing it with the vp-compressionbodies $(H_-)_\uparrow$ and $(H_+)_\uparrow$. Observe that $|(\mathcal{H}_2)^{H_-}_\uparrow| = |\mathcal{H}^H_\uparrow| + 1$. Hence, using Lemma 6.8(4),

$$\begin{aligned} I_\uparrow(H_-) &= I_\uparrow(H) - \mu_\uparrow(H) + \mu_\uparrow(H_-) + \mu_\uparrow(H_+) - 6 \\ &= I_\uparrow(H) + 6 - 6 \\ &= I_\uparrow(H). \end{aligned}$$

This proves claim (2).

On the other hand, $(\mathcal{H}_2)^{H_+}_\uparrow$ is obtained from $\mathcal{H}^H_\uparrow$ by removing $H_\uparrow$ and replacing it with $(H_+)_\uparrow$. From Lemma 6.8(3) we have $\mu_\uparrow(H_+) < \mu_\uparrow(H)$. Hence $I_\uparrow(H_+) < I_\uparrow(H)$, proving claim (4). □

Corollary 6.10 *If $\mathcal{K}$ is obtained from $\mathcal{H}$ by an elementary thinning sequence, then $\vec{c}(\mathcal{K}) < \vec{c}(\mathcal{H})$.*

Proof Let $\mathcal{H}_1$, $\mathcal{H}_2$, and $\mathcal{H}_3 = \mathcal{K}$ be the multiple vp-bridge surfaces created during the elementary thinning sequence. Corollary 6.9 shows that $I(H_\pm) < I(H)$.

If $J \subset \mathcal{H}^+ \setminus H$, then both $I_\uparrow(J)$ and $I_\downarrow(J)$ are unchanged in passing from $\mathcal{H}$ to $\mathcal{H}_2$. To see this, consider the possible locations of J. If J is neither above nor below H, then J is neither above nor below either of H_- nor H_+ and so $I(J)$ is unchanged. If J is below H, then J is below both H_- and H_+, as any flow line from J to H extends to a flow line from J to H_+ passing through H_-. In this case, passing from $\mathcal{H}$ to $\mathcal{H}_2$ increases the number of vp-compressionbodies above J by one and does not change the number of vp-compressionbodies below J. In the calculation of $I_\uparrow(J)$, we replace $\mu_\uparrow(H)$ with $\mu_\uparrow(H_-) + \mu_\uparrow(H_+)$. By Lemma 6.8, this increases the sum of the indices of the upper vp-compressionbodies above J by 6. It does not change the sum of the indices of the lower vp-compressionbodies below J. Thus, $I(J)$ does not increase when passing from $\mathcal{H}$ to $\mathcal{H}_2$. The analysis when J is above H is nearly identical.

We may therefore conclude that $\vec{c}(\mathcal{H}_2) < \vec{c}(\mathcal{H})$. Finally, either $\mathcal{H}_3 = \mathcal{H}_2$ or $\mathcal{H}_3$ is obtained from $\mathcal{H}_2$ by one or two consolidations. Thus, by Lemma 6.7,

$$\vec{c}(\mathcal{H}_3) \le \vec{c}(\mathcal{H}_2) < \vec{c}(\mathcal{H})$$

as desired. □

6.3 Index is nonnegative

The next lemma will help ensure that our oriented complexity guarantees that we cannot perform an infinite sequence of simplifying moves on an oriented vp-compressionbody.

Lemma 6.11 *Suppose that no component of $\mathcal{H}^-$ is a sphere intersecting T exactly once. Then for any thick surface $H \subset \mathcal{H}^+$, both $I_\uparrow(H)$ and $I_\downarrow(H)$ are nonnegative.*

Proof We prove the statement for $I_\uparrow(H)$; the proof of the statement for $I_\downarrow(H)$ is nearly identical. We may assume that T has no interior vertices (drill them out if necessary). We will also work under the assumption that no vp-compressionbody of $(M,T)\setminus\mathcal{H}$ is a product adjacent to a component of $\mathcal{H}^-$. To see that we may do this, recall from the proof of Lemma 6.7 that consolidation leaves $I_\uparrow(H)$ unchanged if H is not consolidated. If H is consolidated, and $H_\downarrow$ is the product region, then it is easy to see that either $I_\uparrow(H)=\mu(H_\uparrow)\geq 0$ or $I_\uparrow(H)$ is equal to $\mu(H_\uparrow)+I_\uparrow(J)\geq I_\uparrow(J)$ for some thick surface J above H. Finally, if $H_\uparrow$ is the product region, then there exists a thick surface J above H such that $\partial_- H_\uparrow \subset \partial_- J_\downarrow$. We may calculate $I_\uparrow(H)$ from $I_\uparrow(J)$ by subtracting 6 since $\mathcal{H}^H_\uparrow = \mathcal{H}^J_\uparrow \cup H_\uparrow$ and also adding 6 since $\mu(H_\uparrow)=6$. Thus, if $I_\uparrow(J)$ is nonnegative, so is $I_\uparrow(H)$. Henceforth, we assume that $(M,T)\setminus\mathcal{H}$ has no product regions adjacent to a component of $\mathcal{H}^-$.

We can express the definition of $I_\uparrow(H)$ as

$$I_\uparrow(H) = 6 + \sum_{J_\uparrow \in \mathcal{H}^H_\uparrow} (\mu(J_\uparrow)-6).$$

By Lemma 6.3, $\mu(J_\uparrow)\geq 6$ unless $J_\uparrow$ is $(B^3,\varnothing)$ or (B^3,arc) in which cases $\mu(J_\uparrow)=0$ and $\mu(J_\uparrow)=4$, respectively. Thus, if no element of $\mathcal{H}^H_\uparrow$ is a trivial ball compressionbody, then $I_\uparrow(H)\geq 0$. Assume, therefore that at least one element of $\mathcal{H}^H_\uparrow$ is a trivial ball compressionbody. We induct on the number $N(H,\mathcal{H})$ of trivial ball compressionbodies in $\mathcal{H}^H_\uparrow$.

If $H_\uparrow$ is $(B^3,\varnothing)$ or (B^3,arc), then $|\mathcal{H}^H_\uparrow|=1$ and $I_\uparrow(H)=\mu(H_\uparrow)\in\{0,4\}$ as desired. We may therefore assume that $|\mathcal{H}^H_\uparrow|\geq 2$.

We will call vp-compressionbodies $C_1, C_2 \in \mathcal{H}^H_\uparrow$ *adjacent in* $\mathcal{H}^H_\uparrow$ if there is a vp-compressionbody V such that $\partial_+ C_1 = \partial_+ V$ and $\partial_- C_2 \cap \partial_- V \neq \varnothing$ or vice versa. See Figure 15 for an example. If $C_1 \in \mathcal{H}^H_\uparrow$ is a trivial ball compressionbody, then there must be a vp-compressionbody $C_2 \in \mathcal{H}^H_\uparrow$ adjacent in $\mathcal{H}^H_\uparrow$ to C_1 as there is a flow

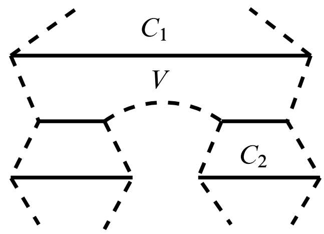

Figure 15: The vp-compressionbodies C_1 and C_2 are adjacent in $\mathcal{H}^H_\uparrow$ for $H = \partial_+ C_2$ or any thick surface H below $\partial_+ C_2$. In this example, there is another possible choice for C_2.

line from H to $\partial_+ C_1 \neq H$. Observe that in such a situation, C_2 is not a trivial ball compressionbody (since $\partial_- C_2 \neq \varnothing$).

Furthermore, if C_1 is a trivial ball compressionbody adjacent in $\mathcal{H}^H_\uparrow$ to C_2, with V the lower compressionbody incident to both, then $\partial_+ V$ is a unpunctured or twice-punctured sphere. Consequently $\partial_- V$ is the union of spheres. Let Γ be the graph with vertices the components of $\partial_- V$ and edges corresponding to the ghost arcs in V. Since $\partial_+ V$ is a sphere, Γ is the union of isolated vertices and trees. Since no component of $\partial_- V$ is a once-punctured sphere, each component P of $\partial_- V$ is a unpunctured or twice-punctured sphere. In particular, if $C_1 \cap T = \varnothing$, then P is unpunctured.

Suppose now that $A \in \mathcal{H}^H_\uparrow$ is adjacent in $\mathcal{H}^H_\uparrow$ to a trivial ball compressionbody $C \in \mathcal{H}^H_\uparrow$. Choose a single component P of $\partial_- A$ such that a flow line from H to $\partial_+ C$ passes through P. This implies that $P \subset \partial_- V$ where V is the lower vp-compressionbody incident to both A and C. By the remarks of the previous paragraph (with $C_1 = C$ and $C_2 = A$), P is a unpunctured or twice-punctured sphere (as are all components of $\partial_- V$).

Cut (M, T) open along all components of $\partial_- V \setminus P$, turning those components into components of ∂M which are unpunctured or twice-punctured spheres. Let $\mathcal{H}'$ be the components of $\mathcal{H}$ which are not now components of ∂M. Observe that $\mathcal{H}'^+ = \mathcal{H}^+$ and that now every flow line from H to $\partial_+ C$ must pass through P. We have not, however, changed $I_\uparrow(H)$ since any compressionbody of $(M, T) \setminus \mathcal{H}$ which was above H is still a compressionbody of $(M, T) \setminus \mathcal{H}'$ above H and we have not created any new vp-compressionbodies above H. We may have disconnected M; however, any vp-compressionbodies not in the component of M containing H were not above H before the cut and we can ignore them for the purposes of the calculation. For convenience of notation, use $\mathcal{H}$ instead of $\mathcal{H}'$ and assume that every flow line from H to $\partial_+ C$ must pass through P.

Now cut open M along P. This cuts M into two components M_1 and M_2 with M_1 containing H and M_2 containing C. Cap off the components of ∂M_1 and ∂M_2 corresponding to P with $(B^3, \varnothing)$ or (B^3, arc) corresponding to whether or not P is a unpunctured or twice-punctured sphere. Let $(\widehat{M}_i, \widehat{T}_i)$ for $i = 1, 2$ be these new (3–manifold, graph) pairs. Observe that $\widehat{\mathcal{H}} = \mathcal{H} \setminus (P \cup \partial_+ C)$ is a multiple vp-bridge surface for $(\widehat{M}_1, \widehat{T}_1)$.

The only upper vp-compressionbody affected by this is A, and we obtain a new vp-compressionbody $\widehat{A}$. If $\widehat{A}$ is a trivial ball compressionbody, then $P = \partial_- A$, A contains no bridge arcs, and $\partial_+ A$ is a sphere. This is enough to guarantee that A is a product compressionbody, contrary to hypothesis. Thus, with respect to $\widehat{\mathcal{H}}$ there is one fewer trivial ball compressionbody above H than with respect to $\mathcal{H}$. Let $\widehat{I}$ be $I_\uparrow(H)$ with respect to $\widehat{\mathcal{H}}$ and let I be $I_\uparrow(H)$ with respect to $\mathcal{H}$. By our inductive hypothesis, we have $\widehat{I} \geq 0$.

We have

$$\mu(A) = \mu(\widehat{A}) + 6 - 2|P \cap T|.$$

Thus

$$I = \widehat{I} + 6 - 2|P \cap T| + (\mu(C) - 6) \geq \mu(C) - 2|P \cap T|.$$

Recalling that $\mu(C) \in \{0, 4\}$ and $|P \cap T| \in \{0, 2\}$, we need only realize that $\mu(C) = 0$ implies $|P \cap T| = 0$ to conclude that

$$I \geq 0. \qquad \square$$

Remark 6.12 By Lemma 6.11, each term of $\vec{c}(\mathcal{H})$ is nonnegative. Thus, any sequence of multiple vp-bridge surfaces $\mathcal{H}$ with $\vec{c}(\mathcal{H})$ strictly decreasing must terminate.

6.4 Extended thinning moves

In this section, we formalize the fact that oriented complexity forbids an infinite sequence of simplifying moves to an oriented multiple vp-compressionbody.

Definition 6.13 An oriented multiple vp-bridge surface $\mathcal{H}$ is *reduced* if it does not contain a generalized stabilization, a perturbation or a removable arc, and if no component of $(M, T) \setminus \mathcal{H}$ is a trivial product compressionbody adjacent to a component of $\mathcal{H}^-$.

Definition 6.14 Suppose that $\mathcal{H} \in \overrightarrow{\text{vp}\mathbb{H}}(M, T)$ is reduced and that T is irreducible. An *extended thinning move* applied to $\mathcal{H}$ consists of the following steps in the following order:

(1) Perform an elementary thinning sequence.

(2) Destabilize, unperturb, and undo removable arcs until no generalized stabilizations, perturbations, or removable arcs remain.

(3) Consolidate all components of $\mathcal{H}^-$ and $\mathcal{H}^+$ cobounding a trivial product compressionbody in $(M,T)\setminus\mathcal{H}$.

(4) Repeat (2) and (3) as much as necessary until $\mathcal{H}$ does not have a generalized stabilization, perturbation, or removable arc or product region adjacent to $\mathcal{H}^-$.

Remark 6.15 Corollary 6.10, Lemma 6.6, and Lemma 6.7 show that each of the steps (1), (2) and (3), if applied nonvacuously, strictly decrease oriented complexity. Thus, by Remark 6.12 they can occur only finitely many times, until either we cannot (nonvacuously) perform any of the steps of an extended thinning move or until we have a multiple vp-bridge surface having a thin level which is a sphere intersecting T exactly once.

We have phrased the steps as we have in order to guarantee that if $\mathcal{H}$ is reduced, then an extended thinning move applied to $\mathcal{H}$ results in a reduced multiple vp-bridge surface. If $\mathcal{H}\in\overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$ is not reduced, we may perform a sequence of consolidations, generalized destabilizations, unperturbings, and undoings of removable arcs to make it reduced. (Such a sequence is guaranteed to terminate because each of those operations strictly decreases oriented complexity.)

Definition 6.16 If $\mathcal{H},\mathcal{K}\in\overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$ then we write $\mathcal{H}\to\mathcal{K}$ if either

- $\mathcal{H}$ is reduced and $\mathcal{K}$ is obtained from $\mathcal{H}$ by an extended thinning move, or
- $\mathcal{H}$ is not reduced, $\mathcal{K}$ is reduced and $\mathcal{K}$ is obtained from $\mathcal{H}$ by a sequence of consolidations, generalized destabilizations, unperturbings, and undoing of removable arcs.

We then extend the definition of $\to$ so that it is a partial order on $\overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$. In particular, if $\mathcal{H}$ is reduced, then $\mathcal{H}\to\mathcal{K}$ means that $\mathcal{K}$ is obtained from $\mathcal{H}$ by a (possibly empty) sequence of extended thinning moves.

Recall that in a poset, a least element is an element x with the property that no element is strictly less than x. In our context, we say that an element $\mathcal{K}\in\overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$ is a *least element* or *locally thin* if it is reduced and if $\mathcal{K}\to\mathcal{K}'$ implies that $\mathcal{K}=\mathcal{K}'$.

The following result follows immediately from our work above. The hypothesis that T is irreducible guarantees that in a sequence of extended thinning moves we never have a thin surface which is a sphere intersecting T exactly once.

Theorem 6.17 *Let (M,T) be a (3–manifold, graph) pair with T irreducible. Suppose that no component of ∂M is a sphere intersecting T two or fewer times. Then for all $\mathcal{H} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$, there is a least element (ie locally thin) $\mathcal{K} \in \overrightarrow{\mathrm{vp}\mathbb{H}}(M,T)$ such that $\mathcal{H} \to \mathcal{K}$.*

7 Sweepouts

Sweepouts, as in most applications of thin position, are the key tool for finding disjoint compressing discs on two sides of a thick surface. In this section, we will use $X - Y$ to denote the set-theoretic complement of Y in X, as opposed to $X \setminus Y$ which indicates the complement of an open regular neighborhood of Y in X.

Definition 7.1 Suppose that (C,T) is a vp-compressionbody and that $\Sigma \subset C$ is a trivalent graph embedded in C such that the following hold:

- $(C,T) \setminus \Sigma$ is homeomorphic to $(\partial_+ C \times I, \text{vertical arcs})$.
- Σ contains the ghost arcs of T, and no interior vertex of Σ lies on a ghost arc.
- Each boundary vertex of Σ lies on T or on $\partial_- C$.
- Any edge of T which is not a ghost arc and which intersects Σ is a bridge arc intersecting Σ in a boundary vertex.

Then Σ is a *spine* for (C,T). See Figure 16 for an example.

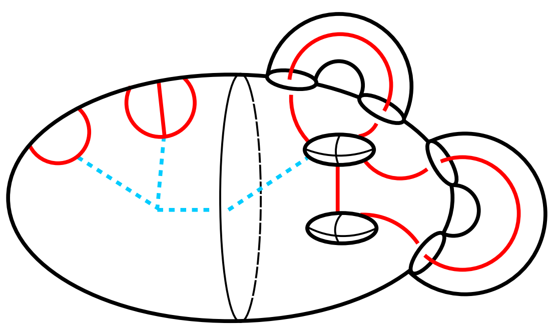

Figure 16: A spine for the vp-compressionbody from Figure 3 consists of the dashed blue graph together with the edges of T that are disjoint from $\partial_+ C$.

Suppose that $H \in \mathrm{vp}\mathbb{H}(M,T)$ is connected and that $\Sigma_\uparrow$ and $\Sigma_\downarrow$ are spines for $(H_\uparrow, T \cap H_\uparrow)$ and $(H_\downarrow, T \cap H_\downarrow)$. The manifold $M - (\Sigma_\uparrow \cup \Sigma_\downarrow)$ is homeomorphic to $H \times (0,1)$ by a map taking $T - (\Sigma_\uparrow \cup \Sigma_\downarrow)$ to vertical edges. We may extend the homeomorphism to a map $h\colon M \to I$ taking $\Sigma_\downarrow$ to -1 and $\Sigma_\uparrow$ to $+1$. The map h is called a *sweepout* of M by H. Note that $H_t = h^{-1}(t)$ is properly isotopic in $M \setminus T$ to $H \setminus T$ for each $t \in (0,1)$, that $h^{-1}(-1) = \partial_- H_\downarrow \cup \Sigma_\downarrow$, and that $h^{-1}(1) = \partial_- H_\uparrow \cup \Sigma_\uparrow$. If we perturb h by a small isotopy, we also refer to the resulting map as a sweepout.

Theorem 7.2 *Let (M,T) be a (3–manifold, graph) pair. Suppose that $F \subset (M,T)$ is an embedded surface and assume that $H \in \mathrm{vp}\mathbb{H}(M,T)$ is connected and doesn't bound a trivial vp-compressionbody on either side. Then, H can be isotoped transversely to T such that, after the isotopy, H and F are transverse and one of the following holds:*

1. $H \cap F = \varnothing$.
2. *$H \cap F \neq \varnothing$, every component of $H \cap F$ is essential in F and no component of $H \cap F$ bounds an sc-disc for H.*
3. *H is sc-weakly reducible.*

Remark 7.3 The essence of this argument can be found in many places. It originates with Gabai's original thin position argument [3] and is adapted to the context of Heegaard splittings by Rubinstein and Scharlemann [11]. A version for graphs in S^3 plays a central role in [4].

Proof Let h be a sweepout corresponding to H, as above. Perturb the map h slightly so that $h|_F$ is Morse with critical points at distinct heights. Let

$$0 = v_0 < v_1 < v_2 < \cdots < v_n = 1$$

be the critical values of $h|_F$. Let $I_i = (v_{i-1}, v_i)$. Label I_i with $\downarrow$ (resp. $\uparrow$) if some component of $F \cap H_t$ bounds a sc-disc below (resp. above) H_t for some $t \in I_i$.

Observe that, by standard Morse theory, the label(s) on I_i are independent of the choice of $t \in I_i$.

Case 1 (some interval I_i is without a label) Let $t \in I_i$. If $H_t \cap F = \varnothing$, then we are done, so suppose that $H_t \cap F \neq \varnothing$.

Suppose that some component $\zeta \subset H_t \cap F$ is inessential in F. This means that ζ bounds an unpunctured or once-punctured disc in F. Without loss of generality, we may assume that ζ is innermost in F. Let $D \subset F$ be the disc or once-punctured disc it bounds. Since ζ does not bound an sc-disc for H_t, the disc D is properly isotopic in $M \setminus T$ relative to ∂D into H_t. Let $B \subset M$ be the 3–ball bounded by D and the disc in H_t. By an isotopy supported in a regular neighborhood of B, we may isotope H_t to eliminate ζ (and possibly some other inessential curves of $H_t \cap F$). Repeating this type of isotopy as many times as necessary, we may assume that no curve of $H_t \cap F$ is inessential in F. If $H_t \cap F = \varnothing$, we have the first claim. If $H_t \cap F \neq \varnothing$, then we have claim (2).

Suppose, therefore, that each I_i has a label.

Case 2 (some I_i is labeled both $\uparrow$ and $\downarrow$) Since H_t is transverse to F for each $t \in I_i$, we have claim (3).

Case 3 (there is an i such that I_i is labeled $\downarrow$ and I_{i+1} is labeled $\uparrow$, or vice versa) The labels cannot change from I_i to I_{i+1} at any tangency other than a saddle tangency. Let $\epsilon > 0$ be smaller than the lengths of the intervals I_i and I_{i+1}. Since H_t is orientable, under the projections of $H_{v_i-\epsilon}$ and $H_{v_i+\epsilon}$ to H, the 1–manifold $H_{v_i-\epsilon} \cap F$ can be isotoped to be disjoint from $H_{v_i+\epsilon} \cap F$. Since some component of the former set bounds an sc-disc on one side of H and some component of the latter set bounds an sc-disc on the other side of H, we have claim (3) again.

Case 4 (all I_i are labeled $\downarrow$ and not $\uparrow$, or all I_i are labeled $\uparrow$ and not $\downarrow$) Without loss of generality, assume that each I_i is labeled $\uparrow$ and not $\downarrow$. In particular, I_1 is labeled $\uparrow$ and not $\downarrow$. Fix $t \in I_1$ and consider H_t. Since H does not bound a trivial vp-compressionbody to either side, the spine for $(H_\downarrow, T \cap H_\downarrow)$ has an edge e. Since I_i is below the lowest critical point for $h|_F$, the components of $F \cap (H_t)_\downarrow$ intersecting e are a regular neighborhood in F of $F \cap e$. Let $D_\downarrow$ be a meridian disc for e with boundary in H_t and which is disjoint from $F \cap (H_t)_\downarrow$. Since I_1 is labeled $\uparrow$, there is a component $\zeta \subset H_t \cap F$ such that ζ bounds an sc-disc $D_\uparrow$ for H_t in $(H_t)_\uparrow$. The pair $\{D_\uparrow, D_\downarrow\}$ is then a weak reducing pair for H_t, giving claim (3). □

Remark 7.4 Observe that in claim (3), we can only conclude that H is sc-weakly reducible, not that H is c-weakly reducible. This arises in case 4 of the proof, when we use an edge of the spine to produce an sc-disc. This is one reason for allowing semicompressing and semicut discs in weak reducing pairs.

Corollary 7.5 *Suppose that $H \in \mathrm{vp}\mathbb{H}(M,T)$ is connected and sc-strongly irreducible. If a component S of ∂M is c-compressible, then the component of $(M,T) \setminus H$ containing S is a trivial product compressionbody.*

Proof Let $F \subset (M, T)$ be a c-disc for S. Let (C, T_C) and (E, T_E) be the components of $(M, T) \setminus H$, with $S \subset \partial_- C$. If (E, T_E) is a trivial compressionbody, we may isotope F out of E to be contained in C. This contradicts the fact that $\partial_- C$ is c-incompressible in C. Hence, (E, T_E) is not a trivial product compressionbody. Since $S \subset \partial_- C$ is c-compressible, it either has positive genus or intersects T at least three times. In particular, (C, T_C) is not a trivial ball compressionbody. Suppose, for a contradiction, that (C, T_C) is not a trivial product compressionbody. Then by Theorem 7.2 H can be isotoped transversely to T such that after the isotopy one of the following holds:

(1) $H \cap F = \varnothing$.

(2) $H \cap F \neq \varnothing$, every component of $H \cap F$ is essential in F and no component of $H \cap F$ bounds an sc-disc for H.

Since $\partial_- C$ is c-incompressible in C, by Lemma 3.5, the first conclusion cannot hold. Since no curve in a disc or once-punctured disc is essential, the second conclusion is also impossible. Thus, (C, T_C) is a trivial product compressionbody. □

Theorem 7.6 (properties of locally thin surfaces) *Suppose (M, T) is a (3–manifold, graph) pair with T irreducible. Let $\mathcal{H} \in \overrightarrow{\text{vp}\mathbb{H}}(M, T)$ be locally thin. Then the following hold:*

(1) *$\mathcal{H}$ is reduced.*

(2) *Each component of $\mathcal{H}^+$ is sc-strongly irreducible in the complement of $\mathcal{H}^-$.*

(3) *No component of $(M, T) \setminus \mathcal{H}$ is a trivial product compressionbody between $\mathcal{H}^-$ and $\mathcal{H}^+$.*

(4) *Every component of $\mathcal{H}^-$ is c-essential in (M, T).*

(5) *If (M, T) is irreducible and if $\mathcal{H}$ contains a 2–sphere disjoint from T, then $T = \varnothing$ and $M = S^3$ or $M = B^3$.*

Proof Without loss of generality, we may assume that T has no vertices (drilling them out to turn them into components of ∂M if necessary). Claims (1) and (3) are immediate from the definition of locally thin. If some component of $\mathcal{H}^+$ is sc-weakly reducible in $(M, T) \setminus \mathcal{H}^-$, then since T is irreducible, we could perform an elementary thinning sequence, contradicting the definition of locally thin. Thus, (2) also holds.

Next we show that each component of $\mathcal{H}^-$ is c-incompressible. Suppose, therefore, that $S \subset \mathcal{H}^-$ is a thin surface. We first show that S is c-incompressible and then that it is not ∂–parallel. Suppose that S is c-compressible by a c-disc D. By an innermost disc argument, we may assume that no curve of $D \cap (\mathcal{H}^- \setminus S)$ is an essential curve in $\mathcal{H}^-$. By passing to an innermost disc, we may also assume that $D \cap (\mathcal{H}^- \setminus S) = \varnothing$. Let (M_0, T_0) be the component of $(M, T) \setminus \mathcal{H}^-$ containing D. Let $H = \mathcal{H}^+ \cap M_0$ and recall that H is connected. By Corollary 7.5 applied to H in (M_0, T_0), the vp-compressionbody between S and H is a trivial product compressionbody. This contradicts property (3) of locally thin multiple vp-bridge surfaces. Thus, each component of $\mathcal{H}^-$ is c-incompressible.

We now show no sphere component of $\mathcal{H}^-$ bounds a 3–ball in $M \setminus T$. Suppose that $S \subset \mathcal{H}^-$ is such a sphere and let $B \subset M \setminus T$ be the 3–ball it bounds. By passing to an innermost such sphere, we may assume that no component of $\mathcal{H}^-$ in the interior of B is a 2–sphere. If there is a component of $\mathcal{H}^-$ in the interior of B, that component would be compressible, a contradiction. Thus the intersection H of $\mathcal{H}$ with the interior of B is a component of $\mathcal{H}^+$. The surface H is a Heegaard splitting of B. If H is a sphere, it is parallel to S, contradicting (3). If H is not a sphere, then by [22] it is stabilized, contradicting (1). Thus, each component of $\mathcal{H}^-$ is c-incompressible in (M, T) and not a sphere bounding a 3–ball in $M \setminus T$. In particular, if (M, T) is irreducible no component of $\mathcal{H}^-$ is a sphere disjoint from T.

We now show that no component of $\mathcal{H}^-$ is ∂–parallel. Since T may be a graph and not simply a link, this does not follow immediately from our previous work. Suppose, to obtain a contradiction, that a component F of $\mathcal{H}^-$ is boundary parallel in the exterior of T. An analysis (which we provide momentarily) of the proof of [19, Theorem 9.3] shows that $\mathcal{H}$ either has a perturbation or a generalized stabilization or is removable. We elaborate on this:

As in [19, Lemma 3.3], since all components of $\mathcal{H}^-$ are c-incompressible, we may assume that the product region W between F and $\partial(M \setminus T)$ has interior disjoint from $\mathcal{H}^-$. (That is, F is innermost.) Observe that W is a compressionbody with $F = \partial_+ W$ and the component $H = \mathcal{H}^+ \cap W$ is a vp-bridge surface for $(W, T \cap W)$.

If there is a component of $T \cap W$ with both endpoints on F, then F must be a 2–sphere, and W is a 3–ball with $T \cap W$ a ∂–parallel arc. In this case, if $T \cap W$ is disjoint from H, then H is a Heegaard surface for the solid torus obtained by drilling out the arc $T \cap W$. Since H is not stabilized, it must then be a torus. In particular, since W is a 3–ball, this implies that H is meridionally stabilized, a contradiction. Thus, in particular, if $T \cap W$ has a component with both endpoints on F, then H

intersects each component of $T \cap W$. We now perform a trick to guarantee that this is also the case when $T \cap W$ does have a component with endpoints on F.

Let $(\overline{W}, \overline{T})$ be the result of removing from $(W, T \cap W)$ an open regular neighborhood of all edges of $T \cap W$ which are disjoint from H. (By our previous remark, all such edges have both endpoints on $\partial W \setminus F$.) Then $\overline{T}$ is a 1–manifold properly embedded in $\overline{W}$ with no edge disjoint from H.

If $\overline{T}$ has at least one edge, then by [19, Theorem 3.5] (which is a strengthening of [18, Theorem 3.1]), one of the following occurs:

- (i) $H \in \mathrm{vp}\mathbb{H}(\overline{W}, \overline{T})$ is stabilized, boundary-stabilized along $\partial \overline{W} \setminus F$, perturbed or removable.
- (ii) H is parallel to F by an isotopy transverse to T.

Consider possibility (i). If $H \in \mathrm{vp}\mathbb{H}(\overline{W}, \overline{T})$ is stabilized, boundary-stabilized along $\partial \overline{W}$, or removable, then $H \in \mathrm{vp}\mathbb{H}(W, T)$ would have a generalized stabilization or it would be removable with removing discs disjoint from the vertices of T, an impossibility. If $H \in \mathrm{vp}\mathbb{H}(\overline{W}, \overline{T})$ is perturbed, $H \in \mathrm{vp}\mathbb{H}(W, T)$ has a perturbation (since $\overline{T}$ is a 1–manifold), also an impossibility. Thus, (i) does not occur. Possibility (ii) does not occur since none of the vp-compressionbodies of $(M, T) \setminus \mathcal{H}$ are trivial product compressionbodies adjacent to $\mathcal{H}^-$.

We may therefore assume that $\overline{T}$ has no edges. Then by [14], H is either parallel to F by an isotopy transverse to T or is boundary-stabilized along $\partial \overline{W}$. The former situation contradicts the assumption that $(M, T) \setminus \mathcal{H}$ contains no trivial product compressionbody adjacent to $F \subset \mathcal{H}^-$. In the latter situation, since H is boundary-stabilized in $\overline{W}$, it has a generalized stabilization as a surface in $\overrightarrow{\mathrm{vp}\mathbb{H}}(W, T)$, contradicting the assumption that $\mathcal{H}$ is reduced. Thus, once again, F is not boundary-parallel in $M \setminus T$. We have shown therefore that no component of $\mathcal{H}^-$ is ∂–parallel, and thus that each component of $\mathcal{H}^-$ is c-essential in (M, T).

It remains to show that if (M, T) is irreducible and if some component of $\mathcal{H}$ is a sphere disjoint from T, then $T = \varnothing$ and $M = B^3$ or $M = S^3$. Assume that (M, T) is irreducible. We have already remarked that since each component of $\mathcal{H}^-$ is c-essential, no component of $\mathcal{H}^-$ is a sphere disjoint from T. We now show that no component H of $\mathcal{H}^+$ is a sphere disjoint from T, unless $T = \varnothing$ and M is S^3 or B^3. Suppose that there is such a component $H \subset \mathcal{H}^+$. Let (C, T_C) and (D, T_D) be the vp-compressiobodies on either side of H. By the definition of vp-compressionbody, the surfaces $\partial_- C$ and $\partial_- D$ are the unions of spheres, and T_C

and T_D are the unions of ghost arcs. Consider the graphs $\Gamma_C = \partial_- C \cup T_C$ and $\Gamma_D = \partial_- D \cup T_D$ (thinking of the components of $\partial_- C$ and $\partial_- D$ as vertices of the graph). By the definition of vp-compressionbody since H is a sphere disjoint from T, the graphs Γ_C and Γ_D are the union of trees. If either Γ_C or Γ_D has an edge, then a leaf of Γ_C or Γ_D is a sphere intersecting T exactly once. This contradicts the irreducibility of (M, T). Consequently, both T_C and T_D are empty. Since no spherical component of $\mathcal{H}^-$ is disjoint from T, this implies that $\partial_- C \cup \partial_- D$ is a subset of ∂M. Since $M \setminus T$ is irreducible, this implies that $\partial_- C \cup \partial_- D$ is either empty or a single sphere. Consequently, M is either S^3 or B^3 and $T = \varnothing$. □

8 Decomposing spheres

The goal of this section is to show that if we have a bridge surface for a composite knot or graph, we can untelescope it so that a summing sphere shows up as a thin level.

We start with a simple observation (likely well known) that compressing essential twice and thrice-punctured spheres results in a component which is still essential. The proof is straightforward and similar to that of Lemma 3.3, so we leave it to the reader.

Lemma 8.1 *Assume that (M, T) is a 3–manifold graph pair with T irreducible. Suppose that $P \subset (M, T)$ is an essential sphere with $|P \cap T| \le 3$. Let P' be the result of compressing P along an sc-disc D. Then at least one component of $P' \subset (M, T)$ is an essential sphere intersecting T at most three times.*

Theorem 8.2 *Suppose that (M, T) is a (3–manifold, graph) pair with T irreducible. Suppose that there is an essential sphere $P \subset (M, T)$ such that $|P \cap T| \le 3$. If $\mathcal{H} \in \overrightarrow{\text{vp}\mathbb{H}}(M, T)$ is locally thin, then some component F of $\mathcal{H}^-$ is an essential sphere with $|F \cap T| \le 3$. Furthermore, $|F \cap T| \le |P \cap T|$ and F can be obtained from P by a sequence of compressions using sc-discs.*

Proof Let $P \subset (M, T)$ be an essential sphere such that $|P \cap T| \le 3$. As in the proof of Lemmas 3.3 and 8.1, since T is irreducible, if P_0 is a sphere resulting from an sc-compression of P, then $|P_0 \cap T| \le |P \cap T|$.

Without loss of generality, we may assume that the given P was chosen so that no sequence of isotopies and sc-compressions reduces $|P \cap \mathcal{H}^-|$. The intersection $P \cap \mathcal{H}^-$ consists of a (possibly empty) collection of circles. We show it is, in fact, empty. Suppose, for a contradiction, that γ is a component of $|P \cap \mathcal{H}^-|$. Without loss of generality, we may suppose it is innermost on P. Let $D \subset P$ be the unpunctured disc or once-punctured disc which it bounds. Since $\mathcal{H}^-$ is c-incompressible, γ must bound a unpunctured or once-punctured disc E in $\mathcal{H}^-$. Thus, if $|P \cap \mathcal{H}^-| \neq 0$, then there is a component of the intersection which is inessential in $\mathcal{H}^-$.

Let $\zeta \subset P \cap \mathcal{H}^-$ be a component which is inessential in $\mathcal{H}^-$ and which, out of all such curves, is innermost in $\mathcal{H}^-$. Let $E \subset \mathcal{H}^-$ be the unpunctured or once-punctured disc it bounds. Observe that ζ also bounds an unpunctured or once-punctured disc on P. If E is not an sc-disc for P, then we can isotope P to reduce $|P \cap \mathcal{H}^-|$, contradicting our choice of P. Thus, E is an sc-disc. By Lemma 8.1, compressing P along E creates two spheres, at least one of which intersects T no more than three times and is essential in the exterior of T. Since this component intersects $\mathcal{H}^-$ fewer times than does P, we have contradicted our choice of P. Hence $P \cap \mathcal{H}^- = \varnothing$.

We now consider intersections between P and $\mathcal{H}^+$. Since P is disjoint from $\mathcal{H}^-$, we may apply Theorem 7.2 to the component (W, T_W) of $(M, T) \setminus \mathcal{H}^-$ containing P. We apply the theorem with $H = \mathcal{H}^+ \cap W$ and $F = P$. If some component of $\mathcal{H}^-$ is a once-punctured sphere, we are done, so assume that no component of $\mathcal{H}^-$ is a once-punctured sphere. By cutting open along $\mathcal{H}^-$ and replacing (M, T) with the component containing P, we may assume that $\mathcal{H}^- = \varnothing$ and that $H = \mathcal{H}$ is connected. Apply Theorem 7.2 to P (in place of F) to see that we can isotope H transversely to T in $M \setminus \mathcal{H}^-$ so that one of the following occurs:

(1) $H \cap P = \varnothing$.

(2) $H \cap P$ is a nonempty collection of curves, each of which is essential in P.

(3) H is sc-weakly reducible.

Since $\mathcal{H}$ is locally thin in $\overrightarrow{\mathrm{vp}\mathbb{H}}(M, T)$, (3) does not occur. Since P contains no essential curves, (2) does not occur. Thus, $H \cap P = \varnothing$.

Let (C, T_C) be the component of $M \setminus \mathcal{H}$ containing P. By Lemma 3.3, after some sc-compressions, P is parallel is a component of $\mathcal{H}^- \cup \partial M$. By Lemma 8.1, P is parallel to a component of $\mathcal{H}^-$, and we are done. □

References

[1] **D Bachman**, *Thin position with respect to a Heegaard surface*, preprint (2001) Available at `http://pzacad.pitzer.edu/~dbachman/Papers/tnbglex.pdf`

[2] **A J Casson**, **C M Gordon**, *Reducing Heegaard splittings*, Topology Appl. 27 (1987) 275–283 MR

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

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

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

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

[7] **H Howards**, **Y Rieck**, **J Schultens**, *Thin position for knots and 3–manifolds: a unified approach*, from "Workshop on Heegaard Splittings" (C Gordon, Y Moriah, editors), Geom. Topol. Monogr. 12, Geom. Topol. Publ., Coventry (2007) 89–120 MR

[8] **J Johnson**, *Calculating isotopy classes of Heegaard splittings*, preprint (2010) arXiv

[9] **F Manjarrez-Gutiérrez**, *Circular thin position for knots in* S^3, Algebr. Geom. Topol. 9 (2009) 429–454 MR

[10] **Y Rieck**, **E Sedgwick**, *Thin position for a connected sum of small knots*, Algebr. Geom. Topol. 2 (2002) 297–309 MR

[11] **H Rubinstein**, **M Scharlemann**, *Comparing Heegaard splittings of non-Haken 3–manifolds*, Topology 35 (1996) 1005–1026 MR

[12] **M Scharlemann**, *Heegaard splittings of compact 3–manifolds*, from "Handbook of geometric topology" (R J Daverman, R B Sher, editors), North-Holland, Amsterdam (2002) 921–953 MR

[13] **M Scharlemann**, **J Schultens**, *Comparing Heegaard and JSJ structures of orientable 3–manifolds*, Trans. Amer. Math. Soc. 353 (2001) 557–584 MR

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

[15] **M Scharlemann**, **A Thompson**, *Thin position and Heegaard splittings of the 3–sphere*, J. Differential Geom. 39 (1994) 343–357 MR

[16] **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., Providence, RI (1994) 231–238 MR

[17] **M Scharlemann**, **M Tomova**, *Uniqueness of bridge surfaces for 2–bridge knots*, Math. Proc. Cambridge Philos. Soc. 144 (2008) 639–650 MR

[18] **S A Taylor**, **M Tomova**, *Heegaard surfaces for certain graphs in compressionbodies*, Rev. Mat. Complut. 25 (2012) 511–555 MR

[19] **S A Taylor**, **M Tomova**, *C-essential surfaces in* (3*–manifold, graph*) *pairs*, Comm. Anal. Geom. 21 (2013) 295–330 MR

[20] **S A Taylor**, **M Tomova**, *Additive invariants for knots, links and graphs in* 3*–manifolds*, preprint (2016) arXiv

[21] **A Thompson**, *Thin position and bridge number for knots in the* 3*–sphere*, Topology 36 (1997) 505–507 MR

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

Department of Mathematics and Statistics, Colby College
Waterville, ME, United States

Department of Mathematics, University of Iowa
Iowa City, IA, United States

scott.taylor@colby.edu, maggy-tomova@uiowa.edu

Received: 16 July 2016 Revised: 30 November 2017