Geometry & Topology

Volume 29 (2025)

Tunnel number one knots satisfy the Berge conjecture

Tao Li
Yoav Moriah
Tali Pinsky

msp

msp

Geometry & Topology **29**:6 (2025) 3271–3343
DOI: 10.2140/gt.2025.29.3271
Published: 22 September 2025

Tunnel number one knots satisfy the Berge conjecture

TAO LI
YOAV MORIAH
TALI PINSKY

Let K be a tunnel number one knot in M with irreducible knot exterior, where M is either S^3, or a connected sum of $S^2 \times S^1$ with any lens space. (In particular, this includes $M = S^2 \times S^1$.) We prove that if a nontrivial Dehn surgery on K yields a lens space, then K is a doubly primitive knot in M. For $M = S^3$ this resolves the tunnel number one Berge conjecture. For $M = S^2 \times S^1$ this resolves a conjecture of Greene and Baker, Buck and Lecuona for tunnel number one knots.

57M99, 57K10, 57K30

Introduction

One of the main goals of low-dimensional topology is to determine which manifolds are obtained by Dehn fillings on knots in the 3-sphere. D Gabai [15] proved the Property R conjecture, ie if a Dehn surgery on a knot in S^3 yields $S^2 \times S^1$, then the knot must be a trivial knot. Another major achievement in the 1980s was the proof by C Gordon and J Luecke [17] that if nontrivial surgery on a knot yields S^3, then the knot must be the trivial knot. Both $S^2 \times S^1$ and S^3 are lens spaces. So a natural question is:

Question 1 Which knots in S^3 other than the unknot have nontrivial surgery resulting in a lens space which is not S^3 or $S^2 \times S^1$?

This question was first raised by L Moser [25] in 1971. She determined the surgeries on torus knots yielding lens spaces. Subsequently, Bleiler and Litherland [11], Wang [35] and Wu [36], independently, characterized the surgeries on satellite knots in S^3 which result in lens spaces. The question is still unresolved for hyperbolic knots.

In his celebrated notes, W Thurston [34] proved that each hyperbolic knot has only finitely many surgery slopes so that the manifolds obtained by Dehn fillings along these slopes fail to be hyperbolic (this number is uniformly bounded over all hyperbolic knots; see Agol [1], Bleiler and Hodgson [10], Lackenby [23] and Lackenby and Meyerhoff [24]). There has been a huge effort by many mathematicians to precisely determine these fillings. As lens spaces are nonhyperbolic, the above question can be viewed as a special case of this huge task.

J Berge [7] observed that if a knot K is doubly primitive[1] (see definitions below), then a Dehn surgery on K yields a lens space. He compiled a list of twelve families of doubly primitive knots $K \subset S^3$ including known cases of torus and satellite knots and asked whether a knot $K \subset S^3$ is on this list if and only if K is doubly primitive? (This question was later formally conjectured by C Gordon [22, Problem 1.78].)[2] The question became to be known as the *Berge conjecture*:

Conjecture 2 (the Berge conjecture) Let $K \subset S^3$ be a nontrivial knot which has Dehn surgery resulting in a lens space. Then K is doubly primitive.

A simple closed curve on the boundary of a genus-two handlebody is *primitive* if the handlebody has a compressing disk that transversely intersects this curve in a single point. A knot $K \subset M$ is *doubly primitive* if M has a genus-two Heegaard splitting and K can be isotoped onto the Heegaard surface so that K is primitive in both handlebodies of the Heegaard splitting. Note that if one adds a two-handle to a genus-two handlebody along a primitive curve, then the resulting manifold is a solid torus. Thus a doubly primitive knot always has a Dehn surgery that results in a lens space.

[1] Berge calls these knots "double primitive".

[2] Ultimately the proof that Berge's list for knots in S^3 is complete is a consequence of the lens space realization paper of J Greene [19].

There is a stronger version of the Berge conjecture, which says that if a Dehn surgery on K produces a lens space, then K has a doubly primitive presentation such that the framing given by the Heegaard surface is the surgery slope. Note that a knot may have more than one doubly primitive presentation. By Baker, Doleshal and Hoffman [4, Theorem 10], nontorus doubly primitive knots in S^3 or $S^2 \times S^1$ have no unexpected lens space surgery.

Given a knot K in a closed orientable 3-manifold M, a *tunnel system* for K is a collection of disjoint properly embedded arcs $\{t_1, \dots, t_n\}$ in the knot exterior $E(K) = M \setminus \mathcal{N}(K)$, where $\mathcal{N}(\cdot)$ means a regular neighborhood of $\{\cdot\}$, such that $E(K) \setminus \mathcal{N}\big(\bigcup_{i=1}^{n} t_i\big)$ is a handlebody. The *tunnel number* of the knot K is the minimal number of arcs in a tunnel system for K. Note that a knot K has tunnel number n if and only if $M \setminus \mathcal{N}(K)$ has Heegaard genus $n+1$.

Let $K \subset M$ be a doubly primitive knot. If one pushes K into the interior of either genus-two handlebody of the corresponding Heegaard splitting, then after removing $\mathcal{N}(K)$, this handlebody becomes a compression body and the genus-two Heegaard surface for M becomes a Heegaard surface of $M \setminus \mathcal{N}(K)$. Moreover, the one-handle in the compression body determines an unknotting tunnel for K. This means that all doubly primitive knots have tunnel number one.

Therefore, the Berge conjecture can be divided into two parts:

Conjecture 3 (the tunnel number one Berge conjecture) If $K \subset S^3$ is a tunnel number one knot which admits a Dehn surgery resulting in a lens space, then K is doubly primitive.

Conjecture 4 (the lens space Dehn surgery) If $K \subset S^3$ is a nontrivial knot which admits a Dehn surgery resulting in a lens space, then K is a tunnel number one knot.

As the notion of doubly primitive knots is not limited to S^3, one can ask whether a conjecture similar to the Berge conjecture holds for other manifolds with Heegaard genus at most 2:

Question 5 Let K be a knot in a 3-manifold M such that a Dehn surgery on K yields a lens space. Is K necessarily doubly primitive with respect to a genus-two Heegaard splitting of M?

The answer to Question 5 is known to be negative when M is the Poincaré homology sphere (see Baker [2]) or a lens spaces in general (see Baker, Buck and Lecuona [3]). Namely, there are knots in the Poincaré homology sphere and in nontrivial lens spaces such that surgery on them yield lens spaces but the knots are not doubly primitive with respect to a genus two Heegaard splitting of these spaces.

The next conjecture, which was made by J Greene [19, Conjecture 1.8] and by Baker, Buck and Lecuona [3, Conjecture 1.1], says that the answer to Question 5 is expected to be true if M is $S^2 \times S^1$.

Conjecture 6 (Berge conjecture for $S^2 \times S^1$) If K is a knot in $S^2 \times S^1$ which admits a lens space Dehn surgery, then K is doubly primitive.

In Theorem 7 we prove Conjecture 3. Theorem 7 also implies that Conjecture 6 also holds for knots with tunnel number one.

Theorem 7 *Let $K \subset M$ be a tunnel number one knot with irreducible knot exterior, where M is either S^3 or $(S^2 \times S^1) \# L(r,s)$, where $L(r,s)$ is any lens space. If a nontrivial Dehn surgery on K yields a lens space, then K is doubly primitive.*

Note that if $M = (S^2 \times S^1) \# L(r,s)$ and K is a core curve of a Heegaard solid torus in the lens space summand $L(r,s)$, then a Dehn surgery on K may change $L(r,s)$ to S^3, thereby transforming M to the lens space $S^2 \times S^1$. However, K may not be doubly primitive on the genus-two Heegaard surface of M. This special case is the only reason we require the knot exterior $M \setminus K$ to be irreducible in Theorem 7. In particular, this hypothesis has no impact when M is S^3 or $S^2 \times S^1$, since $S^3 \setminus K$ is irreducible for all knots, and if $(S^2 \times S^1) \setminus K$ is reducible, then K lies in a 3-ball and no nontrivial Dehn surgery produces a lens space. The only point in our proof where this hypothesis is required is in Lemma 1.1.3.

Theorem 7 suggests that the answer to Question 5 might be positive in the case where M is a connected sum of $S^2 \times S^1$ and a lens space.

Let K be a knot in S^3. We say that K admits *genus-reducing surgery* if $S^3 \setminus \mathcal{N}(K)$ has a nontrivial surgery resulting in a manifold M so that $g(M) \leq g(S^3 \setminus \mathcal{N}(K)) - 1$, where $g(X)$ denotes the Heegaard genus of the manifold X. It will be called *strongly genus-reducing surgery* if $g(M) < g(S^3 \setminus \mathcal{N}(K)) - 1$. In this context, the second author would like to make the following conjecture:

Conjecture 8 (strong genus-reducing surgery) Knots in S^3 do not have strongly genus-reducing surgery slopes other than the slope of the meridian.

The following is a weaker conjecture:

Conjecture 9 For any knot $K \subset S^3$, no integer slope is a strongly genus-reducing surgery slope.

It is a consequence of the cyclic surgery theorem (see Culler, Gordon, Luecke and Shalen [14]) that if $K \subset S^3$ is a hyperbolic knot only integer slope surgery can yield a lens space. Thus Conjecture 9, if true, together with Theorem 7 implies the Berge conjecture.

The Berge conjecture has been studied extensively, eg by Berge himself [7; 8; 9], also by D Gabai [16], K Baker, E Grigsby and M Hedden [6], P Ozsváth and Z Szabó [29], J Rasmussen [30], T Saito [31], M Tange [32; 33], Y Ni [27], J Greene [19] and others. There is some strong evidence for the Berge conjecture. For example, Yi Ni [27] proved that if a knot in S^3 admits lens space surgery, then the knot must be fibered (all Berge knots are fibered). Greene [19, Theorem 1.3] proved that if a lens space can be obtained by surgery on a knot in S^3 then such a lens space can be obtained by a knot in the Berge list

However, very little is currently known on either of Conjectures 8 or 9. Some results in this direction were obtained by K Baker, C Gordon and J Luecke [5, Corollary 1.1].

The main idea of the proof and outline of the paper Assume that M is a 3-manifold and $K \subset M$ is a knot. Let $\{U, W\}$ be a Heegaard splitting for $M \setminus \mathcal{N}(K)$ along a genus-g Heegaard surface Σ, where U is a compression body and W is a handlebody, see Definitions 1.1.1. Then Σ is also a Heegaard surface of the manifold $M(r)$ obtained by r-surgery on $M \setminus \mathcal{N}(K)$. If the Heegaard splitting of $M(r)$ along Σ is stabilized, then the handlebody W must contain an essential disk D and the compression body U must contain a planar surface P such that $P \cap D$ is a single point in Σ and, after the Dehn surgery on K, P becomes an essential disk $\widehat{P}$. So $(\widehat{P}, D)$ is a genus-reducing destabilizing disk pair; see Definitions 1.1.1 and 1.1.2 below. The pair (P, D) is called a $(\mathcal{P}, \mathcal{D})$-pair.

As described at the beginning of Section 1.2, we are dealing with a genus-two Heegaard splitting of $M \setminus \mathcal{N}(K)$, where M is either S^3 or $(S^2 \times S^1) \# L(r, s)$, which becomes, after integer surgery, a reducible genus-two Heegaard splitting of a lens space. Thus the genus-two Heegaard splitting must contain a $(\mathcal{P}, \mathcal{D})$-pair (P, D).

Consider a Heegaard diagram of the genus-two Heegaard splitting of M, namely $\widehat{V} = \{\alpha, \gamma\}$ and $\widehat{W} = \{\delta, \varepsilon\}$, where $\delta = \partial D$, ε is a meridian of the handlebody W disjoint from δ and ∂P, α is a boundary curve of a vertical annulus in the compression body U, and γ is the boundary of a disk in U. See the beginning of Section 1.3, Definitions 1.3.1 and 1.3.6 for more details.

The intuition behind the proof is the following: We do not know what the curves $\alpha, \gamma, \delta, \varepsilon$ and ∂P look like. However, since the Heegaard surface Σ is of genus two and since $\gamma \cap \alpha = \delta \cap \varepsilon = \varepsilon \cap \partial P = \varnothing$ and $\partial P \cap \partial D$ is a single point, the possible configurations of these curves should be sufficiently limited to give the required result.

Main idea The name of the game is to analyze all possible configurations for these curves. Note that if we cut the Heegaard surface Σ along $\partial_+ P$, the curves $\widehat{V} = \{\gamma, \alpha\}$ and $\widehat{W} = \{\delta, \varepsilon\}$ are cut into a collection of arcs on $\Sigma \setminus \mathcal{N}(\partial_+ P)$. Using this information define a notion of *complexity* $c(P, D, \alpha, \gamma)$ of a $(\mathcal{P}, \mathcal{D})$-pair; see Definition 1.2.1. In order to reduce this complexity, we perform *wave moves* (see Definition 1.3.2) on the above curves which yield new curves and meridional disks so that one of the following holds:

(1) We find a new $(\mathcal{P}, \mathcal{D})$-pair or a new α with smaller complexity $c(P, D, \alpha, \gamma)$. If this is the case then, after iterating the argument if need be, we use Proposition 1.2.2 to conclude that K is a Berge knot.

(2) The knot K is a Berge–Gabai knot, see [8; 16], and hence doubly primitive, ie K is a Berge knot.

(3) The new Heegaard diagram, which is not standard, has no wave. In this case Theorems 1.3.7 and 1.3.8 are contradicted when M is one of S^3 or $(S^2 \times S^1) \# L(r, s)$, where $L(r, s)$ is a lens space.

Outline of the paper In Sections 1.1–1.3, we define the necessary notions and prove some of the lemmas that are needed throughout the paper. In particular, we define a notion of *complexity* and show that if there is a $(\mathcal{P}, \mathcal{D})$-pair and an annulus with very small complexity, then the knot K must be doubly

primitive; see Proposition 1.2.2. We also study the Whitehead graphs of genus-two Heegaard splittings of M. The edges of a Whitehead graph of M are restricted by Theorems 1.3.7 and 1.3.8. For example if the Heegaard diagram is not "standard" (as in Definition 1.3.6), the edges cannot be *blocking-edges* as defined in Definition 1.3.10.

The next crucial step, done in Section 2.1, is to first cut the Heegaard surface Σ open along the curve $\partial_+ P$, denoting the resulting compact surface by $\widehat{\Sigma}$, and then fix a "standard" decomposition for $\widehat{\Sigma}$. The surface $\widehat{\Sigma}$ is composed of two annuli, denoted by $\mathcal{A}_r$ and $\mathcal{A}_l$, and two rectangles denoted by $\mathcal{R}^u$ and $\mathcal{R}^d$. The arcs $\gamma \setminus \partial_+ P$, $\alpha \setminus \partial_+ P$, $\delta \setminus \partial_+ P$, and the curve ε in $\widehat{\Sigma} = \mathcal{A}_r \cup \mathcal{R}^u \cup \mathcal{R}^d \cup \mathcal{A}_l$ have certain properties and restrictions, which are discussed in lemmas, remarks and propositions numbered 2.1.1–2.1.9.

In Section 2.2 we concentrate on the $\{\varepsilon, \delta\}$ and $\partial_+ P$ curves. In order to describe the curves we define the notion of *paths*, more specifically *short* and *long* paths, which the curves $\{\varepsilon, \delta\}$ must follow (take); see Definition 2.2.3. Each of the curves can take a path many times. Since the curves $\{\varepsilon, \delta\}$ after being cut along $\partial_+ P$ can have many parallel arcs which take the same path, it is convenient to think of them as *train tracks*. This is done in the discussion below Definition 2.2.3. These train tracks will play a crucial role in Section 3.

Section 3, which is the essence of the proof, is a careful and detailed analysis of the curve δ in $\widehat{\Sigma}$ depending on the paths that δ takes. It turns out that there are four possibilities:

(1) δ takes one short path in $\mathcal{A}_l$ and one short path in $\mathcal{A}_r$.

(2) δ takes one short path in $\mathcal{A}_l$ or $\mathcal{A}_r$ and two short paths in the other annulus.

(3) δ takes two short paths in $\mathcal{A}_l$ and two short paths in $\mathcal{A}_r$.

(4) δ takes a long path in one (or both) of the annuli.

This analysis gives the three possibilities as in the main idea above, which proves the theorem.

Acknowledgements We thank John Berge for important remarks regarding $(1,1)$-knots and for pointing out the importance of the papers by T Homma, M Ochiai and M Takahashi [21] and M Ochiai [28] to the Berge conjecture. Li thanks Ken Baker and Alex Zupan for helpful communications. Moriah and Pinsky thank Dale Rolfsen for his support. We thank the referee for numerous corrections and suggestions.

Li is partially supported by an NSF grant. Moriah thanks the University of British Columbia for its generous hospitality.

1 Preliminaries

Throughout the paper, we assume that M is either S^3 or $(S^2 \times S^1) \# L(r,s)$, where $L(r,s)$ is a lens space. Since S^3 is a lens space, the case $M = S^2 \times S^1$ is included. For any space Y, we use $|Y|$ to denote the number of components of Y, use $\mathrm{int}(Y)$ to denote the interior of Y and use $\mathcal{N}(Y)$ to denote a small regular neighborhood of Y. Finally, throughout the paper the hyperelliptic involution will be denoted by π.

1.1 Primitive and doubly primitive

We begin with some definitions:

Definition 1.1.1 A *compression body* is a connected 3-manifold U obtained by adding two-handles to a product $\Sigma \times I$ ($I = [0,1]$) along $\Sigma \times \{0\}$, where Σ is a closed and orientable surface, and then capping off resulting 2-sphere boundary components by 3-balls. The surface $\Sigma \times \{1\}$ is denoted by $\partial_+ U$ and $\partial U \setminus \partial_+ U$ is denoted by $\partial_- U$. If $\partial_- U = \varnothing$ then U is a handlebody. One can also view a compression body with $\partial_- U \neq \varnothing$ as a manifold obtained by adding one-handles on the same side of $\partial_- U \times I$.

A *Heegaard splitting* of a 3-manifold Y, denoted by (U, W), is a decomposition of Y into two compression bodies U and W along a closed surface $\Sigma = \partial_+ U = \partial_+ W$, and Σ is called a *Heegaard surface*.

Let $K \subset M$ be a knot in M. Throughout the paper, we will use (U, W) to denote a Heegaard splitting for $M \setminus \mathcal{N}(K)$ with $\partial_- U = \partial(M \setminus \mathcal{N}(K))$ being the boundary torus, and we use Σ to denote the Heegaard surface. Note that a trivial Dehn filling along $\partial_- U$ extends U to a handlebody V, and hence Σ is also a Heegaard surface of M. Throughout the paper, we use (V, W) to denote this Heegaard splitting of M, with $\Sigma = \partial V = \partial W$, $K \subset V$ and $U = V \setminus \mathcal{N}(K)$.

For any surface S properly embedded in a compression body U, we use $\partial_- S$ to denote $S \cap \partial_- U$ and $\partial_+ S$ to denote $S \cap \partial_+ U \subset \Sigma$.

Throughout this paper P will denote a planar surface which is not a disk in U. We require further that P has a single boundary component, denoted by $\partial_+ P$ on $\partial_+ U$ and $N \geq 1$ boundary components on $\partial_- U$.

An annulus A properly embedded in the compression body U is called a *vertical annulus* if A is incompressible and has one boundary component in $\partial_+ U$ and the other boundary component in $\partial_- U$. Note that there are many vertical annuli in U even with the same $\partial_- A$ curve, by taking band sums of one such A with compressing disks in U.

Definition 1.1.2 Let (U, W) be a Heegaard splitting of $M \setminus \mathcal{N}(K)$ as above. Let P be a planar surface properly embedded in U and D a compressing disk in W. Suppose $\partial_+ P$ is a single curve transversely intersecting ∂D in Σ. We say that (P, D) is a $(\mathcal{P}, \mathcal{D})$-*pair* with respect to the Heegaard splitting with slope r if $\partial D \cap \partial_+ P$ is a single point and $\partial_- P$ consists of essential curves of slope r in the boundary torus. If P is an annulus then we say that (P, D) is an $(\mathcal{A}, \mathcal{D})$-*pair*.

In this paper we assume that some nontrivial Dehn surgery on K yields a lens space.

Lemma 1.1.3 *Let M be either S^3 or $(S^2 \times S^1) \# L(r, s)$, and let K be a knot in M with irreducible knot exterior. Suppose K admits a nontrivial Dehn surgery resulting in a lens space. Then either K is doubly primitive, or the surgery slope is an integer slope.*

Proof If $M = S^3$, then since K admits lens space surgery, by the cyclic surgery theorem (see [14]), either the surgery slope r is an integer or K is a torus knot. Since a torus knot is doubly primitive (see [7]), the lemma holds.

Suppose therefore that M is $S^2 \times S^1$ or $(S^2 \times S^1) \# L(r,s)$. Since $M \setminus \mathcal{N}(K)$ is irreducible, by [13, Corollary 1.4], either the surgery slope is an integer slope or $M \setminus \mathcal{N}(K)$ is a simple Seifert fibered space. If $M \setminus \mathcal{N}(K)$ is a simple Seifert fibered space, then it contains an essential annulus A which divides $M \setminus \mathcal{N}(K)$ into two solid tori. Since M is $(S^2 \times S^1) \# L(r,s)$ and since $M \setminus \mathcal{N}(K)$ is irreducible, $M \setminus \mathcal{N}(K)$ admits an essential planar surface with ∞-slope, ie $S^2 \setminus \mathcal{N}(K)$, where S^2 is an essential 2-sphere in $(S^2 \times S^1) \# L(r,s)$. By [18], the boundary slope of A and the boundary slope of the planar surface $S^2 \setminus \mathcal{N}(K)$ has geometric intersection one. Hence the boundary slope of A is an integer slope. Since A divides $M \setminus \mathcal{N}(K)$ into two solid tori and ∂A has integer slope, K must lie on a torus which divides M into two solid tori, which implies that $M = S^2 \times S^1$. Similar to the argument about torus knots (see [7]), K is doubly primitive with respect to a genus-two Heegaard surface of M. □

Proposition 1.1.4 *Let $K \subset M$ be a tunnel number one knot with irreducible knot exterior, where M is S^3 or $(S^2 \times S^1) \# L(r,s)$. Suppose K admits a lens space surgery, Then, if there is an $(\mathcal{A}, \mathcal{D})$-pair with respect to a genus-two Heegaard splitting of $M \setminus \mathcal{N}(K)$ where $\partial_- A$ has the surgery slope, then K is doubly primitive.*

Proof Suppose (A, D) is such a pair, ie $A \subset U$ is a vertical annulus meeting the disk $D \subset W$ in a single point in $\partial_+ A$. If we perform Dehn filling along the slope r of $\partial_- A$, A extends to a disk and the (A, D)-pair extends to a stabilizing pair of disks in the manifold $M(r)$ obtained by Dehn filling along the slope r. Since the Heegaard splitting has genus two, this implies that $M(r)$ has Heegaard genus at most one and hence is a lens space.

As in Lemma 1.1.3, this means that the slope r of $\partial_- A$ is an integer or K must be doubly primitive. Suppose r is an integer, then $\partial_- A$ and hence $\partial_+ A$ are isotopic to the knot K in M. So we may isotope K to the curve $\partial_+ A \subset \Sigma$. Since $\partial_+ A$ transversely intersects ∂D in one point, K is primitive in W. Since U is a compression body, U contains a vertical annulus A' such that $\partial_- A'$ has meridional slope. Hence $\partial_+ A'$ bounds a disk D' in the handlebody $V = U \cup \mathcal{N}(K)$ intersecting K in one point. So K is primitive in V as well. □

Remark 1.1.5 A tunnel number one knot may have multiple different unknotting tunnels. Thus, a priori, it is possible that a knot K is not doubly primitive with respect to some genus-two Heegaard surface of $S^3 \setminus \mathcal{N}(K)$ but is doubly primitive with respect to a different genus-two Heegaard surface of $S^3 \setminus \mathcal{N}(K)$.

In fact, the $(-2,3,7)$-pretzel knot K is a tunnel number one knot with four unknotting tunnels; see [20]. John Berge pointed out that the compression body in one of these genus-two Heegaard splittings can be obtained by adding a two-handle to the boundary of a genus-two handlebody H along a simple closed

curve $\rho \subset \partial H$ which represents the word $x^3y^2x^2y^2x^3y^3$ in $\pi_1(H) = F(x, y)$, the free group on two generators. In this case, there is only one simple closed curve disjoint from ρ which represents a primitive in H, namely the curve xy, which is the meridian of $H(\rho) = S^3 \setminus \mathcal{N}(K)$. So clearly K cannot be doubly primitive with respect to this Heegaard splitting. However, the $(-2, 3, 7)$-pretzel knot is on the Berge list.

This example implies that the proof of the main theorem, Theorem 7, will have to deal with this possible situation. This will be done in the next section.

1.2 A notion of complexity for a $(\mathcal{P}, \mathcal{D})$-pair

Let $K \subset M$ be a tunnel number one knot and (U, W) a genus-two Heegaard splitting of $M \setminus \mathcal{N}(K)$, where W is a handlebody, and let Σ be the Heegaard surface. The surface Σ also determines a Heegaard splitting (V, W) of M with $V = U \cup \mathcal{N}(K)$. Suppose the Dehn filling along a slope r yields a lens space. Note that Σ is also a genus-two Heegaard surface of the lens space. By Bonahon and Otal [12], the genus-two Heegaard splitting of the lens space is stabilized by a stabilizing pair of disks (D, E). Assume $D \subset W$ and $E \subset V$. Hence $E \setminus \mathcal{N}(K) \subset V \setminus \mathcal{N}(K) = U$ is a planar surface $P \subset U$ which intersects $D \subset W$ in a single point. Thus, a stabilizing pair of disks gives rise to a $(\mathcal{P}, \mathcal{D})$-pair (P, D), where $\partial_- P$ has slope r. By Lemma 1.1.3, we may suppose that r is an integer slope. If P is compressible in U, then by compressing P, we obtain a new planar surface P' with $\partial_+ P' = \partial P_+$. So, after finitely many compressions, we obtain a new $(\mathcal{P}, \mathcal{D})$-pair with incompressible planar surface. Assume throughout the paper that P is incompressible in U. By Proposition 1.1.4, we shall assume P is not an annulus (as otherwise we are done).

We now proceed to define a notion of complexity for $(\mathcal{P}, \mathcal{D})$-pairs. As the Heegaard splitting is of genus two, there is a unique, up to isotopy, nonseparating disk in U which is denoted by C. There is a vertical annulus A in U such that $\partial_- A$ has the ∞-slope (ie the meridional slope of K in M). Note that A is not unique; in fact, one can construct infinitely many such annuli by band-summing A with C in different ways. Suppose $C \cap A = \varnothing$. Let $\gamma = \partial C$ and $\alpha = \partial_+ A$. Note that γ and α lie on the Heegaard surface Σ, and both γ and α bound nonseparating disks in the handlebody $V = U \cup \mathcal{N}(K)$.

Definition 1.2.1 Given a $(\mathcal{P}, \mathcal{D})$-pair (P, D) in a Heegaard splitting of $M \setminus \mathcal{N}(K)$ and a collection of disjoint simple closed curves $\{\alpha_1, \dots, \alpha_n\} \subset \Sigma$ in general position with respect to $(\partial P_+, \partial D)$, define the complexity $c_0(P, D, \alpha_1, \dots, \alpha_n)$ to be the number of segments of $\{\alpha_1 \cup \cdots \cup \alpha_n\} \setminus \partial_+ P$ which intersect D.

In the current situation, for a given (P, D), consider the complexity $c_0(P, D, \alpha, \gamma)$, where $\gamma = \partial C$ and $\alpha = \partial_+ A$ and where A is a vertical annulus in U such that $\partial_- A$ has the ∞-slope. Further, define $c_1(D, \alpha, \gamma)$ to be the number of intersection points of ∂D and $\alpha \cup \gamma$. Now, define the pair

$$c(P, D, \alpha, \gamma) = (c_0(P, D, \alpha, \gamma), c_1(D, \alpha, \gamma))$$

with the lexicographical order to be the complexity of (P, D).

As mentioned above, the general idea of the proof is to show one can reduce the complexity of the diagram. The following proposition shows that if one reaches a diagram with small complexity, then K is a Berge knot. It is thus a main ingredient of the proof. This proposition also deals with the situation described in Remark 1.1.5, where a knot space has multiple Heegaard splittings in which the knot K is not doubly primitive with respect to some and is doubly primitive with respect to others.

Proposition 1.2.2 *Let (P, D) be a $(\mathcal{P}, \mathcal{D})$-pair, and $\alpha = \partial_+ A$ as above. Suppose $|P \cap A|$ is minimal up to isotopy on P and A. If $c_0(P, D, \alpha) \leq 1$, then K is doubly primitive with respect to a (possibly different) genus-two Heegaard splitting of M.*

Proof We have assumed at the beginning of Section 1.2 that P is incompressible. This implies P does not have a ∂-compressing disk at the boundary torus $\partial\mathcal{N}(K)$. To see this, if P has such a ∂-compressing disk, after boundary compressing P along this disk we obtain a new surface P' with a boundary curve bounding a disk on the boundary torus $\partial\mathcal{N}(K)$. If this curve bounds a disk in P' as well, then since $M \setminus \mathcal{N}(K)$ is irreducible, P must be a boundary parallel annulus. If this curve is essential in P', then the disk in $\partial\mathcal{N}(K)$ gives rise to a compressing disk for P. Both cases contradict the choice of P.

By Lemma 1.1.3, we may assume that the curves in $\partial_- P$ have integer slope. Hence each component of $\partial_- P$ intersects $\partial_- A$ in exactly one point. Since P is incompressible and $|P \cap A|$ is minimal we may assume that, perhaps after isotopy, $P \cap A$ contains no closed curve that is trivial in A. Since P does not have ∂-compressing disk at the boundary torus, after isotopy, $P \cap A$ contains no arc with both endpoints in $\partial_- A$. Further, each component of $A \cap P$ must be either an essential arc in A or an arc in A with both endpoints in $\alpha = \partial_+ A$.

Let $k_1, \ldots, k_N$ be arcs in $A \cap P$ that are essential in A. Thus $k_1, \ldots, k_N$ divide A into N rectangles $R_1, \ldots R_N$. By Proposition 1.1.4, we may assume P is not an annulus and hence $N \geq 2$.

Suppose $c_0(P, D, \alpha) \leq 1$. Then the intersection $\alpha \cap \partial D$, if nonempty, must lie in a single arc of $\alpha \setminus P$ and thus in the boundary of a single rectangle R_i, and without loss of generality we may assume that $\alpha \cap \partial D \subset \partial R_1$. Set $\widehat{R} = R_2 \cup \cdots \cup R_N$. If $P \cap \widehat{R}$ contains arcs that are ∂-parallel in A, we can perform ∂-compressions on P along these arcs. Since $\widehat{R} \cap D = \varnothing$, such ∂-compressions do not create new intersection points with D. Hence, after the ∂-compressions, one of the resulting planar surfaces still intersects D in a single point. Thus we may assume that no arc in $P \cap \widehat{R}$ is ∂-parallel in A; in other words, $P \cap \widehat{R} = \bigcup_{i=1}^{N} k_i$.

Let K^* denote the dual knot to K in the corresponding lens space L, where $L \setminus \mathcal{N}(K^*) \cong M \setminus \mathcal{N}(K)$. We view P, D and A in the exterior of K^*. The planar surface P can be viewed as an N-punctured disk in $L \setminus \mathcal{N}(K^*)$. The arcs $k_1, \ldots, k_N$, when viewed in P, are arcs connecting the N components of $\partial_- P$ to $\partial_+ P$. Let μ_i be the component of $\partial_- P$ that contains an endpoint of k_i, where $i = 1, \ldots, N$, and let N_{μ_i} be the closure of a small neighborhood of $\mu_i \cup k_i$ in P. So N_{μ_i} is an annulus of which μ_i is a

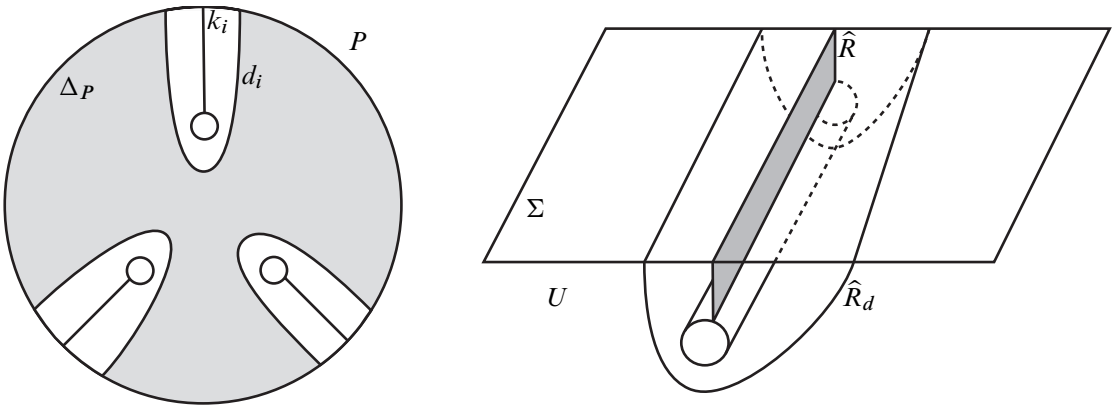

Figure 1: The surfaces Δ_P, left, and $\widehat{R}_d$, right, used to modify the Heegaard splitting.

boundary component. We may view the annulus N_{μ_i} as a subsurface of P cut off by an arc d_i, where $d_i \subset \partial N_{\mu_i}$ and d_i is properly embedded in P; see Figure 1, left. The arc d_i is the arc in P that goes around $\mu_i \cup k_i$ as indicated in Figure 1, left.

Since $k_1, \dots, k_N$ are vertical arcs of the rectangle $\widehat{R}$ there is a rectangle $\widehat{R}_d = I \times I$ "following" $\widehat{R}$ so that each d_i is a vertical arc in $\widehat{R}_d$ of the form of $\{x\} \times I$ and $\widehat{R}_d \cap \Sigma = I \times \partial I$; see Figure 1, right (the shaded region denotes $\widehat{R}$). The rectangle $\widehat{R}_d$ can be obtained by first taking two parallel copies of $\widehat{R}$ and then connecting them around K^*. Note that $\{d_1, \dots, d_N\}$ is a collection of disjoint vertical arcs in $\widehat{R}_d$. Since $\widehat{R}$ is disjoint from D we have $\widehat{R}_d \cap D = \varnothing$.

Next we modify the Heegaard splitting (U, W) of $M \setminus \mathcal{N}(K) = L \setminus \mathcal{N}(K^*)$. Let $U' = U \setminus \mathcal{N}(\widehat{R}_d)$. Note that $\widehat{R}_d$ deformation retracts to d_i for any i. So $U' \cong U \setminus \mathcal{N}(d_i)$ and one may view $\mathcal{N}(\widehat{R}_d)$ as a "fat" tunnel in U. As U is a compression body, there is a vertical annulus $A_{\mu_i} \subset U$ between any μ_i and a curve in Σ. Notice that, after isotopy, $A_{\mu_i} \cap A$ is a single arc k which is vertical in both A_{μ_i} and A. Hence k_i is isotopic to k in the annulus A. As $k \subset A_{\mu_i}$, we can isotope $\mu_i \cup k_i$ into A_{μ_i}. Hence the annulus N_{μ_i} can be isotoped to a subsurface of A_{μ_i}. In particular, d_i can be isotoped into A_{μ_i}.

As d_i has both end points on Σ, d_i is ∂-parallel in A_{μ_i} and hence ∂-parallel in U. Thus drilling out a tunnel in U with the arc d_i as its core results in a compression body isotopic to $U' = U \setminus \mathcal{N}(\widehat{R}_d)$. The complement of U' is a genus-three handlebody W'. Thus (U', W') is a stabilized genus-three Heegaard splitting of $M \setminus \mathcal{N}(K) = L \setminus \mathcal{N}(K^*)$.

Note that N_{μ_i} is a vertical annulus in U'. Let $\Delta \subset W'$ be a cocore (meridional) disk of the (fat) tunnel $\mathcal{N}(\widehat{R}_d)$ that we drill out from U. So $\Delta \cap N_{\mu_i}$ is a single point and (N_{μ_i}, Δ) is an $(\mathcal{A}, \mathcal{D})$-pair.

Let Δ_P be the closure of $P \setminus \bigcup_{i=1}^{N} N_{\mu_i}$; see the shaded region in Figure 1, left. Since $\widehat{R}$ connects all the components of $\partial_- P$, Δ_P is a disk which can be viewed as a properly embedded disk in U'. Since $\widehat{R}_d$ is disjoint from D, the operation of drilling this tunnel does not affect D. Hence we may view D

as a disk in W', and $D \cap \Delta_P = D \cap P$ is a single point. Thus (D, Δ_P) is a stabilizing pair of disks for the genus-three splitting (U', W'). By cutting W' along D and adding $\mathcal{N}(D)$ as a two-handle to U' along ∂D, we destabilize the splitting (U', W') into a new genus-two Heegaard splitting of $M \setminus \mathcal{N}(K)$. As D is disjoint from both N_{μ_i} and Δ, the pair (N_{μ_i}, Δ) described above remains an $(\mathcal{A}, \mathcal{D})$-pair in the resulting genus-two Heegaard splitting. By Proposition 1.1.4, K is doubly primitive with respect to this new genus-two Heegaard splitting. □

By Definition 1.2.1, $c_0(P, D, \alpha) \leq c_0(P, D, \alpha, \gamma)$. Hence K is doubly primitive if $c_0(P, D, \alpha, \gamma) \leq 1$. In light of Proposition 1.2.2 we adopt the following:

Assumption 1.2.3 Fix a $(\mathcal{P}, \mathcal{D})$-pair (P, D) and suppose P is incompressible in U. By Proposition 1.1.4, we may assume P is not an annulus. Throughout the paper, we assume P, D and α are chosen so that $c(P, D, \alpha, \gamma)$ is minimal among all the $(\mathcal{P}, \mathcal{D})$-pairs and all such annuli A $(\alpha = \partial_+ A)$.

We finish this section with the following lemma.

Lemma 1.2.4 *Let P and γ be as in Assumption 1.2.3. Then $\partial_+ P \cap \gamma \neq \varnothing$.*

Proof Suppose on the contrary that $\partial_+ P \cap \gamma = \varnothing$. Recall that $\gamma = \partial C$, where C is a compressing disk in U. Since P is incompressible in U and since $\partial_+ P \cap \gamma = \varnothing$, we can perform an isotopy so that $P \cap C = \varnothing$.

If we cut U open along the disk C, then the resulting manifold is homeomorphic to $T^2 \times I$ and $P \subset T^2 \times I$. Note that if P is ∂-compressible in $T^2 \times I$, then P is compressible (one can construct a compressing disk using two parallel copies of a ∂-compressing disk). Since P is assumed to be incompressible, P must be both incompressible and ∂-incompressible and hence π_1-injective in $T^2 \times I$. This means that P must be an annulus, contradicting Assumption 1.2.3. □

1.3 Meridional disks and Whitehead graphs

In this section we define Whitehead graphs and *waves*, which are the main tools in the proof of the theorem.

Let (V, W) be a genus-two Heegaard splitting of a closed 3-manifold M with a Heegaard surface Σ. Let $\widetilde{V} = \{V_1, V_2\}$ and $\widetilde{W} = \{W_1, W_2\}$ be *complete meridian sets* respectively. That is, each is a pair of essential disks such that $V \setminus \widetilde{V}$ and $W \setminus \widetilde{W}$ are 3-balls. Let $v_i = \partial V_i$ and $w_i = \partial W_i$ for $i = 1, 2$. To simplify notation, we also call v_i and w_i meridians of V and W, respectively. If we cut Σ open along $v_1 \cup v_2$, we get a 4-punctured sphere whose boundary consists of four components $\{v_1^+, v_1^-, v_2^+, v_2^-\}$, and $w_1 \cup w_2$ is cut into a collection of properly embedded curves in the 4-punctured sphere.

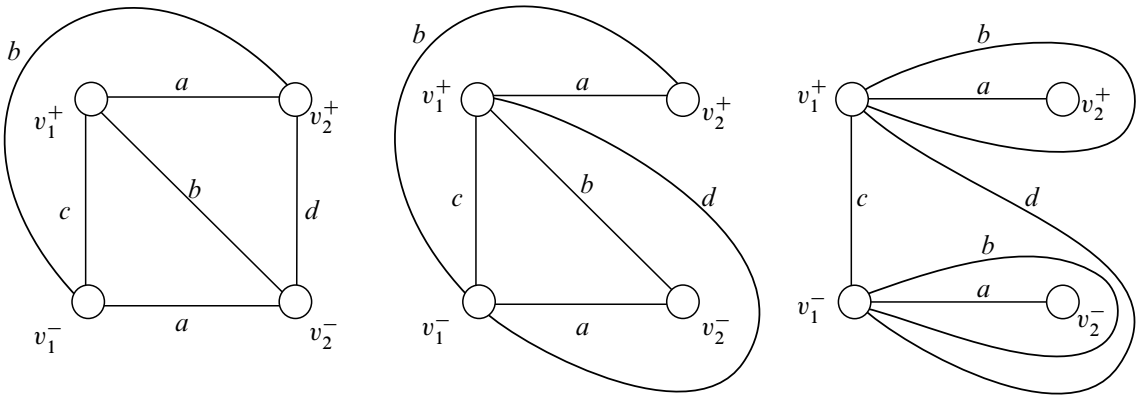

Figure 2: The three graphs given in [28].

Let $\widehat{V} = \{v_1, v_2\}$ and $\widehat{W} = \{w_1, w_2\}$. A Heegaard diagram $(\Sigma, \widehat{V}, \widehat{W})$ for the 3-manifold M is said to be *normal* if no domain in $\Sigma \setminus \{\widehat{V} \cup \widehat{W}\}$ is a bigon; see [28] for more detailed definitions. We can always isotope the curves v_1, v_2, w_1 w_2 in Σ to eliminate all the bigons and obtain a normal Heegaard diagram. Throughout the paper, we assume all the Heegaard diagrams are normal.

Definition 1.3.1 (genus-two Whitehead graphs) The graph $\Gamma(\{v_1, v_2\})$ (or $\Gamma(\widehat{V})$) obtained, as above, by setting $\{v_1^+, v_1^-, v_2^+, v_2^-\}$ to be vertices and the arcs of $(w_1 \cup w_2) \setminus (v_1 \cup v_2)$ to be edges is called the *Whitehead graph* corresponding to $\{v_1, v_2\}$. As $\Sigma \setminus \widehat{V}$ is a 4-hole sphere, $\Gamma(\{v_1, v_2\})$ is a planar graph in the sphere. Similarly we have the Whitehead graph $\Gamma(\{w_1, w_2\})$ (or $\Gamma(\widehat{W})$) corresponding to $\{w_1, w_2\}$.

Ochiai [28, Theorem 1] states that any Whitehead graph of a genus-two Heegaard diagram of a 3-manifold is isomorphic as a planar graph to one of the three graphs in Figure 2, where the integers a, b, c and d represent the number, which can be zero, of parallel arcs corresponding to each graph edge. Note that it is easy to understand this theorem and Figure 2 using the hyperelliptic involution π: Each nonseparating simple closed curve in a genus-two surface is invariant (up to isotopy) under π with orientation reversed. So the involution interchanges v_i^+ and v_i^- for $i = 1, 2$, and induces an involution on the Heegaard diagram. For example, as w_i is invariant under the involution, the numbers of arcs connecting v_1^+ to v_2^+ and v_1^- to v_2^- are the same (indicated by a in Figure 2). Furthermore, this theorem is about the graphs only. If one interchanges v_1 and v_2 or switches the $\pm$-signs of v_i, then the diagram is of the same type. Note that if w_i is parallel to v_i (ie M has an $S^2 \times S^1$ summand) then there are no w_i arcs in $\Gamma(\{v_1, v_2\})$.

Definition 1.3.2 (1) A *standard wave* ζ with respect to a closed curve κ on an orientable surface Σ is a simple arc such that $\zeta \cap \kappa = \partial\zeta$, the arc ζ intersects κ from the same side and ζ is not homotopic into a subarc of κ rel $\partial\zeta$.

(2) When Σ is the boundary of a genus-two handlebody V with a fixed meridian system $\{v_1, v_2\}$, a standard wave ζ with respect to v_i is called an *s-wave* if ζ is disjoint from v_j for $j \neq i$, and $\Sigma \setminus (v_i \cup \zeta)$ is connected.

(3) In a Heegaard diagram $(\Sigma, \widehat{V}, \widehat{W})$ of some 3-manifold M, where $\widehat{V}$ and $\widehat{W}$ are complete sets of meridians, an s-wave $\eta \subset \Sigma$ with respect to a meridian $v_i \subset \widehat{V}$ is called a *wave* with respect to v_i if $\eta \cap \widehat{W} = \varnothing$ or η is a subarc of $\widehat{W}$.

For example, in the Whitehead graph depicted in Figure 2, right, the edges labeled b correspond to subarcs of $w_1 \cup w_2$ that are waves with respect to v_1.

Remark 1.3.3 The configurations in Figure 2 are symmetric under the hyperelliptic involution π. Therefore, if there is a wave η with endpoints in v_i^+, there must be a wave $\eta' = \pi(\eta)$ with endpoints in v_i^-; for example, see the two edges labeled b in Figure 2, right. Moreover, η and η' are attached to opposite sides of v_i and hence are distinct waves.

Definition 1.3.4 Let $\widehat{V} = \{v_1, v_2\}$ be as above. If there exists an s-wave η with respect to v_i, each of the three boundary components of $\mathcal{N}(v_i \cup \eta)$ bounds a disk in the handlebody V. Since V is of genus two, we necessarily have two of the boundary components isotopic to v_1 and v_2 respectively, while the third component is the boundary of a new nonseparating disk. We call the operation of replacing the disk V_i in the complete meridian set $\widetilde{V}$ by the new disk the *wave move* of v_i along η. In this paper we often view the wave move as a two-step operation: The first step is a surgery on v_i which connects the endpoints of $v_i \setminus \mathcal{N}(\partial\eta)$ using two parallel copies of η, thus obtaining two curves. Note that one of these curves is parallel to v_j where $j \neq i$. In the second step we delete the curve parallel to v_j where $j \neq i$, and the remaining curve is the boundary of the new disk obtained from the wave move.

Remark 1.3.5 $\mathcal{N}(v_i \cup \eta)$ in Definition 1.3.4 is a pair of pants in Σ with two of its boundary components isotopic to v_1 and v_2, respectively. So the new curve obtained by the wave move can also be obtained by a band sum of v_1 and v_2 along an arc in $\mathcal{N}(v_i \cup \eta)$.

Definition 1.3.6 Meridian sets $\widehat{V} = \{v_1, v_2\}$ and $\widehat{W} = \{w_1, w_2\}$ for a genus-two Heegaard splitting of S^3 are *standard* if $|v_1 \cap w_1| = |v_2 \cap w_2| = 1$ and $v_1 \cap w_2 = v_2 \cap w_1 = \varnothing$. Meridian sets $\widehat{V} = \{v_1, v_2\}$ and $\widehat{W} = \{w_1, w_2\}$ for a genus-two Heegaard splitting of $(S^1 \times S^2) \# L(r, s)$ are *standard* if

(1) the curves v_1 and w_1 are parallel and disjoint from $v_2 \cup w_2$, and

(2) there is a separating essential simple closed curve in Σ disjoint from $v_1 \cup v_2 \cup w_1 \cup w_2$.

Note that a separating curve in (2) bounds disks in both handlebodies and the two disks form a separating essential 2-sphere in M. A genus-two Heegaard diagram of S^3 or $(S^1 \times S^2) \# L(r, s)$ is *standard* if the corresponding meridian sets are standard.

We are now ready to state the following theorems, which play a crucial role in the argument. The main theorem of Homma, Ochiai and Takahashi in [21] states:

Theorem 1.3.7 *Any genus-two Heegaard diagram for S^3 either is standard or contains a wave.*

For $M = (S^1 \times S^2) \# L(r,s)$, we have the following theorem of Negami and Okita [26]:

Theorem 1.3.8 *Any genus-two Heegaard diagram of $(S^1 \times S^2) \# L(r,s)$ is either standard or contains a wave.*

Lemma 1.3.9 *Suppose $\widetilde{V} = \{V_1, V_2\}$ and $\widetilde{W} = \{W_1, W_2\}$ are complete meridian sets for a genus-two Heegaard splitting (V, W) for a 3-manifold M, and the Whitehead graph of $\Gamma(\widehat{V})$ has edges connecting the vertices as in one of the following cases*:

(1) *There are edges connecting v_1^+ to v_1^- and v_2^+ to v_2^-.*

(2) *There are edges connecting both v_2^+ and v_2^- to the same vertex v_1^+ or v_1^-.*

(3) *There are edges connecting both v_1^+ and v_1^- to the same vertex v_2^+ or v_2^-.*

Then the Heegaard diagram $(\widehat{V}, \widehat{W})$ contains no waves with respect to $\widehat{V}$.

Proof Since this is a Heegaard diagram of a 3-manifold, Theorem 1 of [28] applies and the Whitehead graph is as in Figure 2. For case (1), the Heegaard diagram must be of the type in Figure 2, left, with $c > 0$ and $d > 0$. Note that a wave with respect to v_j must separate v_i^+ from v_i^- for $i \neq j$; see the edge labeled b in Figure 2, right, for an example of a wave. So an edge connecting v_i^+ to v_i^- prevents any wave with respect to v_j, for $i \neq j$ with $i, j \in \{1, 2\}$. Hence case (1) holds.

Cases (2) and (3) are symmetric. For case (2), the Heegaard diagram must be of the type shown at left and center of Figure 2, with $a > 0$ and $b > 0$. Similar to case (1), since the edges labeled a and b in Figure 2, left and center, connect v_1^+, v_1^-, v_2^+ and v_2^- together, there cannot be any wave with respect to v_1 and v_2. □

Definition 1.3.10 Throughout the paper, for any different vertices $v_i^\pm$ and $v_j^\pm$, we use $[v_i^\pm, v_j^\pm]$ to denote an edge of the Whitehead graph $\Gamma(\widehat{V})$ connecting $v_i^\pm$ to $v_j^\pm$. If $[v_i^\pm, v_j^\pm]$ is an edge in either case (1), (2) or (3) of Lemma 1.3.9, we say this edge is a $[v_i^\pm, v_j^\pm]$ *blocking-edges* with respect to V. Lemma 1.3.9 says that, if the Whitehead graph $\Gamma(\widehat{V})$ has both types of blocking-edges in either case (1), (2) or (3) of Lemma 1.3.9, then the Heegaard diagram $(\widehat{V}, \widehat{W})$ contains no waves with respect to $\widehat{V}$. The phrase "blocking-edge" is used in the paper whenever we are searching for edges in case (1), (2) or (3) of Lemma 1.3.9. We would like to emphasize that one type of blocking-edge alone does not "block" waves, and we need both types in case (1), (2) or (3) of Lemma 1.3.9. The corresponding statement of Lemma 1.3.9 also holds for the Whitehead graph $\Gamma(\widehat{W})$ and we can similarly define *blocking-edges* with respect to W.

For example, if two adjacent intersection points of $w_i \cap (v_1 \cup v_2)$ along w_i both belong to v_1 and have the same sign of intersection, then the subarc of w_i between these two points is a $[v_1^+, v_1^-]$ blocking-edge. Moreover, if there is a similar subarc of w_j with endpoints in v_2 and having the same sign of intersection, then we also have a $[v_2^+, v_2^-]$ blocking-edge. By case (1) of Lemma 1.3.9, this means that the Heegaard diagram has no wave with respect to V.

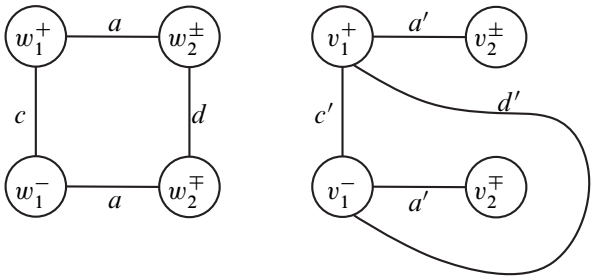

Figure 3: Dual Whitehead graphs, case (1): $\Gamma(\widehat{W})$, left, and $\Gamma(\widehat{V})$, right.

We finish this section with the following lemma on the Whitehead graphs.

Lemma 1.3.11 *Let $\widehat{V} = \{v_1, v_2\}$ and $\widehat{W} = \{w_1, w_2\}$ be complete sets of meridians. Suppose $v_1 \cup v_2$ contains a $[w_1^+, w_1^-]$ blocking-edge and a $[w_2^+, w_2^-]$ blocking-edge. Suppose there is a wave with respect to $\{v_1, v_2\}$ and suppose $\Gamma(\widehat{W})$ is a connected graph.*

(1) *If no subarc of $w_1 \cup w_2$ is a wave with respect to $\{v_1, v_2\}$, then the Whitehead graphs $\Gamma(\widehat{W})$ and $\Gamma(\widehat{V})$ are as shown in Figure 3, where all the labels are nonzero.*

(2) *If a subarc of w_1 or w_2 is a wave with respect to $\{v_1, v_2\}$, then the Whitehead graphs $\Gamma(\widehat{W})$ and $\Gamma(\widehat{V})$ are as shown in Figure 4, where all the labels are nonzero.*

Proof Since $v_1 \cup v_2$ contains a $[w_1^+, w_1^-]$ blocking-edge and a $[w_2^+, w_2^-]$ blocking-edge, the Whitehead graph $\Gamma(\widehat{W})$ must be of the type in Figure 2, left, with $c \geq 1$ and $d \geq 1$. If $a = b = 0$ in Figure 2, left, then the Whitehead graph $\Gamma(\widehat{W})$ is disconnected, a contradiction to the hypothesis. Hence at least one of a and b in Figure 2, left, is nonzero. If $b = 0$ and $a \neq 0$ in Figure 2, left, then $\Gamma(\widehat{W})$ is exactly the left picture of Figure 3. If $a = 0$ and $b \neq 0$ in Figure 2, left, then by interchanging the locations of v_2^+ and v_2^- in Figure 2, left, $\Gamma(\widehat{W})$ is isotopic to the left picture of Figure 3. If neither a or b is zero in Figure 2, left, then $\Gamma(\widehat{W})$ is the left picture of Figure 4.

If $\Gamma(\widehat{W})$ is as in the left picture of Figure 3, then $\Sigma \setminus (v_1 \cup v_2 \cup w_1 \cup w_2)$ consists of two octagons and a collection of quadrilaterals (ie the regions between parallel arcs of the same edge type). If $\Gamma(\widehat{W})$ is as in the left picture of Figure 4, then $\Sigma \setminus (v_1 \cup v_2 \cup w_1 \cup w_2)$ consists of four hexagons and a collection of quadrilaterals.

Next, consider the Whitehead graph $\Gamma(\widehat{V})$. Since there is a wave with respect to $\{v_1, v_2\}$ and since the diagram is connected, the Whitehead graph $\Gamma(\widehat{V})$ must be of type shown in Figure 2, right, with $a \neq 0$ and at least one of c and d nonzero. Without loss of generality, suppose $c \neq 0$.

If $b = d = 0$ in Figure 2, right, then the picture of Figure 2(iii) indicates that $\Sigma \setminus (v_1 \cup v_2 \cup w_1 \cup w_2)$ consists of one 12-gon and a collection of quadrilaterals. This configuration is incompatible with the two possible configurations of $\Gamma(\widehat{W})$ discussed above. So this cannot happen and at least one of b and d in Figure 2, right, is nonzero.

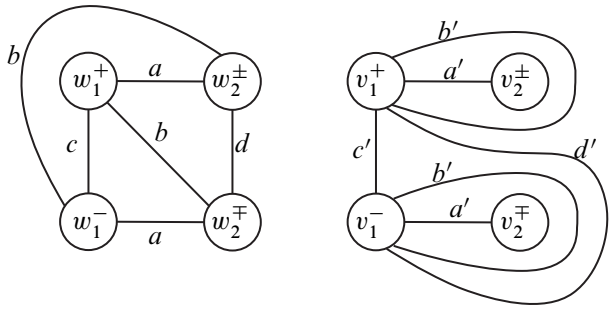

Figure 4: Dual Whitehead graphs, case (2): $\Gamma(\widehat{W})$, left, and $\Gamma(\widehat{V})$, right.

If $b \neq 0$ and $d = 0$ in Figure 2, right, then $\Sigma \setminus (v_1 \cup v_2 \cup w_1 \cup w_2)$ consists of one octagon, two hexagons and a collection of quadrilaterals. This configuration is also incompatible with the two possible configurations of $\Gamma(\widehat{W})$ discussed above. So this cannot happen either.

If $b = 0$ and $d \neq 0$ in Figure 2, right, then $\Gamma(\widehat{V})$ is as in the right picture of Figure 3, in which case $\Sigma \setminus (v_1 \cup v_2 \cup w_1 \cup w_2)$ consists of two octagons and a collection of quadrilaterals. By the description of $\Gamma(\widehat{W})$ above, this means that $\Gamma(\widehat{W})$ must be as in the left picture of Figure 3. Note that an edge in $\Gamma(\widehat{V})$ is a wave if and only if $b \neq 0$ in Figure 2, right. In this configuration, since no edge in $\Gamma(\widehat{V})$ is a wave, no subarc of w_1 and w_2 is a wave with respect to $\{v_1, v_2\}$. This is case (1) of the lemma.

If $b \neq 0$ and $d \neq 0$ in Figure 2, right, then $\Gamma(\widehat{V})$ is as in the right picture of Figure 4, in which case $\Sigma \setminus (v_1 \cup v_2 \cup w_1 \cup w_2)$ consists of four hexagons and a collection of quadrilaterals. By the discussion of $\Gamma(\widehat{W})$ above, this means that $\Gamma(\widehat{W})$ must be as in the left picture of Figure 4. Moreover, since $b \neq 0$ in Figure 2, right, $b' \neq 0$ in Figure 4 and hence a subarc of w_1 or w_2 (ie edge marked b' in Figure 4) is a wave with respect to $\{v_1, v_2\}$. This is case (2) of the lemma. □

2 The planar surface and the curves α, γ, δ and ε

2.1 Fixing the decomposition of $\widehat{\Sigma}$

As in Section 1.2, we have a genus-two Heegaard splitting (U, W) of $M \setminus \mathcal{N}(K)$ along the Heegaard surface Σ. The Heegaard splitting extends to a Heegaard splitting (V, W) of M with $V = U \cup \mathcal{N}(K)$. Given a $(\mathcal{P}, \mathcal{D})$-pair (P, D) as in Section 1.2, the Dehn filling along the slope of $\partial_- P$ yields a lens space L. The union of U and the surgery solid torus is a genus-two handlebody V' in L, and (V', W) is a genus-two Heegaard splitting of the lens space L. Moreover, (P, D) extends to a stabilizing pair of disks for the Heegaard splitting (V', W) of L. This implies that there is a nonseparating disk E in W which is disjoint from $P \cup D$. So $\{D, E\}$ is a complete meridian set for W. Let $\delta = \partial D$ and $\varepsilon = \partial E$. Hence $\varepsilon \cap \partial_+ P = \varnothing$ and $\delta \cap \partial_+ P$ is a single point. In Section 1.2, we denote by γ the curve which bounds the nonseparating disk $C \subset U$, and let $\alpha = \partial_+ A$, where A is a vertical annulus and the slope of $\partial_- A$ is the ∞-slope. Thus

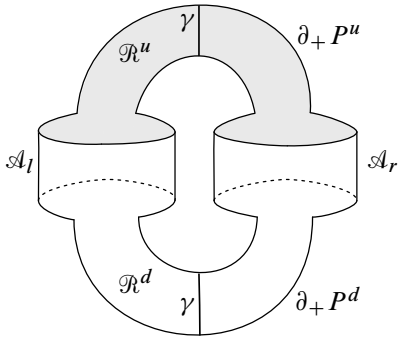

Figure 5: The decomposition of $\widehat{\Sigma}$ into two annuli and two rectangles.

both γ and α bound disks in the handlebody V in M and $\{\gamma, \alpha\}$ is a complete set of meridians for V. In this paper, we will study the Heegaard diagram of M formed by $\widehat{V} = \{\gamma, \alpha\}$ and $\widehat{W} = \{\delta, \varepsilon\}$.

Let $\widehat{\Sigma}$ be the surface obtained by cutting Σ open along $\partial_+ P$, ie $\widehat{\Sigma} = \Sigma \setminus \mathcal{N}(\partial_+ P)$. Throughout the paper, we often view $\widehat{\Sigma}$ as the closure of $\Sigma \setminus \partial_+ P$ under the path metric. As $\partial_+ P$ is nonseparating, $\widehat{\Sigma}$ is a twice-punctured torus. Fix a rectangle-annulus decomposition of this twice-punctured torus as a union of two annuli $\mathcal{A}_r$ and $\mathcal{A}_l$ connected by two rectangles $\mathcal{R}^u$ and $\mathcal{R}^d$, as in Figure 5; see Lemma 2.1.1 below for detailed requirements. The two boundary components of this twice punctured torus are denoted by $\partial_+ P^u$ and $\partial_+ P^d$ (the upper and lower components as depicted in Figure 5). The core curves of the annuli are denoted by $\mathfrak{a}_r$ and $\mathfrak{a}_l$, respectively.

Arcs in the rectangles $\mathcal{R}^u$ and $\mathcal{R}^d$ that are properly embedded and non-∂-parallel in $\widehat{\Sigma}$ will be called *cocore arcs in the rectangles*; eg the arcs marked γ in Figure 5 represent components of $\gamma \cap \widehat{\Sigma}$ that are cocore arcs of rectangles $\mathcal{R}^u$ and $\mathcal{R}^d$. An arc in $\mathcal{R}^u$ or $\mathcal{R}^d$ connecting $\mathcal{A}_l$ to $\mathcal{A}_r$ will be called a *core arc* of the rectangles. An arc properly embedded in $\mathcal{A}_l$ or $\mathcal{A}_r$ connecting $\partial_+ P^u$ to $\partial_+ P^d$ will be called a cocore of the annulus.

In the remaining part of Sections 2.1 and 2.2 we prove, using the above definition of complexity $c(P, D, \alpha, \gamma)$, a series of technical lemmas which are dedicated to a careful analysis of possible intersections of the curves α, γ, δ and ε with the surface $\widehat{\Sigma}$.

Lemma 2.1.1 *The rectangle-annulus decomposition of $\widehat{\Sigma}$ in Figure 5 can be constructed so that each segment of $(\alpha \cup \gamma) \cap \widehat{\Sigma}$ is a cocore arc of a rectangle or an annulus. Furthermore, in each of the rectangles $\mathcal{R}^u$ and $\mathcal{R}^d$, there is at least one γ segment which is a cocore arc for the rectangle.*

Proof Consider the disk $C \subset U$ ($\gamma = \partial C$). Assume $|P \cap C|$ is minimal up to isotopy. Since P is incompressible, $P \cap C$ contains no closed curve. Since P is not an annulus by assumption, $P \cap C \neq \varnothing$ by Lemma 1.2.4. Let ρ be an arc in $P \cap C$ that is outermost in C. The arc ρ cuts off a subdisk Δ_ρ of C with $\Delta_\rho \cap P = \rho$. As $|P \cap C|$ is minimal, Δ_ρ is a ∂-compressing disk for P. Let $\kappa_1 = \partial \Delta_\rho \setminus \mathrm{int}(\rho)$.

So $\kappa_1 \subset \gamma = \partial C$. As M is orientable and since the two endpoints of κ_1 are connected by the arc ρ in C, $\partial\kappa_1$ is a pair of points of $\gamma \cap \partial_+ P$ with opposite signs of intersection. Thus, we may view κ_1 as a properly embedded arc in $\widehat{\Sigma}$ with both endpoints of κ_1 on the same boundary curve, say $\partial_+ P^d$. As γ is a closed curve we are guaranteed to have another arc κ_2 of $\gamma \cap \widehat{\Sigma}$ in $\widehat{\Sigma}$ on the other side of P with both endpoints in $\partial_+ P^u$ (one can also see the existence of κ_2 using the involution π).

For each $i = 1, 2$, consider a regular neighborhood $\mathcal{N}(\kappa_i)$ of κ_i in $\widehat{\Sigma}$. We isotope all the arcs of $(\alpha \cup \gamma) \cap \widehat{\Sigma}$ that are parallel to κ_i into $\mathcal{N}(\kappa_i)$ for $i = 1, 2$, and set $\mathcal{R}^d = \mathcal{N}(\kappa_1)$ and $\mathcal{R}^u = \mathcal{N}(\kappa_2)$. We may set all the arcs of $(\alpha \cup \gamma) \cap \widehat{\Sigma}$ in $\mathcal{R}^d$ and $\mathcal{R}^u$ as cocore arcs of the rectangles.

As $\mathcal{N}(\kappa_1)$ and $\mathcal{N}(\kappa_2)$ are disjoint in the twice-punctured torus $\widehat{\Sigma}$, the complement of $\mathcal{R}^d \cup \mathcal{R}^u$ is a pair of annuli, which we denote by $\mathcal{A}_l$ and $\mathcal{A}_r$. Moreover, each boundary curve of $\mathcal{A}_l$ and $\mathcal{A}_r$ contains exactly one subarc of $\partial\mathcal{R}^d$ and $\partial\mathcal{R}^u$, as depicted in Figure 5. Since all the arcs of $(\alpha \cup \gamma) \cap \widehat{\Sigma}$ that are parallel to κ_i are in $\mathcal{R}^d \cup \mathcal{R}^u$, this implies that the remaining arcs of $(\alpha \cup \gamma) \cap \widehat{\Sigma}$ are nonintersecting arcs in $\mathcal{A}_l$ and $\mathcal{A}_r$ with endpoints in different components of $\partial\mathcal{A}_l$ and $\partial\mathcal{A}_r$. Thus we may set all the arcs of $(\alpha \cup \gamma) \cap \widehat{\Sigma}$ in $\mathcal{A}_l$ and $\mathcal{A}_r$ as cocore arcs of the annuli. This gives a decomposition as required. □

Remark 2.1.2 As in the proof of Lemma 2.1.1, if we boundary compress P along Δ_ρ, we obtain two planar surfaces P_l and P_r. By the construction in the proof, $\partial_+ P_l$ and $\partial_+ P_r$ are isotopic to the two core curves $\mathfrak{a}_l$ and $\mathfrak{a}_r$ of the annuli $\mathcal{A}_l$ and $\mathcal{A}_r$, respectively. In particular, $\partial_+ P_l$ and $\partial_+ P_r$ are nonseparating curves in Σ. If P_l is a disk, then since C is the only nonseparating essential disk in the compression body U (up to isotopy), P_l must be parallel to C, which contradicts the assumption that Δ_ρ is a ∂-compressing disk for P. Thus neither P_l nor P_r is a disk.

Remark 2.1.3 (1) As $\varepsilon \cap \partial_+ P = \varnothing$ the curve ε is contained in the interior of $\widehat{\Sigma}$.

(2) Since the hyperelliptic involution π leaves each simple closed curve in Σ invariant up to isotopy, we may suppose δ and $\partial_+ P$ are invariant under π. Thus $X = \delta \cap \partial_+ P$ is a fixed point of π. Moreover we may suppose the two annuli $\mathcal{A}_l$ and $\mathcal{A}_r$ are invariant under π and hence π interchanges the two rectangles $\mathcal{R}^d$ and $\mathcal{R}^u$.

Lemma 2.1.4 *Each of the intersections $\delta \cap \mathcal{R}^u$ and $\delta \cap \mathcal{R}^d$ contains at least two arcs that are core arcs of the corresponding rectangle.*

Proof Without loss of generality, suppose that at most one component of $\delta \cap \mathcal{R}^d$ is a core arc of $\mathcal{R}^d$. Note that a component of $\delta \cap \mathcal{R}^d$ is either a core arc of $\mathcal{R}^d$ or an arc connecting a boundary edge of $\mathcal{R}^d$ to the intersection point $X = \delta \cap \partial_+ P$. This means that, after an isotopy in $\mathcal{R}^d$, a cocore arc of $\mathcal{R}^d$ intersects δ at most once. In the proof of Lemma 2.1.1, the cocore arc κ_1 of $\mathcal{R}^d$ is a boundary arc of a ∂-compressing disk Δ_ρ for P. So we may isotope κ_1 in $\mathcal{R}^d$ so that κ_1 intersects δ at most once. Since $P \cap \delta = X$ is a single point, the assumption on $\kappa_1 \cap \delta$ implies that the surface obtained by the ∂-compression along Δ_ρ intersects δ in at most three points: the point X and the two copies of $\kappa_1 \cap \delta$ created by the boundary compression.

As in Remark 2.1.2, the two planar surfaces obtained by ∂-compressing P along Δ_ρ are P_l and P_r. So $(P_l \cup P_r) \cap \delta$ consists of at most 3 points. Moreover, $(P_l \cup P_r) \cap \delta$ consists of exactly 3 points if and only if $\kappa_1 \cap \delta \neq \varnothing$, in which case $P_l \cap \delta \neq \varnothing$ and $P_r \cap \delta \neq \varnothing$. This implies that either $P_l \cap \delta$ or $P_r \cap \delta$ is a single point; say, $P_l \cap \delta$. Thus (P_l, D) forms a new $(\mathcal{P}, \mathcal{D})$-pair.

By Remark 2.1.2, $\partial_+ P_l$ is isotopic to the core curve $\mathfrak{a}_l$ of the annulus $\mathcal{A}_l$. We may isotope $\partial_+ P_l$ to the upper boundary component of $\mathcal{A}_l$. After this isotopy, any intersection point of $\alpha \cup \gamma$ with $\partial_+ P_l$ is an intersection point with $\partial_+ P$. Therefore, any segment of $(\alpha \cup \gamma) \setminus \partial_+ P_l$ is composed of a nonzero number of segments of $\alpha \cup \gamma$ in $\widehat{\Sigma}$. It follows that $c(P_l, D, \alpha, \gamma) < c(P, D, \alpha, \gamma)$, contradicting Assumption 1.2.3. □

Lemma 2.1.5 *Let P, D and $\widehat{V} = \{\gamma, \alpha\}$ be as above. Then there is no s-wave with respect to α that is disjoint from δ.*

Proof Suppose, for the sake of contradiction, that there is such an s-wave η with $\partial\eta \subset \alpha$ and $\eta \cap \delta = \varnothing$.

Let α_1 and α_2 be the two components of $\alpha \setminus \mathcal{N}(\partial\eta)$. We define $c_0(\alpha_i)$ for $i = 1, 2$ to be the number of segments of $\alpha_i \setminus \partial_+ P$ that intersect δ. This implies that $c_0(P, D, \alpha) + 2 \geq c_0(\alpha_1) + c_0(\alpha_2)$ (see Definition 1.2.1) because a segment of $\alpha \setminus \partial_+ P$ may be split into two segments by $\partial\eta$,

Recall that the first step of a wave move along η is to connect the two endpoints of α_1 and the two endpoints of α_2 by a pair of arcs parallel to η. Denote the two resulting closed curves by α_1' and α_2' respectively. Since $\eta \cap \delta = \varnothing$ for $i \in \{1, 2\}$ we have $c_0(\alpha_i) \geq c_0(P, D, \alpha_i')$. Therefore, $c_0(P, D, \alpha) + 2 \geq c_0(P, D, \alpha_1') + c_0(P, D, \alpha_2')$.

Because Σ is of genus two, one of the two closed curves α_1' and α_2', obtained after the first step of the wave move, is isotopic to γ. Say it is α_2'. By Lemma 2.1.4 and the construction of $\mathcal{R}^u$ and $\mathcal{R}^d$, $c_0(P, D, \alpha_2') = c_0(P, D, \gamma) \geq 2$. This plus the inequality above implies that $c_0(P, D, \alpha) \geq c_0(P, D, \alpha_1')$. Recall that in the case of a genus-two surface, a curve obtained by a wave move can also be obtained by a band sum. So α_1' can be obtained by a band sum of α and γ. Since $\alpha = \partial_+ A$ and $\gamma = \partial C$, α_1' is a boundary curve of an annulus obtained by a band sum of A and C. Moreover, $|\alpha \cap \delta| = |\alpha_1' \cap \delta| + |\alpha_2' \cap \delta|$, and by Lemma 2.1.4, $\gamma \cap \delta = \alpha_2' \cap \delta \neq \varnothing$. Thus $|\alpha_1' \cap \delta| < |\alpha \cap \delta|$ and $c(P, D, \alpha_1', \gamma) < c(P, D, \alpha, \gamma)$, contradicting Assumption 1.2.3. □

Lemma 2.1.6 *Any s-wave with respect to $\widehat{V} = \{\gamma, \alpha\}$ must intersect $\partial_+ P$.*

Proof As in the proof of Lemma 1.2.4, since P is incompressible and $U \setminus \mathcal{N}(C) \cong T^2 \times I$, each component of $P \setminus \mathcal{N}(C)$ is either a disk or a vertical annulus in $U \setminus \mathcal{N}(C)$. Let ρ' be an arc in the intersection $P \cap C$ which is outermost in P. So ρ' cuts off a subsurface Q' of P with $Q' \cap C = \rho'$. Since $|P \cap C|$ is assumed to be minimal up to isotopy, Q' cannot be a disk. Hence Q' is a vertical annulus in $U \setminus \mathcal{N}(C)$. Now consider the annulus A with $\partial_+ A = \alpha$. As $C \cap A = \varnothing$ and since $\partial_- P$ has integer slope, after isotopy, $Q' \cap A$ is a single vertical arc, and thus $\partial_+ Q' \cap \alpha$ is a single point. Note that $\partial_+ Q'$

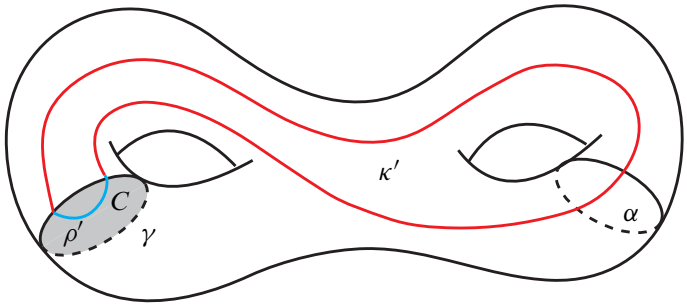

Figure 6: Subarcs of $\partial_+ P$ forming blocking-edges for s-wave of $\{\gamma, \alpha\}$.

is the union of ρ' and a subarc κ' of $\partial_+ P$; see Figure 6. So κ' is an arc properly embedded in $\Sigma \setminus \mathcal{N}(\gamma)$ intersecting α exactly once and having both endpoints in the same boundary component of $\Sigma \setminus \mathcal{N}(\gamma)$.

Now cut Σ open along γ and α and obtain a four-holed sphere. Denote the boundary curves of the four-holed sphere by γ^+, γ^-, α^+ and α^-. The restriction of $\partial_+ P$ to this four-holed sphere is a collection of arcs. It follows that the arc κ' is cut into two subarcs in the four-holed sphere connecting α^+ and α^- to the same curve, say γ^+. By the symmetry induced from the involution π, the curve $\partial_+ P$ also contains subarcs in the four-holed sphere connecting both α^+ and α^- to γ^-. Similar to the proof of cases (2) and (3) of Lemma 1.3.9, these two subarcs of $\partial_+ P$ "block" any s-wave with respect to $\widehat{V} = \{\gamma, \alpha\}$. In other words, any s-wave with respect to $\widehat{V} = \{\gamma, \alpha\}$ must intersect $\partial_+ P$. □

Corollary 2.1.7 *The curves δ and ε do not contain subarcs which are waves with respect to α or γ.*

Proof As above we denote the boundary curves of the four-holed sphere $\Sigma \setminus \mathcal{N}(\alpha \cup \gamma)$ by γ^+, γ^-, α^+ and α^-. If a subarc η of $\delta \cup \varepsilon$ is a wave with $\partial\eta$ in $\gamma^\pm$ (or $\alpha^\pm$), then by the symmetry from the hyperelliptic involution, another subarc η' of $\delta \cup \varepsilon$ is a wave with $\partial\eta'$ in $\gamma^\mp$ (or $\alpha^\mp$). By Lemma 2.1.6, both η and η' must intersect $\partial_+ P$, which means that $\delta \cup \varepsilon$ has at least two intersection points with $\partial_+ P$. This is a contradiction because $\varepsilon \cap P = \varnothing$ and $|\delta \cap P| = 1$. □

Let $\{\delta, \varepsilon'\}$ be a complete set of meridians for W, where ε' may not be the same as ε. Fix an orientation for each of α, γ, δ and ε'. These orientations induce a $\pm$-sign for each point of intersection of the curves.

Proposition 2.1.8 *Let $\alpha, \gamma, \delta, \varepsilon'$ be as above. Consider the Heegaard diagram given by $\{\delta, \varepsilon'\}$ and $\{\alpha, \gamma\}$. Suppose the Heegaard diagram is not a standard Heegaard diagram of S^3 or $(S^2 \times S^1) \# L(r, s)$. Then*

(1) *Suppose there is no wave with respect to $\{\delta, \varepsilon'\}$. Then all the intersection points, of δ with each curve of $\{\alpha, \gamma\}$, have the same sign. Moreover, if no subarc of ε' is a wave with respect to $\{\alpha, \gamma\}$, then the intersection points of ε' with each curve of $\{\alpha, \gamma\}$ all have the same sign.*

(2) *If $\varepsilon' = \varepsilon$, then δ and ε admit orientations so that the intersection points of $\delta \cup \varepsilon$ with each component of $(\alpha \cup \gamma) \setminus \partial_+ P$ all have the same sign.*

Proof Assume for the sake of contradiction that the statement of part (1) of the proposition is false and that δ intersects a curve $v_1 \in \{\alpha, \gamma\}$ with different signs. In particular, δ has a subarc connecting one side of v_1, say the plus side, to the same side. Denote an innermost such subarc by τ, so $\tau \cap v_1 = \partial\tau$ and τ connects the plus side of v_1 to the same side. If τ does not have other intersection points with the other curve $v_2 \in \{\alpha, \gamma\}$ in its interior, then τ is a wave with respect to $v_1 \in \{\alpha, \gamma\}$, contradicting Corollary 2.1.7. Hence τ must intersect the other curve $v_2 \in \{\alpha, \gamma\}$ in its interior. These intersection points of τ with v_2 must have the same sign (otherwise τ contains a subarc that is a wave with respect to $v_2 \in \{\alpha, \gamma\}$ contradicting Corollary 2.1.7).

Denote the two sides of v_i by $v_i^{\pm}$. The arc τ was chosen to connect v_1^+ to v_1^+. The curve v_2 divides τ into a collection of subarcs. Since all the points of $\tau \cap v_2$ have the same sign, the two subarcs of τ containing $\partial\tau$ must be arcs connecting v_1^+ to v_2^+ and v_1^+ to v_2^- respectively, in other words, the two arcs are $[v_1^+, v_2^+]$ and $[v_1^+, v_2^-]$ blocking-edges. Hence, by Lemma 1.3.9, there can be no wave with respect to $\widehat{V} = \{\alpha, \gamma\}$. Since, by the hypothesis, there is no wave with respect to $\{\delta, \varepsilon'\}$, there is no wave in the Heegaard diagram given by $\{\delta, \varepsilon'\}$ and $\{\alpha, \gamma\}$, contradicting Theorems 1.3.7 and 1.3.8. Thus all the intersection points of δ with each curve of $\{\alpha, \gamma\}$ must have the same sign.

If no subarc of ε' is a wave with respect to $\{\alpha, \gamma\}$, then the same proof also works for ε' and hence the intersection points of ε' with each curve of $\{\alpha, \gamma\}$ all have the same sign. This proves part (1) of the proposition.

To prove part (2), we first show that the intersection points of each curve of $\{\delta, \varepsilon\}$ with any component of $(\alpha \cup \gamma) \setminus \partial_+ P$ all have the same sign: Suppose, to the contrary, that a curve $w_1 \in \{\delta, \varepsilon\}$ has opposite signs at intersection points with a segment h of $(\alpha \cup \gamma) \setminus \partial_+ P$. Then a subarc of h connects one side of w_1, say the plus side, to the same side. Let τ' be an innermost such subarc. So $\tau' \cap w_1 = \partial\tau'$.

First suppose that τ' does not intersect the other curve $w_2 \in \{\delta, \varepsilon\}$ in its interior, then τ' is a wave with respect to $w_1 \in \{\delta, \varepsilon\}$ and is disjoint from $\partial_+ P$. If $w_1 = \delta$, since $\tau' \cap \partial_+ P = \varnothing$, a wave move on δ along τ' yields a new meridional curve δ' of W that intersects $\partial_+ P$ in a single point. Let D' be the disk in W bounded by δ'. So (P, D') is a $(\mathcal{P}, \mathcal{D})$-pair. However, it is clear from the wave move that the complexity satisfies $c(P, D', \alpha, \gamma) < c(P, D, \alpha, \gamma)$. This contradicts Assumption 1.2.3. If $w_1 = \varepsilon$ then a wave move on ε along τ' yields a meridional curve of W that is disjoint from $\delta \cup \partial_+ P$ and not parallel to ε, which is impossible in the genus-two surface Σ.

Thus τ' cannot be a wave and hence τ' must contain intersection points with the other curve $w_2 \in \{\delta, \varepsilon\}$ in its interior. By applying the argument used on the arc τ as above to τ', we conclude that the arc τ' contains two subarcs that are $[w_1^+, w_2^+]$ and $[w_1^+, w_2^-]$ blocking-edges. Thus, by Lemma 1.3.9, there can be no wave with respect to $\{\delta, \varepsilon\}$. By Corollary 2.1.7, no subarc of ε is a wave with respect to $\{\alpha, \gamma\}$. Hence, it follows from part (1) of this proposition that the intersection points of w_1 with each curve of $\{\alpha, \gamma\}$ all have the same sign, a contradiction to our assumption on h at the beginning. This proves that the intersection points of each curve of $\{\delta, \varepsilon\}$ with any component of $(\alpha \cup \gamma) \setminus \partial_+ P$ all have the same sign.

This conclusion implies that part (2) of the proposition holds if no component of $(\alpha\cup\gamma)\setminus\partial_+ P$ intersects both δ and ε. Assume therefore, for the sake of contradiction, that h is a component of $(\alpha\cup\gamma)\setminus\partial_+ P$ which intersects both δ and ε. By the conclusion above, we may fix orientations for δ and ε so that the intersection points of $\delta\cup\varepsilon$ with h all have the same sign. So part (2) of the proposition is true unless there is another component h' of $(\alpha\cup\gamma)\setminus\partial_+ P$ such that points in $h'\cap\delta$ and $h'\cap\varepsilon$ have opposite signs. Suppose there is such an arc h'. A subarc of h' has one endpoint in δ and the other endpoint in ε. As they have opposite signs, this subarc is either a $[\delta^+,\varepsilon^+]$ or a $[\delta^-,\varepsilon^-]$ edge. Similarly, since the intersection points of $\delta\cup\varepsilon$ with h all have the same sign, a subarc of h is a $[\delta^-,\varepsilon^+]$ or $[\delta^+,\varepsilon^-]$ edge. Thus $h\cup h'$ contains two subarcs that either connect both δ^+ and δ^- to the same $\varepsilon^\pm$ or connect both ε^+ and ε^- to the same $\delta^\pm$. It now follows from (2) and (3) of Lemma 1.3.9 that the Heegaard diagram contains no wave with respect to $\{\delta,\varepsilon\}$.

By Corollary 2.1.7 no subarc of ε is a wave. So part (1) of this proposition implies that h and h' must belong to different curves of $\{\alpha,\gamma\}$. Without loss of generality, suppose $h\subset\alpha$, $h'\subset\gamma$. Moreover, by part (1) of the proposition, we may choose the orientation of δ and ε so that

(a) points in $\alpha\cap(\delta\cup\varepsilon)$ all have positive signs,

(b) points in $\gamma\cap\delta$ have positive signs, and

(c) points in $\gamma\cap\varepsilon$ have negative signs.

Similar to the argument on h and h' above, this means that a subarc of δ is an $[\alpha^+,\gamma^-]$ or $[\alpha^-,\gamma^+]$ edge and a subarc of ε is an $[\alpha^+,\gamma^+]$ or $[\alpha^-,\gamma^-]$ edge. Again by (2) and (3) of Lemma 1.3.9, the Heegaard diagram contains no wave with respect to $\{\alpha,\gamma\}$. Hence the Heegaard diagram has no wave, contradicting Theorems 1.3.7 and 1.3.8. □

Lemma 2.1.9 *Let α, γ, δ, ε' be as in Proposition 2.1.8. Consider the Heegaard diagram formed by $\widehat{V}=\{\gamma,\alpha\}$ and $\widehat{W}=\{\delta,\varepsilon'\}$. Suppose the Heegaard diagram is not a standard Heegaard diagram of S^3 or $(S^2\times S^1)\,\#\,L(r,s)$. Suppose that $\alpha\cup\gamma$ contains both $[\delta^+,\delta^-]$ and $[\varepsilon'^+,\varepsilon'^-]$ blocking-edges. Furthermore, suppose that no subarc of ε' is a wave with respect to $\widehat{V}=\{\alpha,\gamma\}$. Then:*

(1) *$\Gamma(\widehat{W})$ is of the type in Figure 2, left, where $c\neq 0$, $d\neq 0$ and exactly one of a, b is 0. In other words, $\Gamma(\widehat{W})$ is as in the left picture of Figure 3. See Figure 7, left, for an example of such Heegaard diagram.*

(2) *Suppose that there are more than one $[\delta^+,\delta^-]$-edges in the Heegaard diagram. Then there is a wave with respect to $\{\alpha,\gamma\}$ connecting a $[\delta^+,\delta^-]$-edge to an $[\varepsilon'^+,\varepsilon'^-]$-edge; see the dashed arc in Figure 7, center.*

Proof First note that since $\Gamma(\widehat{W})$ contains both $[\delta^+,\delta^-]$ and $[\varepsilon'^+,\varepsilon'^-]$ blocking-edges, it follows from case (1) of Lemma 1.3.9 that there is no wave with respect to $\{\delta,\varepsilon'\}$. Moreover, it implies that $\Gamma(\widehat{W})$ must be of the type in Figure 2, left, with $c\neq 0$ and $d\neq 0$.

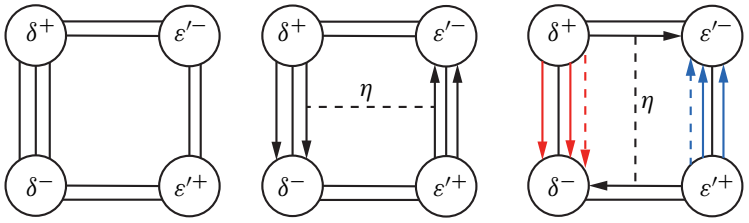

Figure 7: Heegaard diagrams and waves.

If both a and b are 0, then the Whitehead graph is disconnected and there is a circle $\Psi \subset \Sigma$ which is disjoint from $\widehat{V} \cup \widehat{W}$ and which separates the $[\delta^+, \delta^-]$-edges from the $[\varepsilon'^+, \varepsilon'^-]$-edges. Since Ψ is a separating circle in Σ disjoint from $\widehat{V} \cup \widehat{W}$, Ψ bounds separating disks in both handlebodies V and W and the two disks form a separating S^2 in M. This implies that $\Gamma(\widehat{W})$ is a standard Heegaard diagram of a connected sum of two lens spaces $L(c,s) \# L(d,t)$. Since $c > 0$ and $d > 0$, none of $L(c,s)$ and $L(d,t)$ is $S^2 \times S^1$. As $\Gamma(\widehat{W})$ is not the standard Heegaard diagram of S^3 and since M is either S^3 or a connected sum of $S^2 \times S^1$ and a lens space, this is a contradiction.

The argument above implies that the Whitehead graph is connected. Moreover, since there is no wave with respect to $\{\delta, \varepsilon'\}$, by Theorems 1.3.7 and 1.3.8, there must be a wave with respect to $\{\alpha, \gamma\}$.

By Corollary 2.1.7, no subarc of δ is a wave with respect to $\{\alpha, \gamma\}$. Thus, by the hypothesis that no subarc of ε' is a wave with respect to $\{\alpha, \gamma\}$, the conditions in part (1) of Lemma 1.3.11 are satisfied. Hence, $\Gamma(\widehat{W})$ must be as in the left picture of Figure 3. This proves part (1) of the lemma.

Next, we prove part (2) of the lemma. Since no subarc of ε' is a wave, by part (1) of Proposition 2.1.8, we may assign orientations for δ, ε' and γ so that the intersection points of γ with $\delta \cup \varepsilon'$ all have the same sign. In particular, the Heegaard diagram is as in Figure 7, left. Furthermore, the configuration of Figure 7, left, implies that we may choose the orientation for α so that all the intersection points have the same sign.

As illustrated in Figure 7, left, the four curves δ, ε', α and γ cut the Heegaard surface into two octagons and a collections of quadrilaterals. All the $[\delta^+, \delta^-]$-edges are parallel and all the $[\varepsilon'^+, \varepsilon'^-]$-edges are parallel. We say that a $[\delta^+, \delta^-]$-edge or an $[\varepsilon'^+, \varepsilon'^-]$-edge is outermost if it is a boundary edge of an octagon. The hyperelliptic involution π interchanges the two octagons and the two outermost $[\delta^+, \delta^-]$-edges. Furthermore, if there is more than one $[\varepsilon'^+, \varepsilon'^-]$-edge, then π also interchanges the two outermost $[\varepsilon'^+, \varepsilon'^-]$-edges, and if there is a single $[\varepsilon'^+, \varepsilon'^-]$-edge, then π leaves this edge invariant.

Since all the intersection points have the same sign, if the outermost $[\delta^+, \delta^-]$-edges and $[\varepsilon'^+, \varepsilon'^-]$-edges belong to the same curve α or γ, then the dashed arc η in Figure 7, center, is a wave and part (2) of the lemma holds. Next, suppose that they belong to different curves. Without loss of generality, suppose the outermost $[\delta^+, \delta^-]$-edges belong to γ and the outermost $[\varepsilon'^+, \varepsilon'^-]$-edges belong to α.

By Theorems 1.3.7 and 1.3.8, there must be a wave η with respect to $\{\alpha, \gamma\}$. Since no subarc of δ or ε' is a wave, a wave with respect to $\{\alpha, \gamma\}$ must be in an octagon connecting two opposite boundary edges.

Thus, either there is a wave as in Figure 7, center, and part (2) of the lemma holds, or the wave η is as shown in Figure 7, right.

Suppose the wave η is as shown in Figure 7, right. By Lemma 2.1.5, η must be a wave with respect to γ. Perform the wave move along η in two steps: The first step is a surgery on γ along η, resulting in a new meridian γ' and a curve α' parallel to α; see the dashed arcs in Figure 7, right. The second step is to delete α' and obtain a new set of meridians $\{\alpha, \gamma'\}$.

Since α' is parallel to α and since the outermost $[\delta^+, \delta^-]$-edges before the wave move belong to γ, the newly created $[\varepsilon'^+, \varepsilon'^-]$-edge by the first step (see the blue dashed arc in Figure 7, right) must belong to α', and the newly created $[\delta^+, \delta^-]$-edge (see the red dashed arc in Figure 7, right) must belong to γ'. Moreover, this newly created $[\delta^+, \delta^-]$-edge is an outermost $[\delta^+, \delta^-]$-edge. As every essential curve in Σ is invariant under the involution π, the other outermost $[\delta^+, \delta^-]$-edge must also belong to γ'. Since the second step removes α' and leave γ' unchanged, after the wave move, we must have two outermost $[\delta^+, \delta^-]$-edges, both belong to γ'. Furthermore, since α is unchanged by this wave move, the outermost $[\varepsilon'^+, \varepsilon'^-]$-edges still belong to α, and if there is more than one $[\varepsilon'^+, \varepsilon'^-]$-edge before the wave move, there is more than one $[\varepsilon'^+, \varepsilon'^-]$-edge in the new Heegaard diagram. Therefore, we can conclude:

(1) In the new Heegaard diagram after the wave move, the outermost $[\delta^+, \delta^-]$-edges belong to γ' and the outermost $[\varepsilon'^+, \varepsilon'^-]$-edges still belong to α.

(2) Since there is more than one $[\delta^+, \delta^-]$-edge in the diagram before the wave move, there is more than one $[\delta^+, \delta^-]$-edge in new Heegaard diagram. Moreover, if there is more than one $[\varepsilon'^+, \varepsilon'^-]$-edge before the wave move, there is more than one $[\varepsilon'^+, \varepsilon'^-]$-edge in the new Heegaard diagram.

Since there is a $[\delta^+, \delta^-]$-edge and an $[\varepsilon'^+, \varepsilon'^-]$-edge in the new Heegaard diagram, there is no wave with respect to $\{\delta, \varepsilon'\}$ by Lemma 1.3.9. Moreover, γ' has an induced orientation from γ and hence the intersection points of δ and ε' with γ' all have the same sign. So no subarcs of δ or ε' is a wave with respect to $\{\alpha, \gamma'\}$. The argument above implies that a wave with respect to $\{\alpha, \gamma'\}$ must be in an octagon similar to the arc η in Figure 7, right. Thus we can repeat the wave moves described above until the new Heegaard diagram is of the type in Figure 2, left, with $a = b = 0$, which means that M is a connected sum of two lens spaces $L(c, s) \# L(d, t)$ and none of $L(c, s)$ and $L(d, t)$ is $S^2 \times S^1$.

Furthermore, the argument above implies that there is always more than one $[\delta^+, \delta^-]$-edge after each wave move, that is $c \geq 2$. Hence M cannot be S^3. This contradicts the hypothesis that M is either S^3 or a connected sum of $S^2 \times S^1$ with a lens space. □

2.2 Paths, junctions, train tracks and the curves δ and ε

As $\widehat{\Sigma}$ is endowed with the structure from Lemma 2.1.1, it has a decomposition as two annuli and two rectangles $\mathcal{A}_r \cup \mathcal{R}^u \cup \mathcal{A}_l \cup \mathcal{R}^d$. As explained above, we assume that the hyperelliptic involution π maps $\partial_+ P$ to itself. Hence π induces an involution on $\widehat{\Sigma}$ interchanging the two boundary curves $\partial_+ P^u$ and $\partial_+ P^d$.

Remark 2.2.1 (A) The surface, as in Figure 5, has the following three symmetries:

(1) The hyperelliptic involution: a 180°-rotation about the horizontal axis within the projection plane in Figure 5, puncturing each annulus in two points. It interchanges the rectangles, and keeps each annulus in place.

(2) A reflection along a vertical plane cutting through the "middle" of the figure. The vertical plane is disjoint from the two annuli and intersects each rectangle in a cocore arc. The reflection interchanges the two annuli $\mathcal{A}_r$ and $\mathcal{A}_l$.

(3) A reflection along a cylinder which meets each of $\mathcal{R}^u$ and $\mathcal{R}^d$ in a core arc and each of $\mathcal{A}_r$ and $\mathcal{A}_l$ in two cocore arcs. The restriction of this reflection on each rectangle is a reflection along its core arc, and the restriction on each annulus is a reflection of the annulus along a vertical plane that cut the annulus vertically into two halves. This symmetry flips the orientations of the core curves of the two annuli.

(B) The second and the third reflections commute with the hyperelliptic involution.

As it is assumed that π leaves $\delta = \partial D$ and $\partial_+ P$ invariant, π fixes $X = \delta \cap \partial_+ P$. The image of the point X in $\partial_+ P^d$ and $\partial_+ P^u$ is denoted by X^d and X^u, respectively. So the restriction of δ to $\widehat{\Sigma}$ is an arc connecting X^d to X^u, and the restriction of π on $\widehat{\Sigma}$ interchanges X^u and X^d.

First consider the possibility that $X^d \in \partial\mathcal{R}^d$. Hence, by the involution, $X^u \in \partial\mathcal{R}^u$. In this case, let k^d be the arc component of $\delta \cap \mathcal{R}^d$ that contains X^d. If $k^d \cap (\alpha \cup \gamma) = \varnothing$, then by the construction of the rectangles, we may "shrink" the rectangle $\mathcal{R}^d$, by an isotopy so that both k^d and X^d are no longer in $\mathcal{R}^d$. Thus, we may assume that if $X^d \in \partial\mathcal{R}^d$, then $k^d \cap (\alpha \cup \gamma) \neq \varnothing$. Note that we can perform this shrinking operation symmetrically on $\mathcal{R}^u$ so that the symmetry from π is preserved.

Lemma 2.2.2 *No component of $(\delta \cup \varepsilon) \cap \mathcal{A}_l$ and $(\delta \cup \varepsilon) \cap \mathcal{A}_r$ is an arc with both endpoints in the same rectangle $\mathcal{R}^d$ or $\mathcal{R}^u$.*

Proof Assume to the contrary that the statement of the lemma is false. Without loss of generality, suppose $(\delta \cup \varepsilon) \cap \mathcal{A}_l$ has a component ρ with $\partial\rho \subset \mathcal{R}^d$. Note that ρ is a topologically ∂-parallel arc in the annulus $\mathcal{A}_l$. Let k_1 and k_2 be the components of either $\delta \cap \mathcal{R}^d$ or $\varepsilon \cap \mathcal{R}^d$ that contain the two endpoints of ρ. So k_i for $i = 1, 2$ is either a core arc of $\mathcal{R}^d$ or the component k^d of $\delta \cap \mathcal{R}^d$ that contains X^d. By the assumption before Lemma 2.2.2, $k^d \cap (\alpha \cup \gamma) \neq \varnothing$. Since a core arc of $\mathcal{R}^d$ intersects every component of $(\alpha \cup \gamma) \cap \mathcal{R}^d$, k_1 and k_2 must both intersect a component of $(\alpha \cup \gamma) \cap \mathcal{R}^d$. This implies that the two intersection points of $k_1 \cup \rho \cup k_2$ with a component of $(\alpha \cup \gamma) \cap \mathcal{R}^d$ have opposite signs. This contradicts part (2) of Proposition 2.1.8. □

Choose product structures $S^1 \times I$ on each of the annuli $\mathcal{A}_l$ and $\mathcal{A}_r$, and $I \times I$ on each of the rectangles $\mathcal{R}^u$ and $\mathcal{R}^d$, so that each component of $(\alpha \cup \gamma) \cap \widehat{\Sigma}$ is of the form $\{x\} \times I$ in $\mathcal{A}_l$, $\mathcal{A}_r$, $\mathcal{R}^u$ or $\mathcal{R}^d$. We call an arc of the form $\{x\} \times I$ a *vertical* arc in the annuli or rectangles. After performing an isotopy on δ, we may assume each arc in $\delta \cap \mathcal{A}_l$ and $\delta \cap \mathcal{A}_r$ is transverse to all the vertical arcs of $\mathcal{A}_l$ and $\mathcal{A}_r$.

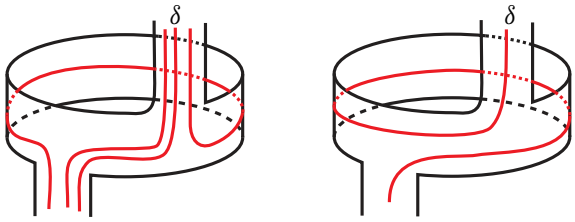

Figure 8: Examples of δ-paths: δ takes 2 short paths, left, and δ takes a long path, right.

Definition 2.2.3 (paths, long and short) Let l be a subarc of δ that is properly embedded in $\mathcal{A}_r$ or $\mathcal{A}_l$. By Lemma 2.2.2, l is an arc connecting $\mathcal{R}^d$ to $\mathcal{R}^u$. By the assumption before Definition 2.2.3, l is transverse to all the vertical arcs. We call l a *regular arc* if one of the following two conditions is satisfied:

(1) Either l is intersects a vertical arc of $\mathcal{A}_r$ or $\mathcal{A}_l$ more than once (in other words, the curve l wraps around the annulus more than once), or

(2) l shares endpoints with two components of $\delta \cap (\mathcal{R}^d \cup \mathcal{R}^u)$ that are core arcs of $\mathcal{R}^d$ and $\mathcal{R}^u$ respectively.

The purpose of condition (2) is to rule out a special case that l is attached to the component of $\delta \cap \mathcal{R}^u$ or $\delta \cap \mathcal{R}^d$ which contains X^u or X^u (eg the arc k^d described before Lemma 2.2.2). A *path* is an isotopy class of regular arcs which setwise preserves $\mathcal{R}^d$ and $\mathcal{R}^u$. Two regular arcs of δ are said to *belong to the same path* if the two arcs are isotopic via an isotopy that setwise preserves $\mathcal{R}^d$ and $\mathcal{R}^u$. We say that δ *takes a path* if it contains a subarc that belongs to that path. A path is called a *long path* in $\mathcal{A}_r$ or $\mathcal{A}_l$ if a regular arc in that path intersects a vertical arc of $\mathcal{A}_r$ or $\mathcal{A}_l$ more than once (ie it satisfies condition (1) above). Otherwise it will be called a *short* path. Note that δ can take at most two different paths in one annulus and that there may be multiple arcs in the same path; see Figure 8 for a picture. We would like to emphasize that, in a short path, both ends of a regular arc must attach to core arcs of $\mathcal{R}^d$ and $\mathcal{R}^u$ in δ.

The decomposition of $\widehat{\Sigma}$ divides δ into a collection of regular arcs and core arcs of $\mathcal{R}^d$ and $\mathcal{R}^u$, as well as two special arcs connecting X^d and X^u to $\partial\mathcal{R}^d$ and $\partial\mathcal{R}^u$. The two special arcs join at X to form a subarc of δ containing X, which we denote by ρ_x. By Definition 2.2.3, each end of ρ_x is attached to either a regular arc in a long path or a core arc in $\mathcal{R}^d$ and $\mathcal{R}^u$ in δ.

Definition 2.2.4 (junctions) The intersections $\partial\mathcal{R}^u \cap \partial\mathcal{A}_l$ and $\partial\mathcal{R}^d \cap \partial\mathcal{A}_l$ are subarcs of different components of $\partial\mathcal{A}_l$; see Figure 5. Choose the product structure on $\mathcal{A}_l$ so that no vertical arc has one endpoint in $\partial\mathcal{R}^u$ and the other end point in $\partial\mathcal{R}^d$.

As illustrated in Figure 9, let $\mathcal{J}_l^u$ and $\mathcal{J}_l^d$ be the unions of all vertical arcs of $\mathcal{A}_l$ with an endpoint in $\partial\mathcal{R}^u$ and $\partial\mathcal{R}^d$, respectively. Similarly let $\mathcal{J}_r^u$ and $\mathcal{J}_r^d$ be the unions of vertical arcs of $\mathcal{A}_r$ with an endpoint in $\partial\mathcal{R}^u$ and $\partial\mathcal{R}^d$, respectively. With this choice $\mathcal{J}_l^u$, $\mathcal{J}_l^d$, $\mathcal{J}_r^u$ and $\mathcal{J}_r^d$ are disjoint squares in $\mathcal{A}_l$ and $\mathcal{A}_r$; see Figure 9. Call these four squares *junctions*.

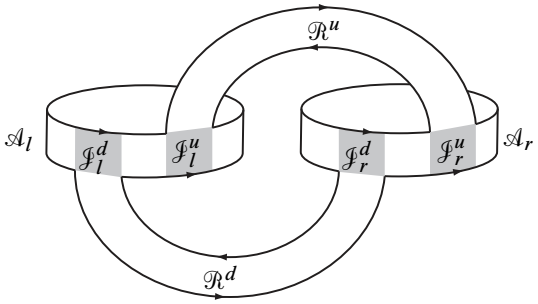

Figure 9: The junctions in $\hat{\Sigma}$.

The purpose of introducing the structure of junctions in Definition 2.2.4 is to construct train tracks, which are our main tools in analyzing curves.

Definition 2.2.5 (train tracks) The four junctions give $\hat{\Sigma}$ a decomposition into ten rectangles: $\mathcal{R}^u$, $\mathcal{R}^d$, $\mathcal{J}_l^u$, $\mathcal{J}_l^d$, $\mathcal{J}_r^u$, $\mathcal{J}_r^d$, $\mathcal{A}_l \setminus (\mathcal{J}_l^u \cup \mathcal{J}_l^d)$ and $\mathcal{A}_r \setminus (\mathcal{J}_r^u \cup \mathcal{J}_r^d)$. When all the arcs of $\delta \setminus \rho_x$ in each of these rectangles with endpoints in the same pair of boundary edges are collapsed into a single arc, we obtain a train track. Now we consider how the special arc ρ_x is affected by the train-track construction. If δ does not take a long path, then the train-track construction pinches the two endpoints of ρ_x onto the train track, and the construction does not affect $\operatorname{int}(\rho_x)$; see Figure 10 for a local picture. If δ takes a long path, then two subarcs of ρ_x at its ends may be affected by the train-track construction. For example, in the left picture of Figure 25, ρ_x contains arcs in the middle of the shaded region where parallel arcs on both sides of ρ_x are collapsed into a single arc in the train track. If this happens, we collapse these two subarcs of ρ_x (which are at the two ends of ρ_x) onto the train track (ie collapse all the arcs in shaded region of Figure 25 into a single segment in the train track). In either case, this train track, together with ρ_x, forms a slightly larger train track, which we denote by τ_D; see Figure 10 and the left picture of Figure 25. The special arc ρ_x becomes a special segment of τ_D, which we still denote by ρ_x.

The switches (or cusp points) of the train track divide the train track into a collection of arcs which are called the *segments* of the train track. The *weight* of δ at a segment of τ_D is the number of δ-arcs which collapse onto this segment. (In other words, the number of times that δ passes this segment.) It follows from our construction that the arc ρ_x is a segment of τ_D that contains the point X; see Figure 10. Clearly the weight of δ at the segment ρ_x is one by the construction. Moreover, by part (2) of Proposition 2.1.8, each segment of τ_D has an induced orientation from δ and the orientations of the segments are compatible at each cusp.

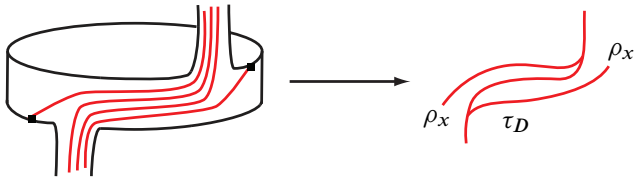

Figure 10: Pinch δ into a traintrack τ_D.

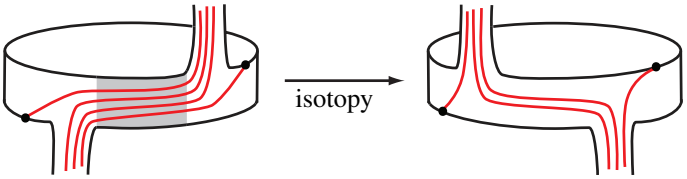

Figure 11: Isotopy on an annulus.

Remark 2.2.6 In the decomposition of $\hat{\Sigma} = \Sigma \setminus \mathcal{N}(\partial_+ P)$, the product structures on the two rectangles $\mathcal{R}^u$ and $\mathcal{R}^d$ depend on their intersections with $\gamma \cup \alpha$; see the γ arcs in Figure 5. So the product structures of $\mathcal{R}^u$ and $\mathcal{R}^d$ are "basically" fixed with respect to δ, and there is no ambiguity in the train track construction.

The product structures on $\mathcal{A}_l$ and $\mathcal{A}_r$ are somewhat flexible on the regions which do not intersect $\gamma \cup \alpha$. For example, if $\gamma \cup \alpha$ does not intersect the shaded region in Figure 11, then we can perform an isotopy on the annulus as in Figure 11 that changes the product structure. The effect of this isotopy on the train track τ_D is performing a splitting move on the train track.

Definition 2.2.7 Consider τ_D as a train track in the Heegaard surface Σ and a small neighborhood of τ_D in Σ, denoted by $\mathcal{N}(\tau_D)$. We can view $\mathcal{N}(\tau_D)$ as an I-bundle over the train track τ_D. The curve $\delta \subset \mathcal{N}(\tau_D)$ is transverse to the I-fibers of $\mathcal{N}(\tau_D)$. Similar to a train track, $\mathcal{N}(\tau_D)$ has a collection of cusps at its boundary. There is a collection of arcs $s_1, \dots, s_k$ in $\mathcal{N}(\tau_D)$ which are disjoint from δ and are transverse to the I-fibers, and which connect the cusps of $\mathcal{N}(\tau_D)$ in pairs. If one cuts open $\mathcal{N}(\tau_D)$ along $s_1, \dots, s_k$, the resulting space is an I-bundle over δ. We call each s_i a *splitting arc*.

3 Proof of the main theorem

In this section we prove Theorem 7 using the terms defined and facts proved in the previous sections.

We begin with a description of the arc $\delta \setminus \partial_+ P \subset \hat{\Sigma}$ connecting X^u to X^d. The possible locations for X^u and X^d in $\partial_+ P^u$ and $\partial_+ P^d$, respectively, are not determined by the construction as the precise gluing map between $\partial_+ P^u$ and $\partial_+ P^d$ is unknown. Nonetheless, since X is assumed to be a fixed point of the hyperelliptic involution π, the points X^u and X^d are symmetric under π.

It follows from Lemma 2.1.4 that δ must take at least one path in each annulus $\mathcal{A}_l$, $\mathcal{A}_r$. The following are the four possible configurations in $\hat{\Sigma}$ for this arc:

- **Configuration 1** δ takes one short path in $\mathcal{A}_l$ and one short path in $\mathcal{A}_r$.
- **Configuration 2** δ takes one short path in $\mathcal{A}_l$ or $\mathcal{A}_r$ and two short paths in the other annulus.
- **Configuration 3** δ takes two short paths in $\mathcal{A}_l$ and two short paths in $\mathcal{A}_r$.
- **Configuration 4** δ takes a long path in one (or both) of the annuli.

A key difference between these four configurations can be seen in the train track τ_D (constructed in Definition 2.2.5).

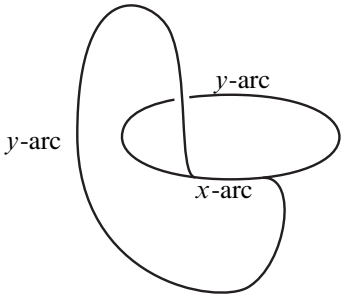

Figure 12: A train track whose complement is a once-punctured torus.

In Configuration 1 the two short paths and a core arc from each rectangle $\mathcal{R}^u$ and $\mathcal{R}^d$ are pinched into a circle, and the train track τ_D is the union of this circle and the special segment ρ_x which contains the point $X = \delta \cap \partial_+ P$, as illustrated in the construction of the train track τ_D in Figure 10. Thus τ_D in the first configuration must be as shown in Figure 12. The x- and y-arcs marked in Figure 12 are the three segments of τ_D, where the x-arc is the segment with cusp directions at its endpoints pointing into this arc. The special segment ρ_x is one of the y-arcs in Figure 12. Note that in this configuration, $\Sigma \setminus \tau_D$ is a once-punctured torus with two cusps at its boundary.

In Configuration 2, see Figure 13, the train track τ_D can be obtained by adding an additional segment to the train track in Figure 12. Hence τ_D in this configuration has four cusps. Since $\partial_+ P \cap \tau_D$ is just a single point X, the train track τ_D is nonseparating in Σ. The complement $\Sigma \setminus \tau_D$ in Configuration 2 can be obtained by cutting the once-punctured torus in Configuration 1 along an essential arc, which means that $\Sigma \setminus \tau_D$ in Configuration 2 is a topological annulus. Moreover, since every essential curve in a genus-two surface is invariant by the hyperelliptic involution, after isotopy, we may assume τ_D is invariant by the involution and this implies that the annulus $\Sigma \setminus \tau_D$ has two cusps at each boundary curve.

Similarly, in Configuration 3, since $\tau_D \cap \partial_+ P$ is a single point, $\Sigma \setminus \tau_D$ is connected. As τ_D has six cusps in this configuration, for Euler characteristic reason, $\Sigma \setminus \tau_D$ is a disk with six cusps at its boundary.

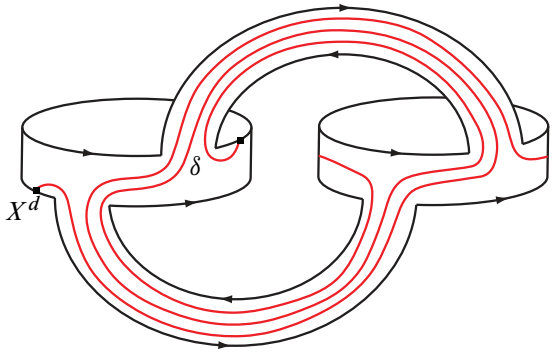

Figure 13: A configuration for δ that takes one short path in $\mathcal{A}_l$ and two short paths in $\mathcal{A}_r$.

In Configuration 4, τ_D takes a long path in one annulus, say $\mathcal{A}_r$. Then τ_D has two possible configurations depending on whether δ takes one short path, two short paths, or another long path in $\mathcal{A}_l$. If τ_D takes one short path in $\mathcal{A}_l$, then τ_D has four cusps, and similar to Configuration 2, the surface $\Sigma \setminus \mathcal{N}(\tau_D)$ is an annulus with two cusps on each boundary curve. If τ_D takes two short paths or another long path in $\mathcal{A}_l$, then τ_D has six cusps, and similar to Configuration 3, $\Sigma \setminus \tau_D$ is a disk with six cusps at its boundary.

We discuss each of these configurations in separate sections. The sections below deal with the configurations above in order 2, 3, 4 and 1 for convenience of the proof.

3.1 The curve δ takes three short paths

In this section, Configuration 2 is considered, where δ takes two short paths in one of the annuli, and a single short path in the other annulus. An example of this situation is given in Figure 13.

Proposition 3.1.1 *If δ takes two short paths in one annulus and a single short path in another annulus, then K is doubly primitive.*

Proof After applying symmetry (2) in Remark 2.2.1 if necessary, assume without loss of generality that δ takes two short paths in $\mathcal{A}_r$ and one short path in $\mathcal{A}_l$. Since symmetry (3) in Remark 2.2.1 interchanges the two short paths in $\mathcal{A}_l$, assume further that δ takes a fixed short path in $\mathcal{A}_l$; see Figure 13.

As explained at the beginning of the section, $\Sigma \setminus \mathcal{N}(\tau_D)$ is an annulus denoted by A_D. The annulus A_D has two cusps in each boundary component of A_D and the cusps correspond to the cusps of the train track τ_D.

By construction, the train track τ_D has a special segment ρ_x which contains the intersection point X and the weight of δ at ρ_x is one. The two ends of ρ_x are two cusps of τ_D, which are called the ρ-cusps. The other two cusps of τ_D are at the two junctions of $\mathcal{A}_r$ where the two short paths meet; see Figure 13. We call these two cusps the j-cusps. As δ is invariant under π and X is a fixed point of π, we may assume ρ_x is invariant under π and π interchanges the two ρ-cusps and the two j-cusps.

A neighborhood $\mathcal{N}(\tau_D)$ of τ_D contains two splitting arcs as in Definition 2.2.7. Denote them by s_ρ and s_j. These splitting arcs connect the cusps of $\mathcal{N}(\tau_D)$ in pairs such that s_ρ and s_j are disjoint from δ, and if one splits $\mathcal{N}(\tau_D)$ along s_ρ and s_j, the resulting surface $\mathcal{N}(\tau_D) \setminus (s_\rho \cup s_j)$ is a product neighborhood of δ.

We claim that one splitting arc, say s_ρ, connects the two ρ-cusps and the other splitting arc s_j connects the two j-cusps. To see this, consider the modified train track τ_D^- obtained by removing ρ_x from τ_D. As in the discussion at the beginning of the section, the train track τ_D^- is as in Figure 12 and $\mathcal{N}(\tau_D^-)$ is a once-punctured torus. We view $\mathcal{N}(\tau_D^-)$ as an I-bundle over τ_D^- and view s_ρ and s_j as arcs in $\mathcal{N}(\tau_D^-)$ transverse to the I-fibers.

Assume, for the sake of contradiction, that s_ρ connects a ρ-cusp to a j-cusp. Then s_ρ connects a cusp of $\mathcal{N}(\tau_D^-)$ to a smooth boundary point of $\mathcal{N}(\tau_D^-)$. Since s_ρ is transverse to the I-fibers, this implies that s_ρ is nontrivial in the once-punctured torus $\mathcal{N}(\tau_D^-)$. Similarly, s_j is also nontrivial in $\mathcal{N}(\tau_D^-)$.

The symmetry induced by π implies that the two endpoints of ρ_x (ie the two ρ-cusps) and the two j-cusps alternate along the boundary curve of $\mathcal{N}(\tau_D^-)$; see Figure 13.

If s_ρ connects a ρ-cusp to a j-cusp, then the endpoints ∂s_ρ and ∂s_j do not alternate along $\partial\mathcal{N}(\tau_D^-)$. However, the endpoints of any pair of disjoint nonparallel essential arcs in a once-punctured torus do alternate along the boundary curve. Thus s_ρ and s_j must be parallel in the once-puncture torus $\mathcal{N}(\tau_D^-)$. Hence, $\mathcal{N}(\tau_D^-)\setminus(s_\rho\cup s_j)$ has two components, one of which is an annulus. Since s_ρ and s_j are transverse to the I-fibers of $\mathcal{N}(\tau_D^-)$, this annular component of $\mathcal{N}(\tau_D^-)\setminus(s_\rho\cup s_j)$ must be an I-bundle over a simple closed curve with each I-fiber a subarc of an I-fiber of $\mathcal{N}(\tau_D^-)$. This is impossible because s_ρ and s_j are the splitting arcs of $\mathcal{N}(\tau_D)$ and $\mathcal{N}(\tau_D)\setminus(s_\rho\cup s_j)$ is an I-bundle over δ, ie we have a contradiction. Thus, one splitting arc, say s_ρ, connects the two ρ-cusps, and the other s_j connects the two j-cusps.

Claim 3.1.2 *Each of the splitting arcs s_ρ and s_j passes through both $\mathcal{R}^u$ and $\mathcal{R}^d$. In other words, each contains subarcs that are core arcs of $\mathcal{R}^u$ and $\mathcal{R}^d$. In particular, each splitting arc intersects γ at least twice.*

Proof Since the splitting arc s_j connects the two j-cusps, as illustrated in Figure 13, s_j must pass through both rectangles $\mathcal{R}^d$ and $\mathcal{R}^u$. Now we consider s_ρ.

Recall the construction of the special arc ρ_x. As discussed before Definition 2.2.4, in general, each end of ρ_x is attached either to a long path or to a component of $\delta\cap(\mathcal{R}^d\cup\mathcal{R}^u)$. Since there is no long path in the configuration of Proposition 3.1.1, the two ends of ρ_x are attached to two components of $\delta\cap(\mathcal{R}^d\cup\mathcal{R}^u)$ that are core arcs of $\mathcal{R}^d$ and $\mathcal{R}^u$, respectively.

In the construction of the train track τ_D, all the core arcs of $\delta\cap\mathcal{R}^d$ (and $\delta\cap\mathcal{R}^u$) are pinched onto the same segment of τ_D. As illustrated in Figure 10, this means that, similar to the two j-cusps, the cusp directions at the two ends of ρ_x must point into $\mathcal{R}^d$ and $\mathcal{R}^u$, respectively. Thus, s_ρ passes through both $\mathcal{R}^d$ and $\mathcal{R}^u$. □

As $\varepsilon\cap\delta=\varnothing$, the intersection $\varepsilon\cap\mathcal{N}(\tau_D)$ is a collection of arcs through the cusps of $\mathcal{N}(\tau_D)$. By part (2) of Proposition 2.1.8, we can choose an orientation on ε so that the induced orientation on all the arcs in $\varepsilon\cap\mathcal{N}(\tau_D)$ is compatible with the orientation of δ near the cusps of $\mathcal{N}(\tau_D)$; see the arrows in Figure 14.

If ε goes through all the cusps of $\mathcal{N}(\tau_D)$, then as shown in Figure 14, there must be an s-wave ρ with respect to ε which is parallel to a δ-arc. Perform a sequence of wave moves on ε along s-waves similar to ρ (as in Figure 14) so that the resulting curve ε' goes into at most one cusp in each boundary component of $A_D=\Sigma\setminus\mathcal{N}(\tau_D)$. This implies that the arcs in $\varepsilon'\cap\mathcal{N}(\tau_D)$, if there are any, are all parallel to either s_ρ or s_j.

Claim 3.1.3 *The intersection $\varepsilon'\cap\mathcal{N}(\tau_D)$ contains at least two arcs.*

Proof The proof is by contradiction and is divided into two cases:

(i) $|\varepsilon'\cap\mathcal{N}(\tau_D)|=0$, ie $\varepsilon'\cap\mathcal{N}(\tau_D)=\varnothing$, or

(ii) $|\varepsilon'\cap\mathcal{N}(\tau_D)|=1$.

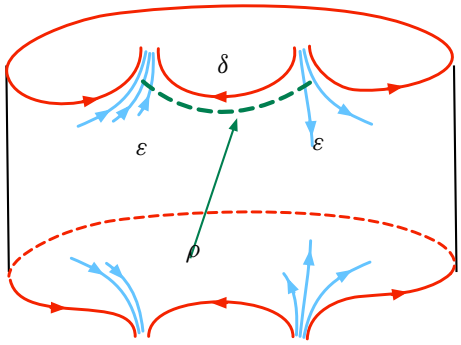

Figure 14: The annulus $A_D = \Sigma \setminus \tau_D$ and the s-wave ρ with respect to ε.

Case (i) ($\varepsilon' \cap \mathcal{N}(\tau_D) = \varnothing$, ie ε' is isotopic to the core curve of the annulus A_D) Since $\partial_+ P \cap \rho_x = \partial_+ P \cap \delta$ is the single point X, the intersection $\partial_+ P \cap A_D$ is an essential arc in A_D. As ε' is isotopic to the core curve of A_D, the intersection $\varepsilon' \cap \partial_+ P$ is a single point. Let E' denote the disk in the handlebody W which is bounded by ε'. Hence (P, E') is a $(\mathcal{P}, \mathcal{D})$-pair.

Showing that the $(\mathcal{P}, \mathcal{D})$-pair (P, E') has a smaller complexity than (P, D) will rule out Case (i). To compute the complexity, consider the arcs in $(\alpha \cup \gamma) \cap \widehat{\Sigma}$, or equivalently arcs in $(\alpha \cup \gamma) \setminus \partial_+ P$. There are two possible subcases:

The first subcase is that when a component of $(\alpha \cup \gamma) \setminus \partial_+ P$ intersects ε' then it also intersects δ. In this subcase, $c_0(P, E', \alpha, \gamma) \le c_0(P, D, \alpha, \gamma)$ by the definition of the complexity.

The second subcase is that there is a component of $(\alpha \cup \gamma) \setminus \partial_+ P$ which intersects ε' but does not intersect δ. Let $\mathscr{C}$ denote the union of the components of $(\alpha \cup \gamma) \cap A_D$ which are not ∂-parallel in A_D. So, in case (i), we have $|(\alpha \cup \gamma) \cap \varepsilon'| = |\mathscr{C}|$.

The arcs in $\mathscr{C}$ can be viewed as vertical arcs of the annulus A_D and $\partial_+ P \cap A_D$ can be viewed as a spiral in A_D. In the second subcase, since $A_D = \Sigma \setminus \tau_D$, a component of $(\alpha \cup \gamma) \setminus \partial_+ P$ which intersects ε' but does not intersect δ must be contained in the interior of A_D. This implies that the arc in $\mathscr{C}$ which contains such a component of $(\alpha \cup \gamma) \setminus \partial_+ P$ must meet $\partial_+ P$ more than once. Hence the curve $\partial_+ P \cap A_D$ wraps around A_D more than once, which means that each arc in $\mathscr{C}$ must intersect $\partial_+ P$ at least once. Thus, for each component σ of $\mathscr{C}$, we can pick one intersection point of $\sigma \cap \partial_+ P$, and so we have a total of $|\mathscr{C}|$ such intersection points. These $|\mathscr{C}|$ intersection points divide $\alpha \cup \gamma$ into $|\mathscr{C}|$ arcs, each of which meets δ. Since these $|\mathscr{C}|$ points are in $\partial_+ P$, each of the these $|\mathscr{C}|$ arcs contains at least one component of $(\alpha \cup \gamma) \setminus \partial_+ P$ which meets δ. By our definition of complexity, this implies that $c_0(P, D, \alpha, \gamma) \ge |\mathscr{C}|$. Moreover, by the definition of the complexity,

$$c_0(P, E', \alpha, \gamma) \le |(\alpha \cup \gamma) \cap \varepsilon'|.$$

As $|(\alpha \cup \gamma) \cap \varepsilon'| = |\mathscr{C}|$, we have $c_0(P, E', \alpha, \gamma) \le c_0(P, D, \alpha, \gamma)$ in both subcases.

Since each arc in $\mathscr{C}$ intersects ε' in one point, Lemma 2.1.4 implies that in fact

$$|(\alpha \cup \gamma) \cap \varepsilon'| < |(\alpha \cup \gamma) \cap \delta|$$

and thus

$$c(P, E', \alpha, \gamma) < c(P, D, \alpha, \gamma).$$

This contradicts Assumption 1.2.3 on the choice of the pair (P, D) and therefore $\varepsilon' \cap \mathcal{N}(\tau_D) \neq \varnothing$.

Case (ii) ($\varepsilon' \cap \mathcal{N}(\tau_D)$ is a single arc) Consider the train track τ_D^- obtained by removing the segment ρ_x from τ_D. Note that $\mathcal{N}(\tau_D^-)$ is a once-punctured torus and $\delta \cap \mathcal{N}(\tau_D^-)$ is a single arc in $\mathcal{N}(\tau_D^-)$. Since $\varepsilon' \cap \rho_x = \varnothing$, we have $\varepsilon' \cap \mathcal{N}(\tau_D) = \varepsilon' \cap \mathcal{N}(\tau_D^-)$. Hence $\varepsilon' \cap \mathcal{N}(\tau_D^-)$ is also a single arc in $\mathcal{N}(\tau_D^-)$. Further, note that the arc $\varepsilon' \cap \mathcal{N}(\tau_D)$ is parallel to one of the two splitting arcs s_ρ or s_j. Similar to the argument before Claim 3.1.2, the endpoints ∂s_ρ (or ∂s_j) and the endpoints $\partial \rho_x$ alternate along the boundary of $\mathcal{N}(\tau_D^-)$. Hence the endpoints of the arc $\varepsilon' \cap \mathcal{N}(\tau_D^-)$ and the endpoints of the arc $\delta \cap \mathcal{N}(\tau_D^-)$ alternate along the boundary of $\mathcal{N}(\tau_D^-)$. As $\mathcal{N}(\tau_D^-)$ is a once-punctured torus, this implies that $\delta \cap \mathcal{N}(\tau_D^-)$ and $\varepsilon' \cap \mathcal{N}(\tau_D^-)$ are nonparallel essential arcs in $\mathcal{N}(\tau_D^-)$. Thus we may assume that $\delta \cap \mathcal{N}(\tau_D^-)$ is of slope $1/0$ and $\varepsilon' \cap \mathcal{N}(\tau_D^-)$ is of slope $0/1$ with respect to a basis of the relative first homology of the once-punctured torus $\mathcal{N}(\tau_D^-)$.

Since δ takes two short paths in $\mathcal{A}_r$, we can isotope the core curve $\mathfrak{a}_r$ of $\mathcal{A}_r$ into $\mathcal{N}(\tau_D^-)$. So there is an arc of slope p/q in $\mathcal{N}(\tau_D^-)$ interesting $\mathfrak{a}_r$ in a single point. Let E' and D be the disks bounded by ε' and δ in W, respectively. We can take p parallel copies of D and q parallel copies of E', and perform a sequence of band sums of these disks along the boundary of $\mathcal{N}(\tau_D^-)$ to obtain a disk D' in W so that $\partial D' \cap \mathcal{N}(\tau_D^-)$ is an arc of slope p/q. Hence $\partial D'$ intersects $\mathfrak{a}_r$ in a single point. Recall that $\mathfrak{a}_r$ is a boundary curve of a planar surface P_r, ie $\mathfrak{a}_r = \partial_+ P_r$, where P_r is obtained by a single ∂-compression on P; see Remark 2.1.2. Thus the pair (P_r, D') is a new $(\mathcal{P}, \mathcal{D})$-pair.

Now compute the complexity of this new (P_r, D') pair. The curve $\mathfrak{a}_r$ can be viewed as the union of two segments of τ_D. Since δ intersects every α- and γ-arc in $\mathcal{A}_r$ and intersects at least two other γ-arcs in the rectangles $\mathcal{R}^d$ and $\mathcal{R}^u$, we have

$$|(\alpha \cup \gamma) \cap \mathfrak{a}_r| < c_0(P, D, \alpha, \gamma).$$

By the definition of the complexity,

$$c_0(P_r, D', \alpha, \gamma) \leq |(\alpha \cup \gamma) \cap \partial_+ P_r| = |(\alpha \cup \gamma) \cap \mathfrak{a}_r|.$$

Thus, when putting the two inequalities together we have $c_0(P_r, D', \alpha, \gamma) < c_0(P, D, \alpha, \gamma)$ again, a contradiction to Assumption 1.2.3. This finishes the proof of the claim. □

Next, consider the Heegaard diagram formed by the $\{\delta, \varepsilon'\}$ and $\{\alpha, \gamma\}$ sets of curves.

Claim 3.1.4 *The curve γ has two distinct subarcs that are $[\varepsilon'^-, \varepsilon'^+]$-edges and also has two distinct subarcs that are $[\delta^-, \delta^+]$-edges.*

Proof As illustrated by the arrows in Figure 14, before the wave moves on ε, the orientations of δ and ε are compatible within $\mathcal{N}(\tau_D)$ (see part (2) of Proposition 2.1.8). The wave moves on ε, that were performed above, are disjoint from $\mathcal{N}(\tau_D)$. Hence there is an induced orientation on ε' and the induced orientations of the arcs $\varepsilon' \cap \mathcal{N}(\tau_D)$ are compatible with the orientation of δ.

In the paragraph before Claim 3.1.3, we concluded that the arcs in $\varepsilon' \cap \mathcal{N}(\tau_D)$ are all parallel to either s_ρ or s_j in $\mathcal{N}(\tau_D)$. By Claim 3.1.3, $\varepsilon' \cap \mathcal{N}(\tau_D)$ contains at least two such parallel arcs. As the orientations of the arcs $\varepsilon' \cap \mathcal{N}(\tau_D)$ are compatible, the intersection points of each arc of $\gamma \cap \mathcal{N}(\tau_D)$ with ε' all have the same sign. Thus each subarc of $\gamma \cap \mathcal{N}(\tau_D)$ between two curves of $\varepsilon' \cap \mathcal{N}(\tau_D)$ is an $[\varepsilon'^-, \varepsilon'^+]$-edge. By Claim 3.1.2, γ intersects each arc in $\varepsilon' \cap \mathcal{N}(\tau_D)$ at least twice, so γ has at least two distinct subarcs that are $[\varepsilon'^-, \varepsilon'^+]$-edges. Similarly, since ε' does not go through all cusps, it follows from Claim 3.1.2 that γ has two distinct subarcs that are $[\delta^-, \delta^+]$-edges. □

In the Heegaard diagram formed by $\{\delta, \varepsilon'\}$ and $\widehat{V} = \{\alpha, \gamma\}$, the existence of both $[\delta^-, \delta^+]$ and $[\varepsilon'^-, \varepsilon'^+]$-edges implies that the Whitehead Graph $\Gamma(\delta, \varepsilon')$ must be of the type in Figure 2, left, with $c \neq 0$ and $d \neq 0$. In fact $c \geq 2$ and $d \geq 2$ by Claim 3.1.4. By part (1) of Lemma 1.3.9, the Heegaard diagram contains no wave with respect to $\{\delta, \varepsilon'\}$. By Theorems 1.3.7 and 1.3.8, it must contain a wave with respect to $\{\alpha, \gamma\}$.

Since every essential simple closed curve is invariant by the involution π after perhaps some isotopy, one may assume that ε', as well as δ, $\partial_+ P$, α and γ, are invariant under π. It also follows from part (1) of Proposition 2.1.8 that we can fix an orientation for α and γ so that the intersection points of α and γ with δ all have the same sign.

Now, assume that η is a wave with respect to $\{\gamma, \alpha\}$ and assume further that $|\eta \cap \partial_+ P|$ is minimal among all such waves. It follows from Lemma 2.1.6 that $\eta \cap \partial_+ P \neq \varnothing$. The curve $\partial_+ P$ divides η into a sequence of subarcs $\kappa_1, \ldots, \kappa_s$ so that the first and last subarcs κ_1 and κ_s contain the endpoints of η. Next, consider the κ_i, with $i = 1, \ldots, s$, as arcs in $\widehat{\Sigma}$. The two arcs κ_1 and κ_s at the ends are arcs connecting $\gamma \cup \alpha$ to $\partial\widehat{\Sigma}$. All other arcs κ_i for $i = 2, \ldots, s-1$ are properly embedded in $\widehat{\Sigma}$.

For each junction $\mathcal{J}$ (see Definition 2.2.4) in the decomposition of $\widehat{\Sigma}$, let $\widehat{\mathcal{J}}$ be a hexagonal neighborhood of $\mathcal{J}$ in $\widehat{\Sigma}$; see Figure 15, left. Any arc σ in $\widehat{\Sigma}$ connecting a component of $(\gamma \cup \alpha) \cap \widehat{\Sigma}$ to $\partial\widehat{\Sigma}$, is called a *junction arc* if after isotopy on $\widehat{\Sigma}$, the arc σ lies in a hexagonal neighborhood $\widehat{\mathcal{J}}$ of a junction $\mathcal{J}$ connecting a pair of opposite edges of $\widehat{\mathcal{J}}$; see the dashed arc in Figure 15, left.

Claim 3.1.5 *Let η and $\kappa_1, \ldots, \kappa_s$ be as above. Then each κ_i for $i = 2, \ldots, s-1$ is isotopic to a vertical arc of $\mathcal{A}_l$ and both κ_1 and κ_s are junction arcs.*

Proof First, consider the arcs κ_i for $i = 2, \ldots, s-1$, which are properly embedded in $\widehat{\Sigma}$. By Lemma 2.1.1, γ intersects both $\mathcal{R}^u$ and $\mathcal{R}^d$; see Figure 5. As η is a wave, $\eta \cap \delta = \varnothing$ and $\kappa_i \cap \gamma = \varnothing$ for $i = 2, \ldots, s-1$. Since δ passes through $\mathcal{R}^u$ and $\mathcal{R}^d$ and since $\kappa_i \cap \gamma = \varnothing$ for $i = 2, \ldots, s-1$, the arc κ_i is either ∂-parallel in $\widehat{\Sigma}$ or an essential arc in $\mathcal{A}_l$ or $\mathcal{A}_r$. If κ_i is boundary parallel, then the bigon disk between κ_i and $\partial_+ P$ can be eliminated by a simple isotopy. This isotopy yields a new wave with fewer intersection points with $\partial_+ P$, contradicting our hypothesis on η. Thus κ_i for $i = 2, \ldots, s-1$ is isotopic to a vertical arc of $\mathcal{A}_l$ or $\mathcal{A}_r$, and since it does not intersect δ it must be in $\mathcal{A}_l$.

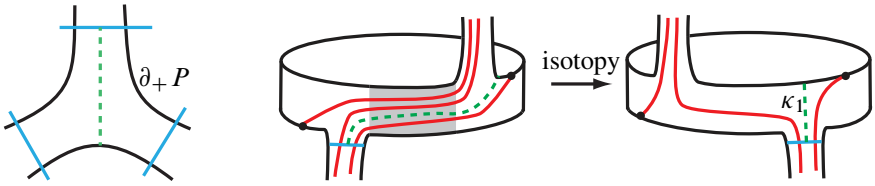

Figure 15: Junction arcs.

Next consider κ_1; the proof for κ_s is identical. Let σ be the component of $(\gamma \cup \alpha) \cap \hat{\Sigma}$ that contains an endpoint Z of κ_1 (which is also an endpoint of the wave η). As illustrated in Figure 5, two γ-arcs in $\mathcal{R}^u$ and $\mathcal{R}^d$ divide $\hat{\Sigma}$ into two topological annuli. Moreover, since $\text{int}(\kappa_1) \cap \gamma = \varnothing$, both σ and κ_i lie in one of the two annuli. Let σ' and σ'' be the two components of $\sigma \setminus Z$. If both $\kappa_1 \cup \sigma'$ and $\kappa_1 \cup \sigma''$ are nontrivial arcs in these two annuli, as illustrated in Figure 15, center and right, perform an isotopy so that κ_1 connects a pair of opposite edges of a hexagonal neighborhood of a junction and hence κ_1 is a junction arc. Note that since η is a wave, if κ_1 is as in Figure 15, center, then the shaded region cannot contain any α- or γ-arc. Hence the isotopy from the center diagram to the right-hand diagram in Figure 15 is in fact the isotopy in Figure 11.

It remains to consider the case that either $\kappa_1 \cup \sigma'$ or $\kappa_1 \cup \sigma''$ is a ∂-parallel arc in $\hat{\Sigma}$. Suppose $\kappa_1 \cup \sigma'$ is ∂-parallel, then $\kappa_1 \cup \sigma'$ and a subarc of $\partial\hat{\Sigma}$ bound a triangle $\Delta \subset \hat{\Sigma}$; see Figure 16, left.

If $\Delta \cap (\delta \cup \varepsilon') = \varnothing$, one can perform an isotopy pushing κ_1 across Δ, as can be seen in Figure 16. This yields a new wave with fewer intersection points with $\partial_+ P$, contradicting our hypothesis on η.

If $\Delta \cap (\delta \cup \varepsilon') \neq \varnothing$, then δ and ε' cut off a subtriangle Δ' in Δ; see the shaded triangle in Figure 17, left. The triangle Δ' is rotated by the involution π into a triangle $\pi(\Delta')$ on the other side of $\partial_+ P$. Since δ, ε' and $\partial_+ P$ are invariant under π, there are two possible configurations, corresponding to whether the vertical arc on the boundary of $\pi(\Delta')$ meets $\partial_+ P$ in Δ or outside it. The two configurations are illustrated in Figures 17, left and right.

Figure 17, left, is the case where an endpoint of $\pi(\sigma)$ lies on the boundary of Δ. In this case the curve α or γ that contains $\pi(\sigma)$ must cut into Δ and meet η. This contradicts the fact that η is a wave and $\eta \cap (\alpha \cup \gamma) = \partial\eta$. So the situation in Figure 17, left, cannot occur.

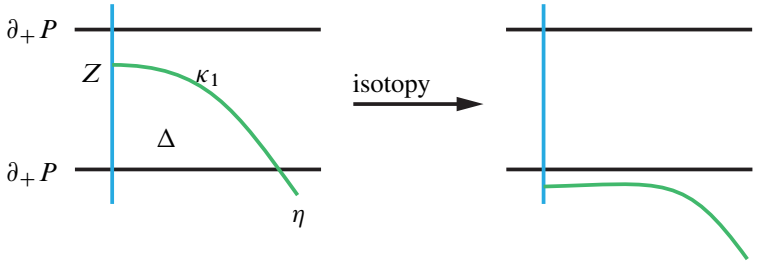

Figure 16: An isotopy on η.

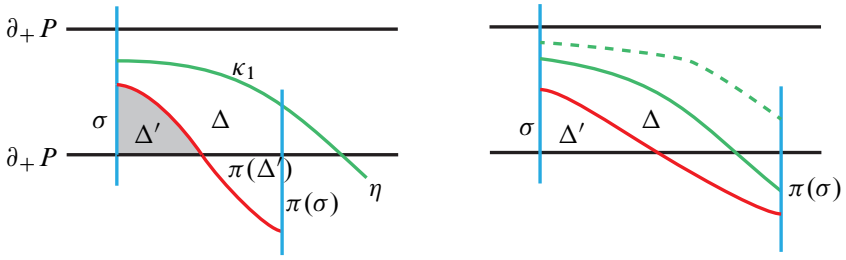

Figure 17: Possible configurations of κ_1.

Figure 17, right, is the configuration where a vertex of Δ lies in the boundary of $\pi(\Delta')$. In this case, the wave η must meet $\pi(\sigma)$. Since $\eta \cap (\alpha \cup \gamma) = \partial\eta$, this means that the union of κ_1 and the extension of κ_1 in $\pi(\Delta')$ — see the green arc in Figure 17, right — must be the whole of η. In this case the wave η can be isotoped to a wave that is disjoint from $\partial_+ P$; see the dashed arc in Figure 17, right, contradicting Lemma 2.1.6. □

It follows from Claim 3.1.5 that, after isotopy, any wave η contains two junction arcs at its two ends. Note that in a hexagonal neighborhood of a junction, two junction arcs connecting different pairs of opposite edges must intersect. Hence the two junction arcs in η belong to two different junctions. For each wave η, the involution π sends η to a disjoint dual wave $\pi(\eta)$; see Remark 1.3.3. Hence the four junction arcs of η and $\pi(\eta)$ must be at different junctions. Since π leaves $\mathscr{A}_l$ and $\mathscr{A}_r$ invariant, π interchanges the two junctions at $\mathscr{A}_l$ and interchanges the two junctions at $\mathscr{A}_r$. As η and $\pi(\eta)$ are distinct waves, this implies that one junction arc of η must be at a junction of $\mathscr{A}_r$ and the other junction arc is at a junction of $\mathscr{A}_l$.

Without loss of generality, suppose κ_1 is a junction arc in $\mathscr{A}_r$ and κ_s is a junction arc in $\mathscr{A}_l$. By assumption, δ takes two short paths in $\mathscr{A}_r$. Therefore, as shown in Figure 18, left, since the junction arc κ_1 is disjoint from δ, up to interchanging η and $\pi(\eta)$, the arc κ_1 must connect an arc in $(\gamma \cup \alpha) \cap \mathscr{R}^u$ to $\partial\mathscr{A}_r$.

Next, perform a wave move along η and study how the new curve, obtained by the wave move, intersects $\hat{\Sigma}$, focusing on how the curve is changed near the arc κ_1. By Lemma 2.1.5, the wave η must be a wave with respect to γ. Perform the wave move along η in two steps as in Definition 1.3.4. The surgery-step

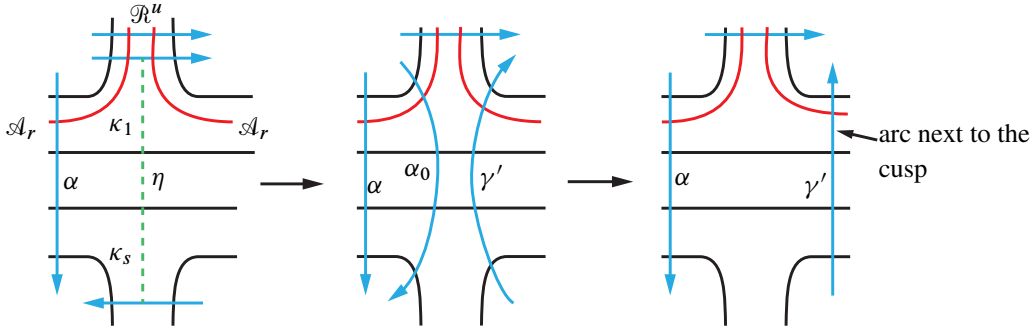

Figure 18: Wave move along η in two steps.

changes γ into two curves γ' and α_0 where α_0 is parallel to α, and the second step is to remove α_0 and obtain a new set of meridians $\{\alpha, \gamma'\}$.

Since κ_1 is a junction arc in $\mathcal{A}_r$, as shown in Figure 18, left and center, the surgery-step of the wave move changes the arc in $\gamma \cap \mathcal{R}^u$ that contains an endpoint of η into two vertical arcs of $\mathcal{A}_r$ next to the junction, one in γ' and the other in α_0.

By Claim 3.1.5, the intersection of both curves γ' and α_0 with $\widehat{\Sigma}$ is a collection of nontrivial arcs with respect to the decomposition of $\widehat{\Sigma}$. Furthermore, after slightly modifying the product structure of $\mathcal{A}_l$ and $\mathcal{A}_r$ if necessary, we may assume both γ' and α_0 meet $\widehat{\Sigma}$ in a collection of vertical arcs in $\mathcal{A}_l$, $\mathcal{A}_r$, $\mathcal{R}^u$ and $\mathcal{R}^d$.

Since α_0 is parallel to α and as shown in Figure 18, center and right, after deleting α_0 in the second step of the wave move, one of two arcs of $(\alpha \cup \gamma') \cap \mathcal{A}_r$ next to a junction belongs to γ' and the other is in α. These two arcs are depicted in Figure 18, right, and are called *arcs in $\mathcal{A}_r$ next to the cusp of δ*. Moreover, since the intersection points of δ with α and γ all have the same sign, the direction of the two arcs next to the cusp must be opposite in $\mathcal{A}_r$; see Figure 18, right.

Let $\mathcal{J}_r^u$ and $\mathcal{J}_r^d$ be the two junctions in $\mathcal{A}_r$ and denote the two components of $\mathcal{A}_r \setminus (\mathcal{J}_r^u \cup \mathcal{J}_r^d)$ by R_α and R_γ. Since the involution π interchanges the junctions $\mathcal{J}_r^u$ and $\mathcal{J}_r^d$, the induced action of π on each of R_α and R_γ is a 180°-rotation. As a consequence of this symmetry, the two outermost arcs of $(\alpha \cup \gamma') \cap R_\alpha$ must belong to the same curve α or γ'. Since the two arcs in $\mathcal{A}_r$ next to the cusp of δ are in different curves, see Figure 18, right, the two outermost arcs of $(\alpha \cup \gamma') \cap R_\alpha$ must be in the same curve, say α, and the two outermost arcs of $(\alpha \cup \gamma') \cap R_\gamma$ are in γ'. Note that it is possible that $(\alpha \cup \gamma') \cap R_\alpha$ or $(\alpha \cup \gamma') \cap R_\gamma$ is a single arc. Thus, after the first wave move along η, the curves γ' and α have the following *two properties*:

(1) $(\alpha \cup \gamma') \cap R_\gamma \neq \varnothing$ and each outermost arc of $(\alpha \cup \gamma') \cap R_\gamma$ is in γ'.

(2) $(\alpha \cup \gamma') \cap R_\alpha \neq \varnothing$ and each outermost arc of $(\alpha \cup \gamma') \cap R_\alpha$ is in α.

If $\gamma' \cup \alpha$ still intersects $\mathcal{R}^u$ and $\mathcal{R}^d$, then similar to Claim 3.1.4, two subarcs of $\gamma' \cup \alpha$ are $[\delta^+, \delta^-]$-edges and two subarcs of $\gamma' \cup \alpha$ are $[\varepsilon'^+, \varepsilon'^-]$-edges, ruling out a wave with respect to $\{\delta, \varepsilon'\}$. The same argument, as above, implies that there must be a wave η' with respect to $\{\alpha, \gamma'\}$.

The next step is to perform a wave move along η'. Before proceeding, notice that η' has similar properties as η. A property of η that was used in the proof of Claim 3.1.5 is that $\eta \cap \partial_+ P \neq \varnothing$, which follows from Lemma 2.1.6. The next claim says that the same is true for the following wave η'.

Claim 3.1.6 *Any wave with respect to $\{\alpha, \gamma'\}$ must intersect $\partial_+ P$.*

Proof We prove this claim using the property that, before the wave move along η, every wave with respect to $\{\alpha, \gamma\}$ must intersect $\partial_+ P$.

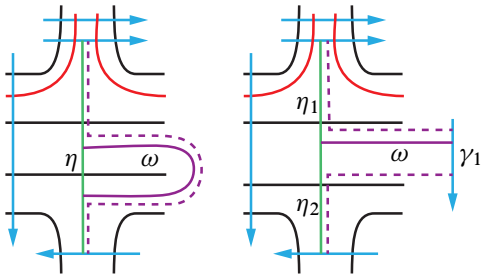

Figure 19: Construct new waves using ω.

Assume for the sake of contradiction that there is a wave ω with respect to $\{\alpha, \gamma'\}$ which is disjoint from $\partial_+ P$. We first show that ω must be a wave with respect to γ'. To see this, consider a pair of pants Q which is a neighborhood of $\gamma \cup \eta$ before the wave move along η. So we may view the three boundary curves of Q as γ, α and γ'. If ω is a wave with respect to α in the Heegaard diagram formed by $\{\alpha, \gamma'\}$ and $\{\delta, \varepsilon'\}$, then $\omega \cap \gamma' = \varnothing$. If ω is also disjoint from γ, then ω is a wave with respect to $\{\alpha, \gamma\}$ and disjoint from $\partial_+ P$, a contradiction to the property that there is no such wave before the wave move along η. Thus, ω must intersect γ and hence ω has subarcs in Q as well as outside Q. As $\omega \cap \gamma' = \varnothing$, this means that a subarc of ω must be a wave with respect to γ in the Heegaard diagram formed by $\{\alpha, \gamma\}$ and $\{\delta, \varepsilon'\}$. This again contradicts the property above since $\omega \cap \partial_+ P = \varnothing$. Therefore, ω must be a wave with respect to γ'.

As described by the surgery-step of the wave move, we may view γ' as the union of η and a subarc γ_1 of γ, ie $\gamma' = \eta \cup \gamma_1$ and $\gamma_1 \subset \gamma$. Since the original γ has no wave disjoint from $\partial_+ P$, there are two possibilities:

(1) $\partial\omega \subset \eta$.

(2) ω has one endpoint in η and the other endpoint in γ_1.

(1) If $\partial\omega \subset \eta$, then the union of ω and two subarcs of η is a wave with respect to γ; see the dashed arc in Figure 19, left. Since $\omega \cap \partial_+ P = \varnothing$, we may view ω as an arc in $\widehat{\Sigma}$. Note that the endpoints of ω cannot be on the same vertical segment of $\eta \setminus \partial_+ P$ as in this case ω must be trivial. Thus, as illustrated in Figure 19, left, this new wave has fewer intersection points with $\partial_+ P$, contradicting our assumption on η.

(2) If ω has one endpoint in η and the other endpoint in γ_1, then the endpoint of ω in η divides η into two subarcs η_1 and η_2. Both $\omega \cup \eta_1$ and $\omega \cup \eta_2$ are waves with respect to γ; see the two dashed arcs in Figure 19, right. Moreover, since $\omega \cap \partial_+ P = \varnothing$, at least one of $\omega \cup \eta_1$ and $\omega \cup \eta_2$ has fewer intersection points with $\partial_+ P$, contradicting our assumption on η.

The two contradictions prove this claim. □

Claim 3.1.6 implies that Claim 3.1.5 also holds for η'. In particular, η' is divided by $\partial_+ P$ into a collection of vertical arcs in $\mathscr{A}_l$ and a pair of junction arcs. Moreover, one of the junction arcs is in $\mathscr{A}_r$ and the other is in $\mathscr{A}_l$; see the dashed arc η in Figure 18, left, for a picture of η'.

Now perform a wave move along η'. If $\partial\eta' \subset \alpha$, call the curve resulting from the wave move an *α-curve*, and if $\partial\eta' \subset \gamma'$, call the resulting curve a *γ-curve*. To simplify notation, we also call α and γ' an α- and a γ-curve, respectively.

Next, apply wave moves repeatedly enabled by Claims 3.1.6 and 3.1.5. To simplify notation, if a wave move is along a wave with respect to an α- or a γ-curve, we always name the resulting curve α' or γ', respectively.

If $(\alpha' \cup \gamma') \cap \mathcal{R}^u \neq \varnothing$ (and $(\alpha' \cup \gamma') \cap \mathcal{R}^d \neq \varnothing$ by symmetry), then similarly to Claim 3.1.4, two subarcs of $\alpha' \cup \gamma'$ are $[\delta^+, \delta^-]$-edges and two subarcs are $[\varepsilon'^+, \varepsilon'^-]$-edges. Moreover, the two properties after the first wave move that were emphasized before Claim 3.1.6 always hold for $\{\alpha', \gamma'\}$. This implies that the Heegaard diagram formed by $\{\delta, \varepsilon'\}$ and $\{\alpha', \gamma'\}$ is not a standard Heegaard diagram of M. Because of the $[\delta^+, \delta^-]$ and $[\varepsilon'^+, \varepsilon'^-]$ blocking-edges, there is no wave with respect to $\{\delta, \varepsilon'\}$. So, Theorems 1.3.7 and 1.3.8 imply that there must be a wave with respect to $\{\alpha', \gamma'\}$. Furthermore, Claims 3.1.6 and 3.1.5 also hold for the new wave with respect to $\{\alpha', \gamma'\}$ and we can continue with the wave moves described above.

Since each wave move reduces the number of intersection points, this process will eventually end. Stop performing the wave moves whenever $(\alpha' \cup \gamma') \cap (\mathcal{R}^d \cup \mathcal{R}^u) = \varnothing$. This means that, at this stage, every arc in $(\alpha' \cup \gamma') \cap \widehat{\Sigma}$ is a vertical arc in either $\mathcal{A}_l$ or $\mathcal{A}_r$.

As above, call each component of $\alpha' \cap \widehat{\Sigma}$ an α'-arc and each component of $\gamma' \cap \widehat{\Sigma}$ a γ'-arc. It follows from our construction that the two properties before Claim 3.1.6 also hold for $\{\alpha', \gamma'\}$. Hence for the rectangles R_α and R_γ we have both that

(1) $(\alpha' \cup \gamma') \cap R_\alpha \neq \varnothing$ and each outermost arc of $(\alpha' \cup \gamma') \cap R_\alpha$ is in α', and

(2) $(\alpha' \cup \gamma') \cap R_\gamma \neq \varnothing$ and each outermost arc of $(\alpha' \cup \gamma') \cap R_\gamma$ is in γ'.

Claim 3.1.7 *If $(\alpha' \cup \gamma') \cap R_\gamma$ (resp. $(\alpha' \cup \gamma') \cap R_\alpha$) contains more than one arc before the wave move, then there is more than one arc in R_γ (resp. R_α), after the wave move.*

If R_γ contains exactly one γ'-arc (resp. R_α contains exactly one α'-arc) before the wave move and if the next wave move is with respect to γ' (resp. α'), then after this wave move, R_γ (resp. R_α) contains more than one arc.

Proof Denote the wave by η''. It follows from Claims 3.1.6 and 3.1.5 that, similarly to the first wave η, the curve $\partial_+ P$ divides η'' into a collection of vertical arcs in $\mathcal{A}_l$ and a pair of junction arcs. As in the argument for η, one of the junction arcs is at a junction in $\mathcal{A}_r$. Again, focus on how curves are changed by this wave move along η'' near this junction.

First, assume that $(\alpha' \cup \gamma') \cap R_\gamma$ contains more than one arc. Then R_γ has two outermost arcs and both are γ'-arcs. As before, perform the next wave move along η'' in two steps:

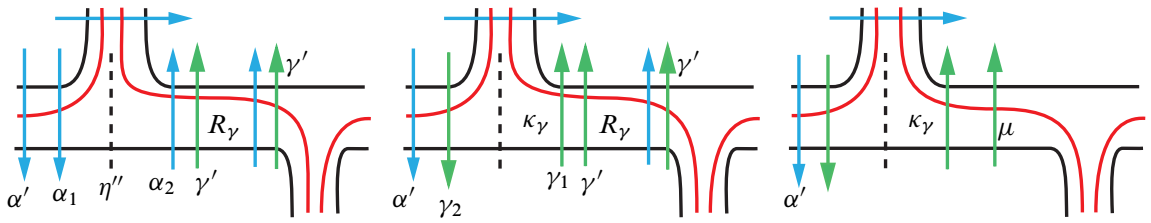

Figure 20: Local pictures of a wave move at a junction in $\mathscr{A}_r$.

If the wave is with respect to α', then the surgery step of the wave move changes α' into two curves α_1 and α_2, where α_2 is parallel to γ'. As illustrated in Figure 20, left, near the junction arc of η'' in $\mathscr{A}_r$, the new arc from the surgery step that lies in R_γ must belong to α_2 since α_2 is parallel to γ'. So after α_2 is deleted in the second step of the wave move, the two outermost γ'-arcs in R_γ remain outermost arcs in R_γ. Hence R_γ still has at least two γ'-arcs, as claimed.

If the wave η'' is with respect to γ', then the surgery step of the wave move changes γ' into two curve γ_1 and γ_2, where γ_2 is parallel to α' and will be removed in the second step of the wave move. Similar to the argument above and as illustrated in Figure 20, center, two new vertical arcs are created near the junction arc of the wave by the surgery step, one in R_α and the other in R_γ. Since γ_2 is parallel to α', the new arc in R_α is in γ_2 and the new arc in R_γ is in γ_1; see Figure 20, center.

Let κ_γ denote this new vertical arc in R_γ next to the junction; see Figure 20, center. The second step of the wave move deletes γ_2 and the curve γ_1 becomes the new γ'-curve. So κ_γ becomes a new outermost γ'-arc in R_γ.

Note that κ_γ cannot be the only arc left because the deleted curve γ_2 is parallel to α', which means that if an arc in R_γ is deleted in the second step, there is at least one α'-arc in R_γ. Hence there are at least two vertical arcs in R_γ after the wave move on η''. In fact, by the symmetry from π, both outermost arcs in R_γ must be γ'-arcs after the wave move.

For the second part of the claim: In the argument above, suppose $(\alpha' \cup \gamma') \cap R_\gamma$ contains only one arc; denote it by μ. So $\mu \subset \gamma'$. As the next wave move is with respect to γ', after the surgery step of the wave move, R_γ contains two arcs κ_γ and μ; see Figure 20, right. Neither κ_γ nor μ will be removed in the second step of the wave move because there is no α'-arc in R_γ and the removed curve is parallel to α'. A similar argument applies to the R_α and α' case. This proves the latter part of the claim. □

Recall that a sequence of wave moves is performed on $\{\alpha', \gamma'\}$ until $(\alpha' \cup \gamma') \cap (\mathcal{R}^d \cup \mathcal{R}^u) = \varnothing$. Once that is achieved, consider the Heegaard diagram given by $\{\delta, \varepsilon'\}$ and $\{\alpha', \gamma'\}$.

As illustrated in Figure 21, the induced orientations (from a fixed orientation on $\partial_+ P$) on the two arcs in $\partial\widehat{\Sigma} \cap \partial\mathscr{A}_l$ are the same along the annulus $\mathscr{A}_l$. Hence the two endpoints of each arc of $(\alpha' \cup \gamma') \cap \mathscr{A}_l$ represent intersection points of $\alpha' \cup \gamma'$ and $\partial_+ P$ with the same sign. Similarly, the two endpoints of each arc of $(\alpha' \cup \gamma') \cap \mathscr{A}_r$ represent points of $(\alpha' \cup \gamma') \cap \partial_+ P$ with the same sign. Since $\alpha' \cup \gamma'$ does not

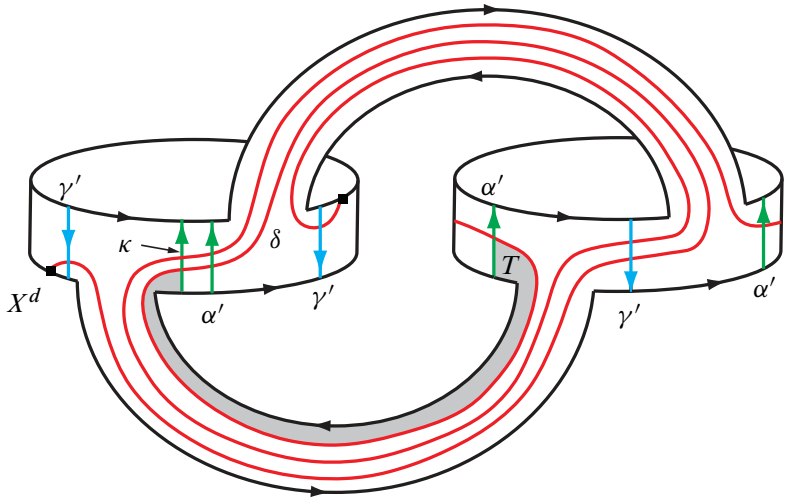

Figure 21: Intersection of α' and γ' with δ and $\widehat{\Sigma}$ after the wave moves.

intersect $\mathcal{R}^d \cup \mathcal{R}^u$ at this stage, this implies that any two points of $\alpha' \cap \partial_+ P$ adjacent along α' must have the same sign. Thus the intersection points of $\alpha' \cap \partial_+ P$ all have the same sign. Similarly, the intersection points of $\gamma' \cap \partial_+ P$ all have the same sign.

As claimed above, after the first wave move along η, R_α contains at least one α'-arc and R_γ contains at least one γ'-arc. Since the intersection points of $\delta \cap (\alpha \cup \gamma)$ all have the same sign before the wave moves, it follows from the wave-move operation that intersection points of $\delta \cap (\alpha' \cup \gamma')$ all have the same sign. Hence the α'-arc and γ'-arc next to the cusp in $\mathcal{A}_r$ must have opposite directions with respect to the core curve of $\mathcal{A}_r$; see Figure 21 and Figure 18, right. This means that the signs of $\alpha' \cap \partial_+ P$ and $\gamma' \cap \partial_+ P$ are opposite. Moreover, since the intersection points of δ with $\alpha' \cup \gamma'$ have the same sign, all the arcs of $R_\alpha \cap (\alpha' \cup \gamma')$ must be α'-arcs and all the arcs in $R_\gamma \cap (\alpha' \cup \gamma')$ must be γ'-arcs, at this stage.

Next, consider $\mathcal{A}_l \cap (\alpha' \cup \gamma')$. Suppose there is a component κ of $\mathcal{A}_l \cap (\alpha' \cup \gamma')$ that meets δ; see Figure 21. The δ-curve from the κ arc to $\mathcal{A}_r$ along $\partial_+ P$ determines a rectangle T, depicted as the shaded region in Figure 21. Two opposite edges of ∂T are subarcs of κ and $(\alpha' \cup \gamma') \cap \mathcal{A}_r$. Another edge of ∂T is a subarc of δ, and the fourth edge is a subarc of $\partial \widehat{\Sigma}$.

Denote by κ' the arc of $(\alpha' \cup \gamma') \cap \mathcal{A}_r$ containing the edge of ∂T opposite to κ. Since the intersection points of δ with $\alpha' \cup \gamma'$ have the same sign, the κ and κ' must have compatible orientation. Since the signs of $\alpha' \cap \partial_+ P$ and $\gamma' \cap \partial_+ P$ are opposite, the arcs κ and κ' are either both α'-arcs or both γ'-arcs; see Figure 21. For the same reason, any arc of $(\alpha' \cup \gamma') \cap \widehat{\Sigma}$ that intersects the rectangle T must also belong to the same curve α' or γ' as κ and κ'.

The discussion above shows that the subarc of δ in ∂T determines an $[\alpha'^-, \alpha'^+]$ blocking-edge if κ is an α'-arc and a $[\gamma'^-, \gamma'^+]$ blocking-edge if κ is a γ'-arc.

Claim 3.1.8 *Either the knot K is doubly primitive and Proposition 3.1.1 holds, or $(\alpha' \cup \gamma') \cap \widehat{\Sigma}$ has two α'-arcs and two γ'-arcs that meet δ and the Heegaard diagram contains both $[\alpha'^-, \alpha'^+]$ and $[\gamma'^-, \gamma'^+]$ blocking-edges.*

Proof Suppose that δ contains no $[\gamma'^-, \gamma'^+]$ blocking-edge. The goal is to show that K must be doubly primitive. The proof for the case that δ contains no $[\alpha'^-, \alpha'^+]$ blocking-edge is similar.

There is no $[\gamma'^-, \gamma'^+]$ blocking-edge, so the argument before the claim implies that no arc of $\mathscr{A}_I \cap (\alpha' \cup \gamma')$ that meets δ can be a γ'-arc. Furthermore, if $R_\gamma \cap (\alpha' \cup \gamma')$ contains more than one arc, since all the arcs of $R_\gamma \cap (\alpha' \cup \gamma')$ are γ'-arcs, two γ'-arcs in R_γ are connected by a subarc of δ in R_γ, which is a $[\gamma'^-, \gamma'^+]$-edge, and this contradicts the assumption. Thus, R_γ contains exactly one γ'-arc.

Since the arcs in $R_\alpha \cap (\alpha' \cup \gamma')$ consist of α'-arcs and since no arc of $\mathscr{A}_I \cap (\alpha' \cup \gamma')$ that meets δ can be a γ'-arc, one concludes that exactly one arc in $\gamma' \cap \widehat{\Sigma}$ intersects δ.

Recall that, after the first wave move, there is at least one arc in R_γ. By Claim 3.1.7, at any stage of the sequence of wave moves, if R_γ contains more than one arc, then R_γ will contain more than one arc regardless of the type of wave move. Since R_γ contains exactly one arc in the end of the process, R_γ must contain exactly one arc after each wave move in the sequence. Furthermore, if any wave move in the sequence, after the first wave move, is with respect to γ', then by the second part of Claim 3.1.7, R_γ will contain more than one arc after that wave move, contradicting the fact that at the last stage there is a single arc. By Lemma 2.1.5 the first wave is with respect to γ. Thus the only possible way to have only one arc in R_γ is if

(a) R_γ contains only one arc after the first wave move on γ, and

(b) all the subsequent wave moves are with respect to the α-curve.

We conclude that the curve γ' will not change and stay fixed after the first wave move. To emphasize this fact, let β denote the image of γ' after the first wave move. Restating the conclusion above: *There is exactly one arc of $\beta \cap \widehat{\Sigma}$ that intersects δ.*

Since the curve β is obtained by a wave move on γ, we can also obtain β by band-summing α with γ. Since $\gamma = \partial C$ and α is the outer boundary of the annulus A, β is the outer boundary of an annulus A' obtained by band-summing A with the disk C. There is exactly one arc of $\beta \cap \widehat{\Sigma}$ that intersects δ. It now follows from Definition 1.2.1 that $c_0(P, D, \beta) = 1$ and thus, by Proposition 1.2.2, K is doubly primitive.

The remaining case is that δ contains no $[\alpha'^-, \alpha'^+]$-edge at the end of the wave-move process. Then by the argument above, we have:

(1) There is exactly one arc of $\alpha' \cap \widehat{\Sigma}$ that intersects δ.

(2) R_α contains only one arc after the first wave move and all the subsequent wave moves are with respect to the γ'.

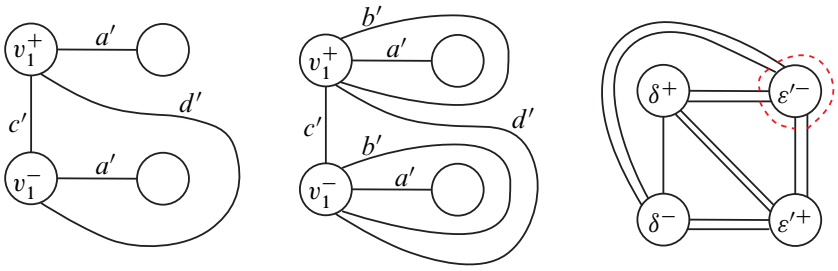

Figure 22: Whitehead graphs of $\Gamma(\{\delta, \varepsilon'\})$.

Since the first wave move is on γ, this means that none of the wave moves in the sequence is on α and the curve α is never changed. So there is exactly one arc of $\alpha \cap \widehat{\Sigma}$ that intersects δ. As before, this means that $c_0(P, D, \alpha) = 1$ and it follows from Proposition 1.2.2 that K is doubly primitive. □

Claim 3.1.9 *Suppose K is not doubly primitive. Then, after the sequence of wave moves, there is no wave with respect to $\{\delta, \varepsilon'\}$.*

Proof First consider the Heegaard diagram formed by curves $\{\delta, \varepsilon'\}$ and $\{\alpha', \gamma'\}$ before the sequence of wave moves has finished. So, at this stage, the curves $\alpha' \cup \gamma'$ still intersect the rectangles $\mathcal{R}^u$ and $\mathcal{R}^d$. As in the argument above and similarly to Claim 3.1.4, this implies that $\alpha' \cup \gamma'$ contains at least two $[\delta^-, \delta^+]$-edges and two $[\varepsilon'^-, \varepsilon'^+]$-edges. Thus, the Whitehead graph $\Gamma(\{\delta, \varepsilon'\})$ is of the type in Figure 2, left, with $c \geq 2$ and $d \geq 2$. If a and b in Figure 2, left, are both zero, then the Heegaard diagram is a standard Heegaard diagram of $L(c, p) \# L(d, q)$. Since $c \geq 2$ and $d \geq 2$, this contradicts our hypothesis that M is either S^3 or $(S^2 \times S^1) \# L(s, t)$. Thus at least one of a and b is nonzero and in particular, the Whitehead graph $\Gamma(\{\delta, \varepsilon'\})$ is connected.

Since the intersection points of δ with α' and γ' all have the same sign, no subarc of δ can be a wave with respect to $\{\alpha', \gamma'\}$. It follows from Lemma 1.3.11 that, if no subarc of ε' is a wave with respect to $\{\alpha', \gamma'\}$, then the Heegaard diagram is as in Figure 7, left, and if a subarc of ε' is a wave with respect to $\{\alpha', \gamma'\}$ then the Heegaard diagram is as in Figure 22, right.

Next, consider the Heegaard diagram determined by $\{\delta, \varepsilon'\}$ and $\{\alpha', \gamma'\}$ after the sequence of wave moves has finished. Assume in contradiction that the claim is false and there is a wave with respect to $\{\delta, \varepsilon'\}$.

Since K is assumed not to be doubly primitive, by Claim 3.1.8, the curve δ contains both $[\alpha'^-, \alpha'^+]$ and $[\gamma'^-, \gamma'^+]$ blocking-edges. Since δ meets both α' and γ' and since the intersection points of δ with α' and γ' all have the same sign, there are subarcs of δ connecting α'^+ to γ'^- and γ'^+ to α'^-. This means that the Whitehead graph $\Gamma(\{\delta, \varepsilon'\})$ is still connected.

Since there is a wave with respect to $\{\delta, \varepsilon'\}$, it follows from Lemma 1.3.11 that $\Gamma(\{\delta, \varepsilon'\})$ must be as in Figure 22, left or center, where v_1 denotes either δ or ε' and all the labels a', b', c', d' are nonzero. In both cases, there are two types of arcs (ie edges marked c' and d' in Figure 22, left and center) connecting v_1^+ to v_1^-, where v_1 is either δ or ε'.

This means that the last wave move (in the sequence of wave moves on $\{\alpha', \gamma'\}$) changes the configuration of $\Gamma(\{\delta, \varepsilon'\})$ from Figure 7, left, or Figure 22, right, to Figure 22, left or center. We will show next that it is impossible to make such a change on the configurations of $\Gamma(\{\delta, \varepsilon'\})$ via wave moves on $\{\alpha', \gamma'\}$. There are two possible cases:

Case (1) Before the last wave move, a subarc of ε' is a wave with respect to $\{\alpha', \gamma'\}$, ie $\Gamma(\{\delta, \varepsilon'\})$ is as shown in Figure 22, right; see Lemma 1.3.11.

In this case, by Lemma 1.3.11, $\Gamma(\{\alpha', \gamma'\})$ is the picture in Figure 4, right. In particular, $\Sigma \setminus (\alpha' \cup \gamma' \cup \delta' \cup \varepsilon')$ consists of four hexagons and a collection of quadrilaterals. Since the intersection points of δ with α' and γ' all have the same sign, no subarc of δ can be a wave. So, the edge labeled b' in Figure 4 denotes b' copies of ε'-arcs. Moreover, all the waves with respect to $\{\alpha', \gamma'\}$ are parallel to these ε'-arcs.

Let ζ_1 and ζ_2 be the two outermost arcs among these b' parallel ε'-arcs. This means that $\zeta_i \subset \varepsilon'$ for $i = 1, 2$, and each ζ_i is an edge of a hexagonal region of $\Sigma \setminus (\alpha' \cup \gamma' \cup \delta' \cup \varepsilon')$. Now we view ζ_i in the Whitehead graph $\Gamma(\{\delta, \varepsilon'\})$. The three dashed arcs in Figure 22, right, are possible pictures of waves parallel and next to ζ_i.

The new meridian after a wave move along ζ_i can be obtained by a surgery using one of the three dashed arcs in Figure 22, right. The effect of this surgery on the configuration of $\Gamma(\{\delta, \varepsilon'\})$ is to merge the two edges of a hexagonal region of $\Sigma \setminus (\alpha' \cup \gamma' \cup \delta' \cup \varepsilon')$ that are adjacent to ζ_i into an arc parallel to the edge opposite to ζ_i in the hexagon. In particular, the wave move along ζ_i does not create any new edge type in $\Gamma(\{\delta, \varepsilon'\})$. However, Figure 22, left or center, has two types of edges (ie edges marked c' and d') connecting v_1^+ to v_1^- but there is only one type of $[\delta^+, \delta^-]$- or $[\varepsilon'^+, \varepsilon'^-]$-edge in Figure 22, right. As the wave move does not create any new edge type in $\Gamma(\{\delta, \varepsilon'\})$, the wave move can never change the configuration of $\Gamma(\{\delta, \varepsilon'\})$ from Figure 22, right, to Figure 22, left or center.

Case (2) No subarc of ε' is a wave with respect to $\{\alpha', \gamma'\}$ before the last wave move.

By Lemma 2.1.9, each wave in the sequence of wave moves must connect a pair of opposite sides of an octagon; see the wave η in Figure 7, center. Similar to Case (1), the effect of this wave move on the configuration of $\Gamma(\{\delta, \varepsilon'\})$ is to merge the two opposite edges of an octagon region of $\Sigma \setminus (\alpha' \cup \gamma' \cup \delta' \cup \varepsilon')$ into another edge of this octagon. In particular, the wave move does not create any new edge type in $\Gamma(\{\delta, \varepsilon'\})$. Thus, similar to Case (1), the wave move can never change the configuration of $\Gamma(\{\delta, \varepsilon'\})$ from Figure 7, left, to Figure 22, left or center.

Therefore, in both cases, $\Gamma(\{\delta, \varepsilon'\})$ cannot be changed to Figure 22, left or center, by the sequence of wave moves and this is a contradiction which proves the claim. □

We now finish the proof of Proposition 3.1.1. By Lemma 1.3.9 and Claims 3.1.8 and 3.1.9 we conclude that there is no wave in the Heegaard diagram after these wave moves. Since the intersection pattern of the curves rules out a standard Heegaard splitting, this is a contradiction to Theorems 1.3.7 and 1.3.8. The only possibility left by Claim 3.1.8 is that K is doubly primitive. □

3.2 The curve δ takes four short paths

This section deals with Configuration 3, where δ takes two short paths in each annulus.

Proposition 3.2.1 *If the curve δ takes two short paths in each of the annuli, then K is doubly primitive.*

Proof The train track τ_D corresponding to this case is obtained from the train track in Proposition 3.1.1 by attaching another segment. Thus, as explained at the beginning of Section 3, the complement $\Sigma \setminus \mathcal{N}(\tau_D)$ is a disk with six cusps on its boundary. So one can think of $\Sigma \setminus \mathcal{N}(\tau_D)$ as a hexagon.

As in Section 3.1, the two cusps at $\partial\rho_x$ are called the *ρ-cusps* and the other four cusps called the *j-cusps*. The j-cusps are located at the four junctions of $\mathcal{A}_l$ and $\mathcal{A}_r$ with the cusp directions pointing into the rectangles $\mathcal{R}^d$ and $\mathcal{R}^u$. The configuration in each annulus is similar to the picture of the right annulus in Figure 13.

The train track τ_D has a special segment ρ_x containing the intersection point X and the weight of δ at ρ_x is one (as in Section 3.1). As before, assume that ρ_x and the hexagon $\Sigma \setminus \mathcal{N}(\tau_D)$ are invariant under the involution π. Similar to Section 3.1, the complement of $\tau_D \setminus \rho_x$ is an annulus as in shown Figure 14 and Figure 23, left, and $\partial_+ P$ must be a core curve of this annulus. Since τ_D and $\partial_+ P$ are invariant under π the intersection point $X = \rho_x \cap \partial_+ P$ is a fixed point of π and hence π interchanges the two ρ-cusps; see Figure 23, left. Moreover, as described in Remark 2.2.1, the involution π interchanges the two junctions in $\mathcal{A}_l$ and the two junctions in $\mathcal{A}_r$.

As indicated by Figure 23, left, the two sides of a small neighborhood $\mathcal{N}(\rho_x)$ are parts of a pair of opposite edges of the hexagon $\Sigma \setminus \mathcal{N}(\tau_D)$. Hence the action of π on the hexagon $\Sigma \setminus \mathcal{N}(\tau_D)$ interchanges these opposite edges. Moreover, as $\partial_+ P$ is invariant under π, the arc $\partial_+ P \cap (\Sigma \setminus \mathcal{N}(\tau_D))$ connects a pair of opposite boundary edges of the hexagon $\Sigma \setminus \mathcal{N}(\tau_D)$.

In $\mathcal{N}(\tau_D)$ there are three splitting arcs (defined in Definition 2.2.7) connecting the cusps of $\mathcal{N}(\tau_D)$ in pairs which are disjoint from δ. When $\mathcal{N}(\tau_D)$ is cut open along these three splitting arcs, the resulting surface is a product neighborhood of δ. Denote the three splitting arcs by s_ρ, s_l and s_r.

By connecting the endpoints of each splitting arc via an arc in the disk $\Sigma \setminus \mathcal{N}(\tau_D)$, we obtain a nontrivial simple closed curve. Since every nontrivial simple closed curve is invariant under the involution π after isotopy, the two cusps at the endpoints of each splitting arc must be invariant under π. Since the involution

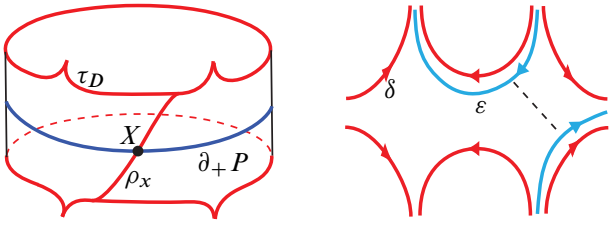

Figure 23: Configurations of $\Sigma \setminus \tau_D$ and s-wave for ε.

interchanges the two ρ-cusps and interchanges the two junctions in each annulus $\mathcal{A}_l$ and $\mathcal{A}_r$, this implies that s_ρ connects the two ρ-cusps, s_l connects the two j-cusps at $\mathcal{A}_l$, and s_r connects the two j-cusps at $\mathcal{A}_r$.

As the cusp directions of the j-cusps point into the two rectangles $\mathcal{R}^d$ and $\mathcal{R}^u$, both s_l and s_r contain subarcs that are core arcs of $\mathcal{R}^d$ and $\mathcal{R}^u$. Similar to Section 3.1, this implies that each of s_l and s_r intersects γ at least twice.

Now we consider s_ρ. Similar to the proof of Claim 3.1.2 in Section 3.1, each end of ρ_x is attached either to a long path or to a component of $\delta \cap (\mathcal{R}^d \cup \mathcal{R}^u)$. Since there is no long path in the current configuration, the two ends of ρ_x are attached to two components of $\delta \cap (\mathcal{R}^d \cup \mathcal{R}^u)$ that are core arcs of $\mathcal{R}^d$ and $\mathcal{R}^u$, respectively. As in the proof of Claim 3.1.2, this implies that s_ρ also contains subarcs that are core arcs of $\mathcal{R}^d$ and $\mathcal{R}^u$, and s_ρ intersects γ at least twice.

By part (2) of Proposition 2.1.8, we may assume that the orientations of δ and ε are compatible in $\mathcal{N}(\tau_D)$; see the arrows of the arcs in Figure 23, right. The surface $\Sigma \setminus \mathcal{N}(\tau_D)$ is a disk with six cusps at its boundary. Similar to the s-wave illustrated in Figure 14, perform wave moves on ε along all possible s-waves in $\Sigma \setminus \mathcal{N}(\tau_D)$; such s-waves are indicated by the dashed arc in Figure 23, right.

Notice that if there are two distinct ε-arcs in the disk $\Sigma \setminus \mathcal{N}(\tau_D)$ that go into two adjacent cusps, then there is an s-wave connecting these two arcs; see Figure 23, right. We perform all possible wave moves along such s-waves on ε and denote the resulting meridian of W by ε'.

Suppose ε' no longer admits any s-wave as in Figure 23, right. In particular, similarly to the argument before Claim 3.1.3, ε' does not pass through all cusps of $\mathcal{N}(\tau_D)$. Consider different types of ε'-arcs in the cusped disk $\Sigma \setminus \tau_D$. The induced action by π on the disk $\Sigma \setminus \tau_D$ is a 180°-rotation, so the ε'-arcs in the cusped disk $\Sigma \setminus \tau_D$ are symmetric with respect to the rotation. Since ε' admits no s-wave in the cusped disk $\Sigma \setminus \tau_D$, the ε'-arcs in $\Sigma \setminus \tau_D$ have three possible configurations, up to rotation and reflection, which are indicated in Figure 24:

(1) There are three types of ε'-arcs in $\Sigma \setminus \tau_D$: diagonals connecting a pair of opposite cusps and arcs parallel to a pair of opposite edges next to the diagonal; see the blue arcs in Figure 24, left.

(2) All the ε'-arcs in $\Sigma \setminus \tau_D$ are parallel to a diagonal connecting a pair of opposite cusps; see the blue arcs in Figure 24, center (ignore the dashed arcs for now).

(3) The ε'-arcs in $\Sigma \setminus \tau_D$ are all parallel to a pair of opposite edges; see the blue arcs in Figure 24, right (ignore the dashed arcs for now).

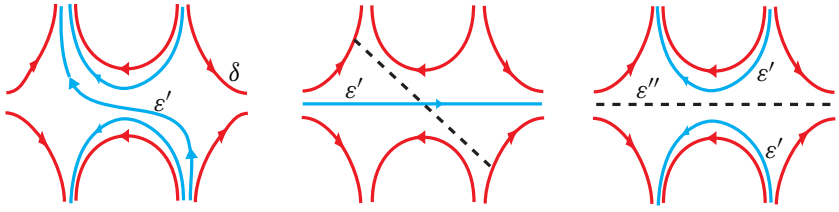

Figure 24: Possible configurations of ε' in $\Sigma \setminus \mathcal{N}(\tau_D)$.

Now consider the Heegaard diagram given by $\{\delta, \varepsilon'\}$ and $\{\alpha, \gamma\}$. There are three cases to discuss.

Case (a) Assume there is a cusp of $\mathcal{N}(\tau_D)$ that contains two ε'-arcs. In other words, two components of $\varepsilon' \cap \mathcal{N}(\tau_D)$ are parallel to the same splitting curve s_ρ, s_l or s_r.

Note that this case includes the configuration in Figure 24, left, since there are multiple curves going into a pair of opposite cusps and since each splitting arc s_ρ, s_l or s_r connects a pair of opposite cusps.

Similar to Claim 3.1.4, since each s_ρ, s_l or s_r intersects γ at least twice and since the orientations of arcs of $\varepsilon' \cap \mathcal{N}(\tau_D)$ are compatible, γ contains two distinct subarcs that are both $[\varepsilon'^-, \varepsilon'^+]$-edges. As explained before Case (a), there is a cusp that contains no ε'-arcs. This means that γ also contains two distinct subarcs that are $[\delta^-, \delta^+]$-edges.

Now the proof for Case (a) is identical to the proof in Section 3.1. Recall that in Claim 3.1.4, we have the same conclusion, which implies that there is no wave with respect to $\{\delta, \varepsilon'\}$ and hence there must a wave η with respect to $\{\alpha, \gamma\}$. As in the argument between Claims 3.1.4 and 3.1.5, we can isotope η to intersect $\partial_+ P$ efficiently. In the current setting, $\mathscr{A}_r$ also has two short paths, which is the same situation as in Section 3.1. Thus Claims 3.1.5–3.1.9 also hold and hence Proposition 3.2.1 holds in Case (a).

Note that the setting in this section is simpler than Section 3.1. The annulus $\mathscr{A}_l$ has two short paths in this setting while $\mathscr{A}_l$ has only one short path in Section 3.1. The additional short path in the current setting could be ignored and the proof can proceed as in Section 3.1.

Next, assume that Case (a) does not occur, ie each cusp of $\mathcal{N}(\tau_D)$ has at most one ε'-arc. In particular, the first possible configuration of ε'-arcs in $\Sigma \setminus \tau_D$, ie Figure 24, left, does not occur. Thus, it remains to discuss the second and third possible configurations, ie Figure 24, center and right.

Case (b) The configuration for $\varepsilon' \cap (\Sigma \setminus \mathcal{N}(\tau_D))$ is as in Figure 24, center.

Since Case (a) does not occur, there is only a single ε'-arc in Figure 24, center, going into the cusp. So the intersection $\varepsilon' \cap (\Sigma \setminus \mathcal{N}(\tau_D))$ is a single arc forming a main diagonal of the hexagon $\Sigma \setminus \mathcal{N}(\tau_D)$. We concluded earlier in the proof of Proposition 3.2.1 that $\partial_+ P \cap (\Sigma \setminus \mathcal{N}(\tau_D))$ is an arc connecting a pair of opposite edges of $\Sigma \setminus \mathcal{N}(\tau_D)$, indicated by the dashed arc in Figure 24, center (also see Figure 23, left), for a picture of how $\partial_+ P$ intersects $\Sigma \setminus \mathcal{N}(\tau_D)$. As shown in Figure 24, center, this conclusion on $\partial_+ P \cap (\Sigma \setminus \mathcal{N}(\tau_D))$ implies that ε' intersects $\partial_+ P$ in exactly one point. Let E' is the disk in W bounded by ε', then (P, E') is a $(\mathcal{P}, \mathcal{D})$-pair. The configuration of ε' implies that

$$|(\alpha \cup \gamma) \cap \varepsilon'| < |(\alpha \cup \gamma) \cap \delta|.$$

Since δ takes two paths in each annulus, every component of $(\alpha \cup \gamma) \cap \widehat{\Sigma}$ must intersect δ and

$$c_0(P, E', \alpha, \gamma) \leq |(\alpha \cup \gamma) \cap \widehat{\Sigma}| = c_0(P, D, \alpha, \gamma).$$

Hence $c(P, E', \alpha, \gamma) < c(P, D, \alpha, \gamma)$, contradicting Assumption 1.2.3 regarding the choice of the pair (P, D).

Case (c) The configuration for $\varepsilon' \cap (\Sigma \setminus \mathcal{N}(\tau_D))$ is as in Figure 24, right.

In this case, we can use a main diagonal of the hexagon $\Sigma \setminus \mathcal{N}(\tau_D)$, see the dashed arc in Figure 24, right, to connect the two ends of the splitting arc of $\mathcal{N}(\tau_D)$ which connects (in $\mathcal{N}(\tau_D)$) the two cusps that ε' does not pass through. This gives a nonseparating simple closed curve ε'' that is disjoint from both δ and ε'; see the dashed arc in Figure 24, right, for a picture of ε'' within $\Sigma \setminus \mathcal{N}(\tau_D)$. As $\{\delta, \varepsilon'\}$ is a complete set of meridians for W, ε'' must also bound a disk in W. Now consider $\{\delta, \varepsilon''\}$. The configuration of ε'' is the same as that of ε' in Case (b). So we can apply the argument in Case (b) to ε'' and obtain the same contradiction to Assumption 1.2.3.

This finishes the proof of the proposition. □

3.3 The curve δ takes a long path

In this section we consider Configuration 4, namely when δ takes a long path in one (or both) of the annuli.

Proposition 3.3.1 *If δ takes a long path in an annulus, then K is doubly primitive or there is an isotopy of the annulus which converts the long path to a short path.*

Proof Without loss of generality, suppose δ takes a long path in $\mathcal{A}_r$.

First note that in some special configurations one can perform an isotopy on $\mathcal{A}_r$ so that δ only takes short paths in $\mathcal{A}_r$. For example, if δ only wraps around $\mathcal{A}_r$ once and if $\alpha \cup \gamma$ does not intersect the shaded region in Figure 25, then after an isotopy in Figure 25, δ only takes short paths in $\mathcal{A}_r$. This isotopy is equivalent to choosing a different product structure for $\mathcal{A}_r$ provided that arcs of $(\alpha \cup \gamma) \cap \mathcal{A}_r$ remain vertical in the new product structure. So assume that δ takes a long path in $\mathcal{A}_r$ after all such isotopies have been performed. In other words, assume that $\delta \cap \mathcal{A}_r$ (similarly $\delta \cap \mathcal{A}_l$ if it contains a long path) cannot be converted to short paths by such an isotopy.

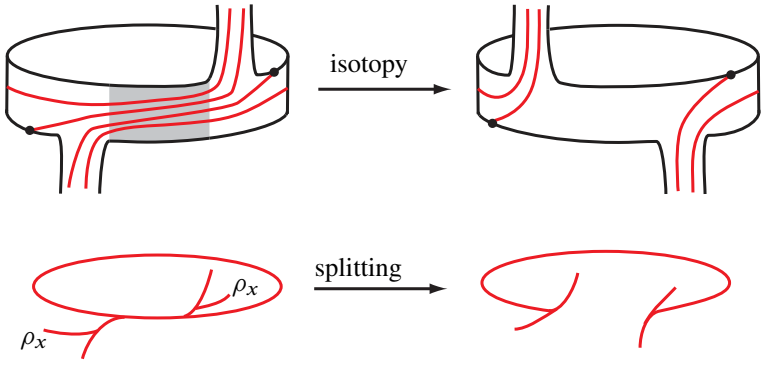

Figure 25: An isotopy converting a long path to a short path.

Claim 3.3.2 *Either Proposition 3.3.1 holds or there is a meridian ε' of W, disjoint from δ, such that there is no wave with respect to $\{\delta, \varepsilon'\}$ in the Heegaard diagram determined by $\{\alpha, \gamma\}$ and $\{\delta, \varepsilon'\}$.*

Proof Consider the possible configurations of $\delta \cap \mathcal{A}_l$, ie the intersection of δ with the other annulus. There are two cases to consider:

Case (1) The curve δ either takes two short paths, or a long path in $\mathcal{A}_l$.

Note that it is still assumed that if δ takes a long path in $\mathcal{A}_l$, and there is no isotopy as in Figure 25 that converts the long path into a short path. Furthermore, as explained at the beginning of Section 3, the surface $\Sigma \setminus \tau_D$ in this case is a disk with six cusps along its boundary. The discussion in this case is similar to Section 3.2. For convenience here is a summary of the setup in Section 3.2:

There are three splitting curves in $\mathcal{N}(\tau_D)$ which are disjoint from δ and which connect the cusps of $\mathcal{N}(\tau_D)$ in pairs. When the surface $\mathcal{N}(\tau_D)$ is cut open along these three splitting curves, the resulting surface is a product neighborhood of δ. As in Section 3.2, each splitting arc is invariant under π. Hence the involution interchanges the two cusps at the endpoints of each splitting arc. The involution interchanges the two ρ-cusps, and also interchanges the other two cusps in $\mathcal{A}_l$ and $\mathcal{A}_r$ respectively. Thus a splitting arc, denoted by s_ρ, connects the two ρ-cusp. The other two splitting arcs, denoted by s_l and s_r, connect two cusps in $\mathcal{A}_l$ and two cusps in $\mathcal{A}_r$, respectively.

The *difference* between the current setting and Section 3.2 is that in the current setting the splitting arcs s_ρ, s_l and s_r may not go through the two rectangles $\mathcal{R}^u$ and $\mathcal{R}^d$. In fact, it is not immediately clear that the three splitting arcs even intersect $\alpha \cup \gamma$.

First consider s_ρ. The proof of Claim 3.1.2 on ρ_x implies that either there is a long path or $(\alpha \cup \gamma) \cap s_\rho \neq \varnothing$. Moreover, it follows from Definition 2.2.5 that the only configuration where $(\alpha \cup \gamma) \cap s_\rho = \varnothing$ is as shown in the left picture of Figure 25, where s_ρ is a short arc cutting through the shaded region of Figure 25 and the shaded region does not contain any α- or γ-arc. Thus, if $(\alpha \cup \gamma) \cap s_\rho = \varnothing$, we can perform an isotopy as shown in Figure 25. The isotopy can be viewed as a change of the product structure of the annulus. Since $(\alpha \cup \gamma) \cap s_\rho = \varnothing$, $\alpha \cup \gamma$ does not intersect the shaded region in Figure 25. Hence, the isotopy does not affect α and γ. Note that the effect of this isotopy on the train track is a splitting along s_ρ; see Figure 25. This isotopy converts a long path to a short path, contradicting our assumption, before the claim, that there is no such isotopy. Thus, $(\alpha \cup \gamma) \cap s_\rho \neq \varnothing$.

Similarly, if $(\alpha \cup \gamma) \cap s_l = \varnothing$ (or $(\alpha \cup \gamma) \cap s_r = \varnothing$), then the configuration of s_l (or s_r) must be a short arc cutting through the shaded region of Figure 25 and the shaded region does not contain any α- or γ-arc. We can perform a similar isotopy (without changing $\alpha \cup \gamma$) which changes the long path into a short path, also contradicting the assumption that there is no such isotopy.

Therefore, s_ρ, s_l and s_r must all meet $\alpha \cup \gamma$. Now the argument is the same as Section 3.2. Performing wave moves on ε in the hexagon $\Sigma \setminus \tau_D$ along s-waves, a meridian ε' is obtained. As in Section 3.2, there are three possible configurations for the ε'-arcs in the hexagon $\Sigma \setminus \tau_D$. If the configuration is as in

Case (a) in Section 3.2, since s_ρ, s_l and s_r all meet $\alpha \cup \gamma$, $\alpha \cup \gamma$ contains subarcs that are $[\delta^-, \delta^+]$ and $[\varepsilon'^-, \varepsilon'^+]$ blocking-edges. By Lemma 1.3.9, there is no wave with respect to $\{\delta, \varepsilon'\}$ and the claim holds. If the configuration is as in Cases (b) and (c) in Section 3.2, then Assumption 1.2.3 is contradicted.

Case (2) The curve δ takes one short path in $\mathscr{A}_l$.

In this case, as explained at the beginning of Section 3, the surface $\Sigma \setminus \mathcal{N}(\tau_D)$ is an annulus with two cusps on each boundary curve. As in the argument before Claim 3.1.3, perform wave moves on ε along s-waves as in Figure 14 to obtain a new meridian ε' of W. In particular, the arcs in $\varepsilon' \cap \mathcal{N}(\tau_D)$, if there are any, are all parallel to one of the two splitting arcs of $\mathcal{N}(\tau_D)$. Furthermore, Claim 3.1.3 also holds for ε' in the current setting, since $\mathcal{N}(\tau_D)$ in Section 3.1 is of the same topological type. So the intersection $\varepsilon' \cap \mathcal{N}(\tau_D)$ contains at least two arcs, all parallel to a splitting arc.

Since it is assumed that a long path cannot be converted to a short path, the argument in Case (1) implies that both splitting arcs of $\mathcal{N}(\tau_D)$ must intersect $\alpha \cup \gamma$. Similarly to Claim 3.1.4, this implies that $\alpha \cup \gamma$ contains subarcs that are $[\delta^-, \delta^+]$ and $[\varepsilon'^-, \varepsilon'^+]$ blocking-edges. By Lemma 1.3.9, this means that there is no wave with respect to $\{\delta, \varepsilon'\}$ as claimed. □

We are now in position to finish the proof of Proposition 3.3.1. By Claim 3.3.2, we may assume that there is no wave with respect to $\{\delta, \varepsilon'\}$. By part (1) of Proposition 2.1.8, the intersection points of $\delta \cap \gamma$ all have the same sign. Next, we consider the core curve $\mathfrak{a}_r$ of $\mathscr{A}_r$. Since δ takes a long path in $\mathscr{A}_r$, an arc in $\delta \cap \mathscr{A}_r$ wraps around the annulus at least once. So the core curve $\mathfrak{a}_r$ is carried by the train track τ_D. Hence $\mathfrak{a}_r$ admits a direction induced from the orientation of δ and τ_D. Since points of $\delta \cap \gamma$ all have the same sign, the intersection points of $\mathfrak{a}_r \cap \gamma$ all have the same sign.

Recall that in our construction (see Remark 2.1.2), an arc in $\gamma \cap \mathcal{R}^u$ determines a ∂-compression disk for P. When P is ∂-compressed along this disk, we obtain two planar surfaces P_l and P_r in the compression body U with $\mathfrak{a}_l = \partial_+ P_l$ and $\mathfrak{a}_r = \partial_+ P_r$. Since γ bounds a disk C in U, if $C \cap P_r \neq \varnothing$, the two endpoints of each arc in $C \cap P_r$ are points in $\gamma \cap \mathfrak{a}_r$ with opposite signs of intersection, a contradiction to the conclusion that the intersection points of $\mathfrak{a}_r \cap \gamma$ all have the same sign. Thus $C \cap P_r = \varnothing$ and hence $\gamma \cap \mathfrak{a}_r = \varnothing$. By Lemma 1.2.4, this means that P_r must be a vertical annulus in the compression body U. Moreover, since the surgery slope is an integer, after isotopy, $P_r \cap A$ is a single vertical arc, where A is the annulus with $\partial_+ A = \alpha$. Hence $\alpha \cap \mathfrak{a}_r$ is a single point. So we have $\mathscr{A}_r \cap \gamma = \varnothing$ and $\mathscr{A}_r \cap \alpha$ is a single arc.

Since there is no wave with respect to $\{\delta, \varepsilon'\}$, by Theorems 1.3.7 and 1.3.8, there must be a wave with respect to $\{\alpha, \gamma\}$ and by Lemma 2.1.5, this wave must be with respect to γ. If a component of $\delta \cap \mathscr{A}_r$ in the long path intersects the arc $\alpha \cap \mathscr{A}_r$ more than once, since $\mathscr{A}_r \cap \gamma = \varnothing$, a subarc of $\delta \cap \mathscr{A}_r$ must be a spiral around $\mathscr{A}_r$ connecting the plus-side of α to its minus-side. This means that a subarc of δ is an $[\alpha^+, \alpha^-]$ blocking-edge. However, by the proof of Lemma 1.3.9, the existence of an $[\alpha^+, \alpha^-]$ blocking-edge means that there is no wave with respect to γ, a contradiction. So each component of $\delta \cap \mathscr{A}_r$ in the long path

intersects the arc $\alpha \cap \mathcal{A}_r$ only once. This means that the shaded region in Figure 25 does not contain any α-arc. As $\mathcal{A}_r \cap \gamma = \varnothing$, the shaded region in Figure 25 intersects neither α or γ. Hence we can perform an isotopy as in Figure 25, changing the long path into a short path with respect to the new product structure. This contradicts our assumption at the beginning of the proof. Therefore, Proposition 3.3.1 holds. □

3.4 The curve δ takes one short path in each annulus, a special case

In this and the next section we rule out the remaining possible configuration, namely, when δ takes only two short paths: One in $\mathcal{A}_l$ and one in $\mathcal{A}_r$. As δ takes *only two short paths* it imposes the least restrictions on the possible configurations and hence this is the most complicated case. We therefore begin by considering a special case in this section and will deal with the general case in the following section. The configurations and proof for the special case are similar to those of Sections 3.1 and 3.3.

As before, we assume δ and $\partial_+ P$ are invariant under π and $X = \delta \cap \partial_+ P$ is a fixed point of π. By symmetry (2) in Remark 2.2.1, we may assume $X^d \cup X^u$ lies in either ∂A_l or $\partial \mathcal{R}^d \cup \partial \mathcal{R}^u$. Moreover, since symmetry (3) in Remark 2.2.1 interchanges the two short paths in $\mathcal{A}_l$, we may assume δ takes a fixed short path in $\mathcal{A}_l$.

Consider the train track τ_D. As explained at the beginning of this section, the shape of τ_D must be as shown in Figure 12. So both $\mathcal{N}(\tau_D)$ and $\Sigma \setminus \mathcal{N}(\tau_D)$ are once-punctured tori. The train track τ_D has a special segment ρ_x which contains the point $X = \delta \cap \partial_+ P$, and the weight of δ at ρ_x is one.

The special case that we discuss in this section is:

Special case There is a component κ of $(\alpha \cup \gamma) \setminus \partial_+ P$ such that $\kappa \cap \tau_D = \kappa \cap \rho_x$ contains more than one point; see Figure 26 where the dashed arc is κ.

Consider the arc κ. By the construction of τ_D, the intersection $\kappa \cap \rho_x$ consists of exactly two points, see Figure 26. There are four possible configurations for $\delta \cap \mathcal{A}_l$ as shown in Figure 26.

Let κ' be the subarc of κ between the two points of $\kappa \cap \rho_x$. By collapsing the arc κ' to a point, we can pinch τ_D to a new train track τ'_D that also fully carries δ. The train tracks $\tau'_D \cap \mathcal{A}_l$ corresponding to the four possible configurations in Figure 26 are shown in Figure 27, respectively. Notice that, if $\mathcal{A}_l$ and $\mathcal{A}_r$ are switched, the train tracks τ'_D in Figure 26(a)–(b) are basically the same as the train track τ_D for two short paths in Section 3.1. Also, the train tracks τ'_D for Figure 26(c)–(d) are basically the same as the train track τ_D for a long path in Section 3.3. The proofs for the two types of configurations will also be the same as the arguments in Sections 3.1 and 3.3, respectively.

Recall the definition of a splitting arc in Definition 2.2.7. The original train track τ_D has two cusps which survive the pinching of κ'. The splitting arc in $\mathcal{N}(\tau_D)$ that connects the two cusps and which is disjoint from δ in $\mathcal{N}(\tau_D)$ was denoted by s_j. Note that Lemma 2.1.4 implies that s_j passes through both $\mathcal{R}^d$ and $\mathcal{R}^u$. Denote the image of s_j in $\mathcal{N}(\tau'_D)$ by s'_j. Clearly, s'_j passes through both rectangles $\mathcal{R}^u$ and $\mathcal{R}^d$.

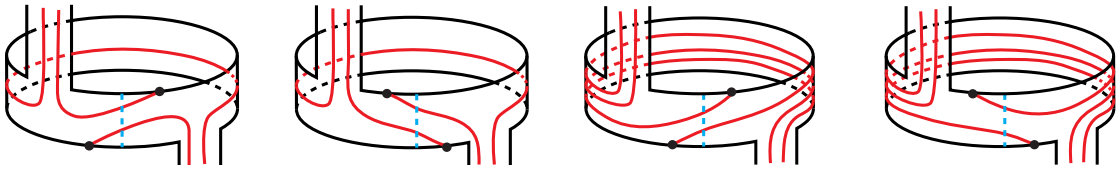

Figure 26: Possible configurations of ρ_x and κ.

For the train track τ_D', we have an additional splitting arc, which we denote by s_ρ', corresponding to the pinching along κ, ie cutting $\mathcal{N}(\tau_D')$ along s_ρ' will undo the pinching operation along κ. Since κ is an arc in $(\alpha\cup\gamma)\cap\hat{\Sigma}$, this new splitting arc s_ρ' intersects $\alpha\cup\gamma$ at least once. Furthermore as $\varepsilon\cap\delta=\varnothing$, the intersection $\varepsilon\cap\mathcal{N}(\tau_D')$ consists of arcs parallel to s_j' and s_ρ'. In all four possible configurations in the Figure 27, the surface $\Sigma\setminus\mathcal{N}(\tau_D')$ is an annulus with two cusps at each boundary curve; see Figure 14 in Section 3.1 for a picture. Now follow the argument in Section 3.1. The first step is to perform a maximal number of wave moves on ε along s-waves in the annulus $\Sigma\setminus\mathcal{N}(\tau_D')$; see Figure 14 for a picture of such s-waves. Denote the resulting meridian by ε'. As in Section 3.1, ε' does not pass through all the cusps; in other words, $\varepsilon'\cap\mathcal{N}(\tau_D')$ consists of arcs all of which are parallel to either s_j' or s_ρ'.

First note that Claim 3.1.3 is also valid for $\mathcal{N}(\tau_D')$ and ε', that is, $\varepsilon'\cap\mathcal{N}(\tau_D')$ contains at least two arcs. The slight difference is that, under the setting of Section 3.1, the proof of Claim 3.1.3 needs the fact that the core curve $\mathfrak{a}_r$ of the annulus $\mathcal{A}_r$ can be isotoped into $\mathcal{N}(\tau_D)$, while in the current setting, the core curve $\mathfrak{a}_l$ of $\mathcal{A}_l$ can be isotoped into $\mathcal{N}(\tau_D')$ in all four configurations of Figure 27.

Similarly to Claim 3.1.4, since both s_j' and s_ρ' intersect $\alpha\cup\gamma$, this means that $\alpha\cup\gamma$ contains both $[\varepsilon'^-,\varepsilon'^+]$- and $[\delta^-,\delta^+]$-edges. So, by Lemma 1.3.9, there is no wave with respect to $\{\delta,\varepsilon'\}$. Moreover, since s_j' passes through both $\mathcal{R}^d$ and $\mathcal{R}^u$, s_j' meets γ at least twice and there are two $[\varepsilon'^-,\varepsilon'^+]$- or $[\delta^-,\delta^+]$-edges. This means that the Heegaard diagram formed by $\{\alpha,\gamma\}$ and $\{\delta,\varepsilon'\}$ is not a standard Heegaard diagram of S^3 or $(S^2\times S^1)\#L(p,q)$. Hence there must be a wave with respect to $\{\alpha,\gamma\}$. By Lemma 2.1.5, this wave must be with respect to γ.

If the configuration of ρ_x is as shown in Figures 26(c) or (d), then the corresponding configuration of the train track τ_D' is Figure 27(c) or (d), respectively. In both configurations, the core curve $\mathfrak{a}_l$ of $\mathcal{A}_l$ is carried by the train track τ_D' and hence has an orientation compatible with the orientation of δ. The argument for these two configurations is the same as in Section 3.3: Since there is no wave with respect to $\{\delta,\varepsilon'\}$, by part (1) of Proposition 2.1.8, the intersection points of $\delta\cap\gamma$ all have the same sign. As in

Figure 27: Possible configurations of τ_D' at $\mathcal{A}_l$.

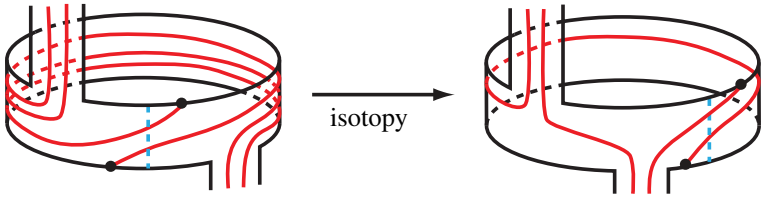

Figure 28: An isotopy converting Figure 26(c) to (b) up to symmetry.

Section 3.3, this implies that $\mathfrak{a}_l$ and hence the planar surface P_l is disjoint from γ and the disk C. As in Section 3.3, this means that P_l is an annulus and $\mathfrak{a}_l \cap \alpha$ is a single point. Thus the arc κ is the only arc in $(\alpha \cup \gamma) \cap \mathscr{A}_l$. As in Section 3.3 and Figure 25, we can perform an isotopy in $\mathscr{A}_l$ which fixes κ and twists $\mathscr{A}_l$ along the short path, such that the cusp directions point into the rectangles $\mathcal{R}^d$ and $\mathcal{R}^u$ after the isotopy; see Figure 28. So this isotopy basically converts the configurations of Figure 26(c) and (d) to the configurations of Figure 26(b) and (a), up to symmetry. Moreover the isotopy does not affect $\alpha \cup \gamma$.

So it remains to consider the configurations in Figure 26(a)–(b). The next step in the proof is to convert the configurations in Figure 26(a)–(b) to the setup in Section 3.1 and apply the arguments used in that section.

Remark 3.4.1 (1) Recall that, for the train track in Section 3.1, both splitting arcs s_j and s_ρ pass through both rectangles $\mathcal{R}^d$ and $\mathcal{R}^u$. In the present case $\mathcal{N}(\tau'_D)$ also has two splitting arcs s'_j and s'_ρ where s'_ρ is the short splitting arc corresponding to the pinching operation. In the current setting for $\mathcal{N}(\tau'_D)$, the splitting arc s'_j indeed passes through both rectangles $\mathcal{R}^d$ and $\mathcal{R}^u$ but the splitting arc s'_ρ is too short and does *not pass* through the rectangles.

(2) The reason we need the property that both splitting arcs pass through both rectangles is the following: In the argument in Section 3.1 (after Claim 3.1.3), we perform a sequence of wave moves on $\{\alpha, \gamma\}$ and obtain a new set of meridians $\{\alpha', \gamma'\}$. Each α'- or γ'-arc in $\mathcal{R}^d$ and $\mathcal{R}^u$ intersects the splitting arcs s_j and s_ρ and hence contain $[\varepsilon'^-, \varepsilon'^+]$- and $[\delta^-, \delta^+]$-edges. Thus we can continue the sequence of wave moves *until* α' *and* γ' *do not intersect* $\mathcal{R}^d$ *and* $\mathcal{R}^u$ *at all.* This implies that the intersection points of α' and γ' with $\partial_+ P$ have consistent signs after these wave moves. So, the property that both splitting arcs pass through the rectangles plays a crucial role in the final part of the argument in Section 3.1.

To overcome this obstacle we next "enlarge" $\mathcal{N}(\tau'_D)$ and s'_ρ so that the new s'_ρ has the desired property. That is, each α- or γ-arc in $\mathcal{R}^d$ or $\mathcal{R}^u$ gives rise to an intersection arc with the new s'_ρ. The "enlarging operation" is as follows:

If there is an α- or a γ-arc that forms a triangle with a cusp of $\mathcal{N}(\tau'_D)$ (see the shaded triangle in Figure 29), then we enlarge $\mathcal{N}(\tau'_D)$ to contain this arc. The operation is illustrated in Figure 29, where the dashed arcs denote s'_ρ. This operation on $\mathcal{N}(\tau'_D)$ is similar in spirit to the pinching operation on a train track.

In order to do that, we will consider the gluing map $\varphi\colon \partial_+ P^u \to \partial_+ P^d$ which glues the two components of $\partial\widehat{\Sigma}$ together to form Σ. The only thing we know for sure about φ is that $\varphi(X^u) = X^d$. Otherwise it

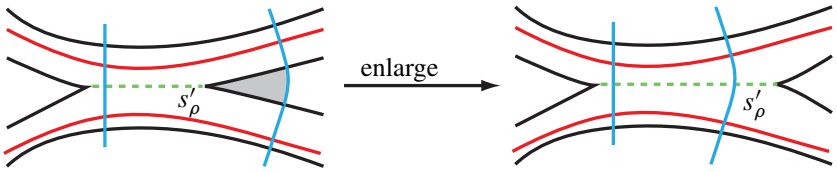

Figure 29: Enlarge $\mathcal{N}(\tau'_D)$.

can stretch and contract various segments when gluing the arc $\partial P^d_+ \setminus X^d$ to the arc $\partial P^d_+ \setminus X^d$. We will consider the various possibilities of the gluing by φ by studying the configuration of δ near the two sides of $\partial_+ P$.

If the configuration of ρ_x is as in Figure 26(a), then consider the shaded triangle region Δ in Figure 30 which contains X^u as a vertex. The gluing map φ sends a boundary edge of Δ to a subarc of $\partial_+ P^d$. We can view how things are identified by φ in the picture of a cyclic cover of Σ dual to $\partial_+ P$, see Figure 30, where the shaded region Δ' is a translation of Δ. The location of Δ' depends on the gluing map φ. Since γ intersects $\mathcal{R}^u$, there must be a γ-arc meeting the interior of Δ; see the top blue arc in Figure 30. So there is a γ-arc meeting the interior of Δ' as illustrated in Figure 30.

Let R be the rectangle in $\mathcal{A}_l \setminus (\mathcal{J}^u_l \cup \mathcal{J}^d_l)$ between the two junctions which contains the points X^u, X^d and the arc κ. Note that R is invariant under the involution π. Moreover, $\pi|_R$ is a 180°-rotation around the center of R, and the vertical arc of R that contains its center is invariant under π. Let r^u be the boundary edge of R in $\partial_+ P^u$ (ie the top edge of R). The involution π rotates r^u to the bottom edge of R.

Consider the subarc of r^u that lies in the boundary of Δ, ie the arc $r^u \cap \partial\Delta$, and the two δ-arcs in R attached to X^u and X^d. If the length of the arc $r^u \cap \partial\Delta$ is less than half of the length of r^u, then the symmetry from π implies that the vertical arc of R which contains its center must separate the two δ-arcs attached to X^u and X^d. In this case no vertical arc of R intersects both δ-arcs, which contradicts the hypothesis that κ intersects both δ-arcs; see Figure 30. Therefore, the length of the arc $r^u \cap \partial\Delta$ must be at least half of the length of r^u and this implies that part of the arc $\varphi(r^u \cap \partial\Delta)$ must be outside R. In other words, $\varphi(r^u \cap \partial\Delta)$ must reach the left junction in Figure 30. By the configuration of Δ and Δ',

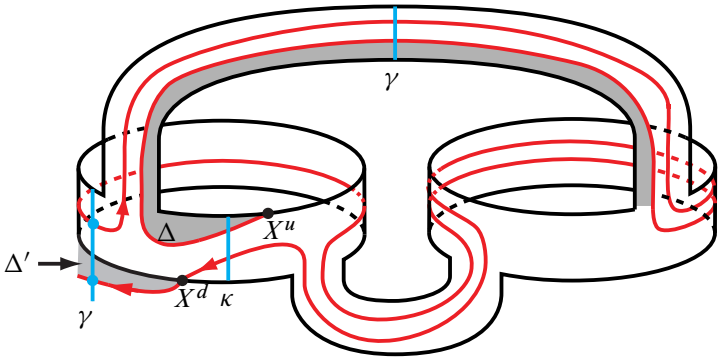

Figure 30: A global picture for Figure 26(a).

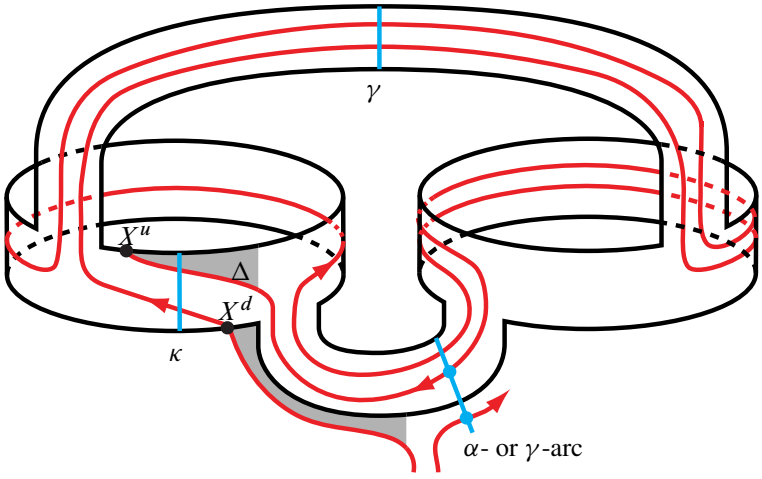

Figure 31: The first possible global picture for Figure 26(b).

this means that the γ-arc that intersects Δ' (the left blue arc in Figure 30) must intersect the δ-curve in $\widehat{\Sigma}$ further to the left of the left junction of $\mathscr{A}_l$.

As illustrated by the two blue dots in the γ-arc that meets Δ' in Figure 30, two intersection points of δ with this γ-arc have opposite signs, contradicting part (1) of Proposition 2.1.8. This means that the configuration in Figure 26(a) cannot happen.

It remains to consider the configuration of ρ_x as in Figure 26(b). Let Δ be the triangular region in Figure 31 and we study the gluing map φ from the translation of Δ in a cyclic cover of Σ dual to $\partial_+ P$. Again there are two possibilities:

The first possibility is that the shaded triangular region marked Δ in Figure 31 plus the junction next to Δ are glued along $\partial_+ P^d$ as illustrated in Figure 31. Figure 31 is the case where the junction next to Δ is glued to $\partial\mathscr{A}_l$ or $\partial\mathscr{R}^d$ and there is an α- or a γ-arc that intersects $\mathscr{R}^d$ and meets a δ-arc on the other side of the junction; see the bottom blue arc in Figure 31. As illustrated by the two blue dots on this arc in Figure 31, this means that two intersection points of δ with this α- or γ-arc have opposite signs, contradicting part (1) of Proposition 2.1.8.

The second possibility is that there is no such α- or γ-arc. Then after an isotopy (fixing $\alpha \cup \gamma$), we may assume that the triangular region Δ is glued to $\partial_+ P^d$ and covers an entire boundary edge of $\mathscr{R}^d$, as illustrated in Figure 32.

Now consider $\mathcal{N}(\tau'_D)$ and the splitting arc s'_ρ described earlier. As illustrated in Figure 29, if an α- or a γ-arc in $\Sigma \setminus \mathcal{N}(\tau'_D)$ is a ∂-parallel arc that cuts off a neighborhood of a cusp, then we can enlarge $\mathcal{N}(\tau'_D)$ to include this neighborhood of the cusp and extend the corresponding splitting arc s'_j or s'_ρ. As the δ-edge of $\partial\Delta$ is parallel to $\partial_+ P$ and since Δ is glued to cover a whole edge of $\partial\mathscr{R}^d$, we can enlarge $\mathcal{N}(\tau'_D)$ to contain Δ, and in particular, contain this boundary edge of $\partial\mathscr{R}^d$. Let $\mathcal{N}(\tau'_D)^+$ be the surface obtained

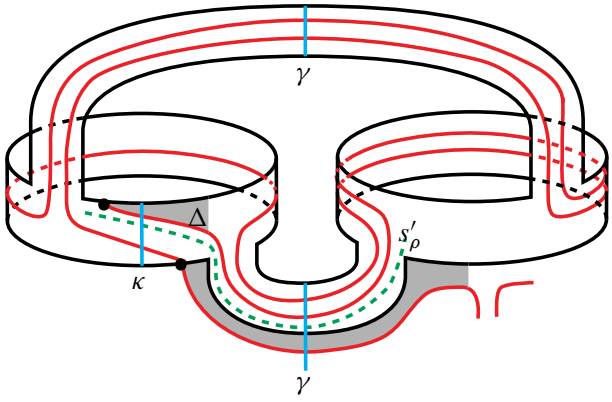

Figure 32: The second possible global picture for Figure 26(b).

by enlarging $\mathcal{N}(\tau'_D)$ as above. As illustrated in Figure 32, the splitting arc s'_ρ can be extended to pass through the rectangle $\mathcal{R}^d$. This picture is symmetric under π, so we can enlarge $\mathcal{N}(\tau'_D)$ and extend the splitting arc s'_ρ in the other direction to pass through $\mathcal{R}^u$ as well. As in Section 3.1, $\Sigma \setminus \mathcal{N}(\tau'_D)^+$ is still an annulus with two cusps at each boundary component.

Now the configuration becomes the same as Section 3.1. In particular, both splitting arcs s'_j and s'_ρ pass through the two rectangles $\mathcal{R}^d$ and $\mathcal{R}^u$. Thus we have all the ingredients needed for the argument in Section 3.1, which can be applied to conclude that K is doubly primitive. This finishes the proof for the special case.

3.5 The curve δ takes one short path in each annulus, the general case

In this section we deal with the remaining general case of Configuration 1. Namely: The curve δ takes only two short paths, one in $\mathcal{A}_l$ and one in $\mathcal{A}_r$, and the configurations in Section 3.4 and Figure 26 do not occur.

As before, we assume δ and $\partial_+ P$ are invariant under the involution π and $X = \delta \cap \partial_+ P$ is a fixed point of π. By symmetry (2) in Remark 2.2.1, we may assume $X^d \cup X^u$ lies in either ∂A_l or $\partial\mathcal{R}^d \cup \partial\mathcal{R}^u$.

Proposition 3.5.1 *Suppose that the curve δ takes one short path in $\mathcal{A}_l$ and one short path in $\mathcal{A}_r$. Then K is doubly primitive.*

Proof Suppose the configuration is not the one in Section 3.4. Hence for each component κ of $(\alpha \cup \gamma) \cap \hat{\Sigma}$ such that $\kappa \cap \tau_D = \kappa \cap \rho_x$, the set $\kappa \cap \rho_x$ contains at most one point.

Consider the train track τ_D. As explained at the beginning of this section, the shape of τ_D must be as in Figure 12. We may assume τ_D is also invariant under π. The train track τ_D consists of three segments, one of which has cusp directions at both endpoints pointing into the segment. Call this segment the x-arc of τ_D. The other two segments have cusps directions pointing away from the segments and we call them

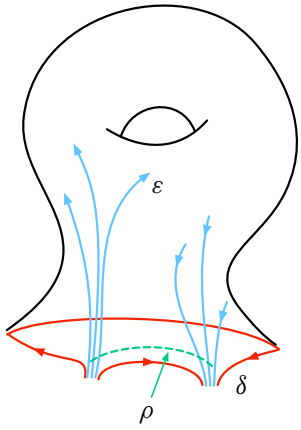

Figure 33: The once-punctured torus T and the s-wave ρ with respect to ε.

the y-arcs; see Figure 12. Note that the special segment ρ_x is a y-arc and, as before, the weight of δ at ρ_x is one.

Let $\mathcal{N}(\tau_D)$ be a small neighborhood of τ_D. It is a once-punctured torus with two cusps on its boundary corresponding to the two cusps of τ_D. We may assume $\varepsilon \cap \mathcal{N}(\tau_D)$ consists of essential arcs in $\mathcal{N}(\tau_D)$. As $\varepsilon \cap \delta = \varnothing$ and δ is a simple closed curve in $\mathcal{N}(\tau_D)$, the set $\varepsilon \cap \mathcal{N}(\tau_D)$ is a collection of parallel arcs going into one cusp of $\mathcal{N}(\tau_D)$ and coming out of the other cusp. It follows from part (2) of Proposition 2.1.8 that the orientations of these parallel arcs $\varepsilon \cap \mathcal{N}(\tau_D)$ (induced from an orientation of ε) are all the same.

Let T denote the complement $\Sigma \setminus \mathcal{N}(\tau_D)$. So T is also a once-punctured torus (with two cusps on the boundary); see Figure 33 for a picture of ε near ∂T.

If ε contains arcs in the cusps of $\mathcal{N}(\tau_D)$, then as shown in Figure 33, there is an s-wave ρ with respect to ε, parallel to one of the δ-arcs connecting the two cusps on ∂T. If there are k_e components of $\varepsilon \cap \mathcal{N}(\tau_D)$, then we can perform k_e consecutive wave moves on ε along such s-waves ρ. The result of the wave moves is a meridional curve ε' of W which is completely contained in T. Let E' be the disk in W bounded by ε'.

The proof will now proceed by the following sequence of claims:

Claim 3.5.2 *There is a properly embedded planar surface $P_{\varepsilon'}$ in U such that $(P_{\varepsilon'}, E')$ forms a $(\mathcal{P}, \mathcal{D})$-pair, and $\partial_- P_{\varepsilon'}$ has the same slope as $\partial_- P$ in the boundary torus $\partial_- U$.*

Proof As explained in Remark 2.1.2, the result of a ∂-compression on P is a pair of planar surface P_l and P_r with $\partial_+ P_l$ and $\partial_+ P_r$ isotopic to the core curves $\mathfrak{a}_l$ and $\mathfrak{a}_r$ of $\mathcal{A}_l$ and $\mathcal{A}_r$ respectively.

Since δ takes one short path in both $\mathcal{A}_l$ and $\mathcal{A}_r$, after isotopy, each of $\mathfrak{a}_l \cap \mathcal{N}(\tau_D)$ and $\mathfrak{a}_r \cap \mathcal{N}(\tau_D)$ is a single nonseparating arc in $\mathcal{N}(\tau_D)$. Note that the two arcs $\mathfrak{a}_l \cap \mathcal{N}(\tau_D)$ and $\mathfrak{a}_r \cap \mathcal{N}(\tau_D)$ cannot be parallel in $\mathcal{N}(\tau_D)$ because this would imply that a band sum of $\mathfrak{a}_l = \partial_+ P_l$ and $\mathfrak{a}_r = \partial_+ P_r$ produces a curve

disjoint from $\mathcal{N}(\tau_D)$ but parallel to $\partial_+ P$. This is a contradiction to the fact that $\delta \cap \partial_+ P$ is a single point. So $\mathfrak{a}_l \cap \mathcal{N}(\tau_D)$ and $\mathfrak{a}_r \cap \mathcal{N}(\tau_D)$ are nonisotopic arcs in the once-punctured torus $\mathcal{N}(\tau_D)$. Hence the endpoints of $\mathfrak{a}_l \cap \mathcal{N}(\tau_D)$ and $\mathfrak{a}_r \cap \mathcal{N}(\tau_D)$ alternate along the boundary of the once-punctured torus (one can also see this by checking the possible configurations geometrically). Now, consider $T = \Sigma \setminus \mathcal{N}(\tau_D)$. The endpoints of $\mathfrak{a}_l \cap T$ and $\mathfrak{a}_r \cap T$ alternate along ∂T, which implies that $\mathfrak{a}_l \cap T$ and $\mathfrak{a}_r \cap T$ are also nonisotopic essential arcs in the once-punctured torus T.

Use the arcs $\mathfrak{a}_l \cap T$ and $\mathfrak{a}_r \cap T$ as representatives of a basis for $H_1(T, \partial T) \cong \mathbb{Z} \oplus \mathbb{Z}$, and suppose they represent elements with slope $1/0$ and $0/1$, respectively. Since $\varepsilon' \subset T$, there is a properly embedded arc in T which intersects ε' in a single point. Suppose this arc has slope p/q, where p and q are coprime. Take p parallel copies of P_l and q parallel copies of P_r, and perform a sequence of band sums of these planar surfaces along ∂T to obtain a planar surface $P_{\varepsilon'}$ so that the intersection of T with the resulting curve $\partial_+ P_{\varepsilon'}$ is an arc of slope p/q. Hence the planar surface $P_{\varepsilon'}$ is such that $\varepsilon' \cap \partial_+ P_{\varepsilon'}$ is a single point. Hence $(P_{\varepsilon'}, E')$ forms a $(\mathcal{P}, \mathcal{D})$-pair. □

Isotope ε' to intersect α and γ minimally. The following is an immediate corollary of Claim 3.5.2.

Corollary 3.5.3 *No subarc of ε' is a wave with respect to $\{\gamma, \alpha\}$.*

Proof The corollary is basically the same as Corollary 2.1.7.

By Claim 3.5.2, $(P_{\varepsilon'}, E')$ forms a $(\mathcal{P}, \mathcal{D})$-pair. If $P_{\varepsilon'}$ is an annulus, then by Proposition 1.1.4, K is doubly primitive and we are done. So we may assume that $P_{\varepsilon'}$ is not an annulus. After replacing the planar surface P in Lemma 2.1.6 with $P_{\varepsilon'}$, one can conclude that any s-wave with respect to $\{\gamma, \alpha\}$ must intersect $\partial_+ P_{\varepsilon'}$: As in the proof of Corollary 2.1.7, if a subarc η of ε' is a wave with respect to $\{\gamma, \alpha\}$, then $\eta \cap \partial_+ P_{\varepsilon'} \neq \varnothing$. By the symmetry induced by π, the curve ε' has another subarc $\eta' = \pi(\eta)$ which is also a wave with respect to $\{\gamma, \alpha\}$ and $\eta' \cap \partial_+ P_{\varepsilon'} \neq \varnothing$. This contradicts the fact that $\varepsilon' \cap \partial_+ P_{\varepsilon'}$ is a single point as $(P_{\varepsilon'}, E')$ is a $(\mathcal{P}, \mathcal{D})$-pair. □

As before, consider the Heegaard diagram formed by $\widehat{W} = \{\delta, \varepsilon'\}$ and $\widehat{V} = \{\gamma, \alpha\}$. Let s_j be the splitting arc in $\mathcal{N}(\tau_D)$ which connects the two cusps, so that $s_j \cap \delta = \varnothing$ and also $\mathcal{N}(\tau_D) \setminus s_j$ is a product neighborhood of δ.

Lemma 2.1.4 implies that s_j passes through both $\mathcal{R}^d$ and $\mathcal{R}^u$. Since γ intersects both $\mathcal{R}^d$ and $\mathcal{R}^u$ and since $\varepsilon' \subset \Sigma \setminus \mathcal{N}(\tau_D)$, our assumption on the orientation of δ implies that γ has two distinct subarcs in $\mathcal{R}^d$ and $\mathcal{R}^u$ that are $[\delta^-, \delta^+]$ blocking-edges.

In a once-punctured torus if two essential embedded curves or arcs intersect minimally then the algebraic intersection number is equal to the geometric intersection number. Hence the intersection points of ε' with each arc of $(\alpha \cup \gamma) \cap T$ all have the same sign.

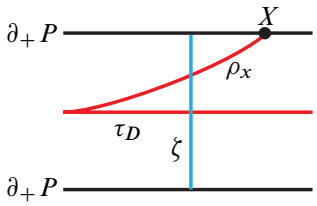

Figure 34: Intersection of ζ with τ_D.

For each arc of $(\alpha \cup \gamma) \cap T$, by collapsing all its intersection points with ε' into one point along this arc, we can construct a train track fully carrying ε'. In fact, we can extend such collapsing/pinching as much as possible to construct a train track τ_E which is either a train track as in Figure 12 or a circle (in which case each arc of $(\alpha \cup \gamma) \cap T$ intersects ε' in at most one point). In either case, each arc of $(\alpha \cup \gamma) \cap T$ intersects the train track τ_E in at most one point.

Claim 3.5.4 $$|(\alpha \cup \gamma) \cap \tau_E| \leq c_0(P, D, \alpha, \gamma).$$

Proof Let ζ be a component of $(\alpha \cup \gamma) \setminus \partial_+ P$ and view ζ as an arc in $(\alpha \cup \gamma) \cap \widehat{\Sigma}$. We first consider the possibility that $\zeta \cap \tau_D$ contains more then one point. In the construction of τ_D, all the parallel δ-arcs, except for the special arc ρ_x, were pinched into a single segment. Thus, ζ intersects $\tau_D \setminus \rho_x$ in at most one point. As $\zeta \cap \tau_D$ contains more then one point, ζ must intersect ρ_x.

It is assumed at the beginning of Proposition 3.5.1 that the configuration is not in the special case in Section 3.4. That is, for each component ζ of $(\alpha \cup \gamma) \setminus \partial_+ P$ such that $\zeta \cap \tau_D = \zeta \cap \rho_x$, the set $\zeta \cap \rho_x$ contains at most one point. Thus, if $\zeta \cap (\tau_D \setminus \rho_x) = \varnothing$, then $\zeta \cap \tau_D = \zeta \cap \rho_x$ contains at most one point. Since $\zeta \cap \tau_D$ contains more then one point, $\zeta \cap (\tau_D \setminus \rho_x) \neq \varnothing$. Figure 34 is a local picture when this happens. Moreover, a subarc of ζ between two points of $\zeta \cap \tau_D$ must connect ρ_x to $\tau_D \setminus \rho_x$.

Recall that in our case the train track τ_D is as in Figure 12. Hence, such a subarc of ζ between two points of $\zeta \cap \tau_D$ is an arc in $\Sigma \setminus \tau_D$ that "cuts off" a cusp, as illustrated in Figure 34. In particular, such an arc corresponds to a ∂-parallel arc in $\Sigma \setminus \mathcal{N}(\tau_D)$ which can be assumed to be disjoint from ε' after isotopy.

Thus, for each component ζ of $(\alpha \cup \gamma) \setminus \partial_+ P$ that meets δ, we can find a subarc $\zeta' \subset \zeta$ so that $\zeta' \cap \delta = \zeta \cap \delta$, ie ζ' contains all the points of $\zeta \cap \delta$, and so that ζ' is disjoint from ε'. Let $\mathscr{C}'$ be the collection of all such subarcs ζ'. Hence, by definition,

$$|\mathscr{C}'| = c_0(P, D, \alpha, \gamma).$$

Since each component of $(\alpha \cup \gamma) \setminus \mathcal{N}(\tau_D)$ intersects τ_E in at most one point, we have

$$|(\alpha \cup \gamma) \cap \tau_E| \leq |(\alpha \cup \gamma) \setminus \mathscr{C}'|.$$

As $|(\alpha \cup \gamma) \setminus \mathscr{C}'| = |\mathscr{C}'| = c_0(P, D, \alpha, \gamma)$, we have $|(\alpha \cup \gamma) \cap \tau_E| \leq c_0(P, D, \alpha, \gamma)$ as stated. □

Claim 3.5.5 *The train track τ_E cannot be a circle.*

Proof Suppose to the contrary that τ_E is a circle. So $\tau_E = \varepsilon'$ and each arc of $(\alpha \cup \gamma) \cap T$ intersects ε' in at most one point. By Claim 3.5.4, we have

$$|(\alpha \cup \gamma) \cap \varepsilon'| \leq c_0(P, D, \alpha, \gamma).$$

Let $P_{\varepsilon'}$ be as in Claim 3.5.2 and let E' be the disk in W bounded by ε'. So $(P_{\varepsilon'}, E')$ is a $(\mathcal{P}, \mathcal{D})$-pair. By the definition of complexity,

$$c_0(P_{\varepsilon'}, E', \alpha, \gamma) \leq |(\alpha \cup \gamma) \cap \varepsilon'|.$$

Hence

$$c_0(P_{\varepsilon'}, E', \alpha, \gamma) \leq c_0(P, D, \alpha, \gamma).$$

Moreover, Lemma 2.1.4 implies that

$$|(\alpha \cup \gamma) \cap \tau_D| < |(\alpha \cup \gamma) \cap \delta|.$$

By the construction of τ_E, we have

$$|(\alpha \cup \gamma) \cap \tau_E| \leq |(\alpha \cup \gamma) \cap \tau_D|.$$

As $\tau_E = \varepsilon'$, we have

$$|(\alpha \cup \gamma) \cap \varepsilon'| = |(\alpha \cup \gamma) \cap \tau_E| \leq |(\alpha \cup \gamma) \cap \tau_D| < |(\alpha \cup \gamma) \cap \delta|.$$

Thus $c(P_{\varepsilon'}, E', \alpha, \gamma) < c(P, D, \alpha, \gamma)$, and this contradicts Assumption 1.2.3. □

Claim 3.5.5 means that there is an arc of $(\alpha \cup \gamma) \cap T$ that intersects ε' in more than one point in T. Hence a subarc of this arc is an $[\varepsilon'^-, \varepsilon'^+]$ blocking-edge. We have concluded before Claim 3.5.4 that γ contains two subarcs in $\mathcal{R}^d$ and $\mathcal{R}^u$ that are $[\delta^-, \delta^+]$ blocking-edges. So by part (1) of Lemma 1.3.9, the Heegaard diagram has no wave with respect to $\{\delta, \varepsilon'\}$. Moreover, since γ contains two distinct $[\delta^-, \delta^+]$-edges, the Heegaard diagram is not a standard Heegaard diagram of S^3 or $(S^2 \times S^1) \# L(r, s)$. Thus, by Theorems 1.3.7 and 1.3.8, there must be a wave with respect to $\{\alpha, \gamma\}$. Since by Lemma 2.1.5, there is no wave with respect to α, there must be a wave with respect to γ.

By part (1) of Proposition 2.1.8, we may assign orientations to α and γ so that the intersection points of δ with $\alpha \cup \gamma$ all have the same sign. By Corollary 3.5.3, the curves α, γ, δ, and ε' satisfy the conditions in Lemma 2.1.9. Thus, by Lemma 2.1.9, the Heegaard diagram must be as depicted in Figure 7, left, and the γ-wave η must connect a $[\delta^-, \delta^+]$-edge to an $[\varepsilon'^-, \varepsilon'^+]$-edge as shown in Figure 7, center,.

As illustrated in Figure 7, left, all the $[\delta^-, \delta^+]$-edges are parallel. Hence we may assume that all the $[\delta^-, \delta^+]$-edges are in the cusp of $\mathcal{N}(\tau_D)$. Similarly, we may assume that all the $[\varepsilon'^-, \varepsilon'^+]$-edges are in the cusp of $\mathcal{N}(\tau_E)$. Figure 35, left, is a picture of the octagon in Figure 7, center, that contains the wave η together with a picture of the cusps of $\mathcal{N}(\tau_D)$ and $\mathcal{N}(\tau_E)$. Since the intersection points of δ with $\alpha \cup \gamma$ all have the same sign, the two edges of the octagon in Figure 35, left, next to the cusp must have opposite orientation; see the orientation of the two arcs marked α and γ in Figure 35, left. This means that these two arcs must belong to different curves of $\{\alpha, \gamma\}$, that is, one α-arc and one γ-arc as labeled

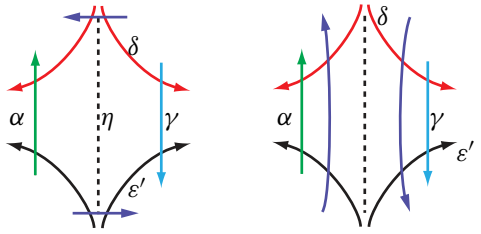

Figure 35: Wave move along η.

in Figure 35, left, since, if they belong to the same curve, then an arc in the octagon connecting these two arcs is a wave and this wave is as shown in Figure 7, right. This contradicts part (2) of Lemma 2.1.9, which claims that the wave should be as in Figure 7, center, and the two are not compatible.

Remark 3.5.6 Let η be the wave as in Figure 35, left (also see Figure 7, center). By Lemma 2.1.5, the wave η must be with respect to γ. As in Definition 1.3.4 and similar to the argument in Section 3.1, a wave move along η can be done in two steps. The first step is to perform a surgery: connect the endpoints of $\gamma \setminus \mathcal{N}(\partial\eta)$ using two parallel copies of η. It can be seen from Figure 35, right, that the resulting two curves do not form any bigon with δ and ε'. One of the resulting curves, which we denote by α_0, is parallel to α, and the other resulting curve, which we denote by γ', is the final curve in the wave move. The configuration of Figure 35, right, implies that the intersection of α_0 (resp. γ') with this octagon is an arc parallel to the edge marked α (resp. γ) in Figure 35, right. The curve α_0 is removed in the second step of the wave move. Thus after the wave move, the two subarcs of $\alpha \cup \gamma'$ next to a cusp of τ_D belong to different curves in $\{\alpha, \gamma'\}$, similar to the original configuration of $\{\alpha, \gamma\}$ before this wave move. Thus, as long as $\{\alpha, \gamma'\}$ contains both $[\delta^-, \delta^+]$ and $[\varepsilon'^-, \varepsilon'^+]$ edges, the component of $\Sigma \setminus (\delta \cup \varepsilon' \cup \alpha \cup \gamma')$ that contains the previous wave remains an octagon; see Figure 7, center, and Figure 35, left. As in Lemma 2.1.9, the next wave with respect to $\{\alpha, \gamma'\}$ must also connect a $[\delta^-, \delta^+]$-edge to an $[\varepsilon'^-, \varepsilon'^+]$-edge as depicted in Figure 35, left, as well.

Remark 3.5.6 implies that there is a sequence of wave moves $\eta_1, \dots, \eta_k$ along waves as shown in Figure 35, left, such that after the last wave move η_k, either

(a) $\Gamma(\delta, \varepsilon')$ contains no more $[\varepsilon'^-, \varepsilon'^+]$ edges, or

(b) $\Gamma(\delta, \varepsilon')$ contains no more $[\delta^-, \delta^+]$ edges.

Note that before the last wave move η_k, the Whitehead graph $\Gamma(\delta, \varepsilon')$ contains *both* $[\delta^-, \delta^+]$ and $[\varepsilon'^-, \varepsilon'^+]$ edges. Hence, we are left with the following three cases regarding τ_E and the Whitehead graph $\Gamma(\delta, \varepsilon')$:

Case 1 The weight of ε' at each segment of the train track τ_E is at least two, and $\Gamma(\delta, \varepsilon')$ contains no more $[\varepsilon'^-, \varepsilon'^+]$-edges after the last wave move η_k.

Case 2 $\Gamma(\delta, \varepsilon')$ contains no more $[\delta^-, \delta^+]$-edges after the last wave move η_k.

Case 3 The weight of ε' at some segment of τ_E is one, and $\Gamma(\delta, \varepsilon')$ contains no more $[\varepsilon'^-, \varepsilon'^+]$-edges after the last wave move η_k.

To simplify notation, we use $\{\alpha', \gamma'\}$ to denote the set of meridional curves of V obtained by the sequence of wave moves, where α' and γ' are derived from the original meridians α and γ, respectively.

Case 1 The weight of ε' at each segment of the train track τ_E is at least two, and $\Gamma(\delta, \varepsilon')$ contains no more $[\varepsilon'^-, \varepsilon'^+]$-edges after the last wave move η_k.

The goal in this case is to show that either the knot K is a *Berge–Gabai knot* (ie K lies in a Heegaard solid torus and the solid torus remains a solid torus after a nontrivial Dehn surgery on K) and hence it is doubly primitive by [8] and [16], or that before the last wave move along η_k, subarcs of ε' are $[\alpha'^+, \alpha^-]$-edges and $[\gamma'^+, \gamma'^-]$-edges. Since $\Gamma(\delta, \varepsilon')$ contains both $[\delta^-, \delta^+]$ and $[\varepsilon'^-, \varepsilon'^+]$ edges before the last wave move η_k, the latter possibility is a contradiction to Theorems 1.3.7 and 1.3.8. So in Case 1 we focus on how ε' intersects $\alpha' \cup \gamma'$ during the sequence of wave moves. This is done by studying the intersection with the train track τ_E.

No subarc of ε' is a wave by Corollary 3.5.3. As in the proof of Lemma 2.1.9 and illustrated in Figure 7, we may assign orientations to ε', α and γ so that all the intersection points have the same sign. Thus the orientation of α and γ are compatible along the train track τ_E. By Claim 3.5.5, τ_E is a train track with two cusps as shown in Figure 12. Recall that the train track τ_E is constructed so that each arc of $(\alpha \cup \gamma) \cap T$ intersects τ_E in at most one point.

Consider the curves α' and γ' during the sequence of wave moves $\eta_1, \dots, \eta_k$ as described above. The curves α' and γ' have an induced orientation from α and γ. We will inductively require that, at any stage of the sequence of wave moves, α' and γ' have the following property: Each arc of $(\alpha' \cup \gamma') \cap T$ intersects the train track τ_E in at most one point and the orientations of α' and γ' are compatible along τ_E. Since the involution π leaves ε', α' and γ' invariant up to isotopy, we may assume that π leaves the train track τ_E invariant. Hence π interchanges the two cusps of τ_E. The two cusp points of the train track τ_E divide τ_E into three segments: one x-arc and two y-arcs, as in Figure 12.

For any segment κ of the train track τ_E, the assumption on π implies that π leaves κ invariant and interchanges the two endpoints of κ. Since π also leaves α' and γ' invariant up to isotopy, this implies that the intersection points of $(\alpha' \cup \gamma') \cap \kappa$ that are outermost in κ are interchanged by π and hence must belong to the same curve α' or γ'.

We first show that $\alpha' \cup \gamma'$ must meet at least two segments of the train track τ_E. Assume to the contrary that $\alpha' \cup \gamma'$ intersects only one segment κ of τ_E. Now consider the region of $\Sigma \setminus (\mathcal{N}(\tau_D) \cup \mathcal{N}(\tau_E))$ corresponding to the octagon of Figure 35, left. We see that the two subarcs of $\alpha' \cup \gamma'$ next to a cusp of τ_E (see the arcs labeled α and γ in Figure 35, left) must correspond to arcs that intersect κ with outermost intersection points along κ and thus belong to the same curve in $\{\alpha', \gamma'\}$, this implies that the two vertical edges of the octagon in Figure 35, left, closest to the cusps of τ_D and τ_E must belong to the same curve of $\{\alpha', \gamma'\}$. This contradicts the conclusion in Remark 3.5.6. Thus we conclude that $\alpha' \cup \gamma'$ must intersect at least two segments of train track τ_E.

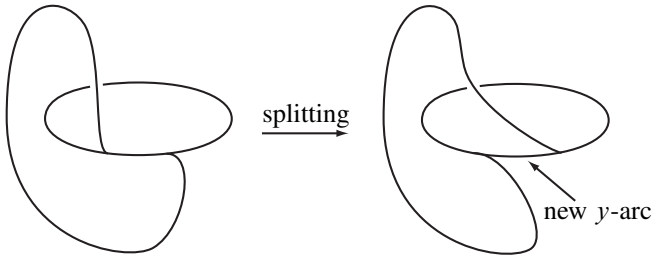

Figure 36: Splitting the train track τ_E.

Let $p\colon \mathcal{N}(\tau_E) \to \tau_E$ be the map collapsing $\mathcal{N}(\tau_E)$ to the train track τ_E. Since the cusp direction of τ_E points into the x-arc, there are two parallel and adjacent ε'-arcs t_1 and t_2 in $\mathcal{N}(\tau_E)$ that collapse onto the x-arc by the map p. If the x-arc intersects $\alpha' \cup \gamma'$, then each intersection point corresponds to an arc of $\alpha' \cup \gamma'$ between the two ε'-arcs t_1 and t_2. As the orientation is compatible, this arc is an $[\varepsilon'^-, \varepsilon'^+]$-edge.

Perform the sequence of wave moves $\eta_1, \dots, \eta_k$, with respect to $\{\alpha', \gamma'\}$, as described in Remark 3.5.6. There are two situations, depending on whether or not α' and γ' meet the x-arc of the train track.

If the x-arc of the train track τ_E intersects $\alpha' \cup \gamma'$, then each intersection point correspond to an $[\varepsilon'^-, \varepsilon'^+]$-edge. As indicated in Figure 35 and Remark 3.5.6, a wave move basically "pushes" an arc out of a cusp of τ_E. Hence after each wave move, the number of intersection points of the x-arc with $\alpha' \cup \gamma'$ decreases by at least one. Moreover, since originally each arc of $(\alpha \cup \gamma) \cap T$ intersects τ_E in at most one point, the configuration of Figure 35, left, implies that after each wave move, each arc of $(\alpha' \cup \gamma') \cap T$ intersects the train track τ_E in at most one point as well. So we can perform wave moves until the x-arc no longer intersects $\alpha' \cup \gamma'$. Note that Remark 3.5.6 says that the two arcs next to a cusp of τ_E are still one α'-arc and one γ'-arc.

If the x-arc of the train track τ_E does not intersect $\{\alpha' \cup \gamma'\}$ and we have not exhausted the sequence of wave moves, then split the train track τ_E as shown in Figure 36. Note that the splitting does not affect α' and γ' as they do not intersect the x-arc. The splitting changes a y-arc of τ_E into the x-arc of the resulting train track, and creates a new y-arc that does not intersect $\alpha' \cup \gamma'$.

Call a y-arc of the train track a *pure* y-arc if it either does not intersect $\alpha' \cup \gamma'$ or intersects only one of the curves in α' or γ'. So the new, after this splitting, y-arc is a pure y-arc. We showed earlier that $\alpha' \cup \gamma'$ must intersect at least two segments of train track. Since the newly created y-arc does not intersect $\alpha' \cup \gamma'$, the new x-arc of the train track must intersect $\alpha' \cup \gamma'$.

To simplify notation, we still use τ_E to denote the train track after the splitting. Now we continue with the sequence of wave moves.

Claim *Each pure y-arc of the train track τ_E remains a pure y-arc after the next wave move.*

Proof There are two cases. The first is when $\kappa \subset \tau_E$ is a pure y-arc which meets one curve in $\{\alpha', \gamma'\}$, say α', before a wave move. Then as in the argument in Remark 3.5.6, the surgery-step of the wave move

"pushes" an arc out of the cusp creating a pair of arcs parallel to the α'- and γ'-edges inside the octagon; see Figure 35, right.

If the wave move is with respect to α', then γ' is unchanged by the wave move. Since κ is a pure y-arc and does not intersect γ' before the wave move, it does not intersect γ' after the wave move either. Hence κ remains a pure y-arc after the wave move.

If the wave move is with respect to γ', then the surgery-step of the wave move "splits" γ' into two curves γ_1 and γ_2. One of γ_1 and γ_2 is parallel to α'. When viewed in the octagon, the surgery-step creates a pair of arcs parallel and next to the α'- and γ'-edges of the octagon as seen in Figure 35, right. Without loss of generality, suppose the arc parallel to the γ'-edge (resp. α'-edge) in the octagon belongs to γ_1 (resp. γ_2). Since the pure y-arc κ only meets α', this means that γ_2 intersects κ while γ_1 does not. Recall that in the construction of a wave move, the set of curves $\gamma' \cup \gamma_1 \cup \gamma_2$ bounds a pair of pants and one of γ_1 or γ_2 is parallel and next to α'. The arc parallel to the γ'-edge in the octagon belongs to γ_1, so γ_1 cannot be the curve parallel to α'. Hence γ_2 is parallel to α' and is deleted in the second step of the wave move and γ_1 is the new γ'-curve resulting from the wave move. Since γ_1 does not intersect κ and γ_2 is removed, this means that after this wave move the pure y-arc κ only meets α' (after the wave move) and hence remains a pure y-arc.

The second case is when a y-arc κ of the train track τ_E does not intersect $\alpha' \cup \gamma'$ before a wave move with respect to γ'. Then similar to the argument above, the surgery-step of the wave move "pushes" an arc out of the cusp and "splits" γ' into two curves γ_1 and γ_2. One of the resulting curves, say γ_1, meets this y-arc κ in exactly one point. Note that γ_1 cannot be the curve parallel to α' because this y-arc κ does not meet α' and two parallel curves must have the same sequence of intersection points with δ and ε' (since there is no bigon intersection by Remark 3.5.6 and since each arc of $(\alpha' \cup \gamma') \cap T$ intersects τ_E in at most one point). So γ_2 is the curve parallel to α' and it is removed in the second step of the wave move. Thus, if a y-arc κ of the train track τ_E does not intersect $\alpha' \cup \gamma'$ before a wave move, then κ intersects the new pair of curves $\alpha' \cup \gamma'$ after the wave move in exactly one point and hence is a pure y-arc.

Therefore, if a y-arc κ is a pure y-arc before a wave move, it remains a pure y-arc after the wave move. □

Continue with the sequence of wave moves and as before, if the x-arc of the train track, after some wave moves, no longer intersects $\alpha' \cup \gamma'$, we split the train track as in Figure 36 and get a new x-arc and a new pure y-arc. Since the splitting in Figure 36 is along an x-arc which does not intersect $\alpha' \cup \gamma'$, when performing the operations of wave moves and splittings, the following two properties are preserved:

(1) The curves $\alpha' \cup \gamma'$ always intersect at least two segments of the train track τ_E.

(2) Each component of $(\alpha' \cup \gamma') \cap T$ intersects the train track τ_E in at most one point.

Since the weight of ε' at each segment of the original train track τ_E is at least two, the splitting arc of $\mathcal{N}(\tau_E)$, see Definition 2.2.7, must "pass through" each segment. This means that, if we proceed with

all the wave moves and splittings, every segment of the train track will eventually become an x-arc. Since the splitting in Figure 36 turns the x-arc into a pure y-arc, after a number of such operations, we reach a configuration where both y-arcs are pure y-arcs and ε' has weight one at each y-arc. This implies that the weight of ε' at the x-arc is 2.

Let y_0 and y_1 be the two y-arcs, such that y_0 is the pure y-arc right after the splitting in Figure 36 and y_1 is the other y-arc which is a pure y-arc *from a previous splitting*. So y_0 does not intersect $\alpha' \cup \gamma'$. Since $\alpha' \cup \gamma'$ always intersects at least two segments of the train track, both y_1 and the x-arc intersect $\alpha' \cup \gamma'$. As y_1 is a pure y-arc, all the points of $y_1 \cap (\alpha' \cup \gamma')$ belong to the same curve. Without loss of generality, suppose $y_1 \cap (\alpha' \cup \gamma') \subset \alpha'$.

As discussed earlier, since $\alpha' \cup \gamma'$ intersects the x-arc at this stage, there is an $[\varepsilon'^-, \varepsilon'^+]$ edge. By the hypothesis of this case, this means that we have not exhausted the sequence of wave moves $\eta_1, \dots, \eta_k$.

Suppose the next wave move is η_m, with $m \le k$. Since y_0 does not intersect $\alpha' \cup \gamma'$ and by the argument above, after the wave move η_m, the y-arc y_0 intersects $\alpha' \cup \gamma'$ in exactly one point. As we have assumed $y_1 \cap (\alpha' \cup \gamma') \subset \alpha'$, after the wave move η_m, then $y_0 \cap (\alpha' \cup \gamma')$ is a point in γ'.

If the x-arc no longer intersects $\alpha' \cup \gamma'$ after the wave move η_m, then since y_1 only intersects α' and since the weight of ε' at y_0 is one, $\gamma' \cap \varepsilon' = \gamma' \cap y_0$ is a single point after the wave move. As α' and γ' are boundary curves of disks in the handlebody V and ε' is the boundary of a disk in W, this means that ε' and γ' are boundary curves of a destabilizing pair of disks for Σ in M (when viewing Σ as a Heegaard surface of M). By Claim 3.5.2, when viewing Σ as a Heegaard surface of the lens space by the Dehn surgery on K, the curves ε' and $\partial_+ P_{\varepsilon'}$ are boundary curves of a destabilizing pair of disks for Σ in the lens space. So after compressing Σ along the disk E' in the handlebody W bounded by ε', the resulting torus is a Heegaard torus for both M and the lens space obtained by the Dehn surgery. This means that K is contained in a Heegaard solid torus in M which remains a solid torus after a nontrivial Dehn surgery. Knots in solid tori with this property are called Berge–Gabai knots. They were classified by Berge [8] and Gabai [16], who also showed that they are all doubly primitive.

Thus, we may assume that, after the η_m wave move, the x-arc still intersects $\alpha' \cup \gamma'$. At this point the situation is that: both y_0 and y_1 are pure y-arcs and all three segments of the train track τ_E intersect $\alpha' \cup \gamma'$.

If the x-arc intersects $\alpha' \cup \gamma'$ in more than one point, then as the surgery-step of the wave move "pushes" only one arc "out of" the x-arc, after the surgery-step, the x-arc still intersects the resulting curves. Recall that each component of $(\alpha' \cup \gamma') \cap T$ intersects the train track τ_E in at most one point. This implies that the second step of the wave move (which deletes a parallel copy of α' or γ') cannot remove all the remaining intersection points of this x-arc with $\alpha' \cup \gamma'$ since the two parallel curves have the same sequence of intersection points with the train track τ_E. Thus, if the x-arc intersects $\alpha' \cup \gamma'$ in more than one point, then this x-arc still intersects $\alpha' \cup \gamma'$ after one additional wave move.

Therefore, we can continue the sequence of wave moves until the x-arc intersects $\alpha' \cup \gamma'$ in a single point. As the x-arc still intersects $\alpha' \cup \gamma'$, $\alpha' \cup \gamma'$ still contains an $[\varepsilon'^+, \varepsilon'^-]$-edge. The hypothesis of Case (1) says that, after we have exhausted the sequence of wave moves $\eta_1, \dots, \eta_k$, there is no more $[\varepsilon'^+, \varepsilon'^-]$-edge. Since we still have an $[\varepsilon'^+, \varepsilon'^-]$-edge at this stage, this means that we have not exhausted the sequence of wave moves. Each wave in the sequence of wave moves connects a $[\delta^+, \delta^-]$-edge to an $[\varepsilon'^+, \varepsilon'^-]$-edge, as explained before Remark 3.5.6. Thus the curves $\alpha' \cup \gamma'$ must still contain a $[\delta^+, \delta^-]$-edge.

As in the earlier discussion, we have $\varnothing \neq y_1 \cap (\alpha' \cup \gamma') \subset \alpha'$ and $\varnothing \neq y_0 \cap (\alpha' \cup \gamma') \subset \gamma'$. Without loss of generality, suppose the intersection point of the x-arc with $\alpha' \cup \gamma'$ belongs to γ'. So there is a subarc of ε' connecting the γ'-arc that intersects the x-arc to a γ'-arc that intersect y_0. As orientation of $\alpha' \cup \gamma'$ is compatible along the train track, this subarc of ε' is a $[\gamma'^+, \gamma'^-]$ edge.

We now have two situations to consider. If $|y_1 \cap (\alpha' \cup \gamma')| \geq 2$, then a subarc of y_1 between two points of $y_1 \cap (\alpha' \cup \gamma')$ corresponds to a subarc of ε' which is an $[\alpha'^+, \alpha'^-]$ edge. By part (1) of Lemma 1.3.9, there is no wave with respect to $\{\alpha', \gamma'\}$. Moreover, we have concluded above that $\alpha' \cup \gamma'$ contains both $[\delta^+, \delta^-]$ and $[\varepsilon'^+, \varepsilon'^-]$ blocking edges at this stage. So by part (1) of Lemma 1.3.9, there is no wave with respect to $\{\delta, \varepsilon'\}$ either. This contradicts Theorems 1.3.7 and 1.3.8.

If $|y_1 \cap (\alpha' \cup \gamma')| = 1$, then since the weight of ε' at y_1 is one, α' intersects ε' in just one point. Hence ε' and α' are boundary curves of disks which are a destabilizing pair for Σ in M (when viewing Σ as a Heegaard surface of M). As before, ε' and $\partial_+ P_{\varepsilon'}$ are boundary curves of a destabilizing pair of disks for Σ in the lens space that was obtained by the Dehn surgery. As in the argument above, this implies that K is a Berge–Gabai knot in a Heegaard solid torus and hence is doubly primitive, by [8] and [16]. This finishes the proof of Case (1).

Case 2 $\Gamma(\delta, \varepsilon')$ contains no more $[\delta^-, \delta^+]$-edges after the last wave move η_k.

The difference between Case 2 and Case 1 is due to the difference in the configuration of τ_D. Recall that the train track τ_D consists of a circle in $\widehat{\Sigma}$ and an arc ρ_x that contains the point $X = \delta \cap \partial_+ P$. By construction, the weight of δ at the segment ρ_x is one, and this is the main difference between τ_D and τ_E. If we simply apply the argument in Case 1 to τ_D, then we may not be able to conclude that both y-arcs of the train track will become pure y-arcs as in Case 1. The way to deal with Case 2 is to use the rectangle-annulus decomposition of $\widehat{\Sigma}$ and the proof below is similar to the proof in Section 3.1.

By Remark 3.5.6, the subarcs of $\alpha' \cup \gamma'$ next to a cusp of τ_D are an α'-arc and a γ'-arc. As before, τ_D is invariant under π. Similar to the discussion on τ_E in Case 1, for any segment of τ_D, the outermost (on the segment) intersection points of $\alpha' \cup \gamma'$ with the segment of τ_D belong to the same curve α' or γ'. As in Case (1), this implies that $\alpha' \cup \gamma'$ intersects at least two segments of the train track τ_D.

Similarly to Case 1, continue with the wave moves and, at each stage, denote the resulting curves by α' and γ'. We first discuss how the curves α' and γ' after the wave move intersect $\partial_+ P$ and $\widehat{\Sigma}$. The goal is

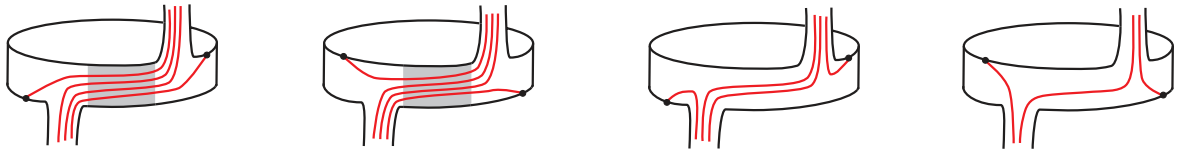

Figure 37: Possible configurations of δ in $\mathscr{A}_l$.

to show that, after some isotopy, the intersection of α' and γ' with $\widehat{\Sigma}$ consists of vertical arcs in $\mathscr{A}_l$, $\mathscr{A}_r$, $\mathscr{R}^u$ and $\mathscr{R}^d$. This will show, in particular, that the claims in Section 3.1 also hold in this setting.

Since every nonseparating simple closed curve in Σ is invariant under π, after isotopy, we may assume that ε', δ and $\partial_+ P$ are all invariant. Consider the wave η. As illustrated in Figure 35, left, η connects an arc in the cusp of $\mathcal{N}(\tau_D)$ to an arc in the cusp of $\mathcal{N}(\tau_E)$. We may suppose η intersects $\partial_+ P$ minimally. Note that the proof of Claim 3.1.5 only uses the minimality of $\eta \cap \partial_+ P$ and the symmetry of δ and ε' under the involution, which means that Claim 3.1.5 also holds for η. That is, $\partial_+ P$ divides η into a pair of junction arcs at the ends and possibly some vertical arcs in $\mathscr{A}_l$ and $\mathscr{A}_r$. Thus, after isotopy, the new curve resulting from the wave move along η intersects $\widehat{\Sigma}$ in a collection of vertical arcs in the rectangles $\mathscr{R}^u$, $\mathscr{R}^d$ and annuli $\mathscr{A}_l$, $\mathscr{A}_r$. Furthermore, Claim 3.1.6 also holds, that is, any wave with respect to the new set of meridians $\{\alpha', \gamma'\}$ must intersect $\partial_+ P$. Therefore, by repeatedly applying Claims 3.1.5 and 3.1.6, we may assume $\{\alpha', \gamma'\}$ always intersect $\widehat{\Sigma}$ in vertical arcs after isotopy.

Before proceeding, we study how these wave moves and isotopies affect the x- and y-arcs of τ_D. Without loss of generality, assume that the cusps of τ_D, which are symmetric under π, lie in $\mathscr{A}_l$.

There are four possible configurations of δ in $\mathscr{A}_l$, as shown in Figure 37. In Figure 37(a, b), the cusp direction of the two cusps of τ_D are horizontal in $\mathscr{A}_l$ and the shaded regions correspond to the x-arc of τ_D. In this case, after a number of wave moves, the shaded regions in Figure 37(a, b) no longer contain any α'- or γ'-arcs. Thus, the next wave arc must pass through the shaded region in Figure 37(a)–(b).

As illustrated in Figure 15, center and right, there is an isotopy so that, after the isotopy, the wave becomes a union of junction arcs and vertical arcs. Notice that the effect of the isotopy from the center diagram to the right-hand diagram in Figure 15 changes the configuration of $\delta \cap \mathscr{A}_l$ from that of Figure 37(a)–(b) to that of Figure 37(c)–(d). Therefore, it remains to consider the configurations in Figure 37(c)–(d) for $\delta \cap \mathscr{A}_l$. In this case the two y-arcs of τ_D are ρ_x and the short path in $\mathscr{A}_l$. Note that in the configurations of Figure 37(c)–(d), the cusp directions of the two cusp of τ_D point into $\mathscr{R}^u$ and $\mathscr{R}^d$.

As in the discussion in Section 3.1, it follows from Lemma 2.1.4, that as long as $(\alpha' \cup \gamma') \cap (\mathscr{R}^u \cup \mathscr{R}^d) \neq \varnothing$, the curves $\alpha' \cup \gamma'$ contain $[\delta^+, \delta^-]$-edges inside $\mathscr{R}^u$ and $\mathscr{R}^d$. Thus, by the hypothesis of Case 2, after a number of wave moves, we reach a situation where $(\alpha' \cup \gamma') \cap (\mathscr{R}^u \cup \mathscr{R}^d) = \varnothing$.

The continuation of the argument is similar in spirit to the argument in the proof of Claim 3.1.8.

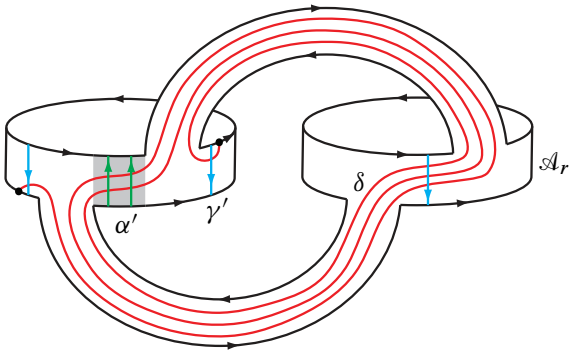

Figure 38: Orientations of $\partial_+ P$, α' and γ'.

Fix an orientation for $\partial_+ P$ and as illustrated in Figure 38, the two endpoints of each vertical arc of $(\alpha' \cup \gamma') \cap \mathscr{A}_l$ and $(\alpha' \cup \gamma') \cap \mathscr{A}_r$ are a pair of intersection points of $(\alpha' \cup \gamma') \cap \partial_+ P$ with the same sign. Since $(\alpha' \cup \gamma') \cap (\mathcal{R}^u \cup \mathcal{R}^d) = \varnothing$ at this stage, the intersection $(\alpha' \cup \gamma') \cap \hat{\Sigma}$ consists of vertical arcs in $\mathscr{A}_l$ and $\mathscr{A}_r$. This implies that all the intersection points of $\alpha' \cap \partial_+ P$ have the same sign and the intersection points of $\gamma' \cap \partial_+ P$ all have the same sign.

Let R_x and R_y be the two rectangles of $\mathscr{A}_l \setminus (\mathcal{J}_l^u \cup \mathcal{J}_l^d)$ and suppose $X \subset \partial R_x$. See the shaded region in Figure 38 for a picture of R_y. Note that ρ_x is a y-arc of the train track τ_D and the arcs of $\delta \cap R_y$ project onto the other y-arc of τ_D when we pinch δ onto τ_D. Moreover, since both $\gamma \cap \mathcal{R}^d \neq \varnothing$ and $\gamma \cap \mathcal{R}^u \neq \varnothing$ before any wave move, at least one wave move must be performed to reach this stage. Similarly to the proof of Case 1, this implies that $\alpha' \cup \gamma'$ must intersect both y-arcs of the train track. In other words, $(\alpha' \cup \gamma') \cap R_y \neq \varnothing$ and $(\alpha' \cup \gamma') \cap \rho_x \neq \varnothing$.

Since the intersection points of δ with α' and γ' all have the same sign, the arcs of $(\alpha' \cup \gamma') \cap R_y$ and the arcs of $(\alpha' \cup \gamma') \cap R_x$ which meet ρ_x must have opposite directions with respect to the core curve $\mathfrak{a}_l$ of $\mathscr{A}_l$; see Figure 38 for a picture. However, the argument above showed that the intersection points of $\alpha' \cap \partial_+ P$ all have the same sign and the intersection points of $\gamma' \cap \partial_+ P$ all have the same sign. These conclusions imply that arcs in $(\alpha' \cup \gamma') \cap R_y$ must belong to only one of the curves α' or γ'. Without loss of generality, suppose that all arcs in $(\alpha' \cup \gamma') \cap R_y$ belong to α'. Then the arcs of $\alpha' \cup \gamma'$ that meet ρ_x all belong to γ'.

The conclusion above says that one y-arc of τ_D intersects only α' and the other y-arc of τ_D intersects only γ'. Therefore, during the sequence of splittings, once $(\alpha' \cup \gamma') \cap (\mathcal{R}^u \cup \mathcal{R}^d) = \varnothing$, then both y-arcs of τ_D are pure y-arcs and the situation becomes similar to Case 1.

Now we can apply the argument for τ_E and ε' in Case 1 to τ_D and δ (ie we switch the roles of τ_D and τ_E). As indicated in Figure 38, the union of the δ-curves in $\mathcal{R}^u$, $\mathcal{R}^d$ and the short path in $\mathscr{A}_r$ are pinched into the x-arc of τ_D. By the argument in Case 1, after some wave moves, we reach a situation that the curves $\alpha' \cup \gamma'$ intersect the x-arc in a single point. Since the involution π interchanges $\mathcal{R}^u$ and $\mathcal{R}^d$, the symmetry from π implies that, if $\alpha' \cup \gamma'$ intersects $\mathcal{R}^u$, it must also intersect $\mathcal{R}^d$. As $\alpha' \cup \gamma'$ intersects the x-arc in a

single point, at this stage, we have that $(\alpha' \cup \gamma') \cap (\mathcal{R}^u \cup \mathcal{R}^d) = \varnothing$ and $\alpha' \cup \gamma'$ intersects the short path in $\mathcal{A}_r$ exactly once; see the blue arc in $\mathcal{A}_r$ in Figure 38. Now the situation becomes exactly the same as Case 1: by performing a sequence of wave moves with respect to $\{\alpha', \gamma'\}$ and possibly splitting the train track τ_D, we can eventually either conclude that K is a Berge–Gabai knot or obtain a contradiction to Theorems 1.3.7 and 1.3.8.

Case 3 The weight of ε' at some segment of τ_E is one, and $\Gamma(\delta, \varepsilon')$ contains no more $[\varepsilon'^-, \varepsilon'^+]$-edges after the last wave move η_k.

The idea of the proof is to convert Case 3 to the setup of Case 2. By the construction of τ_E, the hypothesis of Case 3 implies that there is a properly embedded arc τ in the once-punctured torus $T = \Sigma \setminus \mathcal{N}(\tau_D)$ such that $\tau \cap \tau_E = \tau \cap \varepsilon'$ is a single point. One can band-sum parallel copies of the planar surfaces P_l and P_r, as in Claim 3.5.2, to obtain a planar surface $P_{\varepsilon'}$ such that $\partial_+ P_{\varepsilon'} \cap T = \tau$. So, $(P_{\varepsilon'}, E')$ is a $(\mathcal{P}, \mathcal{D})$-pair, where E' is the disk in W bounded by ε'. In particular, $\partial_+ P_{\varepsilon'} \cap \tau_E = \partial_+ P_{\varepsilon'} \cap \varepsilon'$ is a single point since $\tau \cap \tau_E = \tau \cap \varepsilon'$.

By Claim 3.5.4, we may assume that

$$|(\alpha \cup \gamma) \cap \tau_E| \leq c_0(P, D, \alpha, \gamma).$$

Since $\partial_+ P_{\varepsilon'}$ intersects both ε' and τ_E in a single point, the definition of the complexity implies that

$$c_0(P_{\varepsilon'}, E', \alpha, \gamma) \leq |(\alpha \cup \gamma) \cap \tau_E|.$$

Thus,

$$c_0(P_{\varepsilon'}, E', \alpha, \gamma) \leq c_0(P, D, \alpha, \gamma).$$

Since $c_0(P, D, \alpha, \gamma)$ is minimal among all $(\mathcal{P}, \mathcal{D})$-pairs and all such curves α, the equality holds. In particular,

$$c_0(P_{\varepsilon'}, E', \alpha, \gamma) = |(\alpha \cup \gamma) \cap \tau_E|.$$

We may assume that $P_{\varepsilon'}$ is minimal in the sense that $|\partial_+ P_{\varepsilon'} \cap (\alpha \cup \gamma)|$ is the smallest among all such planar surfaces $P_{\varepsilon'}$ with $\partial_+ P_{\varepsilon'}$ intersecting both τ_E and ε' in a single point.

Let $\widehat{\Sigma}'_\varepsilon$ be the closure (under the path metric) of $\Sigma \setminus \partial_+ P_{\varepsilon'}$. Similar to the decomposition of $\widehat{\Sigma}$ into rectangles and annuli, we have a decomposition of $\widehat{\Sigma}'_\varepsilon$ into a pair of annuli joined by a pair of rectangles. Moreover, each component of $(\alpha \cup \gamma) \cap \widehat{\Sigma}'_\varepsilon$ is a cocore arc of an annulus or a rectangle in this decomposition. Since $c_0(P_{\varepsilon'}, E', \alpha, \gamma) = |(\alpha \cup \gamma) \cap \tau_E|$, each component of $(\alpha \cup \gamma) \cap \widehat{\Sigma}'_\varepsilon$ intersects τ_E in at most one point. As in the discussion on $\widehat{\Sigma}$ in Section 2.1, a cocore arc of a rectangle in the decomposition is a boundary arc of a ∂-compressing disk for $P_{\varepsilon'}$, and if one performs a ∂-compression on $P_{\varepsilon'}$, one obtains two planar surfaces whose ∂_+-boundaries are the core curves of the two annuli in the decomposition. This, plus the assumption that $|\partial_+ P_{\varepsilon'} \cap (\alpha \cup \gamma)|$ is the smallest among all such planar surfaces, implies that Lemma 2.1.4 is also true for ε' and this decomposition of $\widehat{\Sigma}_\varepsilon$. In other words, ε' passes through

each rectangle in the decomposition of $\widehat{\Sigma}'_\varepsilon$ at least twice. In particular, this implies that the curves $\alpha \cup \gamma$ must contain an $[\varepsilon'^+, \varepsilon'^-]$-edge as long as they intersect the two rectangles in the rectangle-annulus decomposition of $\widehat{\Sigma}'_\varepsilon$.

Since τ_E only has two cusps and since each component of $(\alpha \cup \gamma) \cap \widehat{\Sigma}'_\varepsilon$ intersects τ_E in at most one point, the structure of τ_E in $\widehat{\Sigma}'_\varepsilon$ is the same as the structure of τ_D in $\widehat{\Sigma}$. In particular, ε' must take one short path in each annulus in this decomposition. Since ε' passes through each rectangle in the decomposition of $\widehat{\Sigma}'_\varepsilon$ at least twice (similar to Lemma 2.1.4), we have all the ingredients and the argument in Case 2 also works for ε', $\partial_+ P_{\varepsilon'}$ and $\widehat{\Sigma}'_\varepsilon$ in this setup. Thus Case 3 follows from the argument in Case 2 after switching the roles of δ, τ_D and ε', τ_E respectively and using the rectangle-annulus structure of $\widehat{\Sigma}'_\varepsilon$ instead of $\widehat{\Sigma}$. So, after a sequence of wave moves with respect to $\{\alpha', \gamma'\}$, we can eventually either conclude that K is doubly primitive or obtain a contradiction to Theorems 1.3.7 and 1.3.8. □

3.6 The proof

In this section we prove the main theorem.

Proof of Theorem 7 We begin with the meridional systems $\widehat{W} = \{\delta, \varepsilon\}$ and $\widehat{V} = \{\alpha, \gamma\}$, as chosen in Section 1.3 for the Heegaard diagram for M. As described at the beginning of Section 3, we consider all possible configurations of δ depending on the paths that δ takes in the two annuli.

Now Propositions 3.1.1, 3.2.1, 3.3.1 and 3.5.1 deal with all possible configurations of δ. They show that either K is doubly primitive or that the corresponding Heegaard diagram is not induced by a Heegaard splitting of M. This proves Theorem 7. □

References

[1] **I Agol**, *Bounds on exceptional Dehn filling*, Geom. Topol. 4 (2000) 431–449 MR Zbl

[2] **K L Baker**, *The Poincaré homology sphere, lens space surgeries, and some knots with tunnel number two*, Pacific J. Math. 305 (2020) 1–27 MR Zbl

[3] **K L Baker**, **D Buck**, **A G Lecuona**, *Some knots in $S^1 \times S^2$ with lens space surgeries*, Comm. Anal. Geom. 24 (2016) 431–470 MR Zbl

[4] **K L Baker**, **B G Doleshal**, **N Hoffman**, *On manifolds with multiple lens space fillings*, Bol. Soc. Mat. Mex. 20 (2014) 405–447 MR Zbl

[5] **K L Baker**, **C Gordon**, **J Luecke**, *Bridge number, Heegaard genus and non-integral Dehn surgery*, Trans. Amer. Math. Soc. 367 (2015) 5753–5830 MR Zbl

[6] **K L Baker**, **J E Grigsby**, **M Hedden**, *Grid diagrams for lens spaces and combinatorial knot Floer homology*, Int. Math. Res. Not. 2008 (2008) art. id. rnm024 MR Zbl

[7] **J Berge**, *Some knots with surgeries yielding lens spaces*, preprint (1990) arXiv 1802.09722

[8] **J Berge**, *The knots in $D^2 \times S^1$ which have nontrivial Dehn surgeries that yield $D^2 \times S^1$*, Topology Appl. 38 (1991) 1–19 MR Zbl

[9] **J Berge**, *The simple closed curves in genus two Heegaard surfaces of S^3 which are double-primitives*, unpublished manuscript (2010)

[10] **S A Bleiler**, **C D Hodgson**, *Spherical space forms and Dehn filling*, Topology 35 (1996) 809–833 MR Zbl

[11] **S A Bleiler**, **R A Litherland**, *Lens spaces and Dehn surgery*, Proc. Amer. Math. Soc. 107 (1989) 1127–1131 MR Zbl

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

[13] **S Boyer**, **X Zhang**, *On Culler–Shalen seminorms and Dehn filling*, Ann. of Math. 148 (1998) 737–801 MR Zbl

[14] **M Culler**, **C M Gordon**, **J Luecke**, **P B Shalen**, *Dehn surgery on knots*, Ann. of Math. 125 (1987) 237–300 MR Zbl

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

[16] **D Gabai**, *Surgery on knots in solid tori*, Topology 28 (1989) 1–6 MR Zbl

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

[18] **C M Gordon**, **J Luecke**, *Reducible manifolds and Dehn surgery*, Topology 35 (1996) 385–409 MR Zbl

[19] **J E Greene**, *The lens space realization problem*, Ann. of Math. 177 (2013) 449–511 MR Zbl

[20] **D J Heath**, **H-J Song**, *Unknotting tunnels for $P(-2,3,7)$*, J. Knot Theory Ramifications 14 (2005) 1077–1085 MR Zbl

[21] **T Homma**, **M Ochiai**, **M-o Takahashi**, *An algorithm for recognizing S^3 in 3-manifolds with Heegaard splittings of genus two*, Osaka Math. J. 17 (1980) 625–648 MR Zbl

[22] **R Kirby**, *Problems in low-dimensional topology*, from "Geometric topology" (W H Kazez, editor), AMS/IP Stud. Adv. Math. 2.2, Amer. Math. Soc., Providence, RI (1997) 35–473 MR Zbl

[23] **M Lackenby**, *Word hyperbolic Dehn surgery*, Invent. Math. 140 (2000) 243–282 MR Zbl

[24] **M Lackenby**, **R Meyerhoff**, *The maximal number of exceptional Dehn surgeries*, Invent. Math. 191 (2013) 341–382 MR Zbl

[25] **L Moser**, *Elementary surgery along a torus knot*, Pacific J. Math. 38 (1971) 737–745 MR Zbl

[26] **S Negami**, **K Okita**, *The splittability and triviality of 3-bridge links*, Trans. Amer. Math. Soc. 289 (1985) 253–280 MR Zbl

[27] **Y Ni**, *Knot Floer homology detects fibred knots*, Invent. Math. 170 (2007) 577–608 MR Zbl

[28] **M Ochiai**, *Heegaard diagrams and Whitehead graphs*, Math. Sem. Notes Kobe Univ. 7 (1979) 573–591 MR Zbl

[29] **P Ozsváth**, **Z Szabó**, *On knot Floer homology and lens space surgeries*, Topology 44 (2005) 1281–1300 MR Zbl

[30] **J Rasmussen**, *Lens space surgeries and L-space homology spheres*, preprint (2007) arXiv 0710.2531

[31] **T Saito**, *A note on lens space surgeries: orders of fundamental groups versus Seifert genera*, J. Knot Theory Ramifications 20 (2011) 617–624 MR Zbl

[32] **M Tange**, *Lens spaces given from L-space homology 3-spheres*, Exp. Math. 18 (2009) 285–301 MR Zbl

[33] **M Tange**, *A complete list of lens spaces constructed by Dehn surgery, I*, preprint (2010) arXiv 1005.3512

[34] **W T Thurston**, *The geometry and topology of three-manifolds*, lecture notes, Princeton University (1979) Available at `https://url.msp.org/gt3m`

[35] **S C Wang**, *Cyclic surgery on knots*, Proc. Amer. Math. Soc. 107 (1989) 1091–1094 MR Zbl

[36] **Y Q Wu**, *Cyclic surgery and satellite knots*, Topology Appl. 36 (1990) 205–208 MR Zbl

Department of Mathematics, Boston College
Chestnut Hill, MA, United States

Department of Mathematics, Technion
Haifa, Israel

Department of Mathematics, Technion
Haifa, Israel

`taoli@bc.edu, ymoriah@tx.technion.ac.il, talipi@technion.ac.il`

Proposed: Cameron Gordon
Seconded: Mladen Bestvina, Ian Agol

Received: 2 August 2023
Revised: 17 August 2024

GEOMETRY & TOPOLOGY

msp.org/gt

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

The subscription price for 2025 is US $865/year for the electronic version, and $1210/year (+$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 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.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

GEOMETRY & TOPOLOGY

Volume 29 Issue 6 (pages 2783–3343) 2025