AGT

Algebraic & Geometric Topology

Volume 24 (2024)

On keen weakly reducible bridge spheres

PUTTIPONG PONGTANAPAISAN
DANIEL RODMAN

msp

msp

Algebraic & Geometric Topology **24**:8 (2024) 4201–4235
DOI: 10.2140/agt.2024.24.4201
Published: 17 December 2024

On keen weakly reducible bridge spheres

PUTTIPONG PONGTANAPAISAN
DANIEL RODMAN

A bridge sphere is said to be keen weakly reducible if it admits a unique pair of disjoint compressing disks on opposite sides. In particular, such a bridge sphere is weakly reducible, not perturbed, and not topologically minimal in the sense of David Bachman. In terms of Jennifer Schultens' width complex, a link in bridge position with respect to a keen weakly reducible bridge sphere is distance one away from a local minimum. We give infinitely many examples of keen weakly reducible bridge spheres for links in b bridge position for $b \geq 4$.

57K10, 57K20, 57K30

1 Introduction

Suppose that we have a decomposition of the 3–sphere $S^3 = V_+ \cup_\Sigma V_-$ where V_+ and V_- are 3–balls and Σ is a 2–sphere. A link $L \subset S^3$ intersecting Σ transversely is said to be in *bridge position* with respect to Σ if $L \cap V_+ = \alpha_+$ and $L \cap V_- = \alpha_-$, where α_+ and α_- are b–strand trivial tangles. The punctured sphere $\Sigma_L = \Sigma \backslash L$ is called a *b–bridge sphere*. To each bridge sphere, we can assign a disk complex $\mathcal{D}(\Sigma_L)$, which is a simplicial complex whose vertices are isotopy classes of compressing disks in $S^3 \backslash L$ for Σ_L and whose k simplices are spanned by $k+1$ vertices with pairwise disjoint representatives.

We say that Σ_L is *topologically minimal* if one of the following holds:

(1) $\mathcal{D}(\Sigma_L) = \varnothing$.

(2) There exists $i \in \mathbb{N} \cup \{0\}$ such that the i^{th} homotopy group of $\mathcal{D}(\Sigma_L)$ is nontrivial.

The *topological index* of Σ_L is defined to be 0 if $\mathcal{D}(\Sigma_L) = \varnothing$, or the smallest i such that $\pi_{i-1}(\mathcal{D}(\Sigma_L))$ is nontrivial if $\mathcal{D}(\Sigma_L) \neq \varnothing$. The notion of topological minimality was introduced by David Bachman [2010] as a generalization of useful concepts such as incompressibility and strong irreducibility of surfaces in a 3–manifold. It turns out that topologically minimal surfaces possess desirable properties. For instance, in an irreducible 3–manifold, a topologically minimal surface can be isotoped to intersect an incompressible surface in such a way that any intersection loop is essential in both surfaces. Furthermore, the concept of topological minimality gave rise to examples of 3–manifolds containing arbitrarily many nonminimal genus, unstabilized Heegaard surfaces that are weakly reducible [Bachman 2013]. Moriah [2007] dubbed these examples "the nemesis of Heegaard splittings" as they are difficult to find.

Conjecturally, there is a special and mysterious relationship between topologically minimal surfaces and geometrically minimal surfaces, which are surfaces whose mean curvature is identically zero. Every geometrically minimal surface has a Morse index, which roughly speaking counts the maximal number of directions the surface can be deformed so as to decrease its area. Freedman, Hass and Scott [Freedman et al. 1983] showed that every surface of topological index zero is isotopic to a geometrically minimal surface of Morse index zero. By works of Pitts and Rubinstein [1987] and of Ketover, Liokumovich, and Song [Ketover et al. 2019], a Heegaard surface of topological index one is isotopic to a geometrically minimal surface of Morse index at most one. Campisi and Torres [2020] showed that the genus two Heegaard surface of the 3–sphere has topological index three. By Urbano [1990], this Heegaard surface must have Morse index at least six. Thus, it is not true in general that a surface of topological index k is isotopic to a surface of Morse index at most k, but the precise connection is not well understood.

One can ask the interesting question of which surfaces are topologically minimal. Several authors have given examples of topologically minimal Heegaard surfaces [Bachman and Johnson 2010; Campisi and Rathbun 2018; Campisi and Torres 2020; Lee 2015] and bridge surfaces [Lee 2016; Pongtanapaisan and Rodman 2021; Rodman 2018]. Heegaard surfaces that are not topologically minimal have also been studied by several authors who constructed *keen weakly reducible* Heegaard surfaces. That is, each of these surfaces possesses a unique *weak reducing pair*, a pair of compressing disks on opposite sides of the surface whose boundaries are disjoint. By a result of McCullough [1991], the disk complex of the boundary of a handlebody is contractible. Thus having a unique pair of weak reducing disks on distinct sides of a Heegaard splitting means that in the disk complex, there is a unique edge connecting the two contractible subcomplexes corresponding the two handlebodies, resulting in a contractible disk complex for the Heegaard surface. The examples of keen weakly reducible Heegaard surfaces in the literature with simple descriptions include the canonical Heegaard surface of a surface bundle whose monodromy has sufficiently high translation distance by Johnson [2012], some Heegaard surfaces arising from self-amalgamations by E and Lei [2014], and certain unstabilized genus three Heegaard surfaces in irreducible and orientable 3–manifolds by Kim [2016]. More complicated constructions of keen weakly reducible Heegaard surfaces of genus $g \geq 3$ can also be found in [E 2017; Liang et al. 2018].

The goal of this paper is to provide infinitely many examples of nontopologically minimal bridge spheres, which are lacking in the literature, by verifying that the canonical bridge sphere for certain links in plat position is keen weakly reducible. Such links are obtained by "amalgamating" two types of links whose canonical bridge spheres are topologically minimal. Keen weakly reducible bridge spheres also belong to a family of surfaces with finitely many pairs of disjoint compressing disks [E and Zhang 2023], which is interesting in its own right.

Theorem 1.1 *There exist infinitely many links with keen weakly reducible bridge spheres.*

This paper is organized as follows. In Section 2, we discuss properties of a keen weakly reducible bridge sphere related to perturbations of bridge spheres, thin position of links, and essential surfaces in the link

exterior. In Section 3, we define the notion of a plat position for a link, consider a particular family of links in plat position, and describe useful positions of curves on a punctured sphere with respect to a train track. In Section 4, we characterize the behaviors of curves that bound disks above or below the bridge sphere. In Section 5, we use a criterion presented in [Cho 2008] to show that keen weakly reducible bridge spheres are not topologically minimal.

Acknowledgements

The authors would like to thank Roman Aranda, David Bachman, Ryan Blair, Charlie Frohman, and Maggy Tomova for helpful conversations. We thank the referees for finding a mistake in an earlier draft. Research conducted for this paper is supported by the Pacific Institute for the Mathematical Sciences (PIMS). The research and findings may not reflect those of the institute.

2 Consequences of being keen weakly reducible

In this section, we discuss some consequences of putting a link in bridge position with respect to a keen weakly reducible bridge sphere. We remark that a priori keen weakly reducible bridge spheres are not necessarily canonical bridge spheres for links in plat position.

2.1 Unperturbed bridge spheres

Let L be a link in bridge position with respect to Σ. Then $L \cap V_+ = \alpha_+$ is a collection of disjoint embedded arcs with the property that there exists an isotopy (rel $\partial\alpha^+$) taking α_+ into Σ. For each arc α_+^i of α_+, the trace of such an isotopy is a disk called a *bridge disk* D_+^i. From each bridge disk D_+^i we can obtain a compressing disk dD_+^i called the *frontier* of D_+^i using the construction $dD_+^i = (\partial \overline{N}(D_+^i)) \cap V_+$. Analogous definitions can be made for $L \cap V_- = \alpha_-$.

We say that a bridge sphere Σ_L is *perturbed* if there exist two bridge disks $D_+^1 \subset V_+$ and $D_-^1 \subset V_-$ such that $D_+^1 \cap D_-^1$ is a single point contained in L. It is an interesting problem to search for unperturbed bridge spheres for a link up to isotopy since a perturbed bridge can always be obtained from a bridge sphere that is not perturbed by an isotopy which introduces a maximal point and a minimal point as shown in Figure 1. In some cases, the only destabilized bridge sphere is the one that realizes the bridge number [Otal 1985; Ozawa 2011; Zupan 2011]. Another common way to show that a bridge sphere Σ_L for a nontrivial link L is unperturbed is to show that there is no weak reducing pair for Σ_L. Being keen weakly reducible implies the following.

Proposition 2.1 *If Σ_L is keen weakly reducible, then Σ_L is unperturbed.*

It is well known that a perturbed bridge sphere has a weak reducing pair; we prove that result here, and show that such a bridge sphere in fact has at least two weak reducing pairs.

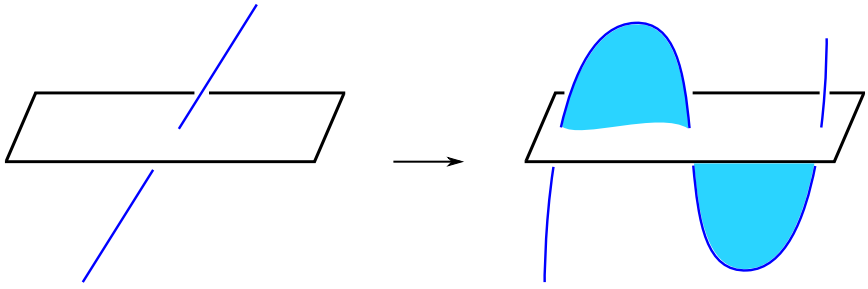

Figure 1: Introducing a canceling pair of critical points.

Proof If Σ_L is a perturbed bridge sphere for a link L in 2–bridge position, then L must be the unknot and there exists a unique compressing disk D above and a unique compressing disk E below. Furthermore, $D \cap E \neq \varnothing$, which implies that Σ_L does not admit a weak reducing pair, and therefore Σ_L cannot be keen weakly reducible. To complete the proof, we consider perturbed bridge spheres for links in b–bridge position, where $b \geq 3$.

Suppose that Σ_L is perturbed. By definition, there exist bridge disks $D^1_+ \subset V_+$ and $D^1_- \subset V_-$ such that $D^1_+ \cap D^1_- = \{p\} \in L$. Let $\mathscr{A}_+$ be a set of b disjoint bridge disks for Σ, each corresponding to one of the components of α_+, and suppose further that $D^1_+ \in \mathscr{A}_+$. Let $\mathscr{A}_-$ be a similarly defined set of bridge disks below Σ with $D^1_- \in \mathscr{A}_-$. (We are able to define $\mathscr{A}_+$ and $\mathscr{A}_-$ *after* D^1_+ and D^1_- by [Scharlemann 2005, Lemma 3.2].) The elements of $\mathscr{A}_-$ may or may not intersect the interior of the arc $D^1_+ \cap \Sigma$. Below, we describe how if they do, we can replace them with another set of b disjoint bridge disks below Σ, each of which is disjoint from the interior of $D^1_+ \cap \Sigma$.

Suppose that the elements of $\mathscr{A}_-$ intersect the interior of $D^1_+ \cap \Sigma$. Consider a point q of intersection closest to p. Let $D'_- \in \mathscr{A}_-$ denote the bridge disk containing q. We perform a surgery on D'_- as depicted in Figure 2, resulting in a new disk D''_-. Notice that D''_- is disjoint from the other elements of $\mathscr{A}_-$, and D'_- and D''_- both correspond to the same bridge arc. In slight abuse of notation, we will replace D'_- with D''_- in the collection $\mathscr{A}_-$. After this replacement, $\mathscr{A}_-$ remains a collection of pairwise disjoint bridge disks for the bridge arcs below Σ. The difference is that now, the elements of $\mathscr{A}_-$ intersect the interior of $D^1_+ \cap \Sigma$ in one fewer point.

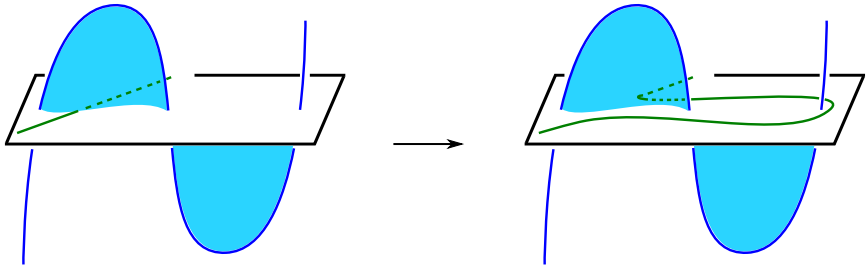

Figure 2: Bridge disks below Σ_K can be isotoped to intersect D^1_+ in two points in K.

We can repeatedly perform such surgeries until the bridge disks of $\mathscr{A}_-$ are all disjoint from the interior of $D^1_+ \cap \Sigma$. It follows that the bridge disks of $\mathscr{A}_-$ intersect D^1_+ only in the two points of $D^1_+ \cap L \cap \Sigma$, each of which intersects a bridge disk of $\mathscr{A}_-$. Since L is in a b bridge position with $b > 2$, these two bridge disks must be distinct. In addition to these two, there must be at least one more bridge disk $D^2_- \in \mathscr{A}_-$ since $b \geq 3$, and so D^2_- is disjoint from D^1_+. Therefore, dD^1_+ and dD^2_- comprise a weak reducing pair for Σ_L.

Now consider dD^1_-. We can mimic the trick in the previous paragraph so that a particular set of b pairwise disjoint bridge disks above Σ intersects D^1_- only in the two points of $D^1_- \cap L \cap \Sigma$. Then there is some bridge disk D^2_+ disjoint from D^1_-, which means that dD^1_- and dD^2_+ comprise another weak reducing pair for Σ_L. Therefore, a perturbed bridge sphere never admits a unique weak reducing pair and can never be keen weakly reducible. □

2.2 Width complex

Suppose that L is a link and $h\colon S^3 \to \mathbb{R}$ is the standard Morse function. Assume also that $h|_L$ is a Morse function. Suppose that $c_1 < \cdots < c_n$ are critical values of $h|_L$. Consider $h^{-1}(r_i)$, where r_i is a regular value between c_i and c_{i+1}. We say that a level sphere $h^{-1}(r_i)$ is a *thin level* if $|h^{-1}(r_{i-1}) \cap L| > |h^{-1}(r_i) \cap L|$ and $|h^{-1}(r_i) \cap L| < |h^{-1}(r_{i+1}) \cap L|$. On the other hand, a level sphere $h^{-1}(r_i)$ is a *thick level* if $|h^{-1}(r_{i-1}) \cap L| < |h^{-1}(r_i) \cap L|$ and $|h^{-1}(r_i) \cap L| > |h^{-1}(r_{i+1}) \cap L|$. We say that a disk $D \subseteq S^3 \backslash L$ is a *strong upper* (*resp. lower*) *disk* with respect to $h^{-1}(r_i)$ if

(1) $\partial D = \alpha \cup \beta$ where $\alpha \subset L$ contains exactly one maximal (resp. minimal) point and β is an arc in $h^{-1}(r_i)$, and

(2) the interior of D contains no critical point with respect to the height function h.

If there exists a strong upper disk and a strong lower disk intersecting in exactly one point lying in L (see Figure 1, for instance), then there is an isotopy that cancels a maximal point and a minimal point. We call such a move a *type I move*. On the other hand, if there exists a strong upper disk and a strong lower disk that are disjoint, then there is an isotopy that interchanges a maximal point and a minimal point. We call such a move a *type II move*.

Schultens [2009] associated to a knot K a graph called the *width complex of* K to understand the structure of the collection of Morse embeddings of a fixed knot K. Two embeddings k and k' of K are considered to be *equivalent* if their thin and thick levels are isotopic. With this definition of equivalence, each vertex of the width complex is an equivalence class of embeddings of K such that $h|_K$ is a Morse function. An edge connects two vertices representing embeddings k and k' if k differs from k' by one of the following moves: a type I move, the inverse of a type I move, a type II move, or the inverse of a type II move. Schultens proved the following interesting result.

Theorem 2.2 [Schultens 2009] *The width complex of a knot is connected.*

The proof of Theorem 2.2 uses the fact that projections of k and k' to the vertical plane differ by a finite number of Reidemeister moves and planar isotopy. Furthermore, each of these local moves either affects an embedding by a type I or a type II move or does not alter the equivalence class at all. As any two projections of a multicomponent link L are also related by Reidemeister moves and planar isotopy, it follows that the width complex of a multicomponent link is also connected.

A vertex that is particularly interesting is one representing an embedding that admits no type I or type II moves. Such an embedding is said to be in *locally thin position*.

Proposition 2.3 *Suppose that l is an embedding of a link L in bridge position with respect to a keen weakly reducible bridge sphere. In the width complex of L, there is an edge between l and an embedding l' of L in a locally thin position.*

Proof Let $D \subseteq V_+$ and $E \subseteq V_-$ be a weak reducing pair for a keen weakly reducible bridge sphere Σ_L.

Claim *∂D and ∂E each cut out a twice punctured disk from Σ_L.*

Proof of claim Suppose that ∂D cuts Σ_L into two components F_1 and F_2, where each component is a punctured disk containing more than two punctures. The loop ∂E is contained in one of the components, say F_1. There exists at least one bridge disk D^1_+ such that $\partial D^1_+ = \alpha \cup \beta$ where $\alpha \subset L$ and $\beta \subset F_2$. Then, $d D^1_+$ and E give rise to a weak reducing pair distinct from D and E, which is a contradiction. The same argument also implies ∂E cuts out a twice-punctured disk from Σ_L. □

Observe that D cuts off a 3–ball containing a unique bridge disk, which is a strong upper disk disjoint from a strong lower disk contained in a 3–ball cut off by E. This pair of disks gives rise to a type II move. After the type II move is performed, there are neither type I nor type II moves left to perform because any pair of strong upper disk and strong lower disk (intersecting in one point of L or mutually disjoint) that emerges after the type II move on D and E will yield a distinct pair of strong upper disk and strong lower disk on Σ_L, and hence Σ_L admits more than one weak reducing pair, which is a contradiction. □

After a type II move is performed along D and E, a thin level emerges. This thin level is incompressible because a compressing disk for this level would imply the existence of another weak reducing pair different from D and E. Thus, we obtain the following corollary.

Corollary 2.4 *A link with a keen weakly reducible bridge sphere contains an essential meridional surface in its exterior.*

3 Setting

In this section, we redevelop and summarize several tools and concepts of Johnson and Moriah [2016]. Specifically, Section 3.1 is a brief summary of Johnson and Moriah's plat links and accompanying tools such as their σ_i and π_y projection maps. Then in Section 3.3, we develop Johnson and Moriah's taos,

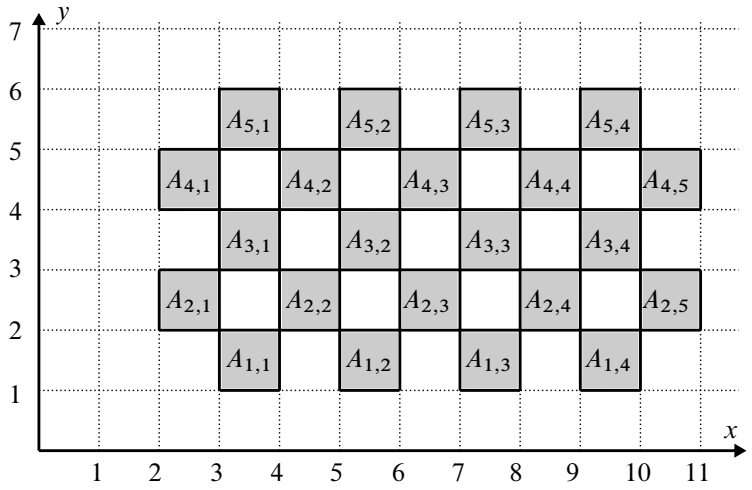

Figure 3: A $(6,5)$–plat structure.

eyelets, and train tracks and adapt them slightly to our situation. Finally in Section 3.4 we define the concepts of carried and almost carried arcs, loops, and graphs in a manner very similar to that of Johnson and Moriah, differing only in some minor ways that suit our purposes.

3.1 Plat positions

Consider the standard Morse function $h\colon S^3 \to \mathbb{R}$ with exactly one maximum, $+\infty$, and one minimum, $-\infty$. Let $\alpha \subset S^3$ be a strictly increasing arc such that $\partial\alpha = \{\pm\infty\}$. We identify $S^3\backslash\alpha$ with $\mathbb{R}^3$ with Cartesian coordinates (x, y, z) in such a way that the xz–plane lies in $h^{-1}(0)$, and more generally, for each $t \in \mathbb{R}$, the plane $y = t$ lies in $h^{-1}(t)$. We orient our perspective so that the x–axis is horizontal, the y–axis is vertical, and the z–axis points towards the reader. (This allows us to use terms like "up", "down", "left", and "right".) We denote $h^{-1}(t)$ by P_t.

For each $y \in \mathbb{R}$, and $k \in \mathbb{Z}$, let $c_{y,k}$ be the circle of radius $\frac{1}{2}$ in P_y, centered at $x = k + \frac{1}{2}$, $z = 0$. The *plat tube* $A_{i,j}$ is defined to be the annulus

$$A_{i,j} = \begin{cases} \bigcup_{y\in[i,i+1]} c_{y,2j} & \text{if } i \text{ is even,} \\ \bigcup_{y\in[i,i+1]} c_{y,2j+1} & \text{if } i \text{ is odd.} \end{cases}$$

For $n, m \in \{2, 3, 4, \ldots\}$, the (n, m)–*plat structure* is the union of the plat tubes $A_{i,j}$ where i ranges from 1 to $n-1$ and j either ranges from 1 to m or 1 to $m-1$ depending upon whether i is even or odd, respectively.

For $n, m \in \{2, 3, 4, \ldots\}$, an (n, m)–*plat braid* is a union of $2m$ pairwise disjoint arcs in $\mathbb{R}^3$ whose projections to the y–axis are monotonic, satisfying the following properties:

(1) One endpoint of each arc lies in P_1 and the other endpoint lies in P_n.

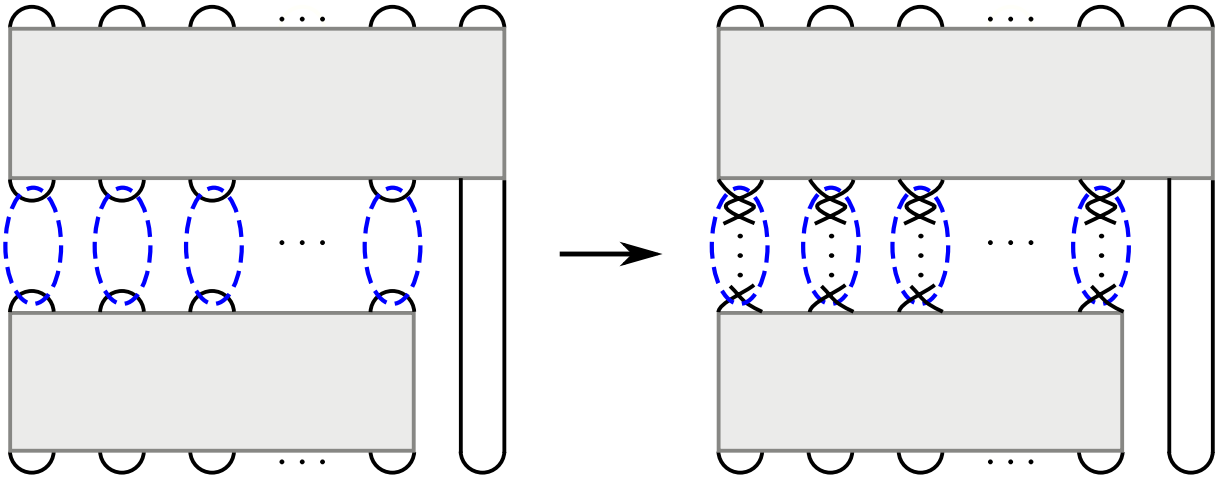

Figure 4: The upper gray box contains D_a, a 4–twisted $(2m-4, m)$–plat braid, where $m \geq 4$. The lower gray box contains D_b a 4–twisted $(n, m-1)$–plat braid.

(2) Each arc can be cut into subarcs, each of which is contained either in the (n, m)–plat structure or in one of the vertical lines $x = 2$ or $x = 2m + 1$ in the xy–plane.

(3) The intersection of the braid with each plat tube consists of a pair of arcs which intersect the plane $z = 0$ in a minimal number of components and whose endpoints lie in $z = 0$.

Observe that the plane $z = 0$ cuts $A_{i,j}$ into two disks. Here we define the *twist number* $a_{i,j}$. If the braid intersects $A_{i,j}$ in vertical arcs, then we define $a_{i,j} = 0$. Otherwise, the disk with nonnegative z–coordinates contains some number of arcs of the plat braid whose projection to the plane $z = 0$ is a set of parallel line segments. We define $|a_{i,j}|$ to be this number of parallel arcs. The sign of $a_{i,j}$ is defined to be the sign of the slope $(\Delta y/\Delta x)$ of the line segments. The integer $a_{i,j}$ is called the *twist number* for $A_{i,j}$.

For our purposes, we will only consider (n, m)–plat braids with n even. In this case, we can obtain a link from an (n, m)–plat braid by first connecting the point $(2j, 1, 0)$ to $(2j + 1, 1, 0)$ for each $1 \leq j \leq m$ with the unique (up to isotopy) arc in the portion of the plane $z = 0$ which lies below the line $y = 1$. Similarly, for each $1 \leq j \leq m$, we also connect the point $(2j, n, 0)$ to the point $(2j + 1, n, 0)$ with the unique arc in the portion of plane $z = 0$ above the line $y = n$. These $2m$ arcs can be isotoped in the plane $z = 0$ (with respect to their endpoints) so that each is injective when projected to the x–axis and each contains either a single maximum or minimum point (with respect to h), with the result that the set of $2m$ arcs is pairwise disjoint. The embedding of a link constructed as the union of the plat braid and these $2m$ arcs in this way is said to be an (n, m)–*plat position of a link*. If a link has an (n, m)–plat position, it is called an (n, m)–*plat link*. A plat link is called k–*twisted* if $|a_{i,j}| \geq k$ for every twist number $a_{i,j}$.

Throughout the rest of the paper, when discussing plat links, $\alpha_+^1 \dots \alpha_+^m$ will refer to the bridge arcs above P_n, labeled from left to right. Likewise, let $\alpha_-^1 \dots \alpha_-^m$ be the bridge arcs below P_1, labeled from left to

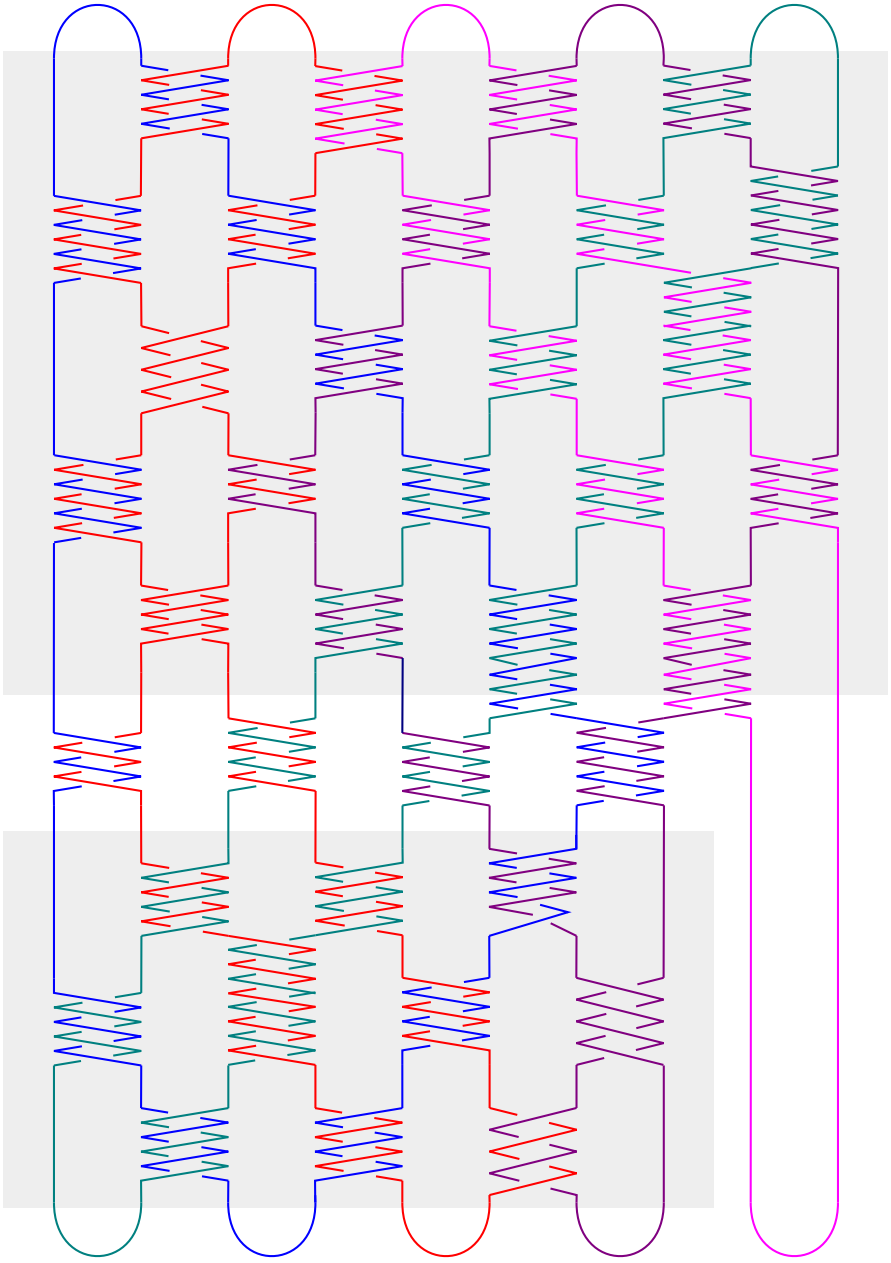

Figure 5: A (10,5)-plat link $L \in \mathcal{L}$.

right. Let L^i be the link component that contains α^i_+. (Of course, since an m–bridge link may have fewer than m components, it may be that the component containing α^i_+ also contains α^j_+ for some $j \neq i$, and so $L^i = L^j$.)

There are two types of projection maps that we will often refer to. The first type of projection map is the Euclidean projection map $\sigma_i : \mathbb{R}^3 \to \mathbb{R}^3$ defined by $\sigma_i(x, y, z) = (x, i, z)$. The second type of

projection map is the map $\pi_y : \mathbb{R} \times [1,n] \times \mathbb{R} \to P_y$, which sends each component of the plat braid to the corresponding point $(j, y, 0)$, and extends to a homeomorphism from $P_{y'}$ to P_y for each $y' \in [1,n]$. (In a slight abuse of notation, we will refer to this homeomorphism as π_y.)

3.2 The family of links we consider

Let D_a be a 4–twisted (n_a, m)–plat position of a link L_a such that $m \geq 4$ and $n_a = 2m-4$, and let D_b be a 4–twisted $(n_b, m-1)$–plat position of a link L_b. Position D_a above D_b as shown on the left of Figure 4. Let D_{ab} be an (n_a+n_b, m)–plat position obtained from $D_a \sqcup D_b$ by replacing each of the 0–tangles in the dashed ovals with vertical half-twists as shown on the right of Figure 4. (Note: the subscripts a and b are used for "above" and "below".) We define $\mathcal{L}$ to be the family of links constructed in this fashion which have the following additional properties.

(1) The rightmost twist regions of every row alternate in sign from row to row. That is, the sign of the rightmost nonzero twist region of each row is opposite to the sign of the rightmost nonzero twist regions of any adjacent rows.

(2) The sign of $a_{n-1,2}$, the twist number for the second twist region in the top row of D_a, is even, and every other twist region that involves L^3 has an odd twist number. (This forces L^3 to be an unknot component containing the bridge arcs α_+^3 and α_-^m.)

(3) The signs of the twist numbers for the twist regions involving L^m are chosen so that L^m contains the lower left bridge arc α_-^1.

(4) The rest of the twist numbers for D_{ab} are chosen so that D_{ab} is an m–component link and so that the bridges α_+^m and α_-^1 are contained in the same link component, namely L^m. (It follows that for each $i, j \in \{1, \dots, m\}$, with $i \neq j$, L^i is a distinct link component from L^j.)

(5) Excluding the pair $\{L^1, L^3\}$, every pair of link components comprises a two-bridge nonsplit sublink. (Note: The sublink $L^1 \cup L^3$ will always be split no matter what set of twist numbers is chosen.)

Below in Proposition 3.1, we will show that $\mathcal{L}$ is a nonempty set. First, observe that $\mathcal{L}$ is a family of links in (n,m)–plat position for $m \geq 4$ and $n = n_a + n_b$ with the following conditions on the twist numbers:

(1) For $i > n_b$, $|a_{i,j}| \geq 4$ for all possible values of j.

(2) If i is odd, $1 \leq i \leq n_b$, and $1 \leq j \leq m-2$, (resp. $j = m-1$), then $|a_{i,j}| \geq 4$ (resp. $a_{i,j} = 0$).

(3) If i is even, $1 \leq i \leq n_b$, and $1 \leq j \leq m-1$, (resp. $j = m$), then $|a_{i,j}| \geq 4$ (resp. $a_{i,j} = 0$).

(4) If $a_{i,*}$ denotes the rightmost nonzero twist number in row i, then $a_{i,*} \cdot a_{i-1,*} < 0$ and $a_{i,*} \cdot a_{i+1,*} < 0$. In other words, the signs of the rightmost nonzero twist numbers alternate from row to row.

Figure 5 shows an example of what $L \in \mathcal{L}$ may look like. In this case, $n = 10$, and $m = 5$. It follows from the definition of the family $\mathcal{L}$ that for any $L \in \mathcal{L}$, L^3 is the component containing the lower right bridge arc α_-^m.

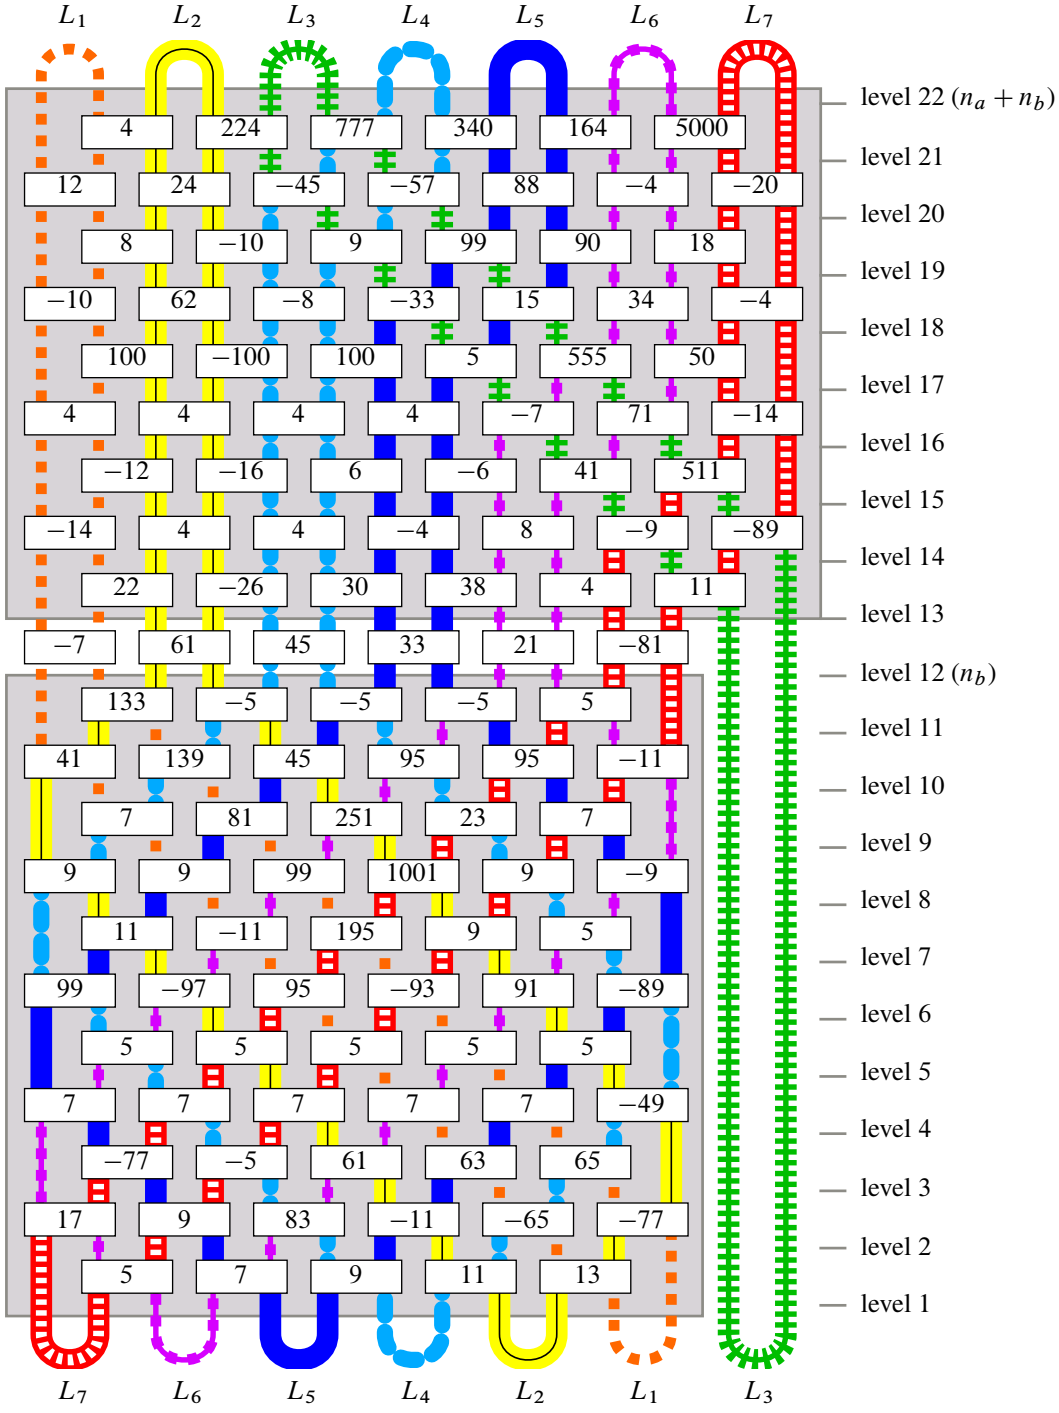

Figure 6: This figure illustrates Proposition 3.1, showing an example a 7–bridge link in $\mathcal{L}$. The seven different colors and line styles represent seven different link components. Each rectangle represents a twist region, and the integer inside each rectangle is the twist number indicating the number and sign of half twists present in that twist region.

Proposition 3.1 *For each integer $m \geq 4$, the family $\mathcal{L}$ contains infinitely many links of bridge number m.*

Proof It is not difficult to satisfy properties (1) and (2). We need to show that properties (3), (4), and (5) can be satisfied. We will do this by constructing an infinite family of examples for each bridge number $m \geq 4$.

Fix $m \geq 4$. Let D_a and D_b be a $(2m-4, m)$–plat link and an $(2m-2, m-1)$–plat link. We will show that the right choices of the parities and magnitudes of the twist numbers of D_a and D_b will allow us to fulfill the conditions given above.

First, after choosing odd numbers for the particular twist numbers in D_a prescribed by property (2), choose all even numbers for the rest of the twist numbers in D_a. Then choose any integers of magnitude at least four (regardless of parity) for the row of twist regions between D_a and D_b. The twist number choices we have made so far guarantee that at level n_b, the punctures of P_{n_b} occur in pairs corresponding to the link components in this order, from left to right: $L^1, L^2, L^4, L^5, \ldots, L^m, L^3$. That is, they are arranged in numerical order from left to right except that the punctures of L^3 appear at the end of the line.

Then for every twist region below P_{n_b} (ie the twist numbers corresponding to D_b), we choose all odd twist numbers. This guarantees that the punctures of P_1 occur in pairs corresponding from left to right to the link components $L^m, L^{m-1}, \ldots, L^5, L^4, L^2, L^1, L^3$. That is, they area arranged in reverse numerical order, except that again, the punctures of L^3 are at the end of the line. Thus L is an m–component link whose lower left bridge arc is contained in L^m, satisfying conditions (3) and (4).

Since every twist number in D_b is odd, it follows that for each pair $\{L^i, L^j\}$ of distinct link components from the set $\{L^1, L^2, L^4, \ldots, L^m\}$ (the set of all link components excluding L^3), there are exactly four twist regions in D_b which involve both L^i and L^j. The choices of twist numbers in D_a guarantees that there is one twist region containing arcs of both L^2 and L^3, and there are exactly four twist regions containing arcs of both L^3 and L^j for each $j \geq 4$. Now let $\{L^i, L^j\}$ be any pair of link components except for the pair $\{L^1, L^3\}$. To satisfy condition (5), simply choose twist numbers such that the linking number of $L^i \cup L^j$ is nonzero. For example here is one way to do so. There will be some positive number N of twist regions that involve strands from both L^i and L^j. For these twist regions, choose twist numbers $t_1, t_2, \ldots, t_N$ such that $|t_1| > \sum_{k=2}^{N} |t_k|$. □

Proposition 3.2 *Each link in the family $\mathcal{L}$ is nonsplit.*

Proof Let $L \in \mathcal{L}$, and assume S is a splitting sphere for L.

Case 1 *The link components L^1 and L^2 are both on the same side of S.*

In this case, let L^j be a link component on the other side of S. Then by condition (5) of the definition of $\mathcal{L}$, $L^2 \cup L^j$ is a nonsplit link which is split by S, a contradiction.

Case 2 *The link components L^1 and L^2 are not both on the same side of S.*

Then S is a splitting sphere for the nonsplit link $L^1 \cup L^2$, another contradiction. □

To say that a given compressing disk C is a *cap* is to say that there exists some bridge disk D such that $C = dD$. If α is the bridge arc corresponding to D, then we say that C is a cap *for* α. It follows that ∂C cuts the bridge sphere into two components, one of which is a twice-punctured disk, where the two punctures are the intersection points of the bridge sphere with α.

Proposition 3.3 *Let D and E be compressing disks above and below P, respectively. If $\{D, E\}$ is a weak reducing pair for P, then D is a cap for α^1_+, and E is a cap for α^m_-.*

Proof The loop $\partial D \subset P$ partitions the link components of L into two nonempty sets, A and A' (based on which side of ∂D the punctures of each link component lie). Let A be the set containing L^1. The loop ∂E also partitions the link components into two nonempty sets, B and B'. Let B be the set containing L^3. If A contains L^i for any $i \neq 1$, then $\{D, E\}$ is a weak reducing pair for the sublink $L^i \cup L^3$, a nonsplit 2–bridge link, a contradiction. Similarly, if B contains L^j for any $j \neq 3$, then $\{D, E\}$ is a weak reducing pair for the sublink $L^1 \cup L^j$, a nonsplit 2–bridge link, another contradiction. Therefore D is a cap for α^1_+, the bridge arc above P contained in L^1, and E is a cap for α^m_-, the bridge arc below P contained in L^3. □

The rest of the paper will be devoted to proving that each $L \in \mathcal{L}$ admits a keen weakly reducible bridge sphere. The reason Proposition 3.3 does not immediately imply this is because for any given bridge arc, there are infinitely many distinct caps for that bridge arc, provided there are at least three bridges on each side of the bridge sphere, which is the case for all of the links in $\mathcal{L}$.

3.3 Plat train tracks

Speaking generally, let Σ_L denote a bridge sphere, and let I denote a closed unit interval. A *train track* τ is a compact subsurface of Σ_L whose interior is fibered by open intervals and the fibration extends to a fibration of τ by closed intervals except for at finitely many intervals called *singular fibers*. Let α be a singular fiber, and denote its closed neighborhood in τ by $\overline{N}(\alpha)$. Then there is a homeomorphism $f : \overline{N}(\alpha) \to (I \times I) \backslash \left(\left(\frac{1}{4}, \frac{3}{4}\right) \times \left(\frac{1}{2}, 1\right] \right)$ such that $f(\alpha) = I \times \left\{\frac{1}{2}\right\}$. We will refer to the inverse image of $\left(I \times \left\{\frac{1}{2}\right\}\right) \backslash \left(\left(I \times \left[0, \frac{1}{4}\right]\right) \sqcup \left(I \times \left[\frac{3}{4}, 1\right]\right) \right)$ under f as a *switch* of τ; see Figure 7.

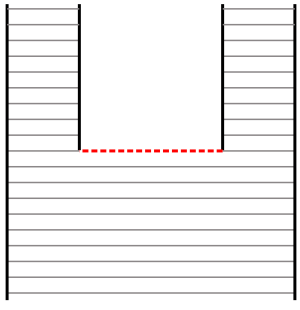

Figure 7: A train track at a singular fiber. The closed red line segment is a switch.

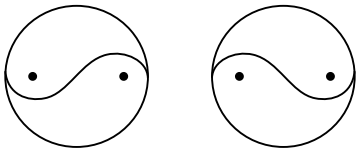

Figure 8: At left, a left-handed tao diagram. At right, a right-handed tao diagram.

In this paper, we will assign a train track τ_i to each bridge sphere P_i for $i = 1, 2, \ldots, n-1$. (There is no need for a train track at the top level P_n.) To this end, we will construct a certain trivalent graph, called a train graph, on each bridge sphere based on the parity of i and the twist numbers $a_{i,j}$ for the row. The train track will then be constructed from the train graph in a natural way.

We define a *train graph* to be a connected trivalent graph with the property that the three edges incident to each vertex are tangent to each other at the vertex, and not all three edges emanate from the vertex in the same direction. (See the left side of Figure 14.) Below, we will construct a specific train graph T_i embedded in P_i for each $i = 1, 2, \ldots, n-1$, and these train graphs will have the property that $P_i \setminus T_i$ consists of $2m$ once-punctured disks and one (nonpunctured) disk. We will informally express this by saying that each puncture is "surrounded by" T_i.

To construct each train graph, there are various cases to consider. Recall from Section 3.2 that $\mathcal{L}$ is a family of links in (n, m)–plat position for $n = n_a + n_b$. If i is odd and $n-1 \geq i \geq n_b+1$ (resp. $i < n_b+1$), we define $\ell_{i,j}$ to be the circle in P_i centered at $\left(2j + \frac{3}{2}, i, 0\right)$ with radius $\frac{3}{4}$ for $j = 1, 2, \ldots, m-1$ (resp. for $j = 1, 2, \ldots, m-2$). If i is even and $n-2 \geq i \geq n_b+1$ (resp. $i < n_b+1$), we define $\ell_{i,j}$ to be the circle in P_i centered at $\left(2j + \frac{1}{2}, i, 0\right)$ with radius $\frac{3}{4}$ for $j = 1, 2, \ldots, m$ (resp. $j = 1, 2, \ldots, m-1$).

Now, each $\ell_{i,j}$ cuts out a twice-punctured disk from P_i. We will distinguish two types of arcs that separate the two punctures. If $\ell_{i,j}$ is directly below a positive twist region, then we draw a *right-handed tao arc* separating the two punctures as shown on the right of Figure 8. In the case where $\ell_{i,j}$ is directly below a negative twist region, we instead draw a *left-handed tao arc*. The union of a left-handed tao arc (resp. right-handed tao arc) with $\ell_{i,j}$ will be called a *left-handed tao diagram* (resp. *right-handed tao diagram*). An important aspect of these tao diagrams is that at a tao arc's endpoints, the circle and the tao arc are tangent to each other as pictured.

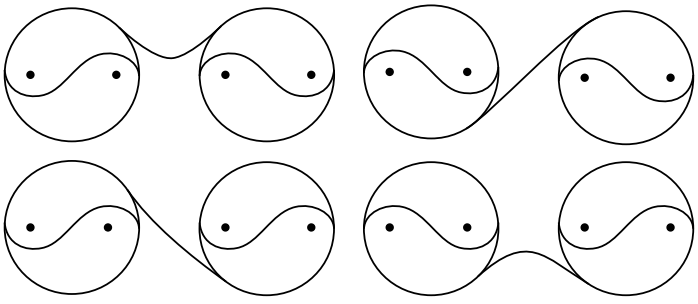

Figure 9: The way we add an edge between two adjacent tao diagrams depends on their handedness.

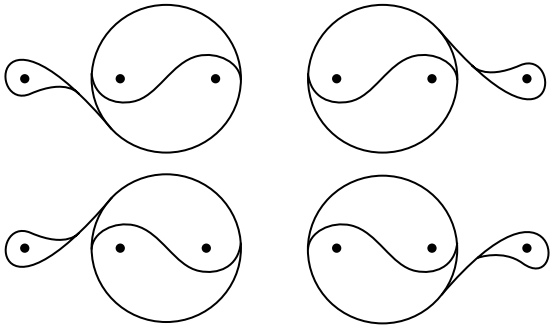

Figure 10: Adding eyelets to "leftover" punctures that are adjacent to a tao diagram.

At this point we have constructed various disconnected tao diagrams in each P_i. We next connect each pair of adjacent tao diagrams with an edge in one of the four ways pictured in Figure 9, depending on the handedness of each tao diagram. If i is even and $i \geq n_b + 1$, the result of this procedure is a train graph which we call T_i.

In all other cases (ie if i is odd and/or $i < n_b + 1$), begin the construction of the train graph as above, combining tao diagrams and connecting edges; however, after doing so, there will be "leftover" punctures that are not surrounded by any tao diagrams. If any such puncture is adjacent to a puncture surrounded by a tao diagram, then we modify our graph according to Figure 10, adding a vertex and two edges to the graph in a way that depends on which side of the tao diagram the puncture is on and the handedness of the tao diagram. The newly added subgraph consists of two edges, one forming a loop around a puncture, and the other connecting the loop to a tao diagram. We refer to such a subgraph as an *eyelet*. If i is odd and $i \geq n_b + 1$, this procedure gives a connected trivalent graph containing two eyelets, surrounding all the punctures. We call this train graph T_i.

For $1 \leq i < n_b + 1$, there are still "leftover" punctures that are not surrounded by a tao diagram or eyelet. Since the sign of the rightmost nonzero twist region of a row is opposite to the sign of the rightmost

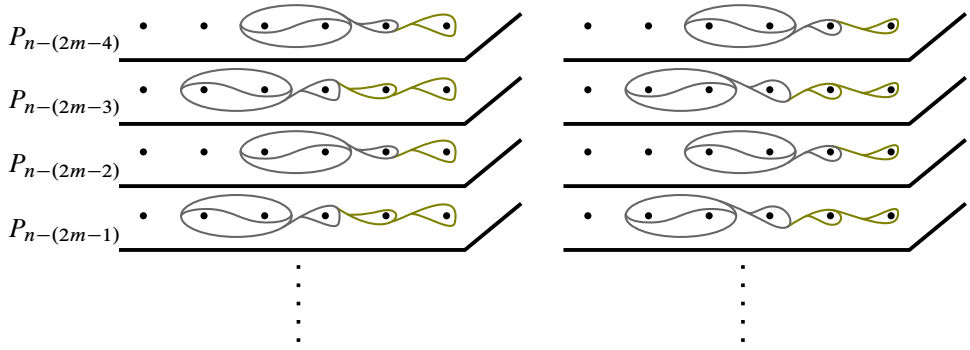

Figure 11: Adding eyelets to "leftover" punctures on P_i for $1 \leq i < n_b + 1$. If the rightmost tao in $P_{n-(2m-4)}$ is left-handed (resp. right-handed), we add eyelets according to the left (resp. right) picture.

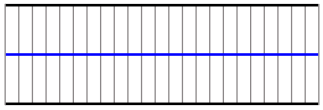

Figure 12: A fibration of $\overline{N}(e)$ by intervals.

nonzero twist region of adjacent rows, there are only two possibilities for the rightmost tao and the eyelet adjacent to it. These are depicted in gray in Figure 11. If $a_{n-1,m-1} < 0$, we add the eyelets according to the left of Figure 11. If $a_{n-1,m-1} > 0$, we add the eyelets according to the right of Figure 11. The newly added eyelets are colored gold. After doing so, we have a train graph that surrounds all the punctures for each P_i for $1 \leq i \leq n-(2m-4)$, and we call this train graph T_i.

We have constructed a train graph T_i on each bridge sphere P_i. Now we will use each train graph T_i to construct a train track τ_i on each sphere P_i. Let V_i and E_i be the vertex set and the edge set for T_i, respectively. For each vertex $v \in V_i$, let $\overline{N}(v)$ be a closed regular neighborhood of v in P_i.

Let e' denote the connected component of $T_i \setminus \bigcup_{v\in V_i} \overline{N}(v)$ corresponding to the edge e. Let $\overline{N}'(e')$ be a closed regular neighborhood of e' in P_i, and then define $\overline{N}(e) = \overline{N}'(e') \setminus \bigcup_{v\in V_i} \overline{N}(v)$. Notice that $\big(\bigcup_{v\in V_i} \overline{N}(v)\big) \sqcup \big(\bigcup_{e\in E_i} \overline{N}(e)\big)$ is a regular neighborhood of T_i which we call $\overline{N}(T_i)$, and the set $\{\overline{N}(e), \overline{N}(v) \mid e \in E_i, v \in V_i\}$ is a partition for $\overline{N}(T_i)$.

We fiber each set $\overline{N}(e)$ with interval fibers, each one intersecting e' transversely exactly once as in Figure 12. Then we impose a singular fibration on each $\overline{N}(v)$ containing exactly one singular fiber, as in Figure 13. This makes $\overline{N}(v)$ into a neighborhood of a switch in a train track. The surface $\overline{N}(T_i)$, together with the singular fibration, is a train track which we call τ_i, constructed from the train graph T_i. This construction process is illustrated in Figure 14.

3.4 Carried and almost carried

We want to isotope certain objects in the bridge sphere P_i to a position that behaves nicely with respect to the train track τ_i.

Definition 3.4 For an arc α (not necessarily properly) embedded in P_i, τ_i is said to *almost carry* α if the following are true.

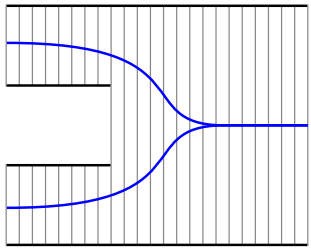

Figure 13: A singular fibration on $\overline{N}(v)$ containing one singular fiber.

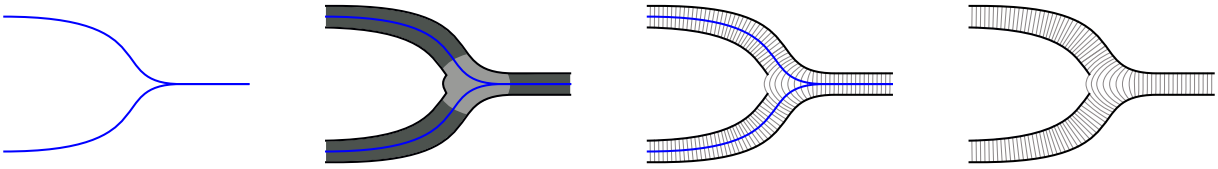

Figure 14: Constructing a train track τ_i from the train graph T_i.

(1) For each point $p \in \alpha$, either $p \notin \tau_i$ or p is a transverse intersection point of α with an interval fiber of τ_i.

(2) No point of α is an endpoint of an interval fiber of τ_i.

(3) No connected component α' of $\alpha \setminus \mathring{\tau}_i$ is parallel (rel $\partial\alpha'$) into a switch.

(4) No connected component α' of $\alpha \setminus \mathring{\tau}_i$ is parallel (rel $\partial\alpha'$) into an arc $\alpha'' \subseteq \partial\tau_i$ with the property that α'' is partitioned into three subintervals: the outer two being subintervals of switches and the middle subinterval of α'' being an interval of fiber endpoints of τ_i (see Figure 15).

Remark 3.5 It follows from Definition 3.4 that if τ_i almost carries an arc α, and an endpoint of α lies in $\partial\tau_i$, then that endpoint lies in the interior of a switch.

Remark 3.6 If an arc α satisfies conditions (1), (2), and (3) of Definition 3.4, then each arc of $\alpha \cap (P_i \setminus \mathring{\tau}_i)$ which is properly embedded in $P_i \setminus \mathring{\tau}_i$ but which does not satisfy condition (4) can be isotoped into the train track, as illustrated in Figure 15. This results in a position of α which *is* now almost carried.

Definition 3.7 A loop $\ell \subseteq P_i$ is said to be *almost carried* by the train track τ_i if every connected component of $\tau_i \cap \ell$ and every connected component of $(P_i \setminus \mathring{\tau}_i) \cap \ell$ is an arc which is almost carried by τ_i.

Remark 3.8 An arc or loop in P_i which is completely disjoint from τ_i still satisfies the definition of being almost carried by τ_i.

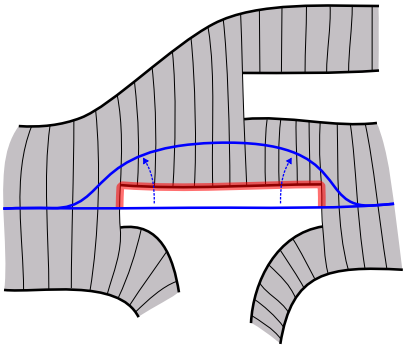

Figure 15: The red shaded arc is α'' from Definition 3.4.

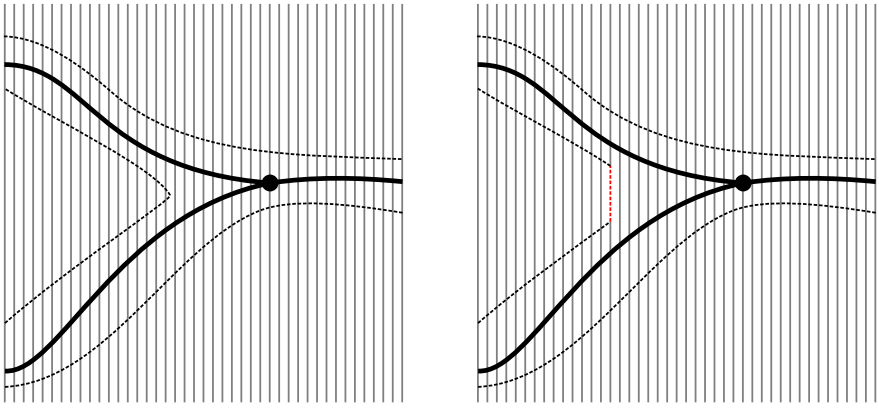

Figure 16: The black graph is T, the vertical gray lines are fibers of σ, and $\overline{N}(T)$, the closed regular neighborhood of T, is outlined with dashed lines. On the right, we see that a slight isotopy of $\overline{N}(T)$ makes it into a train track diagram τ with fibers inherited from σ. That is, every fiber of τ is a subinterval of a fiber of σ.

Definition 3.9 A train graph T is said to be *almost carried* by the train track τ_i if each edge of T is almost carried by τ_i.

As a simple example, for each i, the train graph T_i is almost carried by the train track τ_i.

Definition 3.10 Let T_i' be a subgraph of the train graph T_i (eg a tao), and let $\tau_i' \subseteq \tau_i$ be the sub train track constructed from T_i' following the instruction in Section 3.3. If ℓ is a loop or train graph, then ℓ is said to *cover* T_i' if ℓ is almost carried by τ_i' and ℓ intersects every interval fiber of τ_i'.

Definition 3.11 Let τ and σ be two different train tracks contained in the same bridge sphere P_i. Then σ is said to *almost carry* τ if for each interval fiber I of τ, I is disjoint from σ or I is contained in the interior of some interval fiber of σ.

Proposition 3.12 *Let T be a train graph in the bridge sphere P, corresponding to train track diagram τ, and let σ be another train track diagram in P. If T is almost carried by σ, then τ is almost carried by σ. Furthermore, if p is a point in T which lies in the interval fiber $I \subseteq \tau$, and p also lies in the interval fiber $J \subseteq \sigma$, then $I \subseteq J$.*

Proof Our strategy here is to reexamine the construction of τ and see that it has the desired properties. Let T be a train graph in P, almost carried by σ, and let $\overline{N}(T)$ be a closed regular neighborhood of T.

By definition, every point of T is either disjoint from σ or lies in the interior of a fiber interval of σ. It follows that $\overline{N}(T)$ is disjoint from the fiber endpoints of σ.

Near each vertex v of T which lies in σ, we perform a slight isotopy of $\overline{N}(T)$ (pictured in Figure 16) as follows. We locate an arc λ of $\partial\overline{N}(T)$ located between the two edges of T which emanate from v in the same direction. We isotope $\overline{N}(T)$ so that λ is a subinterval of one of the fiber intervals of σ.

Next we allow the portion of $\overline{N}(T)$ which lies inside σ to inherit a fibering from σ this way: if J is a fiber of σ, then $J \cap \overline{N}(T)$ is a (possibly empty) set of fibers of $\overline{N}(T)$.

After extending this fibration to the rest of $\overline{N}(T)$ which lies outside of σ, $\overline{N}(T)$, endowed with a fibration, is now a train track diagram τ with the desired properties. □

4 How compressing disks meet train tracks

It is desirable to isotope a simple closed curve to intersect a train track in a way over which we can have some control.

For future convenience, we partition the compressing disks into two disjoint sets. Consider α_+^1, the leftmost bridge arc above P_n. A vertical isotopy of α_+^1 into bridge sphere P_n traces out a bridge disk D_+^1. Let $B = dD_+^1$; that is, B is the cap which is the frontier of a regular neighborhood of D_+^1 in V_+. Similarly, consider α_-^m, the rightmost bridge arc below P_1. The bridge arc α_-^m gives rise to a bridge disk D_-^m and a corresponding caps $B' = dD_-^m$. We will refer to these two isotopy classes of caps as *blue* disks. Compressing disks for P_i that are not blue will be referred to as *red* disks.

Proposition 4.1 [Johnson and Moriah 2016, Lemma 8.4] *If D is a compressing disk above P_n, then $\pi_{n-1}(\partial D)$ covers at least one tao of τ_{n-1}, and away from those one or more taos, $\pi_{n-1}(\partial D)$ intersects τ_{n-1} in almost carried arcs or in fiber intervals.*

Definition 4.2 A subgraph T_i' of T_i is called a *mini-graph* of T_i if it has the following properties.

(1) T_i' is a union of taos, connecting arcs, and eyelets of the train graph T_i.

(2) Two adjacent taos of T_i are contained in T_i' if the taos' connecting arc is contained in T_i'.

(3) An eyelet $E \subseteq T_i$ is contained in T_i' only if both the tao T nearest to E in T_i and every other eyelet between E and T are also contained in T_i'.

Definition 4.3 Let T_i' be a mini-graph of T_i. The mini-graph *directly below* T_i' is defined to be the unique mini-graph T_{i-1}' of T_{i-1} with the following properties.

(1) If $T \subseteq T_i'$ is a tao or an eyelet and $\sigma_{i-1}(T)$ intersects a tao T' of T_{i-1}, then $T' \subseteq T_{i-1}'$.

(2) If $T \subseteq T_i'$ is a tao and $\sigma_{i-1}(T)$ intersects two taos of T_{i-1}, say T' and T'', then the connecting arc between T' and T'' is contained in T_{i-1}'.

(3) If $T \subseteq T_i'$ is a tao and $\sigma_{i-1}(T)$ intersects an eyelet E' of T_{i-1}, then $E' \subseteq T_{i-1}'$.

(4) If $T \subseteq T_i'$ is an eyelet and $\sigma_{i-1}(T)$ is an eyelet of T_{i-1}, then $\sigma_{i-1}(T)$ is also an eyelet of T_{i-1}'.

Let $i, i-j \in \{1, 2, \dots, n-1\}$, with $i > i-j$. Let T_i' and T_{i-j}' be mini-graphs of T_i and T_{i-j}, respectively. We say that T_{i-j}' is *below* T_i' if and only if there exists a sequence of mini-graphs $T_i', T_{i-1}', T_{i-2}', \dots, T_{i-j}'$ such that for $k = 1, 2, \dots, j$, the mini-graph T_{i-k-1}' is directly below T_{i-k}. Naturally, we will say that a mini-graph T_i' lies (*directly*) *above* a mini-graph T_{i-j}' if and only if T_{i-1}' lies (directly) below T_i'.

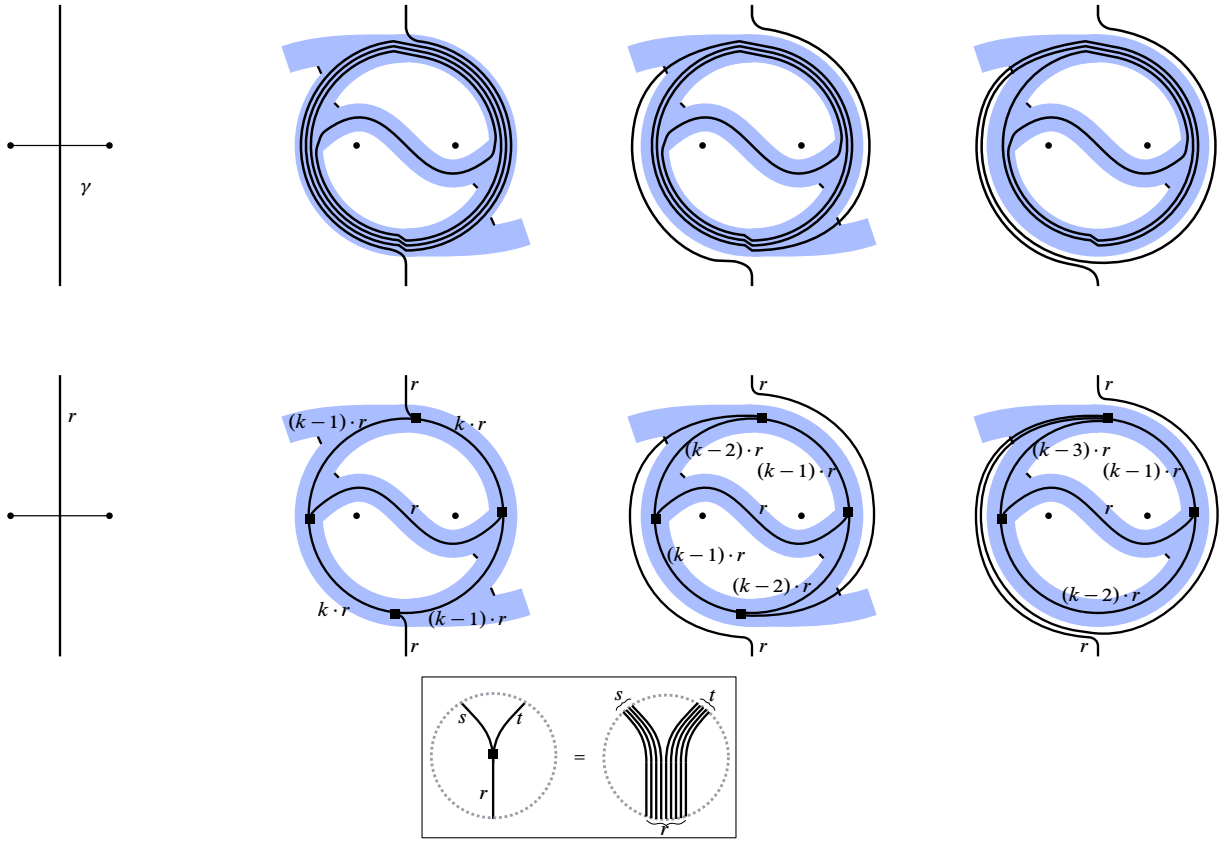

Figure 17: Illustration of Proposition 4.1. In this and following figures, an arc with a label (such as the arc coming out of the bottom of the second picture in the bottom row, marked with an r), represents not just one, but a number of parallel arcs, according to the label. The box at the bottom of the figure illustrates that the small squares placed at various places in this and following figures represents a set of parallel arcs separating into two different sets of parallel arcs. In particular, the square does *not* represent a vertex of a graph.

Observation 4.4 If T_i' and T_i'' are mini-graphs of T_i, then $T_i' \cup T_i''$ is a mini-graph of T_i as well.

Observation 4.5 Suppose that T_i' and T_i'' are mini-graphs of T_i, that T_{i-1}' and T_{i-1}'' are mini-graphs of T_{i-1}, that T_{i-1}' is below T_i', and that T_{i-1}'' is below T_i''. Then $T_{i-1}' \cup T_{i-1}''$ is the mini-graph below $T_i' \cup T_i''$.

Proposition 4.6 *First, if T_{n-1}' is the leftmost tao of T_{n-1}, and if T_1' is the mini-graph of T_1 constructed by excluding from T_1 only the rightmost two eyelets, then T_{n-1}' is above T_1'. Second, if T_{n-1}'' is any other tao of T_{n-1} besides the leftmost tao, then T_{n-1}'' is either above T_1 itself or above some mini-graph T_1'' of T_1 constructed by excluding from T_1 the leftmost eyelet and/or the rightmost eyelet.*

In reading through the following proof, the reader may find it helpful to use the example link in Figure 5 to help locate and visualize the various mini-graphs we discuss.

Proof We will need to define a few more specific mini-graphs. Define T'_{n_b+1} to be the mini-graph constructed from T_{n_b+1} by excluding from it the rightmost eyelet, tao, and connecting arc. Define T'_{n_b} to be the mini-graph of T_{n_b} consisting of all of the taos and connecting arcs of T_{n_b} but excluding both eyelets. Now the reader can observe that by virtue of the dimensions of the link L_a, that is, since $n_a = 2m-4$, the tao T'_{n-1} is above the mini-graph T'_{n_b+1}. Further, observe that T'_{n_b+1} is directly above T'_{n_b}, which is above T'_1. Therefore, T'_{n-1} is above T'_1.

For the next part of the proof, suppose that T''_{n-1} is a tao of T_{n-1}, but not the leftmost one. Appealing again to the dimensions of L_a, the tao T''_{n-1} is above some mini-graph T''_{n_b+1} of T_{n_b+1} which includes at least all of the taos of T_{n_b+1} but the leftmost one. That is, while T''_{n_b+1} *may* contain the leftmost tao of T_{n_b+1}, T''_{n_b+1} *does* contain all the other taos of T_{n_b+1}. It follows that T''_{n_b+1} is directly above some mini-graph T''_{n_b} of T_{n_b} which contains at least all but the leftmost tao of T_{n_b} and also the first of the two eyelets on the right side of T_{n_b}. Going down another level, T''_{n_b} must be directly above a mini-graph T''_{n_b-1} of T_{n_b-1} which contains all the taos and connecting arcs of T_{n_b-1} as well as the first and second (but not necessarily the third) eyelet on the right.

If $n_b = 2$, then $T_{n_b-1} = T_1$, and we can define $T''_1 = T''_{n_b-1}$, in which case the proof is finished, for the tao T''_{n-1} is above T''_1, which has been shown to have the desired properties.

If $n_b > 2$, then observe that at each level from level $n_b - 1$ down to level 1, T''_{n_b-1} will be above a mini-graph consisting of all of the level's taos and connecting arcs as well as all of the level's eyelets, possibly excluding the leftmost eyelet and/or the rightmost eyelet. Therefore T''_{n_b-1} is above some mini-graph T''_1 with the desired properties, which finishes the proof. □

Definition 4.7 It will be helpful to name a few special types of mini-graphs.

(0) If T is a tao, we will call T a *type* 0 mini-graph.

(1) A *type* 1 mini-graph consists of a final eyelet of T_i and an adjacent tao.

(2) A *type* 2 mini-graph consists of a final eyelet E_2, an eyelet E_1 adjacent to E_2, and a tao adjacent to E_1.

(3) A *type* 3 mini-graph consists of a final eyelet E_3, an eyelet E_2 adjacent to E_3, an eyelet E_1 adjacent to E_2, and a tao adjacent to E_1.

(4) Collectively we will refer to mini-graphs of type 0, 1, 2, or 3 as *typed* mini-graphs.

Observation 4.8 If T'_i is a mini-graph of T_i, then for some positive integer k, T'_i can be decomposed into a union $T'_i = t_1 \cup t_2 \cup \cdots \cup t_k \cup c$, where $t_1, t_2, \ldots, t_k$ are typed mini-graphs and c is a (possibly empty) union of connecting arcs.

Observation 4.9 Suppose T'_i is a mini-graph of T_i, and T'_{i-1} is the mini-graph directly below T'_i. Let T'_i be decomposed into a union $T'_i = t_1 \cup t_2 \cup \cdots \cup t_k \cup c$, where $t_1, t_2, \ldots, t_k$ are typed mini-graphs and c is a union of connecting arcs. For each $j \in \{1, 2, \ldots, k\}$, let u_j be the mini-graph directly below t_j. Then $T'_{i-1} = u_1 \cup u_2 \cup \cdots \cup u_k$.

Figure 18: A type 0 mini-graph (ie a tao) covers the mini-graph directly below it. In each picture which includes both red and blue arcs, only either the blue arcs or the red arcs will be present, depending on the handedness of the tao.

Proposition 4.10 *For each $i = 2, 3, \ldots, n-1$, if T_i' is a mini-graph of T_i, then $\pi_{i-1}(T_i')$ covers the mini-graph directly below T_i'.*

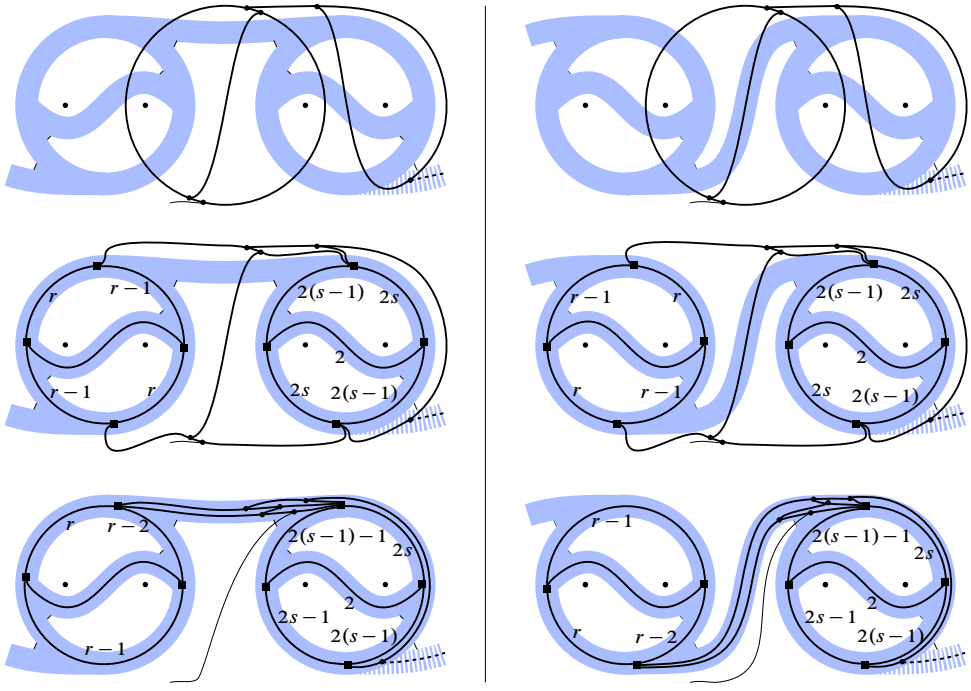

Figure 19: A type 1 mini-graph may lie directly above two taos and their connecting arc. If so, the type 1 mini-graph will cover the two taos and their connecting arc. The dashed lines of the train track diagram and of the train graph are either both present or both absent.

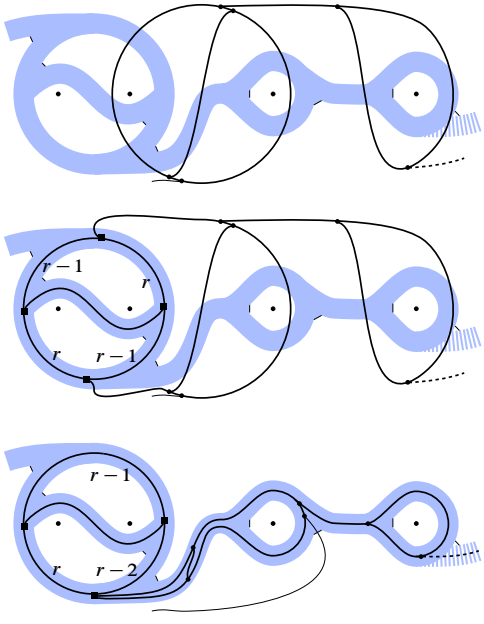

Figure 20: A type 1 mini-graph may lie directly above a type 2 mini-graph, in which case the former will cover the latter.

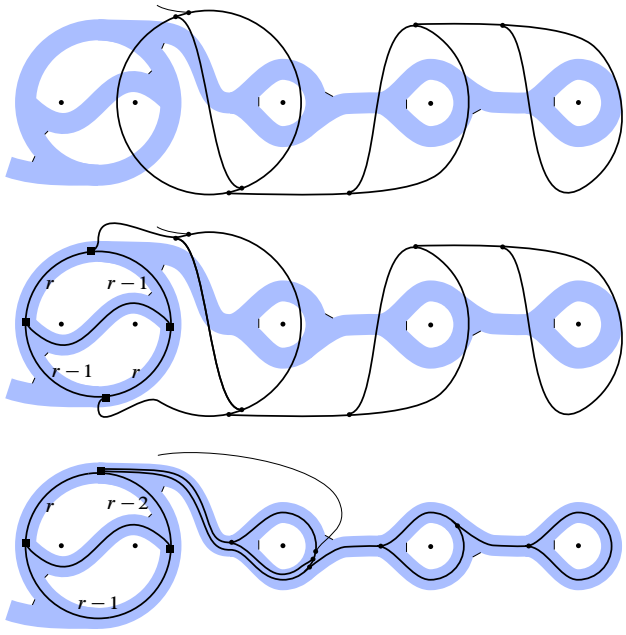

Figure 21: A type 2 mini-graph always lies directly above and covers a type 3 mini-graph.

Proof Fix $2 \leq i \leq n-1$. Let $T_i' \subseteq T_i$ be a mini-graph, and let T_{i-1}' be the mini-graph directly below T_i'. We will prove this proposition by proving several special cases which will lead us to the general result.

To begin, consider the case in which T_i' is a type 0 mini-graph (a tao). The mini-graph T_{i-1}' may consist of two adjacent taos and their connecting arc, as in the first or second column of pictures in Figure 18, or T_{i-1}' may instead be a type 1 mini-graph, as depicted in the third column of pictures in Figure 18. In any case, the result of the isotopy of the bridge sphere from level i to level $i-1$ is shown from the top row of pictures to the second row, or from the fourth row of pictures to the fifth row.

Observe that an isotopy of $\pi_{i-1}(T_i')$ in P_{i-1} (the result of which is shown in the third and sixth rows of pictures in Figure 18) shows how we may push $\pi_{i-1}(T_i')$ into τ_{i-1} so that $\pi_{i-1}(T_i')$ covers T_{i-1}'.

Next, if T_i' is a type 1 mini-graph, then either T_{i-1}' is a pair of taos (pictured in Figure 19) or T_{i-1}' is a type 2 mini-graph (pictured in Figure 20). Either way, the figures illustrate that $\pi_{i-1}(T_i')$ covers T_{i-1}'.

Now suppose T_i' is a type 2 mini-graph. In this case, T_{i-1}' must be a type 3 mini-graph. Figure 21 depicts this case and shows that $\pi_{i-1}(T_i')$ covers T_{i-1}'.

Finally, suppose T_i' is a type 3 mini-graph. It follows that T_{i-1}' is a type 2 mini-graph, as depicted in Figure 22, which shows that as before, $\pi_{i-1}(T_i')$ covers T_{i-1}'.

Observation 4.11 Notice that in each of the cases above, if c is a connecting arc of T_i which is attached to T_i' at vertex v, then the connected component of $c \cap \tau_{i-1}$ which contains v is almost carried by τ_{i-1}.

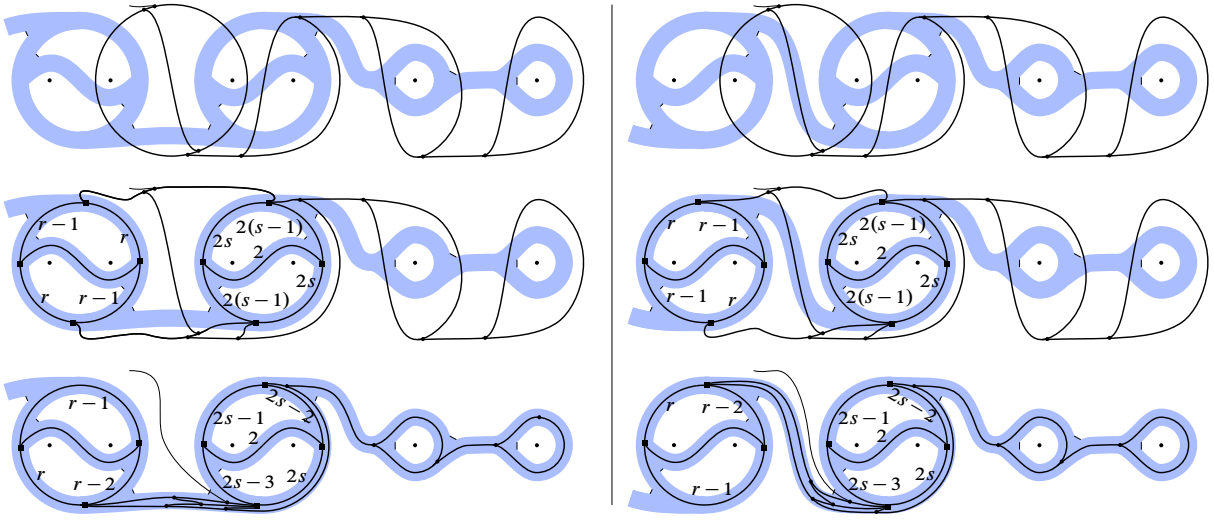

Figure 22: A type 3 mini-graph always lies directly above and covers a type 2 mini-graph (assuming there exists a level below the level of the type 3 minigraph).

Now that we have proven the proposition for cases in which T_i' is a typed mini-graph, we are ready to prove it in the general case where T_i' is an arbitrary mini-graph. According to Observation 4.8, we can view T_i' as a union $T_i' = t_1 \cup t_2 \cup \cdots \cup t_k \cup c$ of typed mini-graphs and connecting arcs. For each $j \in \{1, 2, \ldots, k\}$, define u_j to be the mini-graph of T_{i-1} below t_j. By Observation 4.9, $T_{i-1}' = u_1 \cup u_2 \cup \cdots \cup u_k$. The special cases above demonstrate that for each $j \in \{1, 2, \ldots, k\}$, the mini-graph u_j is covered by t_j, so it follows that $u_1 \cup u_2 \cup \cdots \cup u_k$ is covered by $\pi_{i-1}(t_1 \cup t_2 \cup \cdots \cup t_k)$. Further, if c_0 is one of the connecting arcs of c, then by Observation 4.11, c_0 is also almost carried by τ_{i-1}. Therefore, since $u_1 \cup u_2 \cup \cdots \cup u_k$ is covered by $\pi_{i-1}(t_1 \cup t_2 \cup \cdots \cup t_k)$ and c is almost carried by τ_{i-1}, we can conclude that $u_1 \cup u_2 \cup \cdots \cup u_k$ is covered by $\pi_{i-1}(t_1 \cup t_2 \cup \cdots \cup t_k \cup c)$, or more simply, T_{i-1}' is covered by $\pi_{i-1}(T_i')$. □

Corollary 4.12 *Let ℓ be a loop which covers a mini-graph $T_i' \subseteq T_i$, and let T_{i-1}' be the mini-graph directly below T_i'. The loop $\pi_{i-1}(\ell)$ covers T_{i-1}'.*

Proof Let J be an interval fiber of τ_{i-1}' that T_{i-1}' intersects. By Proposition 4.10, $\pi_{i-1}(T_i')$ covers T_{i-1}', and so by the definition of covering, J is also intersected by $\pi_{i-1}(T_i')$.

Let p be a point of $(\pi_{i-1}(T_i')) \cap J$, and let I be the interval fiber of $\pi_{i-1}(\tau_i)$ which contains p. By Proposition 3.12, $I \subseteq J$. Further, since ℓ covers T_i', ℓ must by definition intersect I. It follows that since $I \subseteq J$, ℓ intersects J. □

Corollary 4.13 *Let $i_1 < i_2$, and let $T_{i_1}' \subseteq T_{i_1}$ be the mini-graph below a mini-graph $T_{i_2}' \subseteq T_{i_2}$. If ℓ is a loop which covers $T_{i_2}' \subseteq T_{i_2}$, then $\pi_{i_1}(\ell)$ covers T_{i_1}'.*

Recall the notation of Proposition 4.6. The leftmost blue disk B above P_n is the only disk whose boundary loop covers T'_{n-1} but no other taos. In contrast, the boundary of every red disk above P_n must cover at least one of the other taos. The next corollary then follows from Proposition 4.6 and Corollary 4.13.

Corollary 4.14 *The boundary of the blue disk B above P_n covers T'_1 (the mini-graph defined in Proposition 4.6), and the boundary of every red disk above P covers either T_1 or some mini-graph T''_1 constructed from T_1 by excluding the leftmost and/or the rightmost eyelet of T_1.*

An almost carried loop ℓ that covers enough taos and eyelets is very beneficial in the sense that its presence allows us to predict the behavior of loops which are disjoint from ℓ.

Remark 4.15 The following is [Johnson and Moriah 2016, Lemma 6.5].

Lemma 4.16 *If ℓ is a loop in P_i that covers a mini-graph T''_i of T_i, and if ℓ' is another loop in P_i disjoint from ℓ, then ℓ' can be isotoped to be almost carried by τ''_i, the train track diagram corresponding to T''_i.*

Proof Let $N(\ell)$ be an open regular neighborhood of ℓ disjoint from ℓ'. Since τ''_i is covered by ℓ, every interval fiber of τ''_i intersects $N(\ell)$.

We perform a small isotopy of τ''_i with the following properties:

(1) The image of each interval fiber of τ''_i at each moment of the isotopy is a subinterval of the original interval fiber.

(2) The endpoints of each interval fiber of τ''_i never intersect ℓ throughout the isotopy.

(3) After the isotopy, both endpoints of every interval fiber of τ''_i lie in $N(\ell)$.

The result of this isotopy is illustrated in Figure 23. The point is that each arc of $\partial\tau''_i$ consisting of endpoints of interval fibers gets pushed into $N(\ell)$. Now each component of $\tau''_i \backslash N(\ell)$ is a band in P_i fibered by intervals (each of which is a subinterval of the original interval fibers of τ''_i). The two interval fibers contained in the boundary of a band will be referred to as *exits*. Note that topologically, each of these bands is a closed disk.

Now we isotope ℓ' in $P_i \backslash N(\ell)$ to intersect these bands minimally. Suppose $\ell' \cap \tau''_i = \varnothing$. Then by Remark 3.8, ℓ' is almost carried by τ''_i. If $\ell' \cap (P_i \backslash \mathring{\tau}''_i)$ is empty (that is, ℓ' lies completely in a band), then we can perform an isotopy of ℓ', pushing it out of the band through an exit, contradicting minimality.

Assume then that both $\ell' \cap \tau''_i$ and $\ell' \cap (P_i \backslash \mathring{\tau}''_i)$ are nonempty. Consider a component α of $\ell' \cap (\tau''_i \backslash N(\ell))$. The component α must be an arc properly embedded in a band. Since $\ell' \cap N(\ell)$ is empty, the endpoints of α must lie in exits. Further, the endpoints of α must lie on *different* exits, for otherwise α could be isotoped out of the band, reducing the number of components of $\ell' \cap (\tau''_i \backslash N(\ell))$, again contradicting minimality. Thus α is an arc that travels through a band from one exit to another, so it follows that the

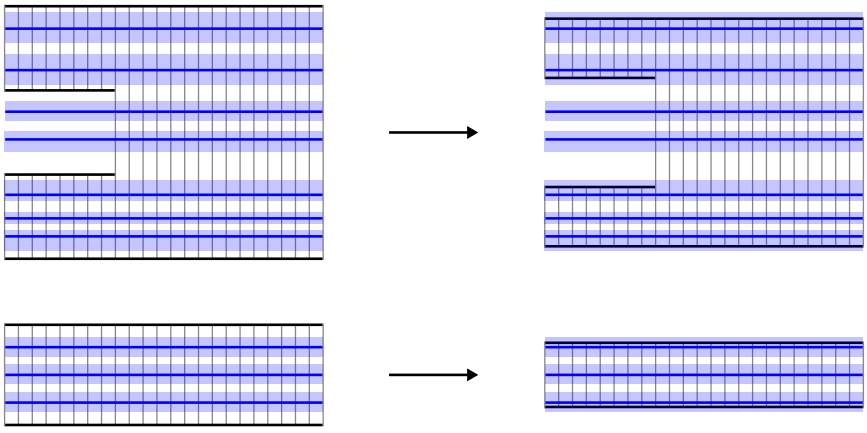

Figure 23: The loop ℓ, shown in blue, covers the train track τ_i''. We perform a small isotopy of τ_i'' which takes each arc of $\partial\tau_i''$ consisting of fiber endpoints into a regular neighborhood $N(\ell)$ of ℓ.

arc α can then be made transverse to each fiber of the band, and thus ℓ' fulfills conditions (1) and (2) of the definition of almost carried. (Note that α also vacuously fulfills condition (3).)

Now let β be a component of $\ell' \cap (P_i \setminus \mathring{\tau}_i{}'')$. The endpoints of β lie on $\partial\tau_i''$. Since the interval fiber endpoints of τ_i'' all lie in $N(\ell)$ and $\ell' \cap N(\ell)$ is empty, the endpoints of β must lie in exits. Clearly then, β satisfies conditions (1) and (2) of the definition of almost carried. The arc β cannot be parallel in $P_i \setminus N(\ell)$ to a subarc of an exit, for the parallelism would guide an isotopy of ℓ' through a band of $\tau_i'' \setminus N(\ell)$, thereby removing two components of intersection between ℓ' and the bands, once again contradicting minimality. Thus ℓ' fulfills condition (3) of the definition of almost carried.

We have shown that each arc of $\ell \cap \tau_i''$ and each arc of $\ell \cap (P_i \setminus \mathring{\tau}_i{}'')$ satisfies conditions (1), (2) and (3) of Definition 3.4. Remark 3.6 tells us that each such arc can be isotoped to be almost carried by τ_i''. Therefore ℓ' is by definition almost carried by τ_i''. □

Henceforth, on P_n, we label the punctures as $p_1, p_2, \ldots, p_{2m}$ in order from left to right. We label the straight arcs connecting the puncture labeled p_{2k-1} to the puncture labeled p_{2k} as β^k. Finally, we label the straight arcs connecting the puncture labeled p_{2k} to the puncture labeled p_{2k+1} as γ^k.

Lemma 4.17 *As above, let T_1' be the mini-graph of T_1 constructed by excluding the rightmost two eyelets, and let τ_1' be the train track diagram corresponding to T_1'. Let T_1'' be a mini-graph of T_1 constructed by possibly excluding the leftmost and/or rightmost eyelet of T_1, and let τ_1'' be the train track diagram corresponding to T_1''. Neither τ_1' nor τ_1'' almost carries the boundary of any red cap for the rightmost bridge arc α_-^m.*

Proof We prove this by contradiction. Let R be a red cap for α_-^m (which implies R is not isotopic to the blue cap $B' = dD_-^m$), and assume ∂R is almost carried by τ_1' or τ_1''. Assume R is in minimal position with respect to the vertical bridge disks below P_1.

We will first establish that R must intersect a lower vertical bridge disk. Let Γ_1 be the unique straight line segment in P_1 which both contains all of the punctures and has p_1 and p_{2m} as endpoints. If ∂R does not intersect Γ_1, then R is trivial, a contradiction. So R must intersect Γ_1. Suppose ∂R passes between the i^{th} and the $(i+1)^{\text{st}}$ vertical bridge disks for $i \in \{1, 2, \dots, m-3, m-2\}$. (This is equivalent to supposing that ∂R intersects $\sigma_1(\gamma_i)$.) Since $\sigma_1(\gamma_i)$ is surrounded by a tao, and since ∂R is almost carried by the train track τ_1' or τ_1'', the loop ∂R is forced to intersect one of the lower vertical bridge disks. Suppose ∂R passes between the $(m-1)^{\text{st}}$ and m^{th} vertical bridge disk. If ∂R does not also pass between some other pair of bridge disks, then R is the blue disk B', a contradiction. Then ∂R must pass between two of the other bridge disks, and so according to the argument above, ∂R will intersect one of the lower bridge disks. Thus we have established that R is not disjoint from the vertical bridge disks.

Let $\Lambda = R \cap \left(\bigcup_{i=1}^{m} D_-^i\right)$, which is nonempty, as shown above. Since R is in minimal position with respect to the bridge disks, Λ contains no loops of intersection, and so Λ is a collection of arcs. Let $\gamma \subseteq \Lambda$ be an outermost arc in R, cutting off an outermost disk R_{out} from R. Define $q = R_{\text{out}} \cap P_1$. Then q is an arc in P_1 with endpoints on β_i for some i. Let β_* be the subarc of β_i which shares its endpoints with q.

We examine what q can look like (and eventually arrive at a contradiction). First if q never crosses the arc Γ_1, then q would define an isotopy of R through which we could decrease the number of components of Λ, contradicting the fact that R is in minimal position with respect to the lower vertical bridge disks. Since the interior of q is by definition disjoint from the lower vertical bridge disks, q must therefore intersect $\sigma_1(\gamma^k)$ for some k. If q passes between the j^{th} and the $(j+1)^{\text{st}}$ vertical bridge disks for $j \in \{1, 2, \dots, m-3, m-2\}$, then as above, since q is almost carried and since the point $q \cap \sigma_1(\gamma_k)$ is surrounded by a tao, the train track τ_1' or τ_1'' (whichever is relevant) will force q to intersect two distinct vertical bridge disks, a contradiction. Therefore q must pass between the $(m-1)^{\text{st}}$ and the m^{th} bridge disks.

Suppose that ∂R is almost carried by τ_1''. There are five ways that q, as an almost carried arc, can pass between the $(m-1)^{\text{st}}$ and the m^{th} bridge disks, and they are illustrated in Figure 24. In cases 1 and 2, since q is almost carried, the endpoints of q must lie on both the $(m-1)^{\text{st}}$ and the $(m-2)^{\text{nd}}$ bridge disks, but that contradicts the definition of q, for both endpoints of q must lie on the same vertical bridge disk. Similarly, in cases 4 and 5, the endpoints of q must lie on both the $(m-1)^{\text{st}}$ and the m^{th} bridge disks, which contradicts the definition of q in the same way. Therefore the only case remaining is case 3, in which we see q must enter a switch of τ_1'' and go on to intersect the $(m-1)^{\text{st}}$ bridge disk. By the definition of q, the other endpoint of q must also intersect the $(m-1)^{\text{st}}$ bridge disk on the same side. Thus q has these properties:

- The arc q has both endpoints on the $(m-1)^{\text{st}}$ bridge disk.
- Of the two bands of τ_1'' going through the $(m-1)^{\text{st}}$ bridge disk, an endpoint of q is contained in the rightmost one.

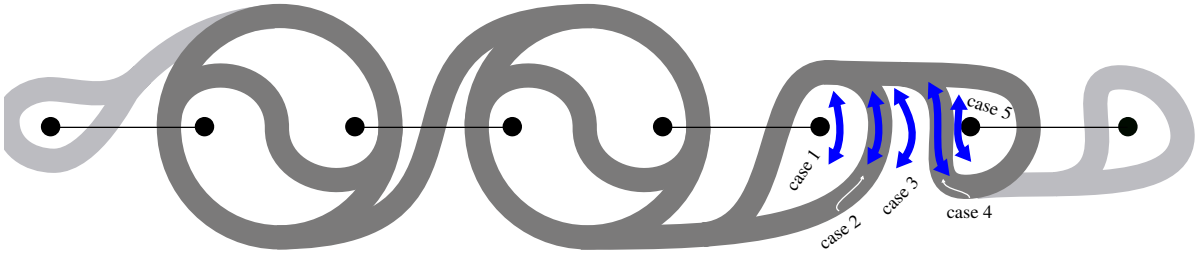

Figure 24: There are only five ways for a loop or arc to pass between the rightmost two bridge arcs while remaining almost carried by τ_1''. The leftmost and rightmost eyelets in this figure are shaded a lighter color to remind the reader that they may or may not be present in τ_1''.

- The arc q is almost carried by τ_1''.
- The interior of q does not intersect any bridge disks below P_1.
- The arc q leaves the $(m-1)^{\text{st}}$ bridge disk in the same direction from both endpoints.

Up to isotopy, there is only one arc that has these properties, and it is depicted in Figure 25. Let a and b be the left and right endpoints, respectively, of q.

The loop $q \cup \beta_*$ cuts P_1 into two punctured disks, one of which contains exactly two punctures: the endpoints of β_m. Call this 2–punctured disk Q (see Figure 25). Observe that ∂R must intersect β_m, or else R would be isotopic to B', contradicting the definition of R. Further, since both endpoints of β_m must be on the same side of ∂R, there must be an even number of points of intersection between β_m and ∂R. It follows that along the interior of β_*, there must be at least four points of intersection with ∂R. Along β_*, let c be the point of $\beta_*^\circ \cap \partial R$ nearest to a, and let d be the point of $\beta_*^\circ \cap \partial R$ nearest to b.

Since R is a cap, R cuts a 2–punctured disk F_R out of P_1 (see Figure 26). Consider the components of $\left(\bigcup_{i=1}^m D_-^i\right) \cap F_R$. There are two components which are arcs that connect a puncture in F_R to ∂R at points we will call x and y. All of the rest of arcs are parallel arcs which separate the two punctures of F_R.

The points x and y cut ∂R into two arcs, one of which must contain both endpoints of q; otherwise q would intersect β_m at point x or at point y, contradicting the definition of q. Now in F_R, a and b are

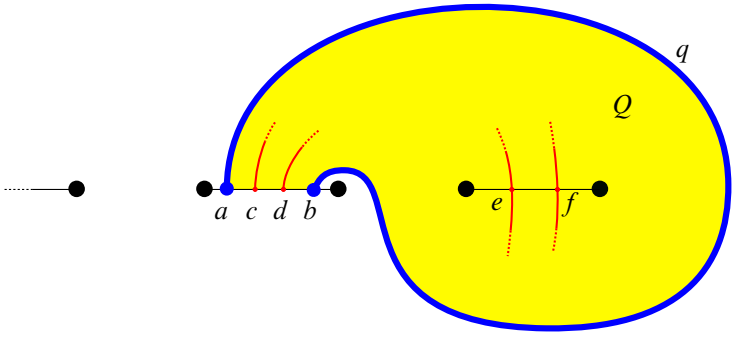

Figure 25: The arc q and the disk Q.

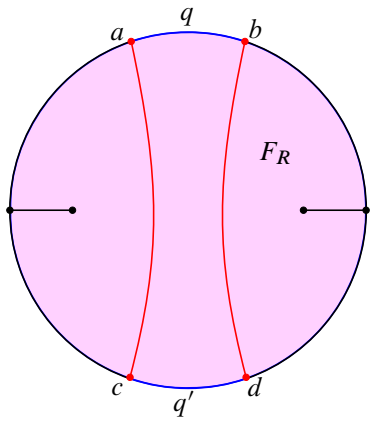

Figure 26: The arc q in relation to the disk F_R. The outer circle is ∂R, and the red vertical arcs are subarcs of β_*.

connected via arcs of ∂R to the points c and d, respectively. These two arcs in F_R with endpoints a, b, c, and d, along with the arc q and another arc of ∂R form a quadrilateral in F_R whose interior is disjoint from $\bigcup_{i=1}^{m} D^i_-$ (clear from Figure 25). Let the side of this quadrilateral whose endpoints are c and d be called q'.

Now q' is parallel to q. But at this point we could repeat this argument, focusing on q' instead of on q, which would lead us to accept the existence of another parallel arc q'', and we could repeat this infinitely many times, each time obtaining another arc of ∂R with endpoints on β_{m-1}, each of which is nested inside the last one. But this contradicts general position; it cannot be the case that β_{m-1} cuts ∂R into infinitely many subarcs. Therefore we conclude that R cannot be almost carried by τ_1''.

Assume then that ∂R is almost carried by τ_1'. In this case, there are more options for what the arc q may look like, but q still must be an arc with endpoints on a β–arc and with interior disjoint from any β–arcs (see Figure 27). The arc q cannot have its endpoints on β_m, as in Figure 28, because that would force

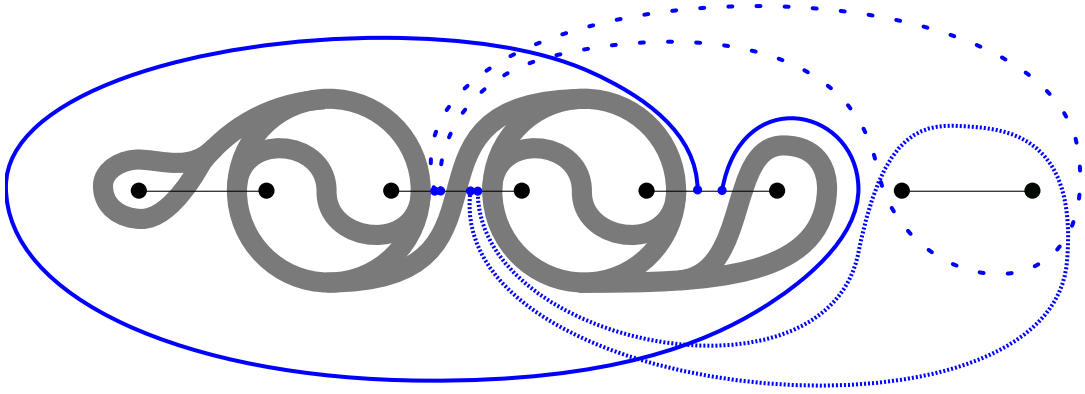

Figure 27: When ∂R is almost carried by τ_1', there are many possibilities for the arc q. Three are illustrated here (each with a different stroke style.)

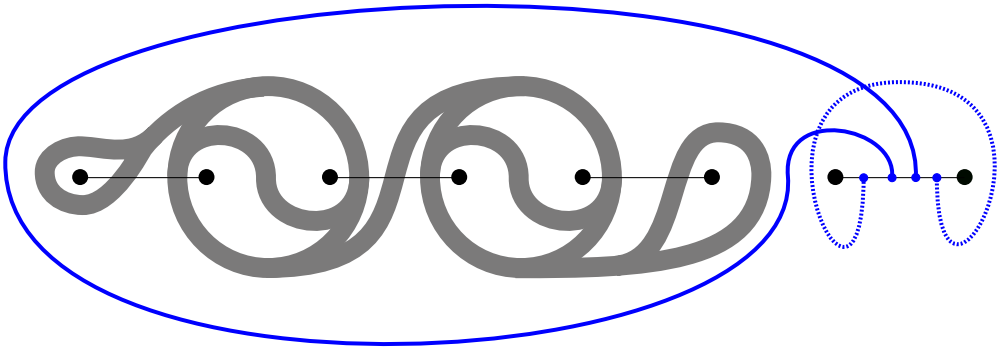

Figure 28: Pictured here are the only nonisotopic possibilities for what q could look like, given that its endpoints lie on β_m. Either way, $\partial R \backslash q$ must contain an arc that cobounds a bigon with β_m, contradicting minimality.

$\partial R \backslash q$ to contain an arc which bounds a bigon with β_m; in other words, ∂R would not be in minimal position with respect to β_m. Then since ∂q lies in β_i for some $i \in \{1, 2, 3, \ldots, m-1\}$, then $q \cup \beta_*$ bounds a disk $Q \subseteq F_R$ which contains β_m. At this point we can apply the same logic as above, leading us to assert the existence of an infinite set $\{q, q', q'', q''', \ldots\}$ of pairwise disjoint arcs of ∂R cut out by β_i, which again contradicts general position. Therefore, we may conclude that ∂R cannot be almost carried by τ_1'' either. □

Lemma 4.18 *The blue disks B and B' are a weak reducing pair for the bridge sphere.*

Proof Observe that for all points along ∂B, the y–coordinates are all less than 4. (We will speak informally this way even though technically we mean that an isotopy class of ∂B in P_n has the property that all the y–coordinates are less than 4.) Moving down a level, $\pi_{n-1}(\partial B)$ is a loop in P_{n-1} whose y–coordinates are all less than 5. Similarly, $\pi_{n-2}(\partial B)$ is a loop in P_{n-2} whose y–coordinates are all less than 6, and so on. In general, for the levels corresponding to D_a, $\pi_{n-k}(\partial B)$ is a loop in P_{n-k} whose y–coordinates are all less than $4+k$. Recall that $n = n_a + n_b$, so $n_b + 1 = n - n_a + 1 = n - (n_a - 1)$. Therefore $\pi_{n_b+1}(\partial B) = \pi_{n-(n_a-1)}(\partial B)$ is a loop in $P_{n-(n_a-1)}$ whose y–coordinates are all less than $4 + n_a - 1$. Since the dimensions of D_a were chosen so that $n_a = 2m-4$, it follows by substitution that the y–coordinates of $\pi_{n_b+1}(\partial B)$ are all less than $4 + (2m-4) - 1 = 2m-1$. This means that $\pi_{n_b+1}(\partial B)$ is completely to the left of the rightmost two punctures of P_{n_b+1}. Thus $\pi_{n_b+1}(\partial B)$ is disjoint from $\sigma_{n_b+1}(\beta_m)$ (the straight line segment connecting those two punctures).

Consider the π–projections of ∂B at consecutively lower levels. For $1 \leq t \leq n_b$, $\pi_t(\partial B)$ will remain disjoint from $\sigma_t(\beta_m)$ because the isotopy π_t from P_{t+1} to P_t fixes the two rightmost punctures of the bridge sphere pointwise. Observe that $\sigma_1(\beta_m) = \beta_m = D_-^m \cap P_1$. Therefore $\pi_1(\partial B)$ is disjoint from β_m, which implies that $\pi_1(\partial B) \cap \partial B' = \varnothing$, and so $\{B, B\}$ are a weak reducing pair. □

Observation 4.19 In general, if we compress the cap R for a bridge arc α along a boundary compressing disk Δ, the result will be two disjoint compressing disks whose boundary loops cut the bridge sphere into

two punctured disks and a twice-punctured annulus, the latter of whose punctures correspond to α. This means that neither of the resulting compressing disks are caps for α.

Lemma 4.20 *The blue cap B for α_+^1 intersects all of the red disks below P_1.*

Proof First, by Proposition 3.3, we already know that B intersects all of the red disks below P_1 which are not caps for α_-^m, so it only remains to show that B also intersects all the red caps for α_-^m.

Observe that $\partial B'$ cuts P_1 into a disk F_1 with two punctures and another disk F_2 with $2m-2$ punctures. By Corollary 4.14, $\pi_1(\partial B)$ covers T_1' (the mini-graph defined in Proposition 4.6 consisting of all of T_1 except the rightmost two eyelets). It follows from Lemma 4.18 that $\pi_1(\partial B)$ is contained in F_2.

Let R be a red cap for α_-^m, and assume by way of contradiction that $\{B, R\}$ is a weak reducing pair. Arrange for R to be in minimal position with respect to the bridge disk D_-^m. An outermost disk Δ on D_-^m cut out by an outermost arc of $D_-^m \cap R$ must be a boundary compressing disk for R, or else D_-^m and R would not be in minimal position. We perform a boundary compression of R along Δ, resulting in a disjoint union $R_1 \sqcup R_2$ of two nonparallel compressing disks for P_1. Since R was disjoint from ∂B, and the boundary compression happened away from ∂B, both R_1 and R_2 are also disjoint from B. Further, by Observation 4.19, neither R_1 nor R_2 are caps for α_-^m, and so we have two weak reducing pairs, $\{B, R_1\}$ and $\{B, R_2\}$ which both contradict Proposition 3.3. □

Lemma 4.21 *If R_a and R_b are red disks above and below the bridge sphere (respectively), then $\{R_a, R_b\}$ is not a weak reducing pair.*

Proof Assume to the contrary that $\{R_a, R_b\}$ is a weak reducing pair of red disks. By Proposition 3.3, R_a and R_b are caps for α_+^1 and α_-^m, respectively. By Corollary 4.14, the loop $\pi_1(\partial R_a)$ covers either T_1 or some mini-graph T_1'' constructed from T_1 by excluding the leftmost and/or the rightmost eyelet of T_1.

Suppose ∂R_b is disjoint from $\pi_1(\partial R_a)$. Then by Lemma 4.16, ∂R_b is isotopic to an almost carried loop, which contradicts Lemma 4.17. □

The following lemma is an immediate corollary of [Pongtanapaisan and Rodman 2021, Theorem 5.10] since the upper braid D_a of our link is a (n_a, m) plat link with $n_a = 2m-4$.

Lemma 4.22 *The cap B' is disjoint from all red disks on the other side of the bridge sphere.*

We have now shown that B and B' are the only weak reducing pair, and so we have proved our main theorem.

Theorem 1.1 *There exist infinitely many links with keen weakly reducible bridge spheres.*

Since a bridge sphere Σ_L of a link L induces a Heegaard surface $\widetilde{\Sigma}_L$ for the 2–fold cover of S^3 branched along L, it is natural to ask whether $\widetilde{\Sigma}_L$ satisfies properties that Σ_L possesses. In our situation, the answer is no.

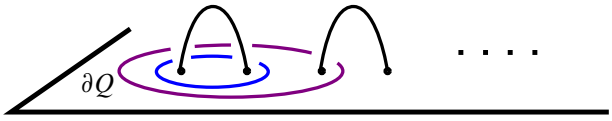

Figure 29: The purple curve bounds a once-punctured disk Q above the bridge sphere and $\partial Q \cap \partial B' = \varnothing$. To visualize Q, imagine a hemisphere shaped disk whose boundary is the purple loop and which is punctured by the second pictured bridge arc.

Proposition 4.23 *Keen weakly reducible bridge spheres in this paper do not lift to keen weakly reducible Heegaard surfaces.*

Proof The blue compressing disk B' is not only disjoint from B, but $\partial B'$ is also disjoint from a curve bounding a once-punctured disk Q above Σ_L (see Figure 29). Such a once-punctured disk lifts to a compressing disk for one of the handlebodies of the Heegaard splitting of the double branched cover. Since B' lifts to a compressing disk in the other handlebody whose boundary is disjoint from lifts of both B and Q, the Heegaard surface $\widetilde{\Sigma}_L$ is not keen. □

5 Nontopologically minimal bridge spheres

One of the main motivations of this article is to search for examples of bridge spheres that are not topologically minimal. The following criterion is needed for our construction of links with nontopologically minimal bridge spheres.

5.1 Cho's criterion

For a link in bridge position, we have that $V_\pm \backslash N(\alpha_\pm)$ is homeomorphic to a handlebody. Therefore, the complex spanned by compressing disks for the bridge sphere for L in $V_\pm \backslash N(\alpha_\pm)$ is a full subcomplex of the disk complex of the handlebody. We recall the following criterion by Cho [2008]:

Theorem 5.1 *If $\mathcal{L}$ is a full subcomplex of the disk complex of the handlebody $\mathcal{H}(V_\pm \backslash N(\alpha_\pm))$ that satisfies the following condition, then $\mathcal{L}$ is contractible*:

Let D and E be disks representing vertices of $\mathcal{L}$ and suppose that $D \cap E \neq \varnothing$. We assume that D intersects E minimally and transversely. If $\Delta \subset D$ is an outermost subdisk cut off by an outermost arc of $D \cap E$, then at least one of the disks obtained from surgery on E along Δ is also a vertex of $\mathcal{L}$.

Proposition 5.2 *The disk complex of $(V_\pm, \alpha_\pm)$ is contractible*

Proof Suppose that compressing disks D and E in $(V_\pm, \alpha_\pm)$ intersect transversely and minimally. Then the boundary of one of the disks that arises from surgery on E along Δ defined as in Theorem 5.1 must enclose at least two punctures. Otherwise, $D \cap E$ would not be minimal. □

Using Cho's criterion and Theorem 1.1, we obtain the following corollary.

Corollary 5.3 *There is an infinite family of nontrivial links with bridge spheres that are not topologically minimal.*

Proof Since the bridge sphere P_1 for each link $L \in \mathcal{L}$ contains a unique pair of disjoint compressing disks on opposite sides of P_1, there is exactly one edge connecting the contractible disk complex of (V_+, α_+) to the contractible disk complex of (V_-, α_-) showing that the disk complex of P_1 is contractible. □

References

[Bachman 2010] **D Bachman**, *Topological index theory for surfaces in* 3*–manifolds*, Geom. Topol. 14 (2010) 585–609 MR Zbl

[Bachman 2013] **D Bachman**, *Stabilizing and destabilizing Heegaard splittings of sufficiently complicated* 3*–manifolds*, Math. Ann. 355 (2013) 697–728 MR Zbl

[Bachman and Johnson 2010] **D Bachman**, **J Johnson**, *On the existence of high index topologically minimal surfaces*, Math. Res. Lett. 17 (2010) 389–394 MR Zbl

[Campisi and Rathbun 2018] **M Campisi**, **M Rathbun**, *Hyperbolic manifolds containing high topological index surfaces*, Pacific J. Math. 296 (2018) 305–319 MR Zbl

[Campisi and Torres 2020] **M Campisi**, **L Torres**, *The disk complex and topologically minimal surfaces in the* 3*–sphere*, J. Knot Theory Ramifications 29 (2020) art. id. 2050092 MR Zbl

[Cho 2008] **S Cho**, *Homeomorphisms of the* 3*–sphere that preserve a Heegaard splitting of genus two*, Proc. Amer. Math. Soc. 136 (2008) 1113–1123 MR Zbl

[E 2017] **Q E**, *On keen weakly reducible Heegaard splittings*, Topology Appl. 231 (2017) 128–135 MR Zbl

[E and Lei 2014] **Q E**, **F Lei**, *Topologically minimal surfaces versus self-amalgamated Heegaard surfaces*, Sci. China Math. 57 (2014) 2393–2398 MR Zbl

[E and Zhang 2023] **Q E**, **Z Zhang**, *On Heegaard splittings with finitely many pairs of disjoint compression disks*, J. Math. Res. Appl. 43 (2023) 496–504 MR Zbl

[Freedman et al. 1983] **M Freedman**, **J Hass**, **P Scott**, *Least area incompressible surfaces in* 3*–manifolds*, Invent. Math. 71 (1983) 609–642 MR Zbl

[Johnson 2012] **J Johnson**, *Mapping class groups of Heegaard splittings of surface bundles*, preprint (2012) arXiv 1201.2628

[Johnson and Moriah 2016] **J Johnson**, **Y Moriah**, *Bridge distance and plat projections*, Algebr. Geom. Topol. 16 (2016) 3361–3384 MR Zbl

[Ketover et al. 2019] **D Ketover**, **Y Liokumovich**, **A Song**, *On the existence of minimal Heegaard surfaces*, preprint (2019) arXiv 1911.07161

[Kim 2016] **J Kim**, *A topologically minimal, weakly reducible, unstabilized Heegaard splitting of genus three is critical*, Algebr. Geom. Topol. 16 (2016) 1427–1451 MR Zbl

[Lee 2015] **J H Lee**, *On topologically minimal surfaces of high genus*, Proc. Amer. Math. Soc. 143 (2015) 2725–2730 MR Zbl

[Lee 2016] **J H Lee**, *Bridge spheres for the unknot are topologically minimal*, Pacific J. Math. 282 (2016) 437–443 MR Zbl

[Liang et al. 2018] **L Liang**, **F Li**, **J Li**, *Reducible handle additions to weakly reducible Heegaard splittings*, Topology Appl. 242 (2018) 33–42 MR Zbl

[McCullough 1991] **D McCullough**, *Virtually geometrically finite mapping class groups of* 3*–manifolds*, J. Differential Geom. 33 (1991) 1–65 MR Zbl

[Moriah 2007] **Y Moriah**, *Heegaard splittings of knot exteriors*, from "Workshop on Heegaard splittings", Geom. Topol. Monogr. 12, Geom. Topol. Publ., Coventry (2007) 191–232 MR Zbl

[Otal 1985] **J-P Otal**, *Presentations en ponts des nœuds rationnels*, from "Low-dimensional topology", Lond. Math. Soc. Lect. Note Ser. 95, Cambridge Univ. Press (1985) 143–160 MR Zbl

[Ozawa 2011] **M Ozawa**, *Nonminimal bridge positions of torus knots are stabilized*, Math. Proc. Cambridge Philos. Soc. 151 (2011) 307–317 MR Zbl

[Pitts and Rubinstein 1987] **J T Pitts**, **J H Rubinstein**, *Applications of minimax to minimal surfaces and the topology of* 3*–manifolds*, from "Miniconference on geometry and partial differential equations, II", Proc. Centre Math. Anal. Austral. Nat. Univ. 12, Austral. Nat. Univ., Canberra (1987) 137–170 MR Zbl

[Pongtanapaisan and Rodman 2021] **P Pongtanapaisan**, **D Rodman**, *Critical bridge spheres for links with arbitrarily many bridges*, Rev. Mat. Complut. 34 (2021) 597–614 MR Zbl

[Rodman 2018] **D Rodman**, *An infinite family of links with critical bridge spheres*, Algebr. Geom. Topol. 18 (2018) 153–186 MR Zbl

[Scharlemann 2005] **M Scharlemann**, *Thin position in the theory of classical knots*, from "Handbook of knot theory", Elsevier, Amsterdam (2005) 429–459 MR Zbl

[Schultens 2009] **J Schultens**, *Width complexes for knots and* 3*–manifolds*, Pacific J. Math. 239 (2009) 135–156 MR Zbl

[Urbano 1990] **F Urbano**, *Minimal surfaces with low index in the three-dimensional sphere*, Proc. Amer. Math. Soc. 108 (1990) 989–992 MR Zbl

[Zupan 2011] **A Zupan**, *Properties of knots preserved by cabling*, Comm. Anal. Geom. 19 (2011) 541–562 MR Zbl

Arizona State University
Tempe, AZ, United States

Mathematics Department, Taylor University
Upland, IN, United States

ppongtan@asu.edu, daniel_rodman@taylor.edu

Received: 24 September 2020 Revised: 6 April 2023

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2024 is US $705/year for the electronic version, and $1040/year (+$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

https://msp.org/

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 24 Issue 8 (pages 4139–4730) 2024