AGT

Algebraic & Geometric Topology

Volume 24 (2024)

Bounding the Kirby–Thompson invariant of spun knots

Román Aranda
Puttipong Pongtanapaisan
Scott A Taylor
Suixin (Cindy) Zhang

msp

Algebraic & Geometric Topology **24**:6 (2024) 3363–3399
DOI: 10.2140/agt.2024.24.3363
Published: 7 October 2024

Bounding the Kirby–Thompson invariant of spun knots

Román Aranda
Puttipong Pongtanapaisan
Scott A Taylor
Suixin (Cindy) Zhang

A bridge trisection of a smooth surface in S^4 is a decomposition analogous to a bridge splitting of a link in S^3. The Kirby–Thompson invariant of a bridge trisection measures its complexity in terms of distances between disk sets in the pants complex of the trisection surface. We give the first significant bounds for the Kirby–Thompson invariant of spun knots. In particular, we show that the Kirby–Thompson invariant of the spun trefoil is 15.

57K45

1 Introduction

Every smooth surface in the 4–sphere S^4 (or indeed any 4–manifold) admits a certain kind of decomposition known as a *bridge trisection*. These bridge trisections are analogous to bridge positions of classical knots in S^3. They give rise to the fundamental notion of the *bridge number* $\mathfrak{b}(S)$ of a knotted smooth surface S. Bridge trisections and bridge number were defined by [Meier and Zupan 2017] and are closely related to the trisections of smooth 4–manifolds of Gay and Kirby [2016]. The major advantage of both bridge trisections and trisections of 4–manifolds is that the handle structure of the knotted surface or 4–manifold is captured using 2–dimensional data on the trisection surface Σ. They also give rise to certain diagrammatic representations of knotted surfaces. In recent years, many authors have connected (bridge) trisections to major open problems in the theory of 2–knots and 4–manifolds [Lambert-Cole 2020; Lambert-Cole et al. 2021; Gay and Meier 2022].

One pressing problem has been to develop new 2–knot or 4–manifold invariants using trisections. Kirby and Thompson [2018] defined a nonnegative integer-valued 4–manifold invariant $\mathscr{L}(M)$ using the cut complex of Σ. Blair et al. [2022] adapted Kirby and Thompson's definition to create an nonnegative integer-valued invariant $\mathscr{L}(S)$ of a smooth surface in S^4. They showed that, for orientable S, if $\mathscr{L}(S) = 0$ then S is an unlink. They also showed that, for a connected, irreducible surface S, $\mathscr{L}(S) > \mathfrak{b}(S) - g(S) - 2$, where $g(S)$ is the genus of S. Using spun knots, Meier and Zupan show that $\mathfrak{b}(S)$ can be arbitrarily large for 2–knots S; consequently, $\mathscr{L}(S)$ can be as well. However, for spun 2–bridge knots, the only previously known lower bound is that $\mathscr{L}(S)$ is nonzero. Calculating $\mathscr{L}(S)$ for specific surfaces remains a challenging problem, as does showing that, for a fixed bridge number, $\mathcal{L}(S)$ can be arbitrarily large. We take steps toward those questions by showing:

Theorem 1.1 *Let $K \subset S^3$ be a 2–bridge knot. If K is the numerator closure of a 2–string trivial tangle with Conway number p/q, then*

$$15 \le \mathcal{L}(S(K)) \le \min\{6d(p/q, 0) + 6, 6d(p/q, \infty) + 9\}.$$

In particular, if K is a trefoil knot $3/1$, then $\mathcal{L}(S(K)) = 15$.

Proof The lower and upper bounds are proven in Corollaries 3.17 and 4.5, respectively. □

More generally, we construct estimates for any spun knot. For a trivial N–tangle T, we define $\mathcal{P}_{\mathrm{comp}}(T)$ and $\mathcal{P}_c(T)$ to be the sets of pants decompositions in the pants complex $p \in \mathcal{P}(\Sigma_{2N})$ such that all loops in p bound compressing disks and c–disks, respectively.

Theorem 1.2 *Let $K = T_K^+ \cup T_K^-$ be a knot in b–bridge position. Let $d \ge 0$ be the distance in $\mathcal{P}(\Sigma_{2b})$ between the sets $\mathcal{P}_c(T_K^+)$ and $\mathcal{P}_{\mathrm{comp}}(T_K^-)$. Then*

$$6b - 8 \le \mathcal{L}(S(K)) \le 6(d + b - 1).$$

Proof The upper bound is proven in Theorem 4.3 for a particular minimal bridge trisection of $S(K)$. Since $\mathcal{L}(S(K))$ is the minimum value of $\mathcal{L}(\mathcal{T})$ along all minimal bridge trisections of $S(K)$ (see Section 2.4), the upper bound holds. The lower bound is [Blair et al. 2022, Theorem 6.3]. □

The invariant $\mathcal{L}(\mathcal{T})$ for a bridge trisection $\mathcal{T}$ with trisection surface Σ is defined using the pants complex of $\mathcal{T}$ and the associated disk complexes (see Section 2.4). Most of the delicate combinatorial work in this paper consists of a careful analysis of paths in the pants complex. Our techniques may, therefore, also be of interest to those working on surface dynamics. In fact, most of our work in Section 3 consists in understanding the combinatorics of $(4, 2)$–bridge trisections. We show:

Theorem 3.16 *Let $\mathcal{T}$ be a $(4, 2)$–bridge trisection for a knotted connected surface F in S^4. Then*

$$L(\mathcal{T}) \ge 15.$$

Meier and Zupan [2017] described bridge trisection diagrams $\mathcal{T}_{\mathrm{MZ}}$ for twist spun knots. Even though (± 1)–twist 2–bridge knots are unknotted, it is unclear whether their bridge trisections $\mathcal{T}_{\mathrm{MZ}}$ are stabilized. They form a family of candidates of nonstabilized nonminimal bridge trisections. In order to disprove this, one could try to build upper bounds for $\mathcal{L}(\mathcal{T}_{\mathrm{MZ}})$ of (± 1)–twist spun knots and use Theorem 3.16 to see they are stabilized.

Acknowledgements Taylor was partially supported by NSF grant DMS-2104022. Aranda and Pongtanapaisan were partially supported by a grant from the Berger Fund and Colby College. Pongtanapaisan acknowledges the Pacific Institute for the Mathematical Sciences for its support. We are grateful to Nathaniel Ferguson for helpful conversations and to Jeffrey Meier for suggesting this project.

2 Preliminaries

In this section, we introduce terminology and recall the definitions of the pants complex, a genus-0 trisection of S^4 and bridge trisections, and the invariant $\mathcal{L}$. For more detailed explanations please refer to [Meier and Zupan 2017; Blair et al. 2022]

2.1 The pants complex

Suppose that Σ is a compact surface with punctures. A simple closed curve $\gamma \subset \Sigma$ is called *essential* if it is disjoint from the punctures, does not bound an unpunctured or once-punctured disk in Σ, and does not cobound an unpunctured annulus in Σ with $\partial\Sigma$. If Σ is a sphere, we define the *inside* of a simple closed curve in Σ to be the sides with the fewest punctures and the *outside* to be a side that is not an inside. Some curves have two inside regions and no outside region. We say that a simple closed curve in a sphere Σ is an *odd curve* if the number of punctures on each side is odd and an *even curve* otherwise.

A *pair of pants* is a sphere with three punctures, an annulus with one puncture, or a disk with two punctures. A *pants decomposition* of Σ is a collection of pairwise disjoint essential curves cutting Σ into pairs of pants. Pants decompositions are considered up to isotopy. If Σ is a sphere with $2b \geq 4$ punctures, then each pants decomposition of Σ has $2b-3$ curves. Define $P(\Sigma)$, the *pants complex*[1] of Σ, as follows. Each pants decomposition of Σ is a vertex of $P(\Sigma)$. Two vertices are connected by an edge if the two corresponding pants decompositions have all but one (isotopy class of) curve in common and the two curves where they differ (have representatives that) intersect minimally in exactly two points. We say that the two endpoints of an edge differ by an *A–move*. The distance $d(x,y)$ between two collections of vertices x and y in $P(\Sigma)$ is the minimum number of edges in a path in $P(\Sigma)$ between a vertex of x and a vertex of y. For a path α in $\mathcal{P}(\Sigma)$, we say that a curve $\gamma \subset \Sigma$ is *unmoved* on α if it (up to isotopy) belongs to every vertex of α. On the other hand, if we have a path from vertex a to vertex b and if c is a curve in a pants decomposition x that is a vertex of the path, then, if the edge of the path leaving x corresponds to an A–move replacing c with c', we say that c is *moved* by the path and write $c \mapsto c'$. Clearly, the length of the path is at least the number of curves moved by the path; some curves may be moved multiple times, so it need not be equal to the number of curves that are moved.

A *trivial tangle* (B_δ, δ) is a 3–ball B_δ containing properly embedded arcs δ such that, fixing the endpoints of δ, we may isotope δ into ∂B_δ. We consider the endpoints of δ on $\Sigma = \partial B_\delta$ to be punctures on Σ. A *c–disk* in (B_δ, δ) is a properly embedded disk $D \subset B_\delta$ transverse to δ, with ∂D essential in the (punctured) surface Σ, and with $|D \cap \delta| \leq 1$. The c–disk D is a *compressing disk* if $|D \cap \delta| = 0$ and a *cut disk* otherwise. The *disk set* $\mathfrak{D}(B_\delta, \delta)$ for (B_δ, δ) consists of the vertices v of $\mathcal{P}(\Sigma)$ such that each curve in the pants decomposition v bounds a c–disk in (B_δ, δ).

Each arc δ_0 of a trivial tangle (B_δ, δ) admits a disk D such that ∂D is the endpoint union of δ_0 with an arc on ∂B_δ and with interior disjoint from δ. Such a disk is called a *bridge disk* and the arc on ∂B_δ is a

[1] It is possible to define higher-dimensional simplices of $\mathcal{P}(\Sigma)$, but we will not make use of them.

shadow arc. There is a collection of pairwise disjoint bridge disks such that each arc of δ belongs to a bridge disk. The union of all the shadow arcs for such a collection of bridge disks is a *complete shadow arc collection*.

For a link $L \subset S^3$, a decomposition $(S^3, L) = (B_\lambda, \lambda) \cup_\Sigma (B_\tau, \tau)$, where each pair (B_δ, δ) is a trivial tangle, is called a *bridge splitting*. The surface $\Sigma = \partial B_i$ for $i = \lambda, \tau$ is the *bridge sphere* of the splitting. An *efficient defining pair* is a pair-of-pants decomposition $(\mathcal{D}_\kappa, \mathcal{D}_\lambda)$ with $x \in \mathcal{D}_\kappa$ and $y \in \mathcal{D}_\lambda$ such that $d(x, y) = d(\mathcal{D}_\kappa, \mathcal{D}_\lambda)$. Zupan [2013] uses this distance to define a knot invariant for knots in S^3. We need the following well-known result (see [Bachman and Schleimer 2005; Zupan 2013]):

Lemma 2.1 *Suppose that Σ is a bridge sphere for an unlink $L \subset S^3$; then:*

(1) *If $|L| \geq 2$, there is a sphere $P \subset S^3$ intersecting Σ in a single essential simple closed curve and separating components of L. Such a sphere is called a* ***reducing sphere*** *for Σ.*

(2) *If L_0 is a component of L such that $|L_0 \cap \Sigma| = 2$, then there is a disk with boundary equal to L_0 and interior disjoint from L such that $L_0 \cap \Sigma$ is a single arc. Furthermore, given a collection of pairwise disjoint reducing spheres, there is such a disk disjoint from them.*

(3) *If L_0 is a component of L such that $|L_0 \cap \Sigma| \geq 4$, then there exist disks D_1 and D_2 on opposite sides of Σ such that:*

  (a) *For $i = 1, 2$, ∂D_i is the endpoint union of a strand of $L \setminus \Sigma$ and an arc on Σ.*

  (b) *For $i = 1, 2$, the interior of D_i is disjoint from $L \cup \Sigma$.*

  (c) *$D_1 \cap D_2$ is a single point (necessarily a puncture of Σ).*

  In this case, we say that L is ***perturbed*** *and call the disks D_1 and D_2 a* ***perturbing pair****. Furthermore, given a collection of pairwise disjoint reducing spheres, there exists a perturbing pair disjoint from them.*

Definition 2.2 For a link L in S^3 with bridge sphere Σ, the intersection of a reducing sphere with Σ is called a *reducing curve* for (S^3, L) on Σ. Notice that an essential curve is a reducing curve if and only if it bounds compressing disks for Σ in both of the trivial tangles on either side of Σ. Similarly, if $\gamma \subset \Sigma$ is a curve bounding cut disks on both sides of Σ, then γ is a *cut-reducing curve* for (S^3, L) on Σ.

2.2 Bridge trisections

Suppose that S is a smooth, closed surface in S^4. A *bridge trisection* $\mathcal{T}$ with trisection surface Σ (a sphere) is defined as follows.[2] Suppose that W_1, W_2 and W_3 are 4–balls in S^4 such that $W_i \cap W_j$ is a 3–ball B_{ij} (for $i \neq j$) and that

$$W_1 \cap W_2 \cap W_3 = B_{12} \cap B_{23} \cap B_{31}$$

[2] It is possible to define higher-genus bridge trisections [Meier and Zupan 2018], but we will not need them in this paper.

is a smooth 2–sphere Σ. Then we say that $S^4 = W_1 \cup W_2 \cup W_3$ is a 0–*trisection* of S^4 [Gay and Kirby 2016]. Suppose also that each of B_{12}, B_{23} and B_{31} are transverse to S and that Σ and S intersect transversally in $2b$ points and that, for each $\{i,j,k\} \in \{1,2,3\}$:

(1) $S \cap W_i$ is a trivial disk system.

(2) In $B_{ij} \cup B_{jk}$, the sphere Σ is a bridge surface for the link $S \cap (B_{ij} \cup B_{jk})$.

(3) The link $S \cap (B_{ij} \cup B_{jk})$ is an unlink of c_j components.

We call $\mathcal{S} = (B_{12}, T_{12}) \cup (B_{23}, T_{23}) \cup (B_{31}, T_{31})$ the *spine* of the bridge trisection and Σ the *bridge surface* of S. The numbers c_1, c_2, c_3 are the *patch numbers* of the bridge trisection. The *bridge number* $\mathfrak{b}(\mathcal{T})$ of the trisection is $\mathfrak{b}(\mathcal{T}) = \frac{1}{2}|S \cap \Sigma|$ and the *bridge number* $\mathfrak{b}(S)$ of S is the minimum of $\mathfrak{b}(\mathcal{T})$ over all bridge trisections $\mathcal{T}$ for S. We say that a trisection $\mathcal{T}$ with bridge number b and patch numbers c_1, c_2 and c_3 is a $(b; c_1, c_2, c_3)$–bridge trisection. As we mentioned, the definitions of bridge trisection and bridge number are due to Meier and Zupan, who also proved that every smooth surface admits a bridge trisection. We let $\mathcal{D}_{ij} \subset \mathcal{P}(\Sigma)$ be the disk set of the tangle (B_{ij}, T_{ij}).

Meier and Zupan [2017] also introduce the notion of a *triplane diagram*, a triple of planar tangle diagrams whose pairwise unions are unlinks. Since a bridge trisection is determined by its spine consisting of a triple of 3–balls (B_{12}, B_{23}, B_{31}) with trivial tangles (T_{12}, T_{23}, T_{31}), one can project each tangle T_{ij} onto a vertical disk in B_{ij} and obtain three planar tangle diagrams. In particular, every knotted surface in S^4 can be represented by a triplane diagram which is unique up to interior Reidemeister moves, bridge sphere braiding, and perturbation and deperturbation. See [Meier and Zupan 2017, Section 2] for details.

Lemma 2.3 *Suppose that $S \subset S^4$ is a topologically knotted sphere with a $(4; c_1, c_2, c_3)$–trisection and $4 = \mathfrak{b}(S)$. Then $c_i = 2$ for all i.*

Proof Since S is topologically knotted, $c_i \geq 2$ for all i by [Meier and Zupan 2017, Corollary 1.12]. The result follows since $2 = \chi(S) = c_1 + c_2 + c_3 - 4$. □

Henceforth, we abbreviate the phrase "$(4; 2, 2, 2)$–trisection" to $(4, 2)$–trisection.

2.3 Spun knots

We now recall a construction of spun knots from a knot $K \subset S^3$ due to Artin [1925]. Let (B^3, K°) be the result of removing a small, open ball centered on a point in K, so that K is a knotted arc with endpoints on the north and south poles, labeled n and s, respectively. Then the spin $S(K)$ of K is the knotted surface given by

$$(S^4, S(K)) = ((B^3, K^\circ) \times S^1) \cup ((S^2, \{n, s\}) \times D^2).$$

Meier and Zupan also showed that every spun b–bridge knot $S(K) \in S^4$ has bridge number at most $3b-2$ by providing an explicit $(3b-2, b)$–bridge trisection, whose corresponding triplane diagram is shown in

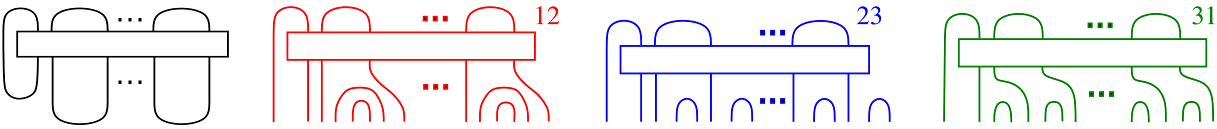

Figure 1: A $(3b-2, b)$–bridge triplane diagram for the spin $\mathcal{S}(K)$ of the b–bridge knot K given in bridge position (left). We will denote the tangles by T_{12}, T_{23}, and T_{31} from left to right.

Figure 1. From now on, we will denote this particular bridge trisection by $\mathcal{T}_{\mathrm{MZ}}$ and, for that trisection, define T_{ij} as indicated for $i, j \in \{1,2,3\}$ with $i \neq j$.

Remark 2.4 For this particular trisection $\mathcal{T}_{\mathrm{MZ}}$ for a spun b–bridge knot, since $\mathfrak{b} = \mathfrak{b}(\mathcal{T}_{\mathrm{MZ}}) = 3b-2$ and $c_i = b$ for all $i \in \{1,2,3\}$, the corresponding bridge sphere is $2\mathfrak{b}$–punctured, and each pants decomposition p^i_{ij} has exactly $2\mathfrak{b}-3 = 2(3b-2)-3 = 6b-7$ curves. Thus, it follows from Lemma 2.7 that there exist $p^i_{ij} \in \mathcal{D}_{ij}$ and $p^i_{ki} \in \mathcal{D}_{ik}$ with $d(p^i_{ij}, p^i_{ki}) = \mathfrak{b} - c_i = (3b-2) - b = 2b-2$.

We note the following:

Theorem 2.5 [Meier and Zupan 2017] *If $K \subset S^3$ has $\mathfrak{b}(K) = 2$, then $\mathfrak{b}(\mathcal{S}(K)) = 4$. Consequently, if $\mathcal{T}$ is a $(4; c_1, c_2, c_3)$–trisection for a spun 2–bridge knot, then each $c_i = 2$.*

Proof We defer to [Meier and Zupan 2017, Section 5] for details. Let $\mathcal{T}$ be a $(b; c_1, c_2, c_3)$–bridge trisection of a spun 2–bridge knot $\mathcal{S}(K)$. By [Meier and Zupan 2017, Corollary 5.3 and Theorem 5.5],

$$\min(c_1, c_2, c_3) \geq \mathrm{mrk}(\mathcal{S}(K)) = \mathrm{mrk}(K),$$

where mrk is the "meridional rank" of the 2–knot or knot. By [Boileau and Zimmermann 1989], $\mathrm{mrk}(K) = 2$, so $c_i \geq 2$ for all i. Also,

$$2 = \chi(\mathcal{S}(K)) = c_1 + c_2 + c_3 - b \geq 6 - b.$$

Thus, $b \geq 4$. Since Meier and Zupan have constructed trisections of spun 2–bridge knots of bridge number 4, $\mathfrak{b}(\mathcal{S}(K)) = 4$. Since the meridional rank of $\mathcal{S}(K) = 2$, $\mathcal{S}(K)$ is topologically knotted. The result follows from Lemma 2.3. □

2.4 The Kirby–Thompson invariant

We now define the Kirby–Thompson invariant of a bridge trisection. For a schematic diagram of the efficient defining pairs for a trisection, see Figure 2.

Definition 2.6 (Kirby–Thompson invariant $\mathcal{L}$) Suppose that $S \subset S^4$ is knotted surface with bridge trisection $\mathcal{T}$ having trisection surface Σ and spine $\mathcal{S} = (B_{12}, T_{12}) \cup (B_{23}, T_{23}) \cup (B_{31}, T_{31})$. For $\{i, j, k\} = \{1,2,3\}$, let (p^j_{ij}, p^j_{jk}) be an efficient defining pair for $(B_{ij}, T_{ij}) \cup_\Sigma (B_{jk}, T_{jk})$. If Σ is a

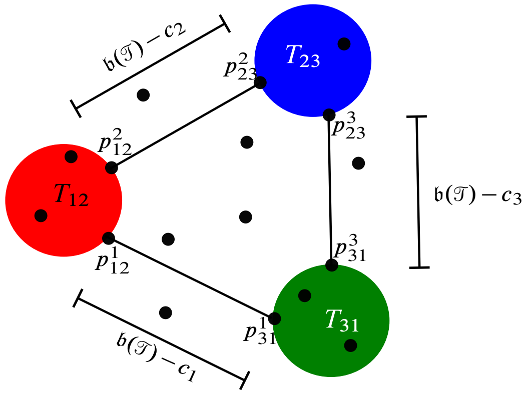

Figure 2: Defining $\mathcal{L}(T)$ via efficient defining pairs. The ellipses represent the disk sets. The line joining p^i_{ij} to p^j_{ij} represents a geodesic path in the pants complex, which has length $\mathfrak{b}(\mathcal{T})-c_i$ for a $(\mathfrak{b}(\mathcal{T}), c_1, c_2, c_2)$–bridge trisection.

sphere with strictly fewer than four punctures, define $\mathcal{L}(\mathcal{T})=0$. Otherwise, define the *Kirby–Thompson invariant* $\mathcal{L}(\mathcal{T})$ to be the minimum of

$$d(p^1_{12}, p^2_{12}) + d(p^2_{23}, p^3_{23}) + d(p^1_{31}, p^3_{31})$$

over all such choices of efficient defining pairs. Define the *Kirby–Thompson invariant* $\mathcal{L}(S)$ to be the minimum of $\mathcal{L}(\mathcal{T})$ over all trisections $\mathcal{T}$ of S with $\mathfrak{b}(\mathcal{T})=\mathfrak{b}(S)$.

The distance between an efficient defining pair in the setting of Definition 2.6 is determined:

Lemma 2.7 [Blair et al. 2022, Lemma 5.6] *If $\mathcal{T}$ is a $(\mathfrak{b}(\mathcal{T}), c_1, c_2, c_3)$–bridge trisection, then every efficient defining pair satisfies*

$$d(p^i_{ij}, p^i_{ik}) = \mathfrak{b}(\mathcal{T}) - c_i.$$

2.5 Reducibility and stabilization of bridge trisection

We provide two related ways in which a bridge trisection may have higher bridge number than necessary: reducibility and stabilization.

Definition 2.8 Given two trisections $\mathcal{T}_i$ for surfaces S_i $(i=1,2)$ in distinct copies of S^4, their *distant sum* is the trisection obtained by taking the connected sum of the two copies of S^4 using a point on each trisection surface disjoint from the surfaces. Their *connected sum* is the trisection obtained by taking the connected sum of the two copies of S^4 using punctures on the two trisection surfaces. For more details, see [Meier and Zupan 2017]. A trisection with trisection surface Σ is *reducible* if there exists an essential simple closed curve in Σ bounding a c–disk in each tangle forming the spine.

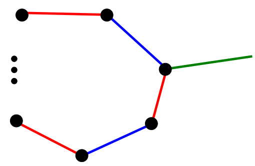

Figure 3: The arrangement of arcs from Lemma 2.12.

Lemma 2.9 *If S is a knotted 2–sphere with $\mathfrak{b}(S) \leq 7$, then no bridge trisection of minimal bridge number is reducible.*

Proof As explained in [Blair et al. 2022], if a trisection $\mathcal{T}$ is a reducible $(4,2)$–bridge trisection for S, then it is the connected sum of two other trisections $\mathcal{T}_1$ and $\mathcal{T}_2$ such that $\mathfrak{b}(\mathcal{T}_1) + \mathfrak{b}(\mathcal{T}_2) = \mathfrak{b}(\mathcal{T}) + 1 \leq 7$ and each has bridge number at least 2. In particular, either $\mathcal{T}_1$ or $\mathcal{T}_2$ has bridge number at most 3, implying that the corresponding surface is unknotted by [Meier and Zupan 2017, Theorem 1.8]. In this case, the other trisection is a trisection for S of smaller bridge number than $\mathcal{T}$. □

Lemma 2.10 *Suppose that $\mathcal{T}$ is a bridge trisection with spine $\bigcup_{i \neq j}(B_{ij}, T_{ij})$. Then $\mathcal{T}$ is reducible or stabilized if and only if there is an essential curve γ bounding a c–disk in each (B_{ij}, T_{ij}). Furthermore, such a curve is a reducing or cut-reducing curve (respectively) for each link $L_j = T_{ij} \cup \overline{T}_{jk}$.*

Proof This follows easily from Lemma 2.1. □

Meier and Zupan [2017, Section 6] define what it means for a bridge trisection to be *stabilized*. This is the analogous to a "perturbed bridge surface" for knots in 3–manifolds or to "stabilized Heegaard splittings" of 3–manifolds. While we do not need the precise definition of stabilization, we need the following two results, both from [Meier and Zupan 2017].

Lemma 2.11 *If $S \subset S^4$, then no stabilized bridge trisection of S has minimal bridge number.*

Lemma 2.12 (stabilization criterion [Meier and Zupan 2017, Lemma 6.2]) *Let $\mathcal{T}$ be a bridge trisection with spine*

$$(B_{12}, T_{12}) \cup (B_{23}, T_{23}) \cup (B_{31}, T_{31}).$$

If, for some $\{i, j, k\} = \{1, 2, 3\}$, there exists a collection of shadow arcs α for (B_{ij}, T_{ij}) and β for (B_{jk}, T_{jk}) and a single shadow arc γ for (B_{ik}, T_{ik}) such that the interiors of all the shadow arcs are disjoint and the following two conditions hold, then $\mathcal{T}$ is stabilized:

(1) *The union $\alpha \cup \beta$ is a simple closed curve (ignoring the punctures).*

(2) *Exactly one endpoint of γ lies on $\alpha \cup \beta$.*

Noting that the union of an arc with an isotopic copy having interior disjoint from the original is a circle produces the following criterion, which we'll use repeatedly:

Lemma 2.13 *Let $\mathcal{T}$ be a bridge trisection with spine*

$$(B_{12}, T_{12}) \cup (B_{23}, T_{23}) \cup (B_{31}, T_{31}).$$

Suppose that there exist $\{i, j, k\} = \{1,2,3\}$ such that there is a shadow arc α for both (B_{ij}, T_{ij}) and (B_{jk}, T_{jk}) and a shadow arc γ for (B_{ik}, T_{ik}) sharing exactly one endpoint with α and with interior disjoint from α. Then $\mathcal{T}$ is stabilized.

We note that Blair et al. [2022] show that,if a $(b; c_1, c_2, c_3)$–bridge trisection $\mathcal{T}$ of a knotted surface S is not reducible, then

$$\mathcal{L}(\mathcal{T}) \geq 2(c_1 + c_2 + c_3) - 8.$$

If $\mathcal{T}$ is a $(4,2)$–bridge trisection, this inequality translates to $\mathcal{L}(\mathcal{T}) \geq 2 \cdot 6 - 8 = 4$. The goal of Section 3 is to further improve this estimate in Theorem 3.16.

3 Combinatorics of (4, 2)–bridge trisections

This section studies relations among pairs-of-pants decompositions of a trisection surface Σ having 8 punctures. For each $\{i, j, k\} = \{1,2,3\}$, the link $L_i = T_{ij} \cup \overline{T}_{ik}$ is a 2–component unlink in 4–bridge position. We define an *inside* of a simple closed curve in Σ to be a side with ≤ 4 punctures and an *outside* to be a side with > 4 punctures. Note that curves with four punctures on each side have two inside regions and no outside region. We say that a puncture or set of punctures is *enclosed* by such a curve if the curve does not separate them and they are all inside the curve. Analyzing which curves in a pants decomposition can enclose which others, produces the next lemma:

Lemma 3.1 *Let (p^i_{ij}, p^i_{ik}) be an efficient defining pair for L_i. Then we may choose notation $p^i_{ij} = \{\gamma_1, \gamma_2, \gamma_3, f_1, f_2\}$ and $p^i_{ik} = \{\gamma_1, \gamma_2, \gamma_3, f'_1, f'_2\}$ so that all of the following hold*:

- *γ_1 is a reducing curve for L_i.*
- *Both γ_2 and γ_3 are cut-reducing curves for L_i.*
- *f_1 and f_2 bound compressing disks for T_{ij} and f'_1 and f'_2 bound compressing disks for T_{ik}.*
- *Every geodesic from p^i_{ij} to p^i_{ik} moves f_1 to f'_1 and f_2 to f'_2, and γ_1, γ_2, and γ_3 are unmoved.*

Proof Recall that Σ has eight punctures, so each pants decomposition has five curves. Let (p^i_{ij}, p^i_{ik}) be an efficient defining pair. By Lemma 2.7, the distance from p^i_{ij} to p^i_{ik} is equal to $\mathfrak{b}(\mathcal{T}) - c_i = 2$. Thus, at least three curves are unmoved by any geodesic in the pants complex joining p^i_{ij} to p^i_{ik}. Let γ_1, γ_2 and γ_3 be three such curves, and let f_1 and f_2 be the other two. Curves in Σ bounding cut disks in one of the tangles in the spine, enclose an odd number of punctures in Σ, while those bounding compressing disks enclose an even number of punctures. Thus, each of γ_1, γ_2 and γ_3 is either a reducing curve or a cut-reducing curve for L_i.

It is impossible for γ_1, γ_2 and γ_3 to all bound cut disks to both sides, because there are only eight punctures and the three curves are pairwise nonparallel. Thus, at least one is a reducing curve. Without loss of generality, we may assume it is γ_1. Since $c_i = 2$, all reducing curves for L_i enclose the same punctures. Thus, γ_2 and γ_3 must be cut-reducing curves. Each encloses exactly three punctures. Since p^i_{ij} is a pants decomposition, all other curves of p^i_{ij} enclose an even number of punctures. Consequently, both f_1 and f_2 must be moved by every geodesic between p^i_{ij} and p^i_{ik}. Thus, each geodesic moves the pair (f_1, f_2) to the pair (f_1', f_2'), which are the curves of p^i_{ik} that are not γ_1, γ_2 or γ_3.

Furthermore, one of γ_2 or γ_3 encloses three punctures as well as either f_1 or f_2. Since no geodesic between p^i_{ij} and p^i_{ik} moves γ_2 or γ_3, there are not two geodesics one of which moves f_1 to f_1' and the other of which moves it to f_2'. Thus, we may assume the notation was chosen so that every such geodesic moves f_1 to f_1' and f_2 to f_2'. □

Remark 3.2 We will often consider efficient defining pairs (p^i_{ij}, p^i_{ik}) and (p^j_{ij}, p^j_{jk}), in which case we choose notation $p^i_{ij} = \{\gamma_1, \gamma_2, \gamma_3, f_1, f_2\}$ and $p^j_{ij} = \{\psi_1, \psi_2, \psi_3, h_1, h_2\}$ as in Lemma 3.1. We refer to any of γ_1, γ_2 or γ_3 as a *γ_n–loop* and any of ψ_1, ψ_2 or ψ_3 as a *ψ_n–loop*.

A *configuration* of either T_{ij}, T_{jk} or L_i is the partition Δ_{ij}, Δ_{jk} or Δ_i (respectively) of the set of the labeled punctures $L = \{1, 2, 3, 4, 5, 6, 7, 8\}$ on Σ built as follows: two punctures are related if they belong to the same connected component of T_{ij}, T_{jk} or L_i, respectively. We will often abbreviate the string "3, 4, 5, 6, 7, 8" as 3–8, and so forth. An element of a configuration with exactly n elements is called an *n–cycle*.

We are interested in the triplet of configurations $(\Delta_1, \Delta_2, \Delta_3)$ for L_1, L_2 and L_3. Up to relabeling, $(4, 2)$–bridge trisection has essentially three options for such triplets. This is formalized in Lemma 3.3:

Lemma 3.3 *Let S be a connected surface in S^4 with a $(4, 2)$–bridge trisection T. Up to permutation of L and choice $\{i, j, k\} = \{1, 2, 3\}$, there are three possible configurations for Δ_i, Δ_j and Δ_k:*

(1) $\Delta_i = \{\{1, 2\}, \{3\text{–}8\}\}$, $\Delta_j = \{\{1\text{–}5, 8\}, \{6, 7\}\}$ *and* $\Delta_k = \{\{3, 4\}, \{1, 2, 5\text{–}8\}\}$.

(2) $\Delta_i = \{\{1, 2\}, \{3\text{–}8\}\}$, $\Delta_j = \{\{1, 2, 6, 7\}, \{3, 4, 5, 8\}\}$ *and* $\Delta_k = \{\{3, 4\}, \{1, 2, 5\text{–}8\}\}$.

(3) $\Delta_i = \{\{1\text{–}4\}, \{5\text{–}8\}\}$, $\Delta_j = \{\{1, 4, 5, 8\}, \{2, 3, 6, 7\}\}$ *and* $\Delta_k = \{\{1, 2, 7, 8\}, \{3\text{–}6\}\}$.

Proof The fact that T is a $(4, 2)$–bridge trisection implies that Δ_1, Δ_2 and Δ_3 each have either one 2–cycle and one 6–cycle or exactly two 4–cycles.

Case 1 (Δ_j has one 2–cycle) After relabeling, we can assume that $\Delta_{ij} = \{\{1, 2\}, \{3, 4\}, \{5, 6\}, \{7, 8\}\}$ and $\Delta_{jk} = \{\{1, 2\}, \{3, 8\}, \{4, 5\}, \{6, 7\}\}$. By connectivity of F, we have that $\{1, 2\} \notin \Delta_{ik}$. We have two cases: either Δ_{ik} shares a common 2–cycle with Δ_{ij} (or Δ_{jk}) or not.

Subcase 1a (Δ_{ij} and Δ_{ik} have a common 2–cycle, say $\{3, 4\} \in \Delta_{ij} \cap \Delta_{ik}$) Suppose $\{6, 7\} \in \Delta_{ik}$. Since $|L_k| = 2$, the labels 5 and 8 must lie in the same component of Δ_{ik} as 1 and 2. This yields option (1) of the statement. Suppose now that $\{6,7\} \notin \Delta_{ik}$; in particular, Δ_{ik} and Δ_{jk} have no common 2–cycle. Focusing on Δ_k, observe that, if $\{5,8\} \notin \Delta_{ik}$, then Δ_{ik} must contain one of $\{1,2\}$ or $\{6,7\}$, which is a contradiction to the previous sentence. Thus, we have $\{5,8\} \in \Delta_{ik}$, concluding that Δ_{ik} must relate the labels 1 and 2 to 6 and 7 somehow. This yields the configuration in option (2) of the statement.

Subcase 1b (Δ_{ik} has no common 2–cycle with either Δ_{ij} and Δ_{jk}) We will see that this case cannot occur. Here, Δ_{ik} is forced to relate 1 and 2 to labels in $\{3–8\}$. After relabeling, we can assume that $\{2,3\} \in \Delta_{ik}$. We have five remaining options for x such that $\{1,x\} \in \Delta_{ik}$. If $x = 4$, in order to have $|\Delta_k| = 2$, it must be that contains $\{7,8\} \in \Delta_{jk}$. Thus, Δ_{jk} and Δ_{ik} have a common 2–cycle, a contradiction. Similarly, we rule out $x = 5, 6, 7$. If $x = 8$, then, as Δ_{ik} does not share a 2–cycle with Δ_{jk}, it must be the case that Δ_{ik} contains either $\{4,6\}$ or $\{4,7\}$. The first possibility implies Δ_i is a single 8–cycle, while the second implies Δ_{ik} and Δ_{ij} share a 2–cycle. Both are impossibilities in this subcase.

Case 2 (Δ_j contains two 4–cycles) We can assume that $\Delta_{ij} = \{\{1,2\},\{3,4\},\{5,6\},\{7,8\}\}$ and $\Delta_{jk} = \{\{1,4\},\{2,3\},\{5,8\},\{6,7\}\}$ without loss of generality. Observe that, if Δ_i or Δ_j has one 2–cycle, then we can permute the symbols $\{i,j,k\}$ and continue as in Case 1, yielding the configurations (1) and (2) in the statement. In particular, if $\{x,y\} \in \Delta_{ik}$, then we must have $\{a,b\},\{c,d\} \in \Delta_{ik}$, where $\{x,a\},\{y,b\} \in \Delta_{ij}$ and $\{x,c\},\{y,d\} \in \Delta_{jk}$.

Subcase 2a (Δ_{ik} relates 1 and 2 to 3 and 4) By the previous paragraph, we are forced to have $\Delta_{ik} = \{\{1,3\},\{2,4\},\{5,7\},\{6,8\}\}$. Thus,

$$\Delta_j = \Delta_k = \Delta_i = \{\{1–4\},\{5–8\}\},$$

which contradicts the fact that F is connected.

Subcase 2b (Δ_{ij} does not relate 1 and 2 to 3 and 4) After relabeling, we can assume that $\{4,5\} \in \Delta_{ik}$. The fact that $|\Delta_k| = |\Delta_i| = 2$ forces $\Delta_{ik} = \{\{4,5\},\{3,6\},\{2,7\},\{1,8\}\}$. This yields configuration (3) in the statement. □

It is easy to see that (MZ–)bridge trisections for (twist) spun 2–bridge knots have configurations as in configuration (2).

Question 3.4 Are there nonstabilized $(4,2)$–bridge trisections of the other types?

A positive answer to Question 3.4 could lead to new examples of $(3,1)$–trisections which have been sought after in the literature.

Remark 3.5 The following combinatorial properties of reducing curves are direct consequences of Lemma 3.3; let ψ_1 and γ_1 be reducing curves in Δ_j and Δ_i, respectively:

- If $\{x,y\}$ are punctures enclosed by γ_1 and if one of them is also enclosed by ψ_1, then both are enclosed by ψ_1.
- Suppose ψ_1 and γ_1 both bound four punctures and that γ_1 bounds $\{x,y,z,w\}$. Then, after relabeling, ψ_1 separates $\{x,y\}$ from $\{z,w\}$.

3.1 Reducing curves

Reducing curves play a special role in trisections. In the case of $(4,2)$–bridge trisections, they restrict the pants decompositions near p^i_{ij} in $\mathcal{P}(\Sigma)$. Lemmas 3.6 and 3.7 show that, in certain circumstances, reducing curves for different links must intersect at least four times. Lemma 3.9 compares the γ_n–curves in p^i_{ij} with the ones (called ψ_n–curves, for convenience) in p^j_{ij}. Lemmas 3.10 and 3.11 imply that A–moves of the form $\gamma_1 \mapsto \psi_n$ and $\gamma_n \mapsto \psi_1$ cannot occur near p^i_{ij}. We rely heavily on theorems of [Lee 2017], governing the relationship between perturbations of a bridge position with bridge disks.

Lemma 3.6 *Suppose L_i has one component intersecting Σ exactly twice and L_j has no such component. Let γ in Σ be a reducing curve for L_i and suppose $\psi \subset \Sigma$ is either a reducing curve or cut-reducing curve for L_j. Then the following hold:*

(1) *If ψ is a reducing curve, then $|\gamma \cap \psi| \geq 4$.*

(2) *If ψ is a cut-reducing curve, and ψ and γ are disjoint, then γ lies inside a 3–punctured disk bounded by ψ.*

Proof Let γ and ψ be as in the statement and assume that they have been isotoped so as to intersect minimally. Let Q be a sphere separating the components of L_j such that $Q \cap \Sigma = \psi$. Let $L_i(1)$ and $L_i(3)$ be the 1–bridge and 3–bridge components of L_i and let L'_j and L''_j be the two components of L_j.

Since γ is a reducing curve for L_i, it is isotopic to the boundary of a regular neighborhood of an arc $\alpha \subset \Sigma$ joining the punctures $L_i(1) \cap \Sigma$. The arc α is the intersection $D \cap \Sigma$ of a disk D such that $\partial D = L_i(1)$ and the interior of D is disjoint from L_i. Observe that there is a shadow arc α' for $(\overline{B}_{ik}, \overline{T}_{ik})$ that is a copy of α.

Suppose that $\gamma \cap \psi = \varnothing$. We may, therefore, assume that D is disjoint from $Q \cap B_{ij}$.

Observe that $E_1 = D \cap B_{ij}$ is a bridge disk for an arc of T_{ij}. Let $K_j \subset B_{ij} \cup \overline{B}_{jk}$ be the link that results from isotoping this arc along E_1 and across Σ. The link K_j is isotopic to L_j, and is, therefore, an unlink of two components. One component is equal to a component of L_j. The result of ∂–reducing (B_{jk}, T_{jk}) along the c–disk $E = Q \cap B_{jk}$ is the disjoint union of two trivial tangles; call them (U_1, τ_1) and (U_2, τ_2). The result of ∂–reducing $(B_{jk}, K_j \cap B_{jk})$ along E is two tangles, one of which is either (U_1, τ_1) or (U_2, τ_2). Without loss of generality, we may assume it is (U_2, τ_2). Call the other one (U'_1, τ'_1). If (U'_1, τ'_1) is a trivial tangle, then so is $(B_{jk}, K_j \cap B_{jk})$. If ψ is a reducing curve, then τ'_1 is a single strand; it must be unknotted, as K_j is an unlink. Otherwise, ψ separates the punctures of Σ into one set with three punctures and the other with five punctures. If γ is on the side with five punctures, we have our theorem, so assume γ is on the side with three punctures. Thus, without loss of generality, (U_1, τ_1) has two strands and (U_2, τ_2) has three strands. Thus, (U'_1, τ'_1) has a single strand and, as before, we see that it is a trivial tangle. Thus, $(B_{jk}, K_j \cap B_{jk})$ is a trivial tangle and Σ is a bridge sphere for K_j.

By [Lee 2017, Theorem 1.1], there is a bridge disk E_2 for a strand of $\overline{T}_{jk}$ in $\overline{B}_{jk}$ such that the arcs α and $\beta = E_2 \cap \Sigma$ intersect in a single point. The three shadow arcs α, α' and β show that Σ is stabilized as in

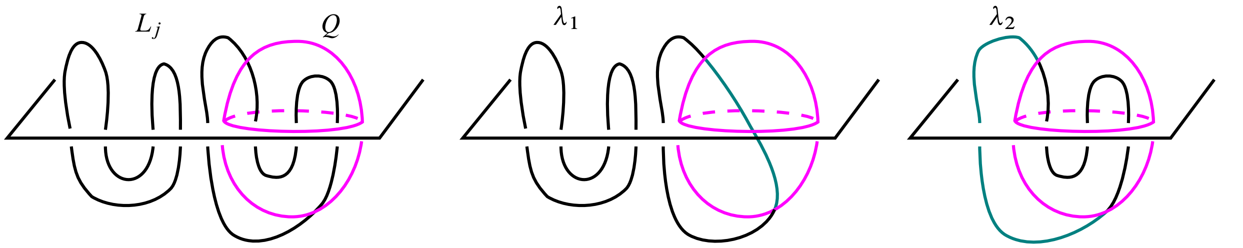

Figure 4: How to build the links λ_1 and λ_2.

Lemma 2.13. This contradicts our assumption on Σ. Thus, $|\gamma \cap \psi| > 0$ when ψ is a reducing curve and γ is on the side with five punctures if ψ is a cut-reducing curve and $|\gamma \cap \psi| = \varnothing$.

Consider the twice-punctured disk $D \subset \Sigma$ bounded by γ. If $|\psi \cap \gamma| > 0$, then $\psi \cap D$ consists of parallel arcs separating the punctures. If ψ is a reducing curve, then it bounds disks in Σ each containing an even number of punctures, in which case $|\psi \cap D|$ is even and $|\psi \cap \gamma|$ is a multiple of 4. Consequently, if ψ is a reducing curve, $|\gamma \cap \psi| \geq 4$. □

Lemma 3.7 *Suppose L_i has one component intersecting Σ exactly twice. That is, L_i is a 2–component link, where one component is in 1–bridge position and the other component is in 3–bridge position. Let $\gamma \subset \Sigma$ be a reducing curve for L_i and suppose $\psi \subset \Sigma$ is a cut-reducing curve for L_j.*

(1) *Suppose that both components of L_j are in 2–bridge position. Then $|\gamma \cap \psi| \neq 2$.*

(2) *Suppose L_j has one component in 3–bridge position. If $|\gamma \cap \psi| = 2$, then the two punctures corresponding to the 1–bridge component of L_j lie inside a 3–punctured disk bounded by ψ.*

Proof Let Q be a cut-reducing sphere such that $Q \cap \Sigma = \psi$. Cut open (S^3, L_j) along Q and glue in (3–ball, unknotted arc) pairs (B^3, α_1) and (B^3, α_2) to obtain (S^3, λ_1) and (S^3, λ_2) (See Figure 4.) In the 3–balls that we glued in we may find once-punctured disks whose boundaries coincide with the images of ψ. Attach those disks to the remnants of Σ to obtain bridge spheres Σ_1 and Σ_2 for (S^3, λ_1) and (S^3, λ_2), respectively. We can recover (S^3, L_j, Σ) by taking the connected sum of the triples $(S^3, \lambda_1, \Sigma_1)$ and $(S^3, \lambda_2, \Sigma_2)$. In particular, λ_1 and λ_2 are unlinks. Since we are decomposing a 2–component unlink L_j via a cut-reducing sphere, we can assume that λ_1 has one component and λ_2 has two components. There are a few cases to consider (see Figure 5). In all of these cases, the strategy is the

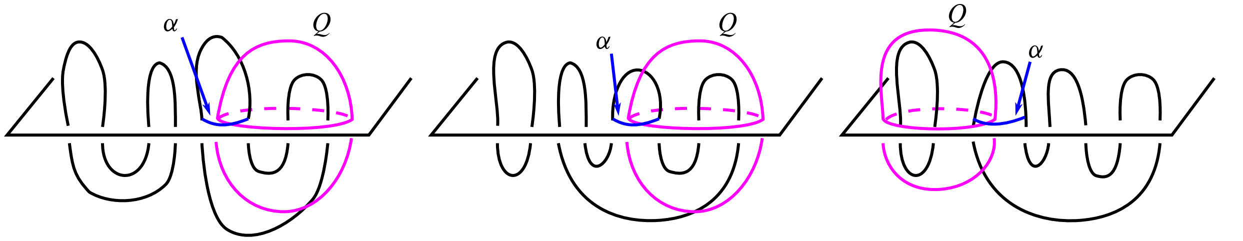

Figure 5: The link $L_j = T_{ij} \cup \overline{T}_{jk}$ in bridge position. The arc α is a shadow for arcs in T_{ij} and T_{ik}.

following. Using the same notation as in Lemma 3.6, there is a shadow arc α' for $(\bar{B}_{ik}, \bar{T}_{ik})$ that is a copy of α for (B_{ij}, T_{ij}). We then use [Lee 2017] to find a shadow in (B_{jk}, T_{jk}) intersecting α only in one endpoint (and no interior points). By Lemma 2.13, this implies that $\mathcal{T}$ is stabilized, contrary to hypothesis.

Let D as in Lemma 3.6. The intersection $D \cap \Sigma$ is a shadow α for arcs in both T_{ij} and T_{ik}. Since $|\gamma \cap \psi| = 2$, the disk $Q_0 = Q \cap B_{ij}$ intersects the disk $E = D \cap B_{ij}$ in a single arc. Thus, E persists to bridge disks E_1 for λ_1 and E_2 for λ_2.

Case 1 (each component of L_j is in 2–bridge position, ie intersects Σ four times) Suppose for the sake of contradiction that $|\gamma \cap \psi| = 2$. Only one component of L_j intersects Q. Without loss of generality, we may assume it is L'_j. Furthermore, all of the punctures $L''_j \cap \Sigma$ must lie in Σ_2 as $|L'_j \cap \Sigma| = 4$. Thus, λ_1 is an unknot intersecting Σ_1 exactly four times. Recall E_1 is a bridge disk for λ_1. Let E'_1 be another bridge disk for λ_1, on the same side of Σ_1 as E_1, but disjoint from E_1. Observe that, in the 4–punctured sphere Σ_1, the frontiers of the arcs $E_1 \cap \Sigma_1$ and $E'_1 \cap \Sigma_1$ are isotopic. Since a reduction along a bridge disk of the 2–bridge unknot is an unknot in 1–bridge position, Theorem 1.2 of [Lee 2017] tells us that each arc of $\lambda_1 \setminus \Sigma_1$ on the opposite side of Σ_1 from E_1 and E'_1 has a bridge disk intersecting both E_1 and E_2 only in one endpoint (and no interior points). Let ϵ be such a disk for the strand of $\lambda_1 \setminus \Sigma_1$ that does not contain α_1. Then ϵ is also a bridge disk for L_j and it intersects α only in one endpoint (and no interior points).

Case 2 (a component of L_j is in 1–bridge position, ie intersects Σ only twice) Suppose that $|\gamma \cap \psi| = 2$. If λ_1 is an unknot intersecting Σ_1 exactly 4 times, then we have the situation with the schematic shown in Figure 5, center. In this case, the shadow we seek for (B_{jk}, T_{jk}) is found as in Case 1. That is, there is a shadow arc α' for $(\bar{B}_{ik}, \bar{T}_{ik})$ that is a copy of a shadow arc α for (B_{ij}, T_{ij}). Since λ_1 is a 2–bridge unknot, Lee [2017] tells us that there is a shadow in $(\bar{B}_{jk}, \bar{T}_{jk})$ intersecting α only in one endpoint (and no interior points). On the other hand, if λ_1 is an unknot intersecting Σ_1 exactly 6 times, we have the second conclusion of our lemma (see Figure 5, right). □

Remark 3.8 Our proofs of Lemmas 3.6 and 3.7 above do not work for higher bridge numbers, as there is a 4–bridge position of the unknot with no complete canceling disk system (see [Lee 2017]).

For the remainder of this section, let p^i_{ij} and p^j_{ij} be pants decompositions belonging to defining pairs for $L_i = T_{ij} \cup \bar{T}_{ik}$ and $L_j = T_{kj} \cup \bar{T}_{ij}$, respectively. Denote their curves by $p^i_{ij} = \{\gamma_1, \gamma_2, \gamma_3, f_1, f_2\}$ and $p^j_{ij} = \{\psi_1, \psi_2, \psi_3, h_1, h_2\}$ as in Lemma 3.1.

Lemma 3.9 *No ψ_n–loop is equal to f_m for any $m \in \{1, 2\}$. Similarly, no γ_n–loop is equal to h_m for any $m \in \{1, 2\}$.*

Proof The second statement follows from the first by reversing the roles in the proof below. We prove the first statement.

By Lemma 3.1, ψ_2 and ψ_3 bound cut disks and f_1 and f_2 bound compressing disks, so the number of punctures they enclose is different modulo 2. Thus, $\psi_n \neq f_1, f_2$ for $n = 2, 3$.

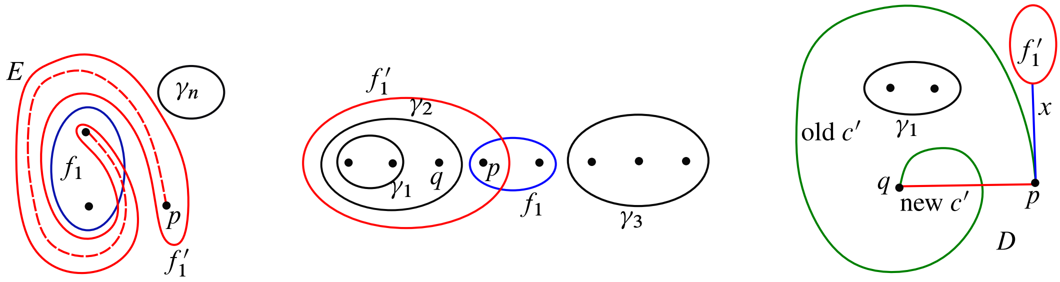

Figure 6: Various subsurfaces of Σ.

Suppose now that $\psi_1 = f_1$. In particular, γ_1 and ψ_1 are disjoint reducing curves. By Lemma 3.6, the number of punctures enclosed by γ_1 and ψ_1 must be the same. For if γ_1 bounds two punctures and ψ_1 bounds four punctures, then the two curves will intersect. But γ_1 and f_1 are distinct curves in the pants decomposition p^i_{ij}, so they cannot both enclose four punctures. We conclude that $\psi_1 = f_1$ and γ_1 enclose two punctures each. Let f_1' and f_2' be simple closed curves such that $p^i_{ik} = \{\gamma_1, \gamma_2, \gamma_3, f_1', f_2'\}$ completes a defining pair (p^i_{ij}, p^i_{ik}) for $T_{ij} \cup \overline{T}_{ik}$. Focus our attention on the A–move corresponding to $f_1 \mapsto f_1'$, which happens inside a 4–holed sphere E. The boundaries of E correspond to boundaries of small neighborhoods of punctures or to some γ_n–curves. Notice that one or two boundaries of E correspond to some γ_n–curves.

Case 1 (∂E has exactly one γ_n–loop) After a surface homeomorphism, we can draw E as in the Figure 6, left. Here, after choosing coordinates for the 4–punctured sphere, f_1 is depicted as a separating curve of slope $1/0$. The conditions $|f_1 \cap f_1'| = 2$ and $f_1' \cap \gamma_n = \varnothing$ imply that f_1' is a separating simple closed curve in E of slope $n/1$ for some $n \in \mathbb{Z}$. In other words, $f_1 = \partial\eta(c)$ and $f_1' = \partial\eta(c')$ for some properly embedded arcs c, c' in E such that c is an arc disjoint from γ_n, and $c \cap c' = \partial c \cap \partial c'$ is exactly one puncture. We pick c' so that the end disjoint from c corresponds to the puncture p on the same side of f_1 as γ_n (see Figure 6, left). Now, recall that f_1' bounds a compressing disk for T_{ik}, and so c' is a shadow for some arc in T_{ik}. Similarly, c is a shadow for arcs in both T_{ij} and T_{kj} because $f_1 = \psi_1$ is a compressing disk for both tangles. By Lemma 2.13, these three shadow arcs with one common endpoint imply that the bridge trisection is stabilized. This concludes Case 1.

Case 2 (∂E has two γ_n–loops) Both must bound cut disks. After a surface homeomorphism, the curves in p^i_{ij} can be depicted as in Figure 6, center. Observe here that f_1' must surround four punctures on each side. Let D be the 4–holed sphere inside Σ cobounded by f_1', γ_1, $\partial\eta(p)$ and $\partial\eta(q)$; see Figure 6, center and right. By construction, there exists an arc x in D with endpoints in p and E such that x is disjoint from $f_1' \cap D$. Since γ_1 and f_1' both bound compressing disks for T_{ik}, it follows that there is an arc in T_{ik} connecting p and q. Furthermore, such arc has a shadow arc c' in Σ with interior disjoint from f_1' and γ_1. Regarded as a subset of D, the arc c' connects E and γ_1. We can slide c' over γ_1 several times and choose a shadow arc c with interior disjoint from x. In particular, c intersects f_1 in one point. This, together with the fact that $f_1 = \psi_1$ bounds reducing curve for $T_{kj} \cup \overline{T}_{ij}$, implies the existence of a shadow arc c for both T_{kj} and T_{ij} with $c \cap c' = \partial c \cap \partial c' = \{p\}$. By Lemma 2.13, we conclude that $\mathcal{T}$ is stabilized. □

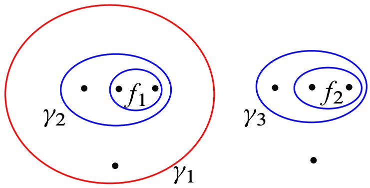

Figure 7: When the reducing curve bounds four punctures, the two cut curves lie on distinct sides.

Lemma 3.10 *If e is an edge in $\mathcal{P}(\Sigma)$ with initial endpoint at p^i_{ij}, then e does not move γ_1 to any ψ_n–loop in p^j_{ij}. Similarly, if e is an edge in $\mathcal{P}(\Sigma)$ with terminal endpoint at p^j_{ij}, then e does not move any γ_n–loop of p^i_{ij} to ψ_1.*

Proof The second statement follows from the first by interchanging the roles of γ_1 and ψ_1, and so we prove only the first statement. Suppose, to establish a contradiction, that γ_1 is moved to some ψ_n–loop by e.

First we show that e does not move γ_1 to ψ_1. Suppose γ_1 bounds a twice-punctured disk D. If e moves γ_1 to ψ_1, then $|\gamma_1 \cap \psi_1| = 2$, so $D \cap \psi_1$ consists of a single arc. It follows that the two punctures of D are on opposite sides of ψ_1, contradicting Remark 3.5. Similarly, ψ_1 does not bound a twice-punctured disk.

Consequently, if e moves γ_1 to ψ_1, then both γ_1 and ψ_1 enclose four punctures. This sets us in the third configuration of Lemma 3.3. First, observe that f_1 and f_2 must be separated by γ_1. This holds since $p^i_{ij} = \{\gamma_1, \gamma_2, \gamma_3, f_1, f_2\}$ is a pants decomposition for Σ, and only γ_1, f_1 and f_2 bound an even number of punctures. Thus, after a surface homeomorphism, we can draw Σ and p^i_{ij} as in Figure 7. We see, therefore, that if e moves γ_1 to ψ_1, then γ_1 and ψ_1 will both bound the same three (out of four) punctures, contradicting Lemma 3.3. Hence, γ_1 cannot be moved first to ψ_1.

We will now see that, due to parity constraints, if e moves γ_1, then γ_1 is moved to a curve bounding an even number of punctures. In particular, γ_1 is not moved to ψ_n for $n = 2, 3$. In order to do this, we focus on the 4–holed sphere, denoted by E, corresponding to the first A–move. The four boundary components of E are loops (or punctures), $\{\partial_1, \partial_2, \partial_3, \partial_4\}$. If γ_1 bounds four punctures, up to surface homeomorphism, then Σ can be depicted as in Figure 7 and we see that each component of ∂E is an odd curve. On the other hand, if γ_1 encloses exactly two punctures, then two components of ∂E are single punctures. The other two boundaries, say ∂_3 and ∂_4, will enclose punctures 1 and 5, 2 and 4, or 3 and 3, respectively. Notice that they cannot enclose punctures 2 and 4, since that will force the existence of a fourth curve in p^i_{ij} bounding an even number of punctures. Thus, in any case, all the components of ∂E are either a single puncture or enclose an odd number of punctures. Consequently, e moves γ_1 to a curve enclosing an even number of punctures, as desired. □

Lemma 3.11 *If e is an edge in $\mathcal{P}(\Sigma)$ with initial endpoint at p^i_{ij}, then e does not move any γ_n–loop of p^i_{ij} to ψ_1. Similarly, if e is an edge in $\mathcal{P}(\Sigma)$ with terminal endpoint at p^j_{ij} then e does not move γ_1 to any ψ_n–loop.*

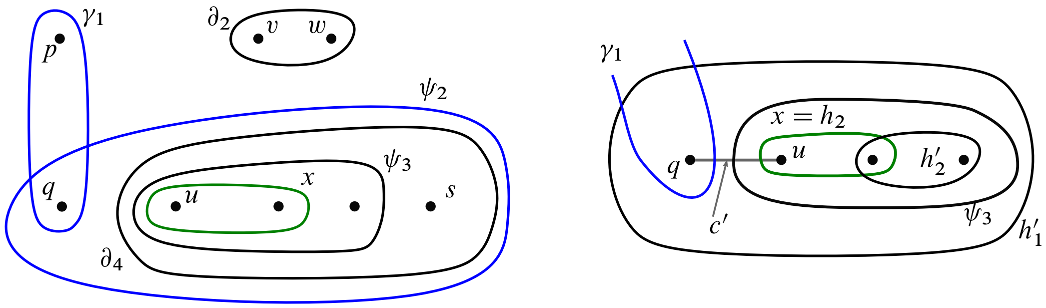

Figure 8: A close look at the A–move $\gamma_1 \mapsto \psi_2$.

Proof As we did in Lemma 3.10, it is enough to show the first statement. The case $\psi_1 \mapsto \gamma_1$ has been discussed in the proof of Lemma 3.10.

We study the case $\gamma_1 \mapsto \psi_2$. In particular, γ_1 and ψ_1 must be disjoint because the endpoint of e is p^j_{ij}. Thus, Lemma 3.6 forces both γ_1 and ψ_1 to bound two punctures each. The 4–holed sphere corresponding to e is drawn in Figure 8, left. Observe that we are forced, by Lemma 3.1, to have one cut curve inside ∂_4 and one compressing curve x. Here, the sets of curves $\{x, \partial_2, \partial_4\}$ and $\{h_1, h_2, \psi_1\}$ agree. Since ψ_1 bounds two punctures, we can assume $\partial_4 = h_1$. Moreover, Lemma 3.7(2) implies that $\psi_1 = \partial_2$, leaving us with $x = h_2$.

Focus on $h'_1 \in p^j_{jk}$. If h'_1 bounds two punctures, we can proceed as in the previous paragraph and conclude that the bridge trisection is stabilized. Thus, h'_1 must bound four punctures. Here, h'_1 bounds q and the curve ψ_3. By focusing on this disk (see Figure 8, right), we see that h'_2 must be disjoint from γ_1 because $(h_2 \cup h'_2) \cap \psi_3 = \varnothing$. This lets us find a shadow c' for T_{jk} connecting q and u, such that c' is disjoint from h'_1 and h'_2. We can slide c' over h'_1 and h'_2 in order to arrange that c' and γ_1 intersect once. Thus, the bridge trisection is stabilized by Lemma 2.13. □

3.2 Improved lower bound

We are ready to prove the lower-bound of Theorem 1.1. The main result of this Section is Theorem 3.16, which states that the Kirby–Thompson invariant of a $(4,2)$–bridge trisection of a knotted sphere in S^4 is at least 15.

As before, let S be a connected surface in S^4 with an unstabilized, irreducible $(4,2)$–bridge trisection $\mathcal{T}$. Fix $\{i, j, k\} = \{1,2,3\}$. Let (p^i_{ij}, p^i_{ik}) and (p^j_{ij}, p^j_{jk}) be defining pairs. Denote the curves in p^i_{ij} and p^j_{ij} by $p^i_{ij} = \{\gamma_1, \gamma_2, \gamma_3, f_1, f_2\}$ and $p^j_{ij} = \{\psi_1, \psi_2, \psi_3, h_1, h_2\}$ as in Lemma 3.1. We know that f_1, f_2, h_1 and h_2 bound compressing disks for T_{ij}; also, each γ_n–curve is a reducing or cut-reducing curve for L_i and each ψ_n–curve is a reducing or cut-reducing curve for L_j; in fact, γ_1 and ψ_1 are reducing curves and the others are cut-reducing curves. Recall that there are essential, simple closed curves f'_1 and f'_2 such that $p^i_{ik} = \{\gamma_1, \gamma_2, \gamma_3, f'_1, f'_2\}$ completes an efficient defining pair (p^i_{ij}, p^i_{ik}). Likewise, there are

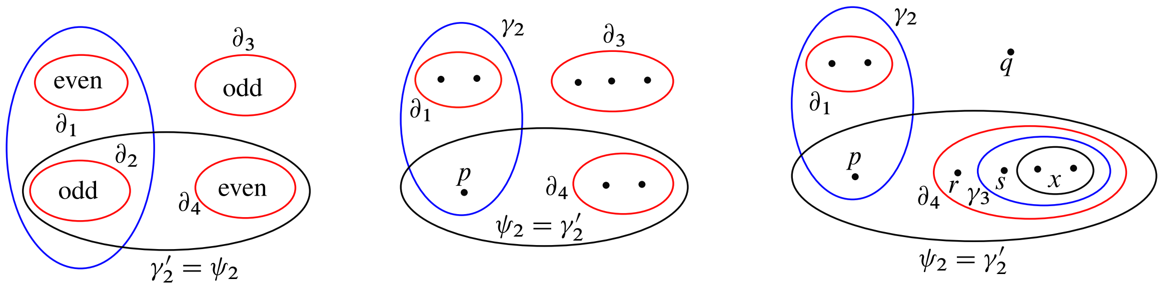

Figure 9: Two subcases, depending on the number of punctures bounded by ∂_3.

essential, simple closed curves h_1' and h_2' such that $p^j_{jk} = \{\psi_1, \psi_2, \psi_3, h_1', h_2'\}$ completes an efficient defining pair (p^j_{ij}, p^j_{jk}).

The proof of Theorem 3.16 will be broken into Propositions 3.13, 3.14 and 3.15, each of them proving that $d(p^i_{ij}, p^j_{ij}) \geq 5$ for each pair, depending on the number of punctures bounded by γ_1 and ψ_1. We begin in Proposition 3.12, showing that such distance is at least 4.

Proposition 3.12 *If $\lambda(ij)$ is a path from p^i_{ij} to p^j_{ij}. The length of $\lambda(ij)$ is at least 4. If it is equal to 4, then at least one of f_1 and f_2 is unmoved by $\lambda(ij)$.*

Proof By Lemma 3.9, no ψ_n–loop is equal to f_1 or f_2 and no γ_n–loop is equal to h_1 or h_2. Thus, if some γ_n–loop is unmoved by $\lambda(ij)$, then it is equal to some ψ_n–loop. But, by Lemma 2.10, this implies that $\mathcal{T}$ is reducible, a contradiction. Thus, $\lambda(ij)$ moves every γ_n–loop, so the length of $\lambda(ij)$ is at least 3. If it is equal to 3, then f_1 and f_2 are unmoved by $\lambda(ij)$ and, if it is equal to 4, at least one of f_1 or f_2 is unmoved by $\lambda(ij)$, as desired. Thus, we simply need to show that the length is not 3.

Assume, for a contradiction, that the length of $\lambda(ij)$ is 3. As f_1 and f_2 are unmoved, by Lemma 3.9, $\{f_1, f_2\} = \{h_1, h_2\}$. By Lemma 2.10, each of the curves $\{\gamma_1, \gamma_2, \gamma_3\}$ moves exactly once. For each $m = 1, 2, 3$, let γ_m' denote the ψ_n–loop to which γ_m is moved by $\lambda(ij)$. Lemmas 3.10 and 3.11 imply that the curves γ_1 and ψ_1 are not involved in the first and third A–moves of $\lambda(ij)$. Thus, $\gamma_1 \mapsto \psi_1$ must be the second A–move in $\lambda(ij)$. We can then assume that γ_2 moves first, $\gamma_2' = \psi_2$ and $\gamma_3' = \psi_3$.

We focus on the 4–holed sphere E where the A–move $\gamma_2 \mapsto \gamma_2'$ occurs. Denote the boundaries of E by ∂_1, ∂_2, ∂_3 and ∂_4. After a surface homeomorphism, we can draw E as in Figure 9, left, which shows the parity of punctures bounded by ∂_n. Since γ_2 is a cut disk, one of its sides contains three punctures. Thus, we may assume that ∂_2 only bounds the puncture p and ∂_1 bounds two punctures. We get two cases, depending on the number of punctures bounded by ∂_3, one or three (see Figure 9).

Case 1 (∂_3 bounds three punctures; in particular, $\partial_3 = \gamma_3$ bounds a cut disk) See Figure 9, center. By the previous paragraph, γ_3 has to be moved in third place and γ_1 in second. Since $\gamma_1 \mapsto \psi_1$ is an A–move, we know that $|\gamma_1 \cap \psi_1| = 2$. This is a contradiction, due to the following argument, also found in

Lemma 3.6. Denote by $D \subset \Sigma$ the twice-punctured disk bounded by γ_1. We have that $\psi_1 \cap D$ consists of parallel arcs separating the punctures. Since ψ_1 is a reducing curve, it bounds disks in Σ each containing an even number of punctures. Therefore, $|\psi_1 \cap D|$ is even and $|\psi_1 \cap \gamma_1|$ is a multiple of 4.

Case 2 (∂_3 bounds one puncture, named q) After a surface homeomorphism, we can draw the curves as in Figure 9, right. Recall that $\gamma_1 \mapsto \psi_1$ is the second A–move in $\lambda(ij)$. It follows that $\gamma_1 \in \{\partial_1, \partial_4, x\}$ and observe that all the possible configurations for the curve $\gamma_1' = \psi_1$ in Figure 9, right, contradict the combinatorial conditions in Remark 3.5. Thus, this case cannot occur. □

Proposition 3.13 *Suppose γ_1 bounds two punctures and ψ_1 bounds four. Then any path $\lambda(ij)$ from p_{ij}^i to p_{ij}^j must be of distance at least 5.*

Proof By Proposition 3.12, it is enough to show the distance from p_{ij}^i to p_{ij}^j is not four. By way of contradiction, let λ be a geodesic path of length four between such pants decompositions. By Lemmas 2.10 and 3.9, each γ_n–curve must move at least once. We have two cases, depending on how many curves of $\{f_1, f_2\}$ are moved.

Case 1 (λ moves one curve of $\{f_1, f_2\}$) Without loss of generality, f_1 is moved and so $f_2 = h_2$ is fixed. In this case, each of $\{\gamma_1, \gamma_2, \gamma_3, f_1\}$ is moved once to one curve among $\{\psi_1, \psi_2, \psi_3, h_1\}$. Denote by x' the image of a loop x under the path λ; ie x and x' differ by one A–move.

First observe that, since h_n and γ_1 are compressing curves for the same tangle, it must happen that, if γ_1 bounds $\{p, q\}$, then they are both on the same side of h_n. Thus, $|\gamma_1 \cap h_n| \equiv 0$ modulo 4. In particular, $\gamma_1' \neq h_n$. Similarly $\gamma_1' \neq \psi_1$. Thus, γ_1' bounds a cut disk, say $\gamma_1' = \psi_2$. In particular, $|\gamma_1 \cap \psi_2| = 2$. This is a contradiction to Lemma 3.7(1). Hence, this case cannot occur.

Case 2 (λ fixes $\{f_1, f_2\}$) We can write $f_1 = h_1$ and $f_2 = h_2$. In this case, one of $\{\gamma_1, \gamma_2, \gamma_3\}$ will move twice and the other γ_n–loops move once along λ. For the curve $\gamma_j \in \{\gamma_1, \gamma_2, \gamma_3\}$ that moves twice, denote by θ the curve γ_j'. We will also refer to θ as the *pivotal curve*.

Subcase 2a (γ_1 moves once along λ) By Lemma 3.6, $|\gamma_1 \cap \psi_1| \geq 4$, so γ_1' must bound a cut disk, say $\gamma_1' = \psi_2$. In particular, $|\gamma_1 \cap \psi_2| = 2$. This is impossible since it contradicts Lemma 3.7(1).

Subcase 2b (γ_1 moves twice along λ) We will first see that $\gamma_n' \neq \psi_1$ for any n. In particular, $\theta' = \psi_1$ and the following property holds: at each vertex of λ, there are at most three pairwise disjoint curves bounding an even number of punctures.

By Lemma 3.6, $\gamma_1' \neq \psi_1$. Suppose, without loss of generality, that $\gamma_2' = \psi_1$. The 4–holed sphere corresponding to the A–move $\gamma_2 \mapsto \psi_1$ has one boundary component bounding one puncture, r, and boundary loops ∂_1, ∂_3 and ∂_4 bounding two, two and three punctures, respectively (see Figure 10, right). Here, there are four pairwise disjoint curves bounding an even number of punctures: $\{\psi_1, \partial_1, \partial_3, x\}$. Since $\gamma_1 \cap \psi_1 \neq \varnothing$ by Lemma 3.6, we know that $\{f_1, f_2, \theta\} = \{\partial_1, \partial_3, x\}$. If $\partial_1 = \theta$, then γ_1 will

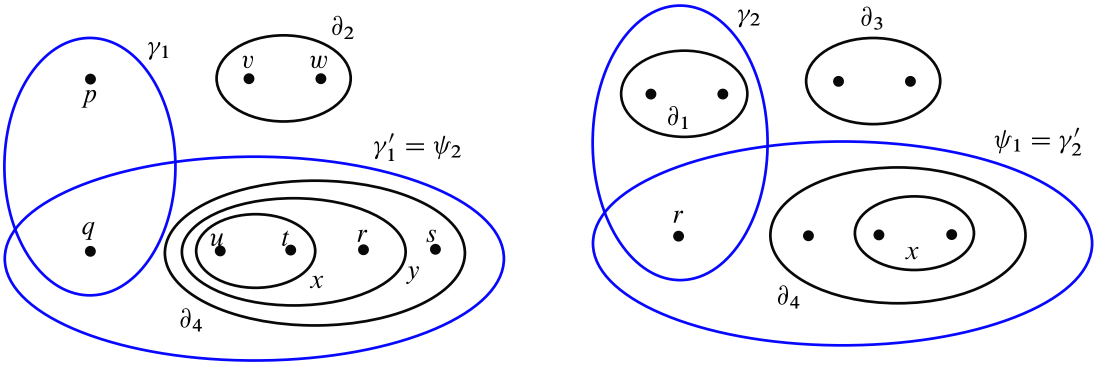

Figure 10: How the curves in Σ look for specific A–moves.

bound r and one of the two punctures bounded by ∂_1. This is impossible since such punctures are on distinct sides of ψ_1. Hence, $\partial_1 = f_1 = h_1$.

Observe that the two punctures bounded by γ_1 must be separated by $\theta = \gamma_1'$; if not, then $|\gamma_1 \cap \theta| \equiv 0$ modulo 4, which makes impossible the A–move $\gamma_1 \mapsto \theta$. We use this to see that, if $\partial_3 = \theta$, then γ_1 would bound one puncture inside ∂_3 with one puncture inside ∂_4. These points are in distinct sides of ψ_1 (see Figure 10, right), which is a contradiction to Remark 3.5. Hence, $x = \theta$, $\partial_3 = f_2$ and $\partial_1 = f_1$. Notice that all the incoming A–moves will occur in the side of ψ_1 containing ∂_4. This forces p_{ij}^j to have at least four curves bounding an even number of punctures, a contradiction to Lemma 3.1. Thus, we conclude that $\gamma_n' \neq \psi_1$, as desired.

By the above, the γ_n cut curves move once along λ to ψ_n cut curves. Without loss of generality, $\gamma_n' = \psi_n$ for $n = 2, 3$. We will assume that $\gamma_3 \mapsto \psi_3$ is not the last A–move in λ; if not, we can relabel the γ_n–curves. We will focus on the 4–holed sphere corresponding to the A–move $\gamma_3 \mapsto \gamma_3'$ (see Figure 11, left). We have two cases, depending on the number of punctures bounded by ∂_2 and ∂_3.

Subcase 2b(i) (both ∂_2 and ∂_3 bound one puncture each) We adopt the notation in Figure 11, center. In this case, we already have three pairwise disjoint curves bounding an even number of punctures, $\{\partial_1, \partial_4, x\}$, so there is a curve y bounding x and one puncture u (see Figure 11, center). Recall that h_n bounds two punctures and $f_n = h_n$ is fixed by λ. This implies that $\partial_1 = f_1$, $x = f_2$ and $\partial_4 \in \{\theta, \psi_1\}$.

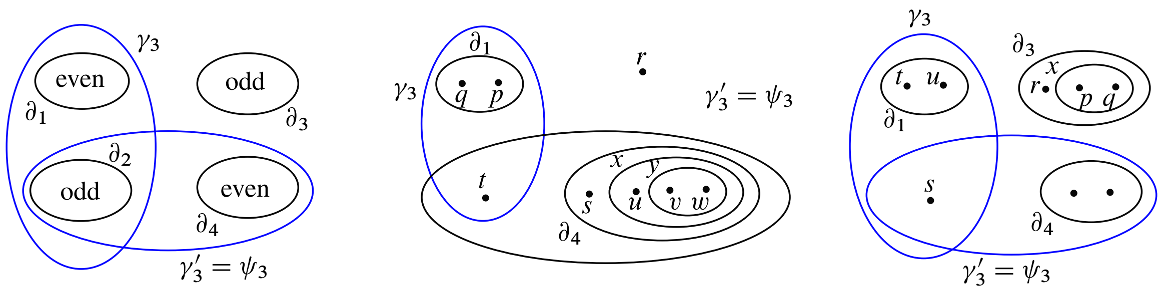

Figure 11: The three possibilities occurring in Subcase 2b.

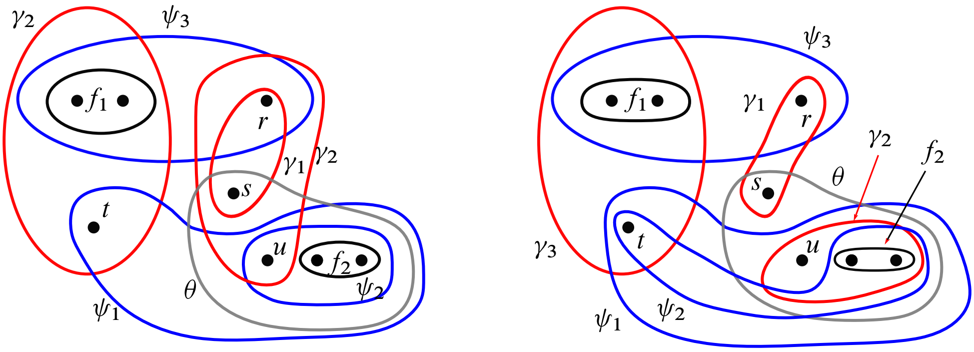

Figure 12: Two paths.

Now, since $\gamma_3 \mapsto \psi_3$ is not the last A–move in λ, there are two possible curves which may move next, y and θ.

Suppose first that y moves before θ does. (The curve θ may or may not move). Then y' must be a cut disk and we get $y' = \psi_2$ and $y = \gamma_2$. Using the notation of Figure 11, center, since γ_1 bounds two punctures and is disjoint from γ_2 and γ_3, we obtain that γ_1 bounds $\{r, s\}$. But ∂_4 separates such punctures, so the only option is $\partial_4 = \theta$. Now, the fact that y' bounds a cut disk implies that it bounds the two punctures inside x and s. The next move $\theta \mapsto \psi_1$ is forced to separate r and s, contradicting Remark 3.5.

It remains to study what happens when ∂_4 moves before y. (The curve y may or may not move). Here, $\partial_4 = \theta$. Focusing on Figure 11, center, we observe that $\psi_1 = \theta'$ bounds the two punctures inside $x = f_2$, together with t and u. By Remark 3.5, γ_1 bounds either $\{r, s\}$ or $\{t, u\}$. The latter is impossible since γ_3 is disjoint from γ_1 and γ_3 separates such punctures. Thus, γ_1 bounds $\{r, s\}$. Since ψ_3 separates r and s, the A–move $\gamma_1 \mapsto \theta$ must appear in λ before $\gamma_3 \mapsto \psi_3$. Moreover, the move $\gamma_2 \mapsto \psi_2 = y$ cannot happen between $\gamma_1 \mapsto \theta$ and $\gamma_3 \mapsto \psi_3$. This claim holds because, if γ_2 moves between γ_1 and γ_3, it would force γ_2 to bound the two punctures inside $x = f_2$ together with s, which implies the contradiction $\gamma_1 \cap \gamma_2 \neq \varnothing$. We are left with two possibilities, depending on the order of the curves moving: $(\gamma_2, \gamma_1, \gamma_3, \theta)$ or $(\gamma_1, \gamma_3, \theta, \gamma_2)$. Figure 12 showcases the two possible paths and which punctures are bounded by each curve.

We focus on the subpath of λ corresponding to the consecutive A–moves $\gamma_1 \mapsto \theta$ followed by $\gamma_3 \mapsto \psi_3$. The second A–move occurs inside a 4–holed sphere with boundaries associated to t, r, f_1 and θ (see Figure 13, left). The fact that γ_1 and γ_3 are disjoint implies that the condition $|\gamma_3 \cap \psi_3| = 2$ is equivalent to $|\gamma_1 \cap \psi_3| = 2$. One can see this claim by noticing that the curves γ_3 and $\partial\eta(\gamma_1 \cup \theta)$ are isotopic in the 4–holed sphere. The condition $|\gamma_1 \cap \psi_3| = 2$ contradicts Lemma 3.7. In other words, Subcase 2b(i) is impossible.

Subcase 2b(ii) (only one of $\{\partial_2, \partial_3\}$ bounds one puncture) Without loss of generality, ∂_2 bounds one puncture and ∂_3 three. This forces the setup in Figure 11, right. The curves along the path λ bounding an even number of punctures are γ_1, ψ_1, $f_1 = h_1$, $f_2 = h_2$ and (possibly) θ. But we have seen that $\theta' = \psi_1$

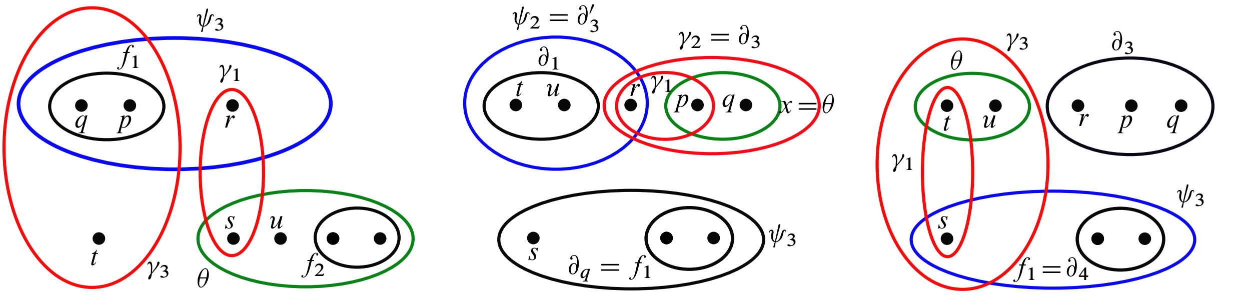

Figure 13: Curves interacting in the consecutive A–moves $\gamma_1 \mapsto \theta$, $\gamma_n \mapsto \psi_n$ for a fixed n.

and $\gamma_1' = \theta$. This implies that $\partial_4 \notin \{\gamma_1, \theta, \psi_1\}$ since all the A–moves starting at ∂_4 will be forced to end at curves bounding two punctures. Thus, we may assume that $\partial_4 = f_1$. Since no curve at this moment bounds four punctures, there should be another A–move after $\gamma_3 \mapsto \psi_3$. Using the notation in Figure 11, right, the curves that might move are $\{\partial_1, \partial_3, x\}$.

Suppose that ∂_3 moves first, then $\partial_3 = \gamma_2$ and $\partial_3' = \psi_2$. Since ψ_2 bounds three punctures, ∂_3' must enclose ∂_1 and the puncture r together. Since ψ_1 separates the cut curves ψ_2 and ψ_3 (Figure 7), it follows that $\partial_1 = f_2$ and ψ_1 separates p and q. Thus, from Remark 3.5, we must have $x = \theta$. Without loss of generality, γ_1 encloses r and p (see Figure 11, right). We now focus on the consecutive A–moves $\gamma_1 \mapsto \theta$ and $\partial_3 = \gamma_2 \mapsto \psi_2$. Observe that $\gamma_2 \mapsto \psi_2$ occurs in a 4–holed sphere with boundaries corresponding to $\psi_3, r, \partial_1 = f_2$ and θ. This local setup in depicted in Figure 13, center. In here, the conditions $\gamma_1 \cap \gamma_2 = \varnothing$ and $|\gamma_2 \cap \psi_2| = 2$ force $|\gamma_1 \cap \psi_2| = 2$. This contradicts Lemma 3.7.

If x moves before ∂_1 and ∂_3, then $\partial_1 = f_2$. In particular, $x = \gamma_1$ and ∂_3 must move so that $\theta' = \psi_1$ can bound four punctures. We can then redefine x to be $\gamma_1' = \theta$ and proceed as if ∂_3 moved first (paragraph above). We get then a contradiction.

The last case to check is when ∂_1 moves before ∂_3 and x. In particular, $x = f_2$ and $\partial_1 \in \{\gamma_1, \theta\}$.

First we see that, if $\partial_1 = \gamma_1$, then ∂_3 will have to move between $\gamma_1 \mapsto \theta$ and $\theta \mapsto \psi_1$. This is true because, if ∂_3 didn't move immediately after, then $(\gamma_1')' = \psi_1$ would separate t and u, contradicting Remark 3.5. In particular, $\partial_3 = \gamma_2$ must move between γ_1 and θ. Moreover, the A–move $\gamma_2 \mapsto \psi_1$ occurs in a 4–holed sphere with boundaries corresponding to θ, $\partial_1 = f_2$ and two boundaries bounding one puncture each. If we switch the labels and redefine γ_2 to be γ_3, we get the situation of Subcase 2b(i). We can then obtain a contradiction.

Therefore, we must have $\partial_1 = \theta$. Since γ_1 is disjoint from γ_3, using the notation in Figure 11, right, we can assume that γ_1 bounds t and s. We obtain the subpath of λ, depicted in Figure 13, right, given by the consecutive A–moves $\gamma_1 \mapsto \theta$ and $\gamma_3 \mapsto \psi_3$. Observe that $\gamma_3 \mapsto \psi_3$ occurs in a 4–holed sphere with boundaries corresponding to s, ∂_3, f_1 and θ. In here, the conditions $\gamma_3 \cap \gamma_1 = \varnothing$ and $|\gamma_3 \cap \psi_3| = 2$ force $|\gamma_1 \cap \psi_3| = 2$, contradicting Lemma 3.7. Hence, Subcase 2b(ii) cannot occur. We have exhausted all the possibilities, thus concluding the proof of the proposition. □

Proposition 3.14 *Suppose that both γ_1 and ψ_1 bound two punctures each. Then any path $\lambda(ij)$ from p^i_{ij} to p^j_{ij} must be of distance at least* 5.

Proof This proof follows the same path as Proposition 3.13. By Proposition 3.12, it is enough to show the distance from p^i_{ij} to p^j_{ij} is not four. By way of contradiction, let λ be a geodesic path of length four between such pants decompositions. By Lemmas 2.10 and 3.9, each γ_n–curve must move at least once. We have two cases, depending on how many curves of $\{f_1, f_2\}$ are moved.

Case 1 (λ moves one curve of $\{f_1, f_2\}$) Without loss of generality, assume $f_2 = h_2$ is fixed. Observe that, since ψ_1 and γ_1 bound two punctures and the curves ψ_1, γ_1, h_1 and f_1 are compressing curves for the same tangle, we obtain that $\gamma_1' \neq h_1, \psi_1$ and $\psi_1 \neq f_1'$. Thus, we can assume that $\gamma_1' = \psi_2$ and $\gamma_2' = \psi_1$. By Lemmas 3.10 and 3.11, the A–moves $\gamma_1 \mapsto \psi_2$ and $\gamma_2 \mapsto \psi_1$ cannot be first nor last in λ.

Subcase 1(a) ($\gamma_1 \mapsto \psi_2$ is second) In particular, $\gamma_2 \mapsto \psi_1$ is third, and there are at most three curves bounding an even number of punctures after the second A–move: $\{f_1, h_1, f_2 = h_2\}$. We focus our attention on the 4–holed sphere corresponding to $\gamma_1 \mapsto \psi_2$. By the previous sentence, we are forced to have an arrangement of curves as in Figure 10, left (compare with Figure 14). In particular, $\{x, \partial_2, \partial_4\} = \{f_1, h_1, f_2\}$ and $y = \gamma_2$. Since ψ_1 is the next curve to appear, ψ_1 must bound $\{r, s\}$. This is already a contradiction since Lemma 3.7(2) implies that ψ_1 bounds two of the three punctures $\{p, v, w\}$. This subcase is impossible.

Subcase 1(b) ($\gamma_1 \mapsto \psi_2$ is third and $\gamma_2 \mapsto \psi_1$ is second in λ) Recall that the only curves bounding an even number of punctures are $\{\gamma_1, \psi_1, f_1, h_1, f_2 = h_2\}$. We need to decide which of the A–moves $\gamma_3 \mapsto h_1$ and $f_1 \mapsto \psi_3$ is first. For us to decide, focus on the 4–holed sphere corresponding to the A–move $\gamma_1 \mapsto \psi_2$. Counting γ_1, there were four or five pairwise disjoint curves bounding an even number of punctures before γ_1 moved (see Figure 14). But every A–move in λ interchanges cut and compressing curves, so the number of even curves after the second A–move will be three or five. Thus, γ_3 moves first, f_1 last, and the curves look like in Figure 14, right. Lemma 3.7(2) implies that $\partial_2 = \psi_1$. Since $\gamma_2 \mapsto \psi_1$ occurs in second place, we can assume that γ_2 bounds $\{p, q, v\}$.

We will focus on ∂_4. First observe that, if $\partial_4 = f_2 = h_2$, then the A–moves in distinct sides of ∂_4 commute. This would let us contradict Lemma 3.11, since we could make $\gamma_2 \mapsto \psi_1$ the first A–move. Suppose now $\partial_4 = f_1$. Since f_1 is the last curve to move, we can assume that $f_1' = \psi_3$ bounds $\{q, u, t\}$. Moreover, because $|\gamma_1 \cap \psi_2| = |\partial_4 \cap \psi_3| = 2$ and ψ_3 is disjoint from x, z and ψ_2, we can see that γ_1 and ψ_3 must intersect in two points. Now, we know that $x = h_a$ for some $a \in \{1, 2\}$. We can use the dual curve $h_a' \in p^j_{jk}$ to find a tuple (c, c') of destabilization shadows as in Lemma 2.13. Thus, $\partial_4 = h_1$ is the remaining option.

If $\partial_4 = h_1$, then we can assume that γ_3 bounds $\{r, s, w\}$ because $\gamma_3 \mapsto h_1$ is the first A–move in λ. Recall that γ_2 bounds $\{p, q, v\}$. By thinking in the 4–holed sphere with boundaries γ_3, ∂_2, z and x, the conditions $|\partial_4 \cap \gamma_3| = |\partial_2 \cap \gamma_2| = 2$ and $\partial_4 \cap \partial_2 = \varnothing$ imply that γ_2 intersects $\partial_2 = \psi_1$ in two points.

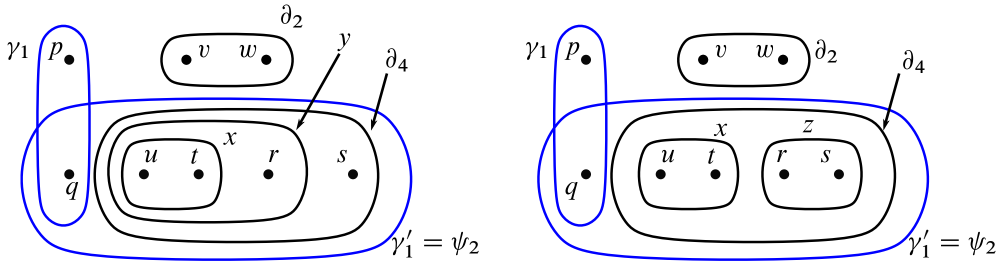

Figure 14: When γ_1 and ψ_2 differ by one A–move, there are either three (left) or four (right) curves disjoint from γ_1 bounding an even number of punctures.

Now, we know that $z = f_a$ for some $a \in \{1,2\}$. We can use the dual curve $f_a' \in p^i_{ik}$ to find a pair of shadows (c, c') as in Lemma 2.13. We have concluded Case 1.

Case 2 (λ fixes $\{f_1, f_2\}$) In this case, one γ_n–curve moves twice and the rest exactly once. We write $f_a = h_a$ and denote by θ the pivotal curve. There are two subcases, depending on how many times γ_1 moves.

Subcase 2a (γ_1 moves once along λ) Recall that γ_1, ψ_1, $f_1 = h_1$ and $f_2 = h_2$ bound compressing disks in T_{ij} and γ_1 and ψ_1 bound two punctures. Thus, $|\gamma_n \cap \alpha|$ and $|\psi_1 \cap \alpha|$ are both divisible by four for all $\alpha \in \{\gamma_1, \psi_1, f_1 = h_1, f_2 = h_2\}$. This implies that γ_1' must bound a cut disk, say $\gamma_1' = \psi_2$. Lemmas 3.10 and 3.11 force γ_1 to move second or third in λ. We can represent the curves in the 4–holed sphere corresponding to $\gamma_1 \mapsto \psi_2$ like in Figure 14. Observe that, before the A–move of γ_1, there are either four or five pairwise disjoint curves bounding an even number of punctures.

We first study ∂_2 in Figure 14. Since ∂_2 bounds two punctures, we have $\partial_2 \in \{f_1 = h_1, f_2 = h_2, \psi_1, \theta\}$. Notice that ∂_2 cannot be θ. If that were the case, θ' would be forced to bound an even number of punctures, say $\{p, v\}$, and $\theta' = \psi_1$. In particular, ψ_1 would separate p and q, which contradicts Remark 3.5. Lemma 3.7 implies that ψ_1 bounds two punctures from $\{p, v, w\}$; thus, $\partial_2 = \psi_1$.

Subcase 2a(i) (there are five even curves) We use the notation in Figure 14, right. We have that the sets of curves $\{x, z, \partial_4\}$ and $\{\theta, f_1, f_2\}$ agree. In particular, by Lemma 3.11, $\gamma_2 \mapsto \psi_1$ must be the second A–move and so $\gamma_3 \mapsto \theta$ is the first one. If ∂_4 is equal to some f_a, then the curves θ and ψ_1 will lie in different sides of ∂_4. We could then permute their corresponding A–moves and obtain $\gamma_2 \mapsto \psi_1$ first in λ, contradicting Lemma 3.11. Thus, we conclude that $\partial_4 = \theta$, $x = f_1 = h_1$ and $z = f_2 = h_2$. Here, we can assume that γ_2 bounds $\{p, q, v\}$ and γ_3 bounds $\{w, r, s\}$. Now, by looking at the 4–holed sphere bounded by γ_2, x, z and $\partial\eta(w)$, we can see that $\gamma_3 \cap \gamma_2 = \varnothing$ and $|\gamma_2 \cap \psi_1| = 2$ imply that $|\gamma_3 \cap \psi_1| = 2$. Then, inside the component of $\Sigma \setminus \gamma_3$ containing w, we can use f_2' to find a tuple of shadows (c, c') satisfying the conditions of Lemma 2.13. Thus, this subcase cannot occur.

Subcase 2a(ii) (before the A–move $\gamma_1 \mapsto \psi_2$, there are four curves bounding an even number of punctures) We can draw the curves in Σ as in Figure 14, left. Since $\partial_2 = \psi_1$, we can assume $x = f_1 = h_1$ and

$\partial_4 = f_2 = h_2$. Now, since ∂_4 is fixed along λ, the A–moves occurring in different sides of ∂_4 can be permuted. Thus, we can assume that $y = \gamma_2$ and so $\gamma_2' \in \{\psi_3, \theta\}$.

Suppose now that $\gamma_2' = \psi_3$. Since $\psi_3 = \gamma_2'$ is forced to bound $\{u, t, s\}$, we can assume that $h_1' \in p_{jk}^j$ bounds $\{t, s\}$. In particular, T_{jk} connects the punctures $\{t, s\}$. On the other hand, since γ_1, f_1 and f_2 bound disks in T_{ij}, we know that T_{ij} connects p, u and r with q, t and s, respectively. The fact that $L_j = T_{ij} \cup \overline{T}_{jk}$ is a 2–component link and ψ_1 is a reducing curve implies that T_{jk} connects the punctures $\{u, r\}$ with $\{p, q\}$. Since γ_2 bounds a cut disk for T_{ik}, we have that T_{ik} must connect r with either u or t. In any case, the fact that $L_k = T_{ik} \cup \overline{T}_{jk}$ is a 2–component link forces v and w to be connected by T_{ik}. Since ψ_1 bounds a compressing disk in both T_{ij} and T_{jk}, we obtain that v and w are connected by the three tangles. This implies the surface S is disconnected, a contradiction.

We are left with $\gamma_2' = \theta$, which forces $\gamma_3' = \psi_1 = \partial_2$ and $\theta' = \psi_3$. Since $\partial_4 = f_2 = h_2$ is fixed along λ, the A–moves on distinct sides of ∂_4 commute. Thus, we can take λ such that $\gamma_3 \mapsto \psi_1$ is the first A–move. This contradicts the conclusion of Lemma 3.11. Hence, this subcase cannot occur.

Subcase 2b (γ_1 moves twice along λ) By symmetry and Subcase 2a, it is enough to study the case that $\theta' = \psi_1$. We write $\gamma_2' = \psi_2$ and $\gamma_3' = \psi_3$. First observe that, since γ_1 and ψ_1 bound disjoint sets of two punctures (Lemma 3.3), the A–moves $\gamma_1 \mapsto \theta$ and $\theta \mapsto \psi_1$ cannot be consecutive in λ. In other words, at least one cut curve must move between those moves. We are left with two options (up to symmetry) for the order of the A–moves along λ: $(\gamma_1, \gamma_3, \gamma_2, \theta)$ and $(\gamma_1, \gamma_3, \theta, \gamma_2)$. We focus on the *second* A–move $\gamma_3 \mapsto \psi_3$. It occurs inside a 4–holed sphere depicted in Figure 11, left.

Subcase 2b(i) (both ∂_2 and ∂_3 bound one puncture each) We use the notation in Figure 11, center, and observe that $y = \gamma_2$. Since $\gamma_1 \mapsto \theta$ and $\gamma_3 \mapsto \psi_3$ are the first two A–moves in λ, we know that the sets of curves $\{x, \partial_1, \partial_4\}$ and $\{\theta, f_1 = h_1, f_2 = h_2\}$ agree. Suppose $\partial_4 = \theta$; then γ_1 is forced to bound $\{r, s\}$. In the 4–holed sphere with boundaries ∂_1, y, $\partial\eta(r)$ and $\partial\eta(t)$, the conditions $\gamma_3 \cap \gamma_1 = \varnothing$ and $|\gamma_3 \cap \psi_3| = 2$ force $|\gamma_1 \cap \psi_3| = 2$. Lemma 3.7 implies that ψ_1 bounds two punctures from $\{q, p, r\}$. This is impossible since $\partial_1 \in \{h_1, h_2\}$ is disjoint from ψ_1. Thus, we conclude that $\partial_4 = f_2 = h_2$.

Suppose now that $\partial_1 = f_1 = h_1$. Since the A–move $\gamma_1 \mapsto \theta$ occurs inside γ_2, we can reuse Figure 11, center, and assume that $x = \gamma_1$ and $\theta = \gamma_1'$ bounds $\{u, v\}$. After $\gamma_1' \mapsto \theta$, the next A–move has to be $\gamma_2 \mapsto \psi_2$. Here, ψ_2 and $\psi_1 = \theta'$ will bound $\{s, u, v\}$ and $\{s, u\}$, respectively. Focus on the 4–holed sphere E corresponding to the A–move $\theta \mapsto \psi_1$. Notice that E has boundaries corresponding to s, u, v and ψ_2. Since $|\gamma_2 \cap \psi_2| = 2$, the intersection $\gamma_2 \cap E$ is an arc with both endpoints on ψ_2 that separates s from $\{u, v\}$ (see Figure 15, left). Since $\theta \cap \gamma_2 = \varnothing$, the condition $|\psi_1 \cap \theta| = 2$ forces ψ_1 to intersect γ_2 in two points.

To end, we study the curve f_2'. For reference, we use the curves and notation from Figure 15, right. We now look at the 4–holed sphere E' with boundaries γ_3, γ_2, $\partial\eta(r)$ and $\partial\eta(t)$. Since $|\psi_1 \cap \gamma_2| = 2$, $\psi_1 \cap E'$

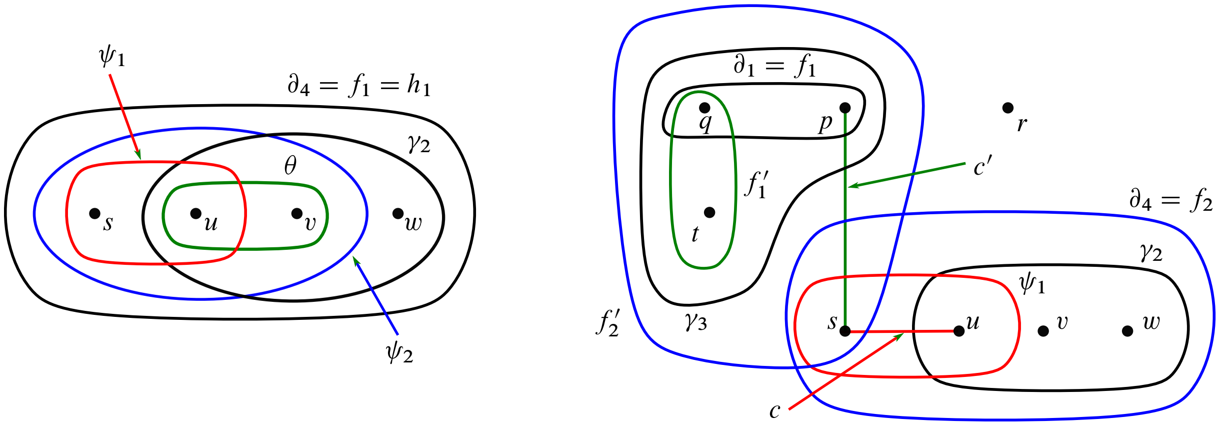

Figure 15: A close-up of some curves in Subcase 2b(i).

is an arc with both endpoints on γ_2 that separates s from r and γ_3. Thus, the conditions $\gamma_2 \cap f_2 = \varnothing$ and $|f_2' \cap f_2| = 2$ imply that ψ_1 intersects f_2' in two points. If f_2' bounds two punctures, we can use the condition $|\psi_1 \cap f_2'| = 2$ to find a tuple (c, c') of shadows satisfying the condition of Lemma 2.13, contradicting the fact that $\mathcal{T}$ is not stabilized.

On the other hand, if f_2' bounds four punctures, we will also find a tuple (c, c') as in Lemma 2.13. The rest of this paragraph explains how to do this. First observe that f_2' will bound γ_3 and s. Since f_1' lies inside γ_3 and intersects f_1 in two points, we can assume that f_1' bounds $\{q, t\}$. Both f_1' and f_2' bound compressing disks in T_{ik}, so we can find a shadow c' of an arc of T_{ik} connecting $\{p, s\}$ such that c' is disjoint from f_1' and f_2'. Inside the disk component of $\Sigma \setminus f_2'$ that contains γ_3, the condition $|\psi_1 \cap f_2'| = 2$ implies that ψ_1 is an arc with both endpoints in f_2' that separates s from f_1' and p. We can slide c' over f_1' and f_2' and assume that $|c' \cap \psi_1| = 1$. The last condition allows us to pick an arc c in Σ connecting $\{s, u\}$ such that $\partial\eta(c) = \psi_1$ and $c \cap c' = \partial c \cap \partial c' = \{s\}$. Notice that c' is a shadow for arcs in T_{ij} and T_{jk}. Hence, the tuple (c, c') satisfies the conditions of Lemma 2.13. This is a contradiction.

We are left with $x = f_1 = h_1$ and $\theta = \partial_1$. Since $\partial_4 = f_2 = h_2$ is fixed along λ, A–moves on distinct sides of ∂_4 commute. Moreover, this setup is equivalent to the previous case ($\partial_1 = f_1 = h_1$): one can reflect Figure 11, center, with respect to ∂_4 and the roles of the curves on each side will reverse. Therefore, this case is impossible.

Subcase 2b(ii) (∂_2 and ∂_3 enclose one and three punctures, respectively) We use the notation of Figure 11, right. One of the curves $\{\partial_1, x, \partial_4\}$ is equal to θ. Observe that, if ρ is a curve such that $\rho \mapsto \partial_4$ is an A–move immediately before $\partial_3 \mapsto \psi_3$, then ρ bounds four punctures. In particular, $\rho \neq \gamma_1$. Thus, $\partial_4 \neq \theta$ and so $\partial_4 = f_1 = h_1$. Suppose now that $x = \theta$. We can assume that γ_1 bounds $\{r, p\}$. By Lemma 3.3, the two punctures bounded by ψ_1 must be distinct than $\{r, p\}$. Here, notice that $\gamma_2' = \psi_2$ is forced to bound $\{t, u, r\}$ and $\theta = x$ must move after γ_2. Moreover, θ' has to bound four punctures, contradicting $\theta' = \psi_1$. Hence, $x = f_2 = h_2$ and $\partial_1 = \theta$.

We are left to discard the case $\partial_1 = \theta$. Since x is fixed along λ and ψ_3 won't move, we see that two out of the three punctures $\{t, u, r\}$ will be bounded by ψ_1. We can assume that γ_1 bounds $\{t, s\}$. By looking at the 4–holed sphere with boundaries γ_3, $\partial\eta(t)$, $\partial\eta(s)$ and $\partial\eta(u)$, we see that the conditions $\psi_3 \cap \theta = \varnothing$, $|\gamma_3 \cap \psi_3| = 2$ and $|\theta \cap \gamma_1| = 2$ imply $|\gamma_1 \cap \psi_3| = 2$. Now, inside the disk of $\Sigma \setminus \psi_3$ containing $\partial_4 = h_1$, one can see that $h_1' \in p_{jk}^j$ must intersect γ_1 in two points. Thus, there is a shadow c' for an arc of T_{jk} with $\partial\eta(c') = h_1'$ and $|c' \cap \gamma_1| = 1$. By taking $c \subset \Sigma$ with $\partial\eta(c) = \gamma_1$ and $c \cap c' = \partial c \cap \partial c' = \{s\}$, we obtain a tuple (c, c') like in Lemma 2.13. Hence, $\overline{T}$ is a stabilization. This finishes the analysis in Case 2. □

Proposition 3.15 *Suppose that both γ_1 and ψ_1 bound four punctures each. Then any path $\lambda(ij)$ from p_{ij}^i to p_{ij}^j must have length at least 5.*

Proof By Proposition 3.12, it is enough to show the distance from p_{ij}^i to p_{ij}^j is not four. By way of contradiction, let λ be a geodesic path of length four between such pants decompositions. By Lemmas 2.10 and 3.9, each γ_n–curve must move at least once.

Notice that, if two pants decompositions differ by the A–move $\gamma_1 \mapsto \psi_1$, then each boundary loop of the 4–holed sphere corresponding to this A–move must bound two punctures. This is true because the curves γ_1 and ψ_1 bound compressing disks for the same tangle T_{ij}. In particular, we know that there are at most five curves bounding an even number of punctures that are involved in λ, say $\{\gamma_1, \psi_1, h_1, f_1, f_2 = h_2\}$ or $\{\gamma_1, \psi_1, \theta, f_1 = h_1, f_2 = h_2\}$, where θ is the pivotal curve. Thus, it cannot contain the edge $\gamma_1 \mapsto \psi_1$.

Case 1 (λ moves one curve of $\{f_1, f_2\}$) Say $f_2 = h_2$ is fixed. Notice that f_1 bounds two punctures since γ_1 bounds four. Also, f_1 and ψ_1 bound compressing disks for the same tangle T_{ij}, so the two punctures bounded by f_1 must be on the same side of ψ_1. Thus, $|f_1 \cap \psi_1|$ is divisible by four. This implies that $f_1' \neq \psi_1$. Similarly, $\gamma_1' \neq h_1$. We can then assume that $\gamma_2 \mapsto \psi_1$ and $\gamma_1 \mapsto \psi_2$ are A–moves along λ. Moreover, by Lemmas 3.10 and 3.11, such A–moves must be in either second or third place. But $\gamma_1 \cap \psi_1 \neq \varnothing$, so $\gamma_1 \mapsto \psi_2$ must be second and $\gamma_2 \mapsto \psi_1$ third.

We now study the 4–holed sphere where the A–move $\gamma_1 \mapsto \psi_2$ occurs. We can assume that the curves look like in Figure 16, left. In particular, $\partial_1 = \gamma_2$ and the sets of curves $\{x, \partial_3, \partial_4\}$ and $\{f_1, h_1, f_2 = h_2\}$ agree. Since the next A–move is $\gamma_2 \mapsto \psi_1$, we obtain that ψ_1 bounds x and ∂_3. From Figure 7, we know that the reducing curve γ_1 (resp. ψ_1) must separate f_1 and f_2 (resp. h_1 and h_2). This implies that $x = f_1$, $\partial_3 = f_2 = h_2$ and $\partial_4 = h_1$.

To end this case, we will analyze the possible shadows of the tangles T_{ij}, T_{ik} and T_{jk}. Figure 17, left, contains the labels of the punctures and the new shadows described throughout this paragraph. Notice that h_1' bounds two punctures, say $\{s, t\}$. By looking at the 4–holed sphere with boundaries ψ_2, s, t and u, we can conclude that h_1' must intersect γ_1 in two points. In particular, there is a shadow c of an arc in T_{jk} connecting $\{s, t\}$ such that $\partial\eta(c) = h_1'$. Since $|h_1' \cap \gamma_1| = 2$, we see that c intersects γ_1 once. Now focus on the disk component of $\Sigma \setminus \gamma_1$ containing γ_2. Since f_1 and γ_1 bound compressing disks for T_{ij}, there

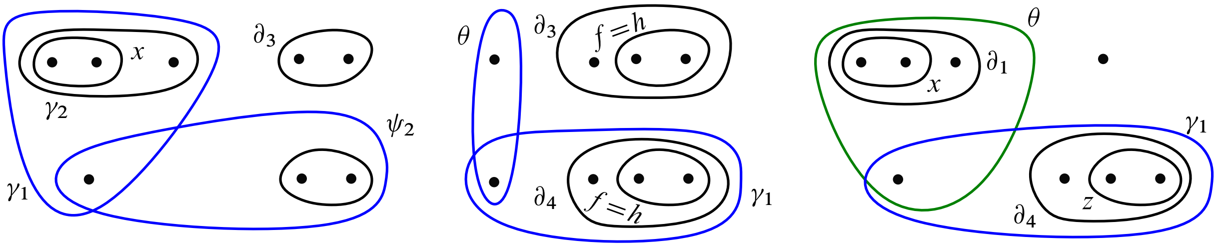

Figure 16: Curve arrangements for specific A–moves.

are shadows a_1 and a_2 for arcs of T_{ij} that are disjoint from $f_1 \cup \gamma_1$ satisfying $\partial\eta(a_1) = f_1$ and such that a_2 connects $\{r, s\}$. Notice that f_1 and h_1' are on opposite sides of γ_2, so $a_1 \cap c = \varnothing$. Moreover, we can think of a_2 as an arc in a 4–holed sphere with boundaries $x = f_1$, γ_1, $\partial\eta(s)$ and $\partial\eta(r)$, where a_2 and c are arcs connecting $\{r, s\}$ and $\{s, \gamma_1\}$, respectively. We can slide a_2 over f_1 and γ_1 and still obtain a shadow arc for T_{ij}. Thus, we can slide a_2 inside this 4–holed sphere and choose a_2 to have interior disjoint from c, ie $a_2 \cap c = \partial a_2 \cap \partial c = \{s\}$. To end, we observe that f_1' bounds two punctures and is inside γ_2. We can assume that f_1' bounds $\{q, r\}$. Since f_1' and γ_1 bound compressing disks for T_{ik}, we can find shadows b_1 and b_2 for arcs in T_{ik} disjoint from f_1' and γ_1 satisfying $\partial\eta(b_1) = f_1'$ and that b_2 connects $\{p, s\}$. Since $|f_1 \cap f_1'| = 2$, we can choose b_1 so that $b_1 \cap a_1 = \partial b_1 \cap \partial a_1 = \{q\}$. As we did with a_2, we can slide b_2 over f_1' and γ_1 until b_2 has interior disjoint from c. We can further slide a_2 and b_2 and see that $a_1 \cup b_1 \cup a_2 \cup b_2$ can be chosen to be a simple closed curve (ignoring the punctures). The tuple $(\alpha, \beta, \gamma) = (\{a_1, a_2\}, \{b_1, b_2\}, c)$ satisfies the conditions of Lemma 2.12, concluding that $\mathcal{T}$ is a stabilization.

Case 2 (λ fixes $\{f_1, f_2\}$) Suppose first that $\gamma_1' = \psi_2$. From Figure 16, left, we note that, before the A–move $\gamma_1 \mapsto \psi_2$, there are four curves bounding an even number of punctures, say $\{\gamma_1, x, \partial_3, \partial_4\}$. Since $\gamma_1 \cap \psi_1 \neq \varnothing$, $f_1 = h_1$ and $f_2 = h_2$, the mentioned A–move is impossible. Thus, $\gamma_1' \neq \psi_2, \psi_3$. Similarly, we see that $\psi_1 \neq \gamma_2', \gamma_3'$. We have already established that γ_1' cannot be equal to ψ_1. Thus, the only option is $\gamma_1' = \theta$ and $\theta' = \psi_1$. In particular, $\gamma_2' = \psi_2$ and $\gamma_3' = \psi_3$.

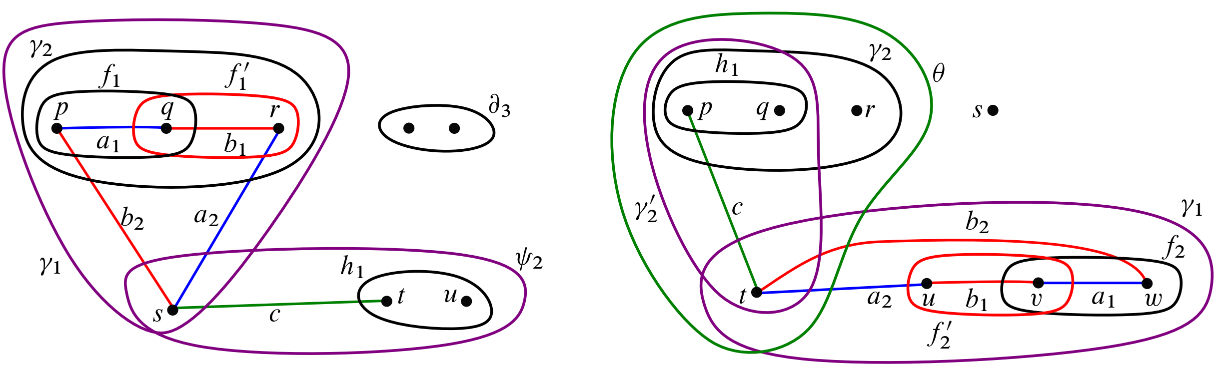

Figure 17: Shadows.

We now study how many punctures θ bounds. First note that θ cannot bound three punctures. This holds because, before the A–move $\gamma_1 \mapsto \theta$, there would be three other curves bounding an even number of punctures (set $\psi_2 = \theta$ in Figure 16, left). This is impossible since only five curves can bound an even number of punctures, $\{\gamma_1, \psi_1, \theta, f_1 = h_1, f_2 = h_2\}$, and ψ_1 and θ intersect γ_1. If θ bounds two punctures, the curves in Σ will look as in Figure 16, center. If θ moves immediately after γ_1, then three out of the four punctures bounded by γ_1 will be on the same side of $\psi_1 = \theta$, contradicting Remark 3.5. If a cut curve moves before θ, we can assume it is $\gamma_2 = \partial_3$. Since γ_2' bounds a cut disk, it is forced to bound θ together with one other puncture. This implies that $\theta' = \psi_1$ will bound two punctures, a contradiction.

The only remaining option is if θ bounds four punctures. Since only $\{f_1 = h_1, f_2 = h_2\}$ are curves disjoint from γ_1 that bound an even number of punctures, we can draw the curves in Σ before the A–move $\gamma_1 \mapsto \theta$ as in Figure 16, right. Moreover, we can assume $x = f_1 = h_1$ and $z = f_2 = h_2$. Recall that $f_1 = h_1$ and $f_2 = h_2$ lie in different sides of both γ_1 and ψ_1 (see Figure 7). Thus, by Remark 3.5, since γ_1 bounds t, u and $f_2 = h_2$, we conclude that ψ_1 bounds r, s and $f_2 = h_2$. But ∂_1 bounds h_1 and r which are on distinct sides of ψ_1. Thus, $\partial_1 \notin \{\psi_2, \psi_3\}$. Similarly, $\partial_4 \notin \{\psi_2, \psi_3\}$. We can then assume that $\partial_1 = \gamma_2$ and $\partial_4 = \gamma_3$. Since θ separates $\{r, t\}$ from $\{s, u\}$, we see that γ_2 moves before θ. Also, $\gamma_2' = \psi_2$ will bound t and $f_1 = h_1$. The A–move $\gamma_2 \mapsto \psi_2$ occurs inside a 4–holed sphere with boundaries $f_1 = h_1$, $\partial\eta(r)$, $\partial\eta(t)$ and θ. Here, γ_1 is an arc with both endpoints in θ that separates t from $f_1 = h_1$ and $\partial\eta(r)$. Thus, since $\gamma_2 \cap \gamma_1 = \varnothing$, the condition $|\gamma_2 \cap \psi_2| = 2$ is equivalent to $|\psi_2 \cap \gamma_1| = 2$. Now, inside ψ_2, we can assume that the curve h_1' bounds $\{p, t\}$. Again, the condition $|\gamma_1 \cap \psi_2| = 2$ implies that $|h_1' \cap \gamma_1| = 2$. In particular, there is a shadow c of an arc in T_{jk} connecting $\{p, t\}$ such that $\partial\eta(c) = h_1'$. The condition $|h_1' \cap \gamma_1| = 2$ implies that c intersects γ_1 once. Focus on the disk component of $\Sigma \setminus \gamma_1$. Here, the arc c is an arc with endpoints in γ_1 and $\{t\}$. We can repeat the argument in Case 1 and find shadows a_1 and a_2 for arcs in T_{ji} and b_1 and b_2 for arcs in T_{ik} as in Figure 17, right. One of the key properties we obtain is that $a_1 \cup b_1 \cup a_2 \cup b_2$ is a simple closed curve (ignoring the punctures) disjoint from γ_1 that intersects c in the puncture $\{t\}$. Then the tuple $(\alpha, \beta, \gamma) = (\{a_1, a_2\}, \{b_1, b_2\}, c)$ satisfies the conditions of Lemma 2.12, concluding that $\mathcal{T}$ is a stabilization. □

Theorem 3.16 *Let $\mathcal{T}$ be a $(4,2)$–bridge trisection for a knotted connected surface S in S^4. Then*

$$L(\mathcal{T}) \geq 15.$$

Proof We first observe that $\mathcal{T}$ is unstabilized and irreducible. If $\mathcal{T}$ is stabilized, then $\mathfrak{b}(S) \leq 3$. By [Meier and Zupan 2017, Theorem 1.8], S is unknotted, contradicting our assumption. If $\mathcal{T}$ is reducible, then, by [Blair et al. 2022], it is either the distant sum or connected sum of two other trisections. In the former case, this would imply that F is disconnected, a contradiction. In the latter case, the two other trisections have bridge numbers $b_1, b_2 \geq 2$ and $b_1 + b_2 - 1 = 4$. Thus, $b_1, b_2 \leq 3$. Again by [Meier and Zupan 2017, Theorem 1.8], this means both surfaces being trisected are unknotted and so S, being their connected sum, is also unknotted.

Let (p^i_{ij}, p^i_{ik}) for $\{i,j,k\} = \{1,2,3\}$ be choices of efficient pairs such that

$$\mathscr{L}(\mathcal{T}) = d(p^1_{12}, p^2_{12}) + d(p^1_{31}, p^3_{31}) + d(p^2_{23}, p^3_{23}).$$

By Lemma 3.3, the reducing curves of p^i_{ij} and p^j_{ij} either

(1) bound two and four punctures each,

(2) both bound two punctures, or

(3) both bound four punctures.

Propositions 3.13, 3.14 and 3.15 state that $d(p^i_{ij}, p^j_{ij}) \geq 5$ in each case. Hence, $\mathscr{L}(\mathcal{T})$ is at least $5+5+5 = 15$. □

Corollary 3.17 *Let $K \neq U$ be a 2–bridge knot in S^3. The spun knot $\mathcal{S}(K)$ satisfies*

$$\mathscr{L}(S(K)) \geq 15.$$

Proof From Theorem 2.5, if $\mathcal{T}$ is a minimal (b, c_1, c_2, c_2)–bridge trisection of $S(K)$, then $b = 4$ and $c_1 = c_2 = c_3 = 2$. By Theorem 3.16, $\mathscr{L}(\mathcal{T}) \geq 15$. □

4 Upper bounds for $\mathscr{L}$–invariant of spun knots

The goal of this section is to build an upper bound for $\mathscr{L}(S(K))$ in terms of the bridge splitting for K. Throughout this section, K will denote a knot in b–bridge position, $K = T^+_K \cup T^-_K$, and $\mathcal{T}_{\mathrm{MZ}}$ is the $(3b-2, b)$–bridge trisection for the spin of K from Section 2.3.

Example 4.1 ($\mathscr{L}$–invariant of spun trefoil) When K is the trefoil knot, the triplane diagrams from Section 2.3 give us the links $L_i = T_{ij} \cup \overline{T}_{ik}$ in Figure 18. In the same figure, we find particular choices for efficient defining pairs (p^i_{ij}, p^i_{ik}) for the link L_i which have bounded distance $d(p^i_{ij}, p^j_{ij}) \leq 5$ (Figure 19). Thus, $\mathscr{L}(S(K)) \leq 15$. One can observe that such paths resemble a particular path in the 4–punctured sphere (Figure 19, right). The main idea of this section is to formalize the resemblance and use it to build a general upper bound in Theorem 4.3.

Recall that a link $L = L_+ \cup L_-$ in bridge position is perturbed if there is a pair of bridge disks (one on each side) intersecting once in one puncture. This notion is equivalent to the existence of a pair of compressing disks (one per tangle) with boundaries f_+ and f_- such that

(1) each $f_\pm$ bounds two punctures,

(2) f_+ and f_- bound one common puncture, and

(3) $|f_+ \cap f_-| = 2$.

Observe that, if $c_\pm$ is the shadow for the bridge disk in the perturbation, then $f_\pm = \partial\eta c_\pm$.

A *perturbation system* is a finite collection of perturbation pairs $\{(c^n_-, c^n_+)\}_n$ with pairwise disjoint interiors such that $\bigcup_n (c^n_+ \cup c^n_-)$ contains no circles in the bridge surface. In other words, it is a collection of

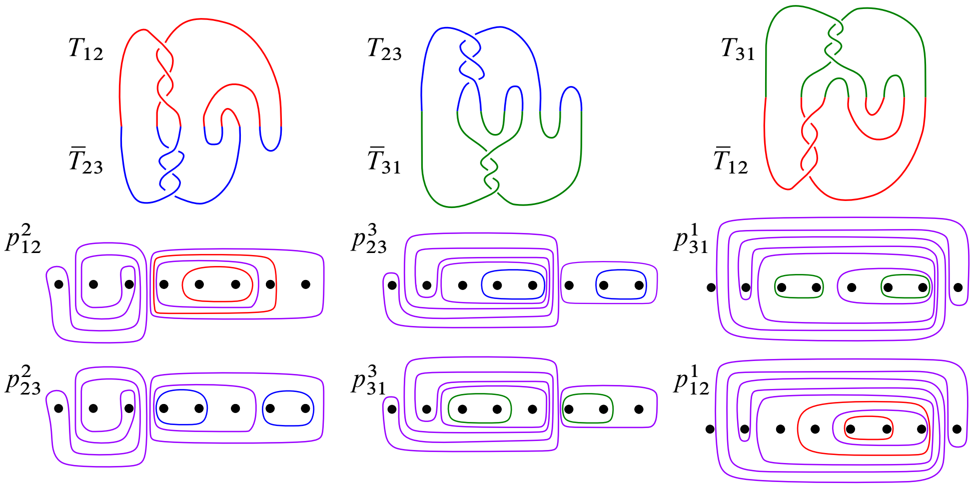

Figure 18: Bridge positions and efficient defining pairs for the links $L_i = T_{ij} \cup \overline{T}_{ik}$.

perturbations that can be undone at the same time. Figure 20 contains examples of perturbation systems. As submanifolds of the bridge surface, the loops $\partial\eta\big(\bigcup_n (c^n_+ \cup c^n_-)\big)$ bound c–disks for L in both sides. We will refer to these curves (resp. spheres) in the bridge surface (resp. S^3) as *sensor* curves (resp. spheres) since they allow us to think of L as a link with lower bridge number.

For the b–bridge links in Figure 20, the perturbation systems will determine two simplicial maps between pants complexes $\mathcal{P}(\Sigma_{2b}) \to \mathcal{P}(\Sigma_{6b-4})$. The main idea of the upper bound for $\mathcal{L}(\mathcal{T}_{\mathrm{MZ}})$ is to induce paths in $\mathcal{P}(\Sigma_{6b-4})$ using information from the splitting of the knot K.

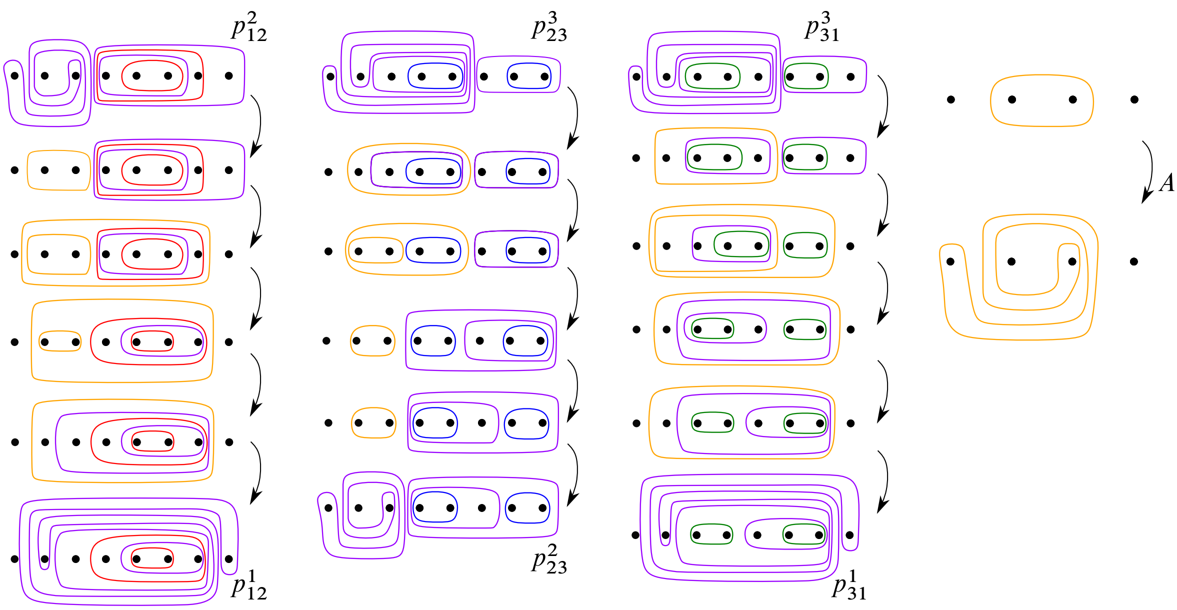

Figure 19: Three paths of length five between p^i_{ij} and p^j_{ij}.

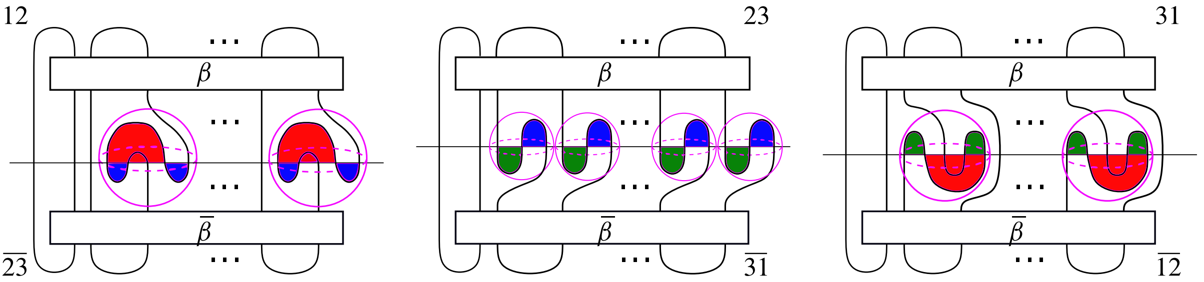

Figure 20: Bridge presentations for the links $L_{\varepsilon,\overline{\delta}} = T_\varepsilon \cup \overline{T}_\delta$.

Fix $(\varepsilon, \delta, \rho)$ to be a cyclic permutation of the labels $(12, 31, 23)$. Focus on the link $L_{\varepsilon,\overline{\delta}} = T_\varepsilon \cup \overline{T_{\overline{\delta}}}$ and the perturbation system in Figure 20. Observe that $L_{\varepsilon,\overline{\delta}}$ is in $(6b-4)$–bridge position. Moreover, if we shrink the sensor spheres to a point by collapsing the 3–ball containing the perturbation disks, we obtain a link isotopic to $L_{\varepsilon,\overline{\delta}}$ in b–bridge position. At the level of the bridge surfaces, this collapsing induces a continuous map between the punctured spheres $g_{\varepsilon,\overline{\delta}}\colon \Sigma_{6b-4} \to \Sigma_{2b}$. Given a pants decomposition $p \in \mathcal{P}(\Sigma_{2b})$, define sets of curves $G^{\pm}_{\varepsilon,\overline{\delta}}(p) = g^{-1}_{\varepsilon,\overline{\delta}}(p) \cup \mu^{\pm}_{\varepsilon,\overline{\delta}} \cup \phi_{\varepsilon,\overline{\delta}}$, where $\mu^{\pm}_{\varepsilon,\overline{\delta}}$ and $\phi_{\varepsilon,\overline{\delta}}$ are collections of curves described in Figure 21. By construction, both $G^{\pm}_{\varepsilon,\overline{\delta}}(p)$ are pants decompositions of Σ_{6b-4}. Furthermore, the functions $\{G^{\pm}_{\varepsilon,\overline{\delta}}\}_{(\varepsilon,\overline{\delta})}$ satisfy several properties described in the following lemma:

Lemma 4.2 *Let $(\varepsilon, \delta, \rho)$ be a cyclic permutation of $(12, 23, 31)$ and let p_0 and p_1 be any two pants decompositions of Σ_{2b}. The following holds:*

(1) *$G^{\pm}_{\varepsilon,\overline{\delta}}\colon \mathcal{P}(\Sigma_{2b}) \to \mathcal{P}(\Sigma_{6b-4})$ is a 1–simplicial map; in other words, if $\lambda \subset \mathcal{P}(\Sigma_{2b})$ is a path from p_0 to p_1, then $G^{\pm}_{\varepsilon,\overline{\delta}}(\lambda)$ is a path connecting $G^{\pm}_{\varepsilon,\overline{\delta}}(p_0)$ and $G^{\pm}_{\varepsilon,\overline{\delta}}(p_1)$.*

(2) *If every loop in p_0 bounds a c–disk in T^+_K, then the tuple $(G^+_{\varepsilon,\overline{\delta}}(p_0), G^-_{\varepsilon,\overline{\delta}}(p_0))$ is an efficient pair for the link $T_\varepsilon \cup \overline{T}_\delta$.*

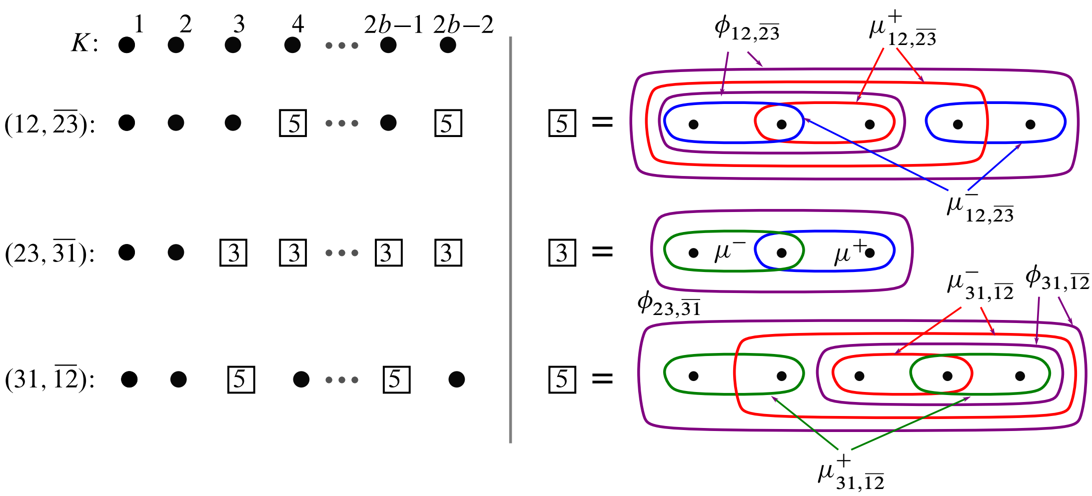

Figure 21: Curves that complete $G^{\pm}_{\varepsilon,\overline{\delta}}$; we removed the indices in the right side of the figure.

(3) *If every loop in p_1 bounds a compressing disk for T_K^-, then the distance in $\mathcal{P}(\Sigma_{6b-4})$ between $G^+_{\varepsilon,\overline{\delta}}(p_1)$ and $G^-_{\rho,\overline{\varepsilon}}(p_1)$ is $2(b-1)$.*

Proof Part (1) follows from the definition of $G^{\pm}_{\varepsilon,\overline{\delta}}$. In order to prove (2), we first observe that $G^+_{\varepsilon,\overline{\delta}}(p_0)$ and $G^-_{\varepsilon,\overline{\delta}}(p_0)$ are pants decompositions with loops bounding c–disks in T_ε and T_δ, respectively. The loops in $\mu^{\pm}_{\varepsilon,\overline{\delta}}$ arise from perturbation pairs and the ones in $\phi_{\varepsilon,\overline{\delta}}$ from sensor loops (see Figure 20). Thus, they bound c–disks. The extra assumption in p_0 implies that $g^{-1}_{\varepsilon,\overline{\delta}}(p_0)$ also bounds c–disks. Next, one can see from Figure 21 that the loops in $\mu^+_{\varepsilon,\overline{\delta}}$ and $\mu^-_{\varepsilon,\overline{\delta}}$ can be paired so that they intersect in two points and are disjoint from the rest. Thus, there is a path in $\mathcal{P}(\Sigma_{6g-4})$ of length $2b-2$. Lemma 2.7 concludes that $(G^+_{\varepsilon,\overline{\delta}}(p_0), G^-_{\varepsilon,\overline{\delta}}(p_0))$ is an efficient pair.

We will now discuss (3). Label the punctures in the bridge sphere for K as in the left side of Figure 21. In particular, since every loop in p_1 bonds a compressing disk for T_K^-, we get that the pairs of punctures $\{2n-1, 2n\}$ belong to the same component of $\Sigma_{2b} \setminus p_1$ for $n = 1, \dots, b$. We denote this collection of loops by $B \subset p_1$. After an isotopy of the bridge surface for K, which changes the surface by a homeomorphism fixing the punctures, we can assume that the loops in B look as in Figure 22. Observe that this isotopy of K does not affect the class of bridge trisection $\mathcal{T}_{\mathrm{MZ}}$; more precisely, it changes the triplane diagrams by a mutual braid transposition by a pure braid [Meier and Zupan 2017, Section 2.5]. We can then consider the pants decompositions $G^+_{\varepsilon,\overline{\delta}}(p_1)$ and $G^-_{\rho,\overline{\varepsilon}}(p_1)$ and see that the loops in $g^{-1}_{\varepsilon,\overline{\delta}}(p_1)$ and $g^{-1}_{\rho,\overline{\varepsilon}}(p_1)$ agree. We also observe that the loops in $\mu^+_{\varepsilon,\overline{\delta}}$ and $\mu^-_{\rho,\overline{\delta}}$ are the same since their corresponding bridge disks agree (see Figure 20). To end, we can perform the length two path of A–moves described by Figure 22 near each loop in B ($b-1$ times), and find a path in $\mathcal{P}(\Sigma_{6b-4})$ replacing the loops $\phi_{\varepsilon,\overline{\delta}}$ by the

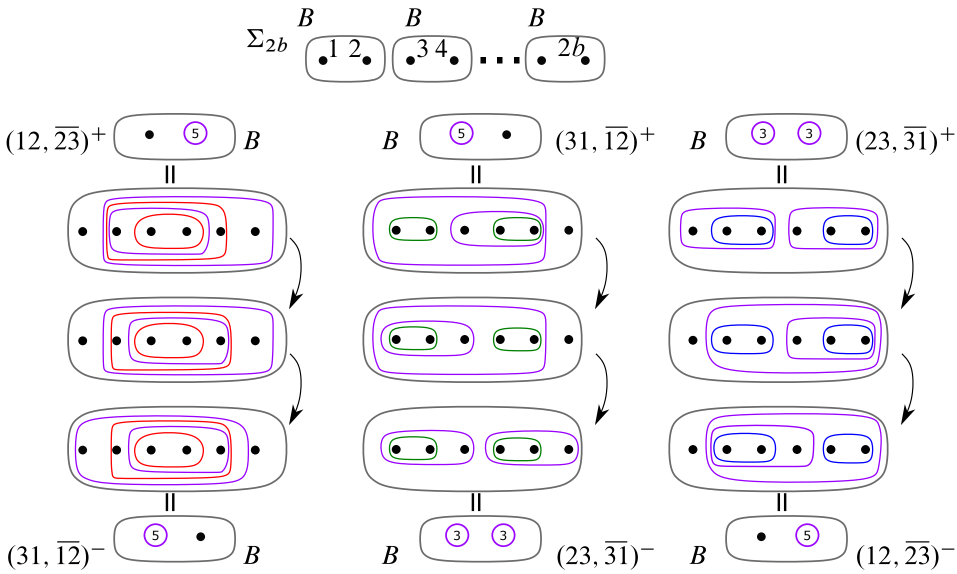

Figure 22: If we perform the sequence of A–moves inside each component of B, we obtain paths of length $2(b-1)$ connecting $\phi_{\varepsilon,\overline{\delta}} \mapsto \phi_{\rho,\overline{\varepsilon}}$.

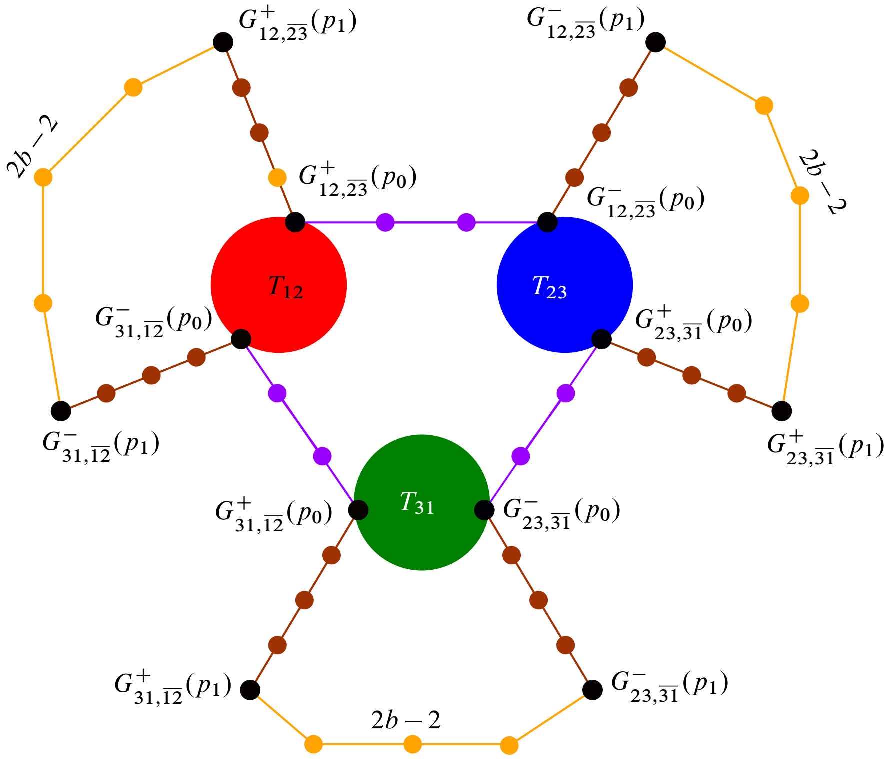

Figure 23: Upper bound for $\mathcal{L}(\mathcal{T})$.

loops $\phi_{\rho,\overline{\varepsilon}}$. Thus, the distance in $\mathcal{P}(\Sigma_{6b-4})$ between $G^+_{\varepsilon,\overline{\delta}}(p_1)$ and $G^-_{\rho,\overline{\varepsilon}}(p_1)$ is at most $2(b-1)$. Since the sets of curves $\phi_{\varepsilon,\overline{\delta}}$ and $\phi_{\rho,\overline{\varepsilon}}$ have no common curve, we conclude that this path is minimal length. □

Motivated by Lemma 4.2, for a trivial N–tangle T, we define $\mathcal{P}_{\mathrm{comp}}(T)$ and $\mathcal{P}_c(T)$ to be the sets of pants decompositions $p \in \mathcal{P}(\Sigma_{2N})$ such that all loops in p bound compressing disks and c–disks, respectively. The upper bound in the following theorem can be summarized in Figure 23.

Theorem 4.3 *Let $K = T^+_K \cup T^-_K$ be a knot in b–bridge position and let $\mathcal{T}_{\mathrm{MZ}}$ be the $(3b-2, b)$–bridge trisection for the spun 2–knot $S(K) \subset S^4$. Let $d \geq 0$ be the distance in $\mathcal{P}(\Sigma_{2b})$ between the sets $\mathcal{P}_c(T^+_K)$ and $\mathcal{P}_{\mathrm{comp}}(T^-_K)$. Then*

$$\mathcal{L}(\mathcal{T}_{\mathrm{MZ}}) \leq 6(d + b - 1).$$

Proof Let $p_0 \in \mathcal{P}_c(T^+_K)$ and $p_1 \in \mathcal{P}_{\mathrm{comp}}(T^-_K)$ be pants decompositions realizing the distance d, and let λ be a geodesic path in $\mathcal{P}(\Sigma_{2b})$ connecting them. In particular, p_0 and p_1 satisfy the conclusions of Lemma 4.2 for any cyclic permutation $(\varepsilon, \delta, \rho)$ of $(12, 23, 31)$. Now, consider the loop in $\mathcal{P}(\Sigma_{6b-4})$ described in Figure 23. By Lemma 4.2, this loop satisfies the conditions in the definition of $\mathcal{L}(\mathcal{T}_{\mathrm{MZ}})$. Since each $G^{\pm}_{\varepsilon,\overline{\delta}}(\lambda)$ is a path of length d, we can conclude the desired inequality. □

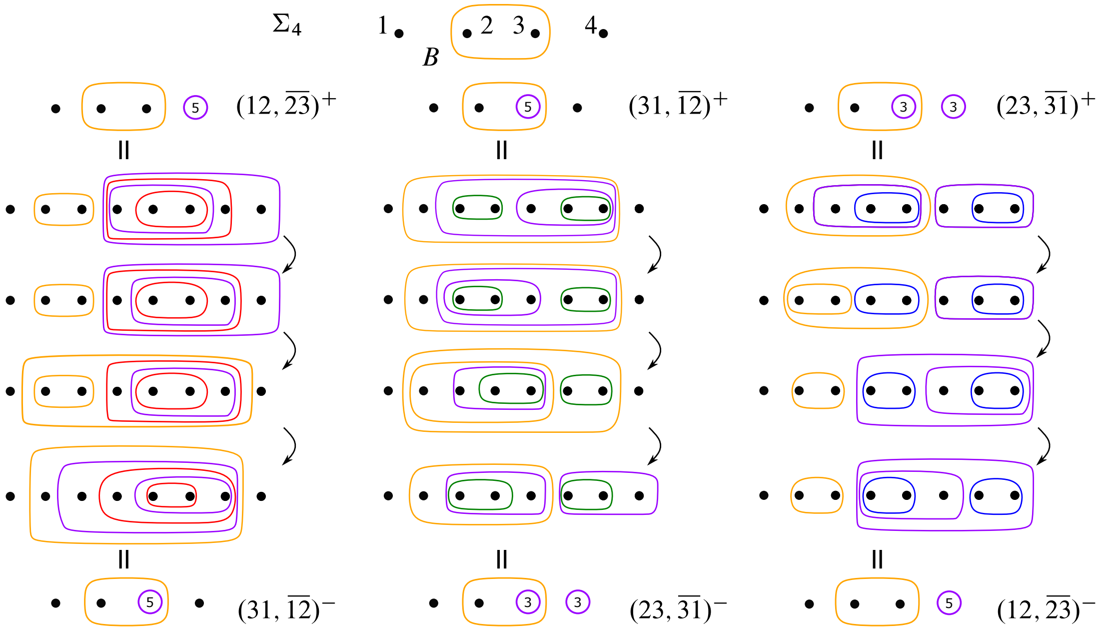

Figure 24: Paths of length three between $G^+_{\varepsilon,\overline{\delta}}(p_1)$ and $G^-_{\rho,\overline{\varepsilon}}(p_1)$.

Remark 4.4 From Figure 23, we can derive a more general upper bound for $\mathcal{L}(\mathcal{T}_{\mathrm{MZ}})$ as follows: if $p_0, p_1 \in \mathcal{P}(\Sigma_{2b})$ are pants decompositions with $p_0 \in \mathcal{P}_c(T_K^+)$, then

$$\mathcal{L}(\mathcal{T}_{\mathrm{MZ}}) \le 6d(p_0,p_1) + d(G^-_{12,\overline{31}}(p_1), G^+_{31,\overline{23}}(p_1)) + d(G^-_{31,\overline{23}}(p_1), G^+_{23,\overline{12}}(p_1)) + d(G^-_{23,\overline{12}}(p_1), G^+_{12,\overline{31}}(p_1)).$$

The following corollary studies the distance between $G^+_{\varepsilon,\overline{\delta}}(p_1)$ and $G^-_{\rho,\overline{\varepsilon}}(p_1)$ for families of pants decompositions other than $\mathcal{P}_{\mathrm{comp}}(T_K^-)$. We use Conway's notation [1970; Kauffman and Lambropoulou 2004; Mulazzani and Vesnin 2001] to describe 2–bridge links. The curve in the top of Figure 22 (resp. Figure 24) bounds a compressing disk on both sides of the 2–bridge link with Conway number 0 (resp. ∞). The distance below can be computed using continued fraction expansions of p/q [Agol 2010]. For details on continued fraction expansions of rational tangles, see [Hatcher 2022].

Corollary 4.5 *Let $K \subset S^3$ be a 2–bridge knot. If K is the numerator closure of a 2–string trivial tangle with Conway number p/q, then*

$$\mathcal{L}(\mathcal{T}_{\mathrm{MZ}}) \le \min\{6d(p/q,0)+6,\, 6d(p/q,\infty)+9\}.$$

Proof For 2–bridge knots, the only curve bounding a compressing disk in T_K^- (resp. T_K^+) is the loop of slope 0 (resp. p/q) in the 4–punctured bridge sphere. Furthermore, there are no cut disks for T_K^+ since b is small. The first inequality, $\mathcal{L}(\mathcal{T}_{\mathrm{MZ}}) \le 6d(p/q,0)+6$, follows from Theorem 4.3.

In order to prove the second inequality, we consider $p_1 \subset \mathcal{P}(\Sigma_4)$ corresponding to the curve $B \subset \Sigma_4$ with slope ∞ in Figure 24. In the same figure, we observe that the distance between the pants decompositions $G^{+}_{\varepsilon,\bar{\delta}}(p_1)$ and $G^{-}_{\rho,\bar{\varepsilon}}(p_1)$ is bounded by three. By Remark 4.4, we conclude $\mathcal{L}(S(K)) \leq 6d(p/q,\infty)+3\cdot 3$, as desired. □

References

[Agol 2010] **I Agol**, *Distance formula in Farey graph?*, reply in MathOverflow thread (2010) Available at https://mathoverflow.net/q/33245

[Artin 1925] **E Artin**, *Zur Isotopie zweidimensionaler Flächen im* $\mathbb{R}_4$, Abh. Math. Sem. Univ. Hamburg 4 (1925) 174–177 MR Zbl

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

[Blair et al. 2022] **R Blair**, **M Campisi**, **S A Taylor**, **M Tomova**, *Kirby–Thompson distance for trisections of knotted surfaces*, J. Lond. Math. Soc. 105 (2022) 765–793 MR Zbl

[Boileau and Zimmermann 1989] **M Boileau**, **B Zimmermann**, *The π–orbifold group of a link*, Math. Z. 200 (1989) 187–208 MR Zbl

[Conway 1970] **J H Conway**, *An enumeration of knots and links, and some of their algebraic properties*, from "Computational problems in abstract algebra" (J Leech, editor), Pergamon, Oxford (1970) 329–358 MR Zbl

[Gay and Kirby 2016] **D Gay**, **R Kirby**, *Trisecting 4–manifolds*, Geom. Topol. 20 (2016) 3097–3132 MR Zbl

[Gay and Meier 2022] **D Gay**, **J Meier**, *Doubly pointed trisection diagrams and surgery on 2–knots*, Math. Proc. Cambridge Philos. Soc. 172 (2022) 163–195 MR Zbl

[Hatcher 2022] **A Hatcher**, *Topology of numbers*, Amer. Math. Soc., Providence, RI (2022) MR Zbl

[Kauffman and Lambropoulou 2004] **L H Kauffman**, **S Lambropoulou**, *On the classification of rational tangles*, Adv. in Appl. Math. 33 (2004) 199–237 MR Zbl

[Kirby and Thompson 2018] **R Kirby**, **A Thompson**, *A new invariant of 4–manifolds*, Proc. Natl. Acad. Sci. USA 115 (2018) 10857–10860 MR Zbl

[Lambert-Cole 2020] **P Lambert-Cole**, *Bridge trisections in $\mathbb{CP}^2$ and the Thom conjecture*, Geom. Topol. 24 (2020) 1571–1614 MR Zbl

[Lambert-Cole et al. 2021] **P Lambert-Cole**, **J Meier**, **L Starkston**, *Symplectic 4–manifolds admit Weinstein trisections*, J. Topol. 14 (2021) 641–673 MR Zbl

[Lee 2017] **J H Lee**, *Reduction of bridge positions along bridge disks*, Topology Appl. 223 (2017) 50–59 MR Zbl

[Meier and Zupan 2017] **J Meier**, **A Zupan**, *Bridge trisections of knotted surfaces in S^4*, Trans. Amer. Math. Soc. 369 (2017) 7343–7386 MR Zbl

[Meier and Zupan 2018] **J Meier**, **A Zupan**, *Bridge trisections of knotted surfaces in 4–manifolds*, Proc. Natl. Acad. Sci. USA 115 (2018) 10880–10886 MR Zbl

[Mulazzani and Vesnin 2001] **M Mulazzani**, **A Vesnin**, *The many faces of cyclic branched coverings of 2–bridge knots and links*, Atti Sem. Mat. Fis. Univ. Modena 49 (2001) 177–215 MR Zbl

[Zupan 2013] **A Zupan**, *Bridge and pants complexities of knots*, J. Lond. Math. Soc. 87 (2013) 43–68 MR Zbl

Department of Mathematical Sciences, Binghamton University
Vestal, NY, United States

Current address: *Department of Mathematics, University of Nebraska – Lincoln*
Lincoln, NE, United States

Department of Mathematical Sciences, University of Saskatchewan
Saskatoon, SK, Canada

Current address: *School of Mathematics and Statistical Sciences, Arizona State University*
Tempe, AZ, United States

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

Department of Mathematics, University of California, Davis
Davis, CA, United States

jarandacuevas2@unl.edu, ppongtan@asu.edu, sataylor@colby.edu, sczhang@ucdavis.edu

Received: 4 January 2022 Revised: 27 May 2022

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 6 (pages 2971–3570) 2024