

ESSENTIAL TWISTED SURFACES IN ALTERNATING LINK COMPLEMENTS

MARC LACKENBY AND JESSICA S. PURCELL

ABSTRACT. Checkerboard surfaces in alternating link complements are used frequently to determine information about the link. However, when many crossings are added to a single twist region of a link diagram, the geometry of the link complement stabilizes (approaches a geometric limit), but a corresponding checkerboard surface increases in complexity with crossing number. In this paper, we generalize checkerboard surfaces to certain immersed surfaces, called twisted checkerboard surfaces, whose geometry better reflects that of the alternating link in many cases. We describe the surfaces, show that they are essential in the complement of an alternating link, and discuss their properties, including an analysis of homotopy classes of arcs on the surfaces in the link complement.

1. INTRODUCTION

Essential surfaces in link complements have played an important role in geometric topology and knot theory. The checkerboard surfaces in alternating links are particularly important. They have been used to analyze volumes [6], to obtain singular structures [2], and to give a polyhedral decomposition [10], among other things. Menasco and Thistlethwaite proved they are incompressible and boundary incompressible [9].

The genus of a checkerboard surface is determined by the crossing number of the diagram. When more and more crossings are added to a single twist region of a diagram of a hyperbolic link, the genus of a corresponding checkerboard surface increases without bound, while the link complement approaches a geometric limit. For this reason, checkerboard surfaces are not always ideally suited for analyzing geometric properties of a hyperbolic link complement.

In this paper we generalize checkerboard surfaces in alternating link complements to another class of surfaces, which we call twisted checkerboard surfaces. These surfaces are immersed in the link complement rather than embedded, but they capture the geometry of the link complement in useful ways when the link has many crossings in some twist regions. These surfaces feature prominently in our recent proof that alternating knots have cusp volume bounded below by a linear function of the twist number of the knot [7].

The main result of this paper is to show that these surfaces are essential. We also analyze geometric and homotopic properties of these surfaces. For example, we determine when two distinct arcs on the surface will be homotopic in the link complement.

To define the surfaces and state our results precisely, we recall some definitions.

1.1. Definitions. A diagram of a link K is said to be *prime* if, for each simple closed curve γ that lies on the plane of projection and meets the diagram transversely exactly

twice in the interiors of edges, the curve γ bounds on one side a portion of the diagram with no crossings.

Menasco showed that any nonsplit, prime, alternating diagram specifies either a hyperbolic link or a $(2, q)$ -torus link [8]. We will be concerned only with hyperbolic alternating links in this paper.

For an alternating link, we define a *twist region* to be a string of bigons arranged end to end in the diagram graph, which is maximal in that there are no additional bigons on either end. A single crossing adjacent to no bigons is also defined to be a twist region.

When we consider diagrams of alternating links, we often want them to have as few twist regions as possible, in the following sense.

Definition 1.1. A diagram is *twist reduced* if any simple closed curve that meets the diagram graph in exactly two vertices and that, at each crossing, runs between opposite regions, encloses a string of bigons of the diagram on one side. (See, for example, [6, Figure 3].)

Suppose a simple closed curve γ in the projection plane meets the diagram in exactly two vertices and that, at each crossing, runs between opposite regions. By sliding γ to contain both vertices on one side, and then applying a flype to the other side, we can either remove both crossings, or move one of the crossings to be in the same twist region as the other. Thus every alternating link has a twist reduced alternating diagram.

Definition 1.2. The *twist number* of an alternating link is the number of twist regions in a twist reduced diagram. We denote the twist number of the link K by $\text{tw}(K)$.

The twist number of an alternating knot is an invariant of the knot, due to work of Dasbach and Lin [3].

1.2. Twisted surfaces. In this subsection, we will define the twisted checkerboard surfaces. First, fix a prime, twist reduced, alternating diagram of the hyperbolic alternating link K . Throughout, we will abuse notation and refer to the link and its diagram by K .

For each twist region of K with at least N_{tw} crossings, where N_{tw} will be determined later, we will *augment* the diagram to obtain a new link diagram, as in Figure 1. That is, to the diagram, add a *crossing circle*, which is a simple closed curve encircling the twist region and bounding a disk in S^3 . We will always ensure that this new crossing circle introduces exactly four new crossings. So, the crossing circle is divided into four arcs. We ensure that two of these arcs are parallel in the diagram, and are as close as possible to the twist region. More precisely, we ensure that there is a square-shaped region of the diagram that includes two of the arcs of the crossing circle, and that there is a triangular region of the diagram that includes a crossing of the twist region and an arc of the associated crossing circle.

Let L be the link consisting of K along with all such crossing circles. By work of Adams [1], the complement of L is hyperbolic. For C a crossing circle, note that $S^3 \setminus C$ is a solid torus. Hence $S^3 \setminus L$ is homeomorphic to $S^3 \setminus \tilde{L}$ where \tilde{L} is obtained from L by removing any even number of crossings from the twist region encircled by C , and the homeomorphism is given by twisting the solid torus $S^3 \setminus C$.




FIGURE 1. An example of link diagrams, left to right, K , L , L_2 and K_2 .

Definition 1.3. Define L_2 to be the diagram in which all except one or two crossings have been removed from each twist region of L encircled by a crossing circle, depending on whether the number of crossings in that twist region in L is odd or even, respectively. When two crossings remain in a twist region, we place the crossing circle so that two of its arcs run through the bigon in the twist region between these two crossings.

Define L_0 to be the diagram in which all except one or zero crossings have been removed. When one crossing remains in a twist region, either in L_0 or L_2 , we still require that the crossing and the crossing circle that encircles it form a triangle in the diagram.

Finally, we let K_i be the (diagram of the) link given by removing the crossing circles from the diagram of L_i , where $i = 0$ or 2 . An example of K , L , L_2 and K_2 is shown in Figure 1.

Definition 1.4. For any crossing encircled by a crossing circle of L_i , for $i = 0$ or 2 , we say the crossing is *associated* with the crossing circle.

We build twisted checkerboard surfaces as follows. Note the diagram of K_i is alternating. Start with its checkerboard surfaces, colored blue and red. Now, when we put the crossing circles of L_i back into the diagram, a small regular neighborhood of each crossing circle intersects a checkerboard surface (either blue or red) in two meridian disks, shown in two shades of blue on the left of Figure 2.




FIGURE 2. Effect of twisting on the intersection of the checkerboard surfaces of K_i with neighborhood of a crossing circle.

Define the *red and blue surfaces*, embedded in the exterior of L_i , to be the punctured red and blue checkerboard surfaces of K_i , respectively, punctured by the crossing circles of L_i . Denote these by R_i and B_i , where $i = 0$ or 2 depending on how many crossings are left in twist regions with an even number of crossings.




FIGURE 3. A cross section of the solid torus added to $S^3 \setminus L$, and how the surface extends into it.

Consider what happens to R_i and B_i under the homeomorphism $(S^3 \setminus L_i) \rightarrow (S^3 \setminus L)$. In a neighborhood of the disk bounded by each crossing circle, the surface is twisted. The two meridian curves go to $1/n$ curves, where $2|n|$ is the number of crossings removed to go from the twist region of L to that of L_i . (The sign on n must be chosen appropriately.) Let R_{tw} denote the minimal number of crossings removed from a twist region. Thus, R_{tw} is an even integer, and $2|n| \geq R_{\text{tw}}$. Note also that crossings are only removed from a twist region if it has at least N_{tw} crossings, and hence

$$(1) \quad R_{\text{tw}} \geq 2\lfloor N_{\text{tw}}/2 \rfloor \text{ if } i = 0 \text{ and } R_{\text{tw}} \geq 2\lceil N_{\text{tw}}/2 \rceil - 2 \text{ if } i = 2.$$

To obtain $S^3 \setminus K$ from $S^3 \setminus L$, we do a meridian Dehn filling on each crossing circle. To construct the twisted checkerboard surfaces, which we continue to color red and blue, do the following. Each cross-sectional meridional disk of a crossing circle in L intersects the punctured blue (or red) surface in $2|n|$ points on the boundary of the disk. Connect opposite points on that disk by attaching an interval that runs through the center of the disk. In other words, attach an I -bundle over S^1 that runs through the center of the Dehn-filling solid torus $|n|$ times. See Figure 3. If n is odd, we attach an annulus. If n is even, each interval has both endpoints on the same curve, and so we attach two Möbius bands. In either case, the result is an immersed surface in $S^3 \setminus K$ which we call the *twisted checkerboard surface*. We continue to color it red or blue, and we denote it by $S_{R,i}$ and $S_{B,i}$, respectively, where $i = 0$ or 2 depending on how many crossings we leave in twist regions containing an even number of crossings.

The following is one of the main results of this paper.

Theorem 1.5. *Let $f: S_{B,i} \rightarrow S^3 \setminus K$ be the immersion of $S_{B,i}$ into $S^3 \setminus K$. Then this immersion is π_1 -injective, provided $N_{\text{tw}} \geq 22$ if $i = 0$, and $N_{\text{tw}} \geq 33$ if $i = 2$.*

Note that by switching the roles of the blue and red surfaces, Theorem 1.5 also implies that the red surface $S_{R,i}$ is π_1 -injective.

There is also a version of this theorem which establishes that, in a suitable sense, the surfaces $S_{B,i}$ and $S_{R,i}$ are boundary-incompressible. Since this term is used in several distinct ways in the literature, we introduce an alternative. We say that a map $f: S \rightarrow M$ between a surface S and a 3-manifold M satisfying $f(\partial S) \subset \partial M$ is *boundary- π_1 -injective* if, for any arc $\alpha: I \rightarrow S$ with endpoints in ∂S , the existence of a homotopy (rel endpoints) of $f \circ \alpha$ into ∂M implies the existence of a homotopy (rel endpoints) of α into ∂S .

Theorem 1.6. *The surface $S_{B,i}$ is boundary- π_1 -injective in $S^3 \setminus \text{int}(N(K))$, provided $N_{\text{tw}} \geq 22$ if $i = 0$ and $N_{\text{tw}} \geq 33$ if $i = 2$.*

Here and throughout, $N(K)$ denotes an embedded regular neighborhood of K in S^3 .

1.3. Acknowledgements. Purcell is supported in part by NSF grant DMS-1252687, and by a Sloan Research Fellowship.

2. GRAPHS FROM SURFACES

In the proofs of Theorems 1.5 and 1.6, the arguments for $i = 0$ and $i = 2$ are slightly different, but use much of the same machinery. In particular, both involve an analysis of a graph in a disk, obtained by the following lemmas.

Lemma 2.1. *If $f: S_{B,i} \rightarrow S^3 \setminus K$ is not π_1 -injective, then there is a map of a disk $\phi: D \rightarrow S^3 \setminus K$ with $\phi|_{\partial D} = f \circ \ell$ for some essential loop ℓ in $S_{B,i}$, such that $\Gamma_B = \phi^{-1}(f(S_{B,i}))$ is a collection of embedded closed curves and an embedded graph in D . The vertices of the graph are those points in D that map to a crossing circle. Each vertex in the interior of D has valence a non-zero multiple of $2n_j$, where $2n_j$ is the number of crossings removed from the twist region at the relevant crossing circle. Each vertex on ∂D has valence $n_j + 1$.*

Proof. If $f: S_{B,i} \rightarrow S^3 \setminus K$ is not π_1 -injective, then there is some essential closed curve γ in $S_{B,i}$ such that $[f(\gamma)] = 0$ in $\pi_1(S^3 \setminus K)$. Consider the curve γ^2 . Note we still have $[f(\gamma^2)] = 0$ in $\pi_1(S^3 \setminus K)$. However, γ^2 lifts to the orientable double cover of $S_{B,i}$. We will use γ^2 for this reason.

The fact that $[f(\gamma^2)] = 0$ in the fundamental group gives a map $\phi: D \rightarrow S^3 \setminus K$ of a disk D into $S^3 \setminus K$ with $\phi|_{\partial D} = f \circ \gamma^2$. However, we need to ensure that ϕ has the correct behaviour near ∂D , and so we construct ϕ in two stages, first near ∂D , and then over the remainder of the interior of D .

Let C denote the crossing circles in $S^3 \setminus K$, and let $N(C)$ denote a small regular neighborhood of C in $S^3 \setminus K$. By construction, $f^{-1}(C)$ is a collection of simple closed curves in $S_{B,i}$, with $f^{-1}(N(C))$ a collection of annuli and Möbius bands. We may ensure that the loop γ^2 in $S_{B,i}$ is transverse to these curves, and so it intersects the annuli and Möbius bands transversely. Each component of $\gamma^2 \cap f^{-1}(N(C))$ is therefore an arc which is mapped into $D^2 \times \{p\}$, for some $p \in S^1$, in a component of $N(C)$ and which runs from one prong of the relevant star to the opposite one. We call this a “sheet” of the star.

Now, γ^2 lifts to the orientable double cover of $S_{B,i}$. We may pick a consistent transverse orientation on this double cover, and using this, we may homotope $f \circ \gamma^2$ in this transverse direction. This homotopy is a map of an annulus into $S^3 \setminus K$, which we take to be the restriction of ϕ to a collar neighbourhood of D . By carefully choosing this homotopy, we can ensure that the image of this homotopy has well-behaved intersection with $N(C)$, as follows.

Consider any arc component of $\gamma^2 \cap f^{-1}(N(C))$, mapping to $D^2 \times \{p\}$ in a component of $N(C)$. Now note $\phi(\partial D)$ runs through $\{0\} \times \{p\}$ in a prong of a star. In a ball around $\{0\} \times \{p\}$, the image of $f(S_{B,i})$ is homeomorphic to the product of an open interval and that star. The homotopy of $\phi(N(\partial D))$ pushes it so that it remains on one side of this sheet of the star. As for an arc component of $\gamma^2 \cap f^{-1}(N(C))$, homotoping a neighborhood of that arc in the direction of the transverse orientation ensures the neighborhood only

intersects $f(S_{B,i})$ on that arc. So, in all cases, near ∂D , $\phi^{-1}(f(S_{B,i}))$ looks like a graph with each vertex on ∂D having valence $n_j + 1$.

We have thus defined ϕ in a collar neighbourhood of ∂D . Since we are assuming that the $[f(\gamma^2)]$ is trivial in $\pi_1(S^3 \setminus K)$, we may extend this to a map $\phi: D \rightarrow S^3 \setminus K$. Using a small homotopy supported away from a neighbourhood of ∂D , make ϕ transverse to all crossing circles, and transverse to $f(S_{B,i})$. Consider $\Gamma_B = \phi^{-1}(f(S_{B,i}))$ on D . Because $S_{B,i}$ is embedded in $S^3 \setminus K$ except at crossing circles, Γ_B consists of embedded closed curves, and embedded arcs (edges) with endpoints corresponding to points of intersection of crossing circles (vertices). Each vertex in the interior of D corresponds to the intersection of $\phi(D)$ with a crossing circle in $S^3 \setminus K$, and the intersection is transverse. Hence it meets the boundary of a neighborhood of that crossing circle in a meridian. The blue surface $f(S_{B,i})$ meets the boundary of the neighborhood of that crossing circle in two curves of slope $\pm 1/n_j$, where $2n_j$ is the number of crossings in K removed from the twist region of that crossing circle. Hence the valence of the vertex in the interior of D is $2n_j$. \square

Note that the vertices in the interior of D have valence precisely $2n_j$ in the above construction. However, later, it will be convenient to permit the existence of vertices that have valence a non-zero multiple of $2n_j$.

A similar result holds when f is not boundary- π_1 -injective.

Lemma 2.2. *Suppose $f: S_{B,i} \rightarrow S^3 \setminus \text{int}(N(K))$ is not boundary- π_1 -injective. Then there is a map of a disk $\phi: D \rightarrow S^3 \setminus \text{int}(N(K))$ with ∂D expressed as a concatenation of two arcs, one mapped by ϕ into $\partial N(K)$, and the other factoring through an essential arc in $S_{B,i}$. Moreover, $\Gamma_B = \phi^{-1}(f(S_{B,i}))$ is a collection of embedded closed curves and an embedded graph on D whose edges have endpoints either at vertices where $\phi(D)$ meets a crossing circle, or on $\phi^{-1}(\partial N(K))$ on ∂D . Each vertex in the interior of D has valence a non-zero multiple of $2n_j$, where $2n_j$ is the number of crossings removed from the twist region at the relevant crossing circle. Each vertex in the interior of the arc in ∂D that maps to $S_{B,i}$ has valence $n_j + 1$. Each vertex on the arc in ∂D that maps to $\partial N(K)$ has valence one.*

Proof. If $f: S_{B,i} \rightarrow S^3 \setminus \text{int}(N(K))$ is not boundary- π_1 -injective, then there is a nontrivial arc on $S_{B,i}$ which is homotopic (rel endpoints) into $\partial N(K)$ in $S^3 \setminus \text{int}(N(K))$. This gives us a map of a disk $\phi: D \rightarrow S^3 \setminus \text{int}(N(K))$ with ∂D consisting of the two arcs required by the lemma. Let α be the sub-arc in ∂D that maps via the essential arc in $S_{B,i}$. Again, we need to control ϕ near ∂D , and so we construct ϕ in two stages. The arc α lifts to the orientable double cover of $S_{B,i}$, which is transversely orientable, and by pushing α in this transverse direction, we obtain the map ϕ in a neighbourhood of α . Similarly, using the fact that $\partial N(K)$ is transversely orientable, we can extend the definition of ϕ over a collar neighbourhood of ∂D . Now extend ϕ over all of D , and then make it transverse to all crossing circles, and transverse to $f(S_{B,i})$. Let $\Gamma_B = \phi^{-1}(f(S_{B,i}))$ on D . Because $S_{B,i}$ is embedded in $S^3 \setminus K$ except at crossing circles, Γ_B consists of embedded closed curves, embedded arcs (edges) with endpoints corresponding to points of intersection of crossing circles (vertices), or with endpoints on $\partial N(K)$.

As in the proof of Lemma 2.1, a vertex in the interior of D corresponds to a transverse intersection of $\phi(D)$ with a crossing circle in $S^3 \setminus K$. Hence the vertex has valence $2n_j$.

Near a vertex in the interior of the arc in ∂D that maps to $S_{B,i}$, the graph looks like half a meridian disk for a crossing circle, and so has valence $n_j + 1$. At a vertex on the arc in ∂D that maps to $\partial N(K)$, the arc in ∂D is transverse to $S_{B,i}$, and so this vertex of Γ_B has valence one. \square

The following well-known result will be central to our proof.

Lemma 2.3. *Let Γ be a connected graph in the 2-sphere that has no bigons and no monogons, and that is neither an isolated vertex nor a single edge joining two vertices. Then Γ contains at least three vertices with valence less than 6.*

Proof. We may add edges to the graph until every complementary region is triangular. Let V , E and F denote the number of vertices, edges and faces. Then $2E = 3F$. Hence,

$$2 = V - E + F = V - E/3 = \sum_v (1 - (d(v)/6)),$$

where the sum runs over each vertex v , and $d(v)$ denotes the valence of a vertex. Since $d(v) > 0$ for each v , we deduce that $1 - (d(v)/6) < 1$, and hence there must be at least three vertices with valence less than 6. \square

Lemma 2.4. *Let Γ be a connected graph in the disk D that includes ∂D , that contains no bigons and no monogons. Then either there is some vertex in the interior of D with valence at most 5, or there are at least three vertices on the boundary with valence at most 3.*


Proof. Double the disk D to form a 2-sphere, and double Γ to form a graph Γ^+ in this 2-sphere. Now Γ^+ contains no bigons or monogons, since this was true of Γ . By Lemma 2.3, Γ^+ must have at least three vertices with valence less than 6. If one of these vertices is disjoint from the copy of ∂D in the 2-sphere, the lemma is proved. On the other hand, if all three vertices lie on ∂D , then their valence in Γ is at most 3. \square

We now focus on the graph Γ_B provided by Lemma 2.1 or Lemma 2.2.

We declare certain edges of Γ_B with at least one endpoint on ∂D to be *trivial*. The precise condition will be given in Definition 4.8, but it does not concern us here. However, if there is a bigon region of Γ_B , then either both of its edges are trivial or neither are. (See Lemma 4.11.) As a result, we say that a *trivial bigon family* is a connected union of trivial bigons, homeomorphic to a disk, and which is maximal, in the sense that none of the bigons are incident along an edge to a trivial bigon not in the family. In addition, when two edges of a triangular region of Γ_B are trivial, then so is the third, again by Lemma 4.11. Another property of trivial edges is that the vertices at their endpoints correspond to the same crossing circle.

Lemma 2.5. *Let Γ_B be the graph in D provided by Lemma 2.1. Suppose that Γ_B has no monogons. Assume also that there are no trivial edges of Γ_B in ∂D . Then Γ_B must have at least $(R_{\text{tw}}/5) - 2$ adjacent non-trivial bigons, where R_{tw} is the minimal number of crossings removed from a twist region.*

Proof. We start by proving a simpler claim: Γ_B must have at least $(R_{\text{tw}}/5)$ adjacent bigons (which may possibly be trivial). Suppose that this claim is not true. Then every collection


FIGURE 4. Trivial arcs and trivial stars in Γ_B

of adjacent bigons has cardinality less than $(R_{\text{tw}}/5)$. Collapse all the adjacent bigons to single edges. Then every interior vertex in the resulting graph has valence more than 5, and every vertex on ∂D has valence more than 3. This contradicts Lemma 2.4. Thus, the main challenge in the proof is to deal with trivial bigons.

Suppose first that there are at least two trivial bigon families that are incident along an interior vertex. Then consider all of the trivial bigon families that are incident to this vertex. These form a star-shaped collection flaring out from this interior vertex. We therefore call such a collection a *trivial star*. In addition, consider all the trivial arcs with both endpoints on ∂D . The complement of these arcs and all the trivial stars is a collection of disks. It is possible that at least one of these disks intersects ∂D in a single point, which occurs when two trivial bigon families have the same pair of endpoints. In this case, we consider this subdisk of D . On the other hand, if all of the disks intersect ∂D in at least one arc, then we consider one that is outermost in D . (See Figure 4.) This separates off a subdisk, which we will still call D . Note that the interior of this subdisk has non-empty intersection with Γ_B , as otherwise the subdisk is a triangle with two of the edges lying on trivial bigons and the third side therefore forming a trivial edge in ∂D , which is contrary to hypothesis, or the subdisk is a collection of trivial bigons, which also implies that some edge of Γ_B in ∂D is trivial. (See Lemma 4.11.)

We will consider the restriction of Γ_B to this subdisk D . If this is not connected, then pass to an innermost connected component. For simplicity, we will still call this graph Γ_B .

Now suppose that Γ_B has an edge loop, in other words, an edge that starts and ends at the same vertex. By hypothesis, this does not bound a monogon, and so there are parts of Γ_B inside the disk that it bounds. Consider an innermost such disk, and focus


 FIGURE 5. Construction of Γ'

on its intersection with Γ_B . Note that in this case, this graph has a single vertex on the boundary of the disk.

We designate certain vertices of Γ_B as *exceptional*. When D is a subdisk bounded by an edge loop, we term the unique vertex on ∂D exceptional. When D is a subdisk separated off by a trivial star, we term the two or three vertices that lay in the trivial star exceptional. When D is a subdisk separated off by a trivial arc, we term the two vertices at the endpoint of this arc exceptional.

Note that the vertices in the interior of D have valence a non-zero multiple of $2n_j$, where $2n_j$ is the number of crossings removed from the twist region at the relevant crossing circle. The vertices on ∂D have valence $n_j + 1$, except possibly the exceptional ones.

We now form a graph Γ' in D as follows. Consider each unexceptional vertex v on ∂D in turn. If v is incident to no trivial bigons, then we leave it untouched. If v is incident to $k \geq 1$ trivial bigon families, then the other endpoints of these bigon families are distinct vertices in the interior of D , by the subdisk replacements we have made above. Replace v with k vertices on ∂D . We view these k vertices as the vertices at the endpoints of the bigon families. Remove the edges meeting v that were part of a trivial bigon family. The edges of Γ_B that used to end at one of these bigon family vertices now end at one of the new vertices. There may have been some edges of Γ_B that ended on v but that were not part of a bigon family. A new endpoint has to be found for such edges. There are two possible choices: we make one arbitrarily. (See Figure 5.) Finally, if Γ_B had any exceptional vertices, then we collapse these to a single vertex.

The resulting graph Γ' has no monogons. This is because any monogon region of Γ' must have come from a monogon, a trivial bigon region of Γ_B , or a trivial arc with both endpoints in the boundary of D , and all of these have been dealt with in our construction above. However, Γ' may have some bigon regions. These arise, for instance, when a triangular region of Γ_B has one of its edges on a trivial bigon with one vertex on ∂D .

In this case, the vertex of the trivial bigon lying on ∂D is removed, and when we assign the endpoints of edges to which it is incident, the triangular region may be collapsed to a bigon. However, we deduce that for any collection of adjacent bigons in Γ' , all but at most two of these bigons came from a bigon in Γ_B . Moreover, if one of these is not a bigon for Γ_B , then it is outermost in the collection.

Note also that every interior vertex of Γ' has valence at least R_{tw} . Moreover, the vertices on the boundary of D have valence at least $(R_{\text{tw}}/2) + 1$, with at most one exception. This is because each unexceptional vertex on ∂D either came from a vertex of Γ_B on ∂D that was untouched, or it came from a vertex in the interior of Γ_B that was at the endpoint of a family of trivial bigons. This family of trivial bigons consisted of at most $n_j - 1$ edges, and so at least $n_j + 1$ edges remain in Γ' . Here, we are using the fact that trivial edges have endpoints on the same crossing circle, and so the values of n_j at the endpoints of a trivial arc are the same.

Now suppose that the lemma is not true, and that every collection of adjacent non-trivial bigons in Γ has cardinality less than $(R_{\text{tw}}/5) - 2$. Then every collection of adjacent non-trivial bigons in Γ' has cardinality less than $(R_{\text{tw}}/5)$. Now collapse all the adjacent bigons of Γ' to single edges. We deduce that every interior vertex of this graph has valence more than 5, and every vertex on the boundary of D , with at most one exception, has valence more than 3. This contradicts Lemma 2.4. \square

Lemma 2.6. *Let Γ_B be the graph in D provided by Lemma 2.2. Suppose that Γ_B has no monogons. Assume also that there are no trivial edges of Γ_B in ∂D . Then Γ_B must have at least $(R_{\text{tw}}/5) - 2$ adjacent non-trivial bigons or there are at least $(R_{\text{tw}}/5) - 2$ adjacent triangles, where one edge of each triangle lies in $\phi^{-1}(\partial N(K))$.*

Proof. We argue as in the proof of Lemma 2.5. Whenever we pass to a subdisk we make sure that it does not contain the arc $\phi^{-1}(\partial N(K))$, and the argument then proceeds exactly as in Lemma 2.5. So, suppose that we do not pass to a subdisk. Then we double D^2 along the arc $\phi^{-1}(\partial N(K))$. The two copies of Γ_B become a single graph, and the vertices that lay on $\phi^{-1}(\partial N(K))$ become the midpoints of edges. In the resulting graph, every vertex on the boundary of the disk has valence $n_j + 1$, where $2n_j$ is the number of crossings removed from the twist region at the relevant crossing circle. We can therefore apply the argument of Lemma 2.5. In all cases, the conclusion is that there are at least $(R_{\text{tw}}/5) - 2$ adjacent non-trivial bigons. In the original disk D , these give at least $(R_{\text{tw}}/5) - 2$ adjacent non-trivial bigons, or at least $(R_{\text{tw}}/5) - 2$ adjacent triangles, where one edge of each triangle lies in $\phi^{-1}(\partial N(K))$. \square

By analyzing properties of alternating diagrams, we will show below that the graph on D coming from $S_{B,0}$ cannot contain three adjacent non-trivial bigons. The graph from $S_{B,2}$ cannot contain five adjacent non-trivial bigons. Assuming these results, we give the proof of Theorem 1.5.

Proof of Theorem 1.5. If $f: S_{B,i} \rightarrow S^3 \setminus K$ is not π_1 -injective, then Lemma 2.1 implies there is a map of a disk $\phi: D \rightarrow S^3 \setminus K$ with $\phi|_{\partial D} = f \circ \ell$ for some essential loop ℓ in $S_{B,i}$, such that $\Gamma_B = \phi^{-1}(f(S_{B,i}))$ is a collection of embedded closed curves and an embedded graph in D . Each vertex in the interior of D has valence a non-zero multiple of $2n_j$, where

$2n_j$ is the number of crossings removed from the twist region at the relevant crossing circle. Each vertex on ∂D has valence $n_j + 1$. In Section 4, we will define a measure of complexity for such maps ϕ , and we choose ϕ to have minimal complexity.

Lemma 4.6 implies that there are no simple closed curves in Γ_B . Lemma 4.7 implies that there are no monogons. Lemma 4.10 states that there are no edges in ∂D that are trivial. So, by Lemma 2.5, Γ_B has at least $(R_{\text{tw}}/5) - 2$ adjacent non-trivial bigons. On the other hand, Proposition 5.8 implies there cannot be three adjacent non-trivial bigons when $i = 0$. If $N_{\text{tw}} \geq 22$, then by equation (1), $R_{\text{tw}} \geq 22$, and there are more than two adjacent bigons. So we deduce that the surface $S_{B,0}$ is π_1 -injective in this case.

For the case $i = 2$, if N_{tw} is at least 33, equation (1) implies that R_{tw} is at least 32, and Lemma 2.5 implies that the graph on D contains more than four adjacent non-trivial bigons. But now, Proposition 5.9 implies that there cannot be five adjacent non-trivial bigons on D when $i = 2$. \square

Similarly, assuming the above results as well as Lemma 6.3, we may prove Theorem 1.6.

Proof of Theorem 1.6. Suppose $S_{B,i}$ is not boundary- π_1 -injective. Then Lemma 2.2 gives a graph on a disk D . This satisfies the hypotheses of Lemma 2.6. So, Γ_B must have at least $(R_{\text{tw}}/5) - 2$ adjacent non-trivial bigons or there at least $(R_{\text{tw}}/5) - 2$ adjacent triangles, where one edge of each triangle lies in $\phi^{-1}(\partial N(K))$. The former case is ruled out by Proposition 5.8 or Proposition 5.9. The latter case is ruled out by Lemma 6.3. \square

3. DIAGRAMS AND PROPERTIES

Our goal is to complete the proofs of the Lemmas referenced in the proofs of Theorems 1.5 and 1.6. The arguments are combinatorial, relying on properties of the diagrams of the links defined in Definition 1.3 above. In this section we discuss these diagram properties.

To simplify the argument, we will consider a modification of the diagrams of L , L_i , and K_i . In particular, recall that we obtained L from K by adding a crossing circle to each twist region of K that had more than N_{tw} crossings. These crossing circles either meet the blue or the red checkerboard surfaces of K . Form a new augmented alternating link L_B from K by only adding those crossing circles that meet the blue checkerboard surface of K . Obtain $L_{B,i}$ by removing pairs of crossings encircled by each crossing circle of L_B , leaving either one or i crossings, where $i = 0$ or $i = 2$. Just as in Definition 1.3, we place the crossing circle in the diagram so that any associated crossing forms a triangle with the crossing circle. In particular, if $i = 2$ and a crossing circle is associated with two crossings, the crossing circle runs through the bigon formed by those two crossings.

Finally, obtain $K_{B,i}$ by removing crossing circles of $L_{B,i}$. Note $K_{B,i}$ differs from K in that red bigon regions have been removed.

Now, notice that the surface B_i embedded in $S^3 \setminus L_i$ can also be embedded in $S^3 \setminus L_{B,i}$. However, now the red surface R_i in $S^3 \setminus L_{B,i}$ is homeomorphic to the red checkerboard surface of K , and to the red checkerboard surface of $K_{B,i}$.

The following lemma discusses the primality of the diagrams of $K_{B,0}$ and $K_{B,2}$.

Lemma 3.1. *The diagrams of $K_{B,2}$ and K_2 are prime.*

The diagram of $K_{B,0}$ may not be prime. However, if γ is a simple closed curve giving a counterexample to primality of $K_{B,0}$, then γ consists of two arcs, one γ_R in the red surface,

and one γ_B in the blue. The arc γ_R can be homotoped (rel endpoints) to run transversely through some crossing disk in the diagram of $L_{B,0}$, intersecting it exactly once. The arc γ_B can be homotoped (rel endpoints) to be disjoint from crossing disks.

Finally, if the diagram of K_0 is not prime, then a simple closed curve giving a counterexample to primality of K_0 must link a crossing circle of L_0 .

Here, a *crossing disk* denotes the twice punctured disk with boundary on the crossing circle, embedded transverse to the plane of projection of the diagram of L_i (or L). There is one crossing disk for each crossing circle. The collection of all disks is embedded.

Proof. Suppose γ is a closed curve in the projection plane meeting the diagram of $K_{B,2}$ twice. Then we claim γ can be isotoped to be disjoint from any crossing disk of $L_{B,2}$. For if γ meets such a disk in the blue surface, since crossing disks intersect blue regions in simple arcs running from the boundary to a point in the interior of the region, the curve γ can simply be pulled off the end of the crossing disk. If γ meets a crossing disk D in the red surface, then it can either be pulled off the disk, or it meets an edge of $L_{B,2}$ running between a crossing associated with D and the crossing disk. Now slide γ along this edge towards D , in a neighborhood of the edge, until γ no longer meets D in the red surface. Then slide γ off of D in the blue surface, as above. Thus in all cases, we may isotope γ to be disjoint from D . Now put back all the crossings to obtain the diagram of K . Because γ misses all crossing disks, it misses all these crossings, hence gives a curve in K meeting the diagram twice. Because the diagram of K is prime, there are no crossings on one side of the curve. Then the same is true for $K_{B,2}$: there are no crossings in the diagram on one side of γ , and $K_{B,2}$ is prime.

Obtain K_2 from $K_{B,2}$ by removing blue bigons in the diagram of $K_{B,2}$. Since $K_{B,2}$ is prime, the same argument as above with $K_{B,2}$ replacing K , and K_2 replacing $K_{B,2}$, and the red surface replacing blue implies that the diagram of K_2 is prime.

Now consider $K_{B,0}$. Let γ be a curve meeting the diagram twice with crossings on either side. Because γ meets the diagram twice in interiors of edges of the diagrams, and because each edge of the diagram meets the red and blue surfaces on either side, we obtain the claim that γ consists of two arcs, γ_R in the red surface and γ_B in the blue. Note that γ_R lies in a single (red) region of the diagram graph, and such a region is a disk. Similarly for γ_B . As above, we may homotope γ_B (rel endpoints) to avoid all crossing disks.

Now, the diagrams of $K_{B,0}$ differs from that of K only in that an even number of crossings in select twist regions have been removed. Homotope γ_R (rel endpoints) in the red region to meet as few crossing disks as possible. If γ_R does not meet a crossing disk, then neither does γ , and we may put back crossings to obtain K without increasing the number of intersections of γ with the diagram. But then γ gives an embedded closed curve in the diagram of K meeting the diagram twice with crossings on either side, contradicting primality of K .

So γ_R must meet a crossing disk. Because γ_R lies in a single red region (which is a disk), γ_R can be homotoped (rel endpoints) to meet the crossing disk exactly once. Because γ_B does not meet the crossing disk, γ must link the corresponding crossing circle exactly once.

Finally, consider K_0 . If γ is a simple closed curve meeting the diagram of K_0 twice with crossings on either side, modify the diagram by putting back crossings bounding

blue bigons to obtain $K_{B,0}$. If after this modification γ still meets the diagram twice with crossings on either side, then the previous paragraph implies that γ links a crossing circle of L_0 . If not, then γ in $K_{B,0}$ must run through a sequence of blue bigons, consisting of crossings in a twist region. As before, γ must link the associated crossing circle. \square

Lemma 3.1 has the following important consequence.

Lemma 3.2. *For $i = 1, 2$, label the regions of the complement of the diagram of $K_{B,i}$ blue or red depending on whether they meet the blue or red surface. Note each crossing circle of $L_{B,i}$ intersects two blue regions.*

- (1) *The blue regions on opposite sides of a crossing of $L_{B,i}$ cannot agree, for $i = 0, 2$.*
- (2) *The red regions on opposite sides of a crossing of $L_{B,2}$ cannot agree. (Note this is not necessarily true for $L_{B,0}$.)*
- (3) *The two blue regions that meet a single crossing circle of $L_{B,i}$ cannot agree. That is, each crossing circle of $L_{B,i}$ meets two distinct blue regions, for $i = 0, 2$.*
- (4) *Suppose the distinct blue regions meeting a single crossing circle meet at the same crossing of the diagram of $K_{B,i}$. Then that crossing is associated with the crossing circle, as in Definition 1.4, for $i = 0, 2$.*

Proof. For (1), if the blue regions on opposite sides of the same crossing do agree, then draw an arc in this region from one side of the crossing to the other. Close this into a simple closed curve by drawing a short arc in a red region meeting that crossing, close enough to the crossing that it does not meet any crossing disks. This simple closed curve contradicts Lemma 3.1. Similarly for item (2), if red regions on opposite sides of a crossing of $L_{B,2}$ agree, then we may connect them into a closed curve with crossings on either side, contradicting Lemma 3.1.

For item (3), if the blue regions meeting one crossing circle agree, we may draw an embedded arc through that region with endpoints on the two intersections of the crossing circle. Attach to this the arc of intersection of the crossing disk D with the projection plane. This gives a simple closed curve γ in the diagram of $K_{B,i}$ meeting the diagram twice. Push slightly off D to ensure that the arc of γ in the red surface meets no crossing disks. By Lemma 3.1, γ must have no crossings on one side. But then we may isotope D away from $K_{B,i}$. This contradicts the fact that the diagram of K is prime.

For (4), draw arcs from either side of the crossing to the intersections of the crossing circle with the blue regions. Connect these by the arc of intersection of the crossing disk with the projection plane. Now twist, yielding K . We have a closed curve of the diagram meeting a crossing coming from the twist region of the crossing circle, and also meeting a single crossing of the diagram. Because K is twist reduced, these crossings belong to the same twist region. Hence the original crossing is associated with that crossing circle. \square

The diagrams $K_{B,i}$, for $i = 0, 2$ may not be twist reduced. However, we will need the fact that they are blue twist reduced, as defined below.

Definition 3.3. A diagram of a link is *blue twist reduced* if every simple closed curve in the projection plane meeting the diagram in exactly two crossings, with sides on the blue checkerboard surface, bounds a string of red bigons.

Lemma 3.4. *For K a knot with prime, twist reduced, alternating diagram, the diagram of $K_{B,i}$, for $i = 0, 2$, is blue twist reduced.*

Proof. Suppose there exists a closed curve γ in the diagram of $K_{B,i}$ meeting the diagram in blue regions, and meeting exactly two crossings. By adding crossings to the diagram, giving red bigons, we obtain the diagram of K . Since γ is disjoint from the red surface, we may isotope it to be disjoint from crossing disks whose boundary crossing circle meets the blue surface. Then when we add back in the crossings and red bigons to obtain the diagram of K , the curve γ remains disjoint from the red surface, and meets the diagram of K in exactly two vertices.

Because K is twist reduced, γ bounds a string of bigons on one side in K , and they must be red bigons. The crossings forming the bigons are all in the same twist region. Either they remain in the diagram of $K_{B,i}$ when we remove crossings from K , or some of them are removed to obtain the diagram of $K_{B,i}$. In the latter case, $i = 2$ and γ bounds a single bigon between the two crossings left in such a twist region in $K_{B,2}$. In either case the lemma is proved. \square

The following lemma will lead to contradictions in particular cases.

Lemma 3.5. *For K a hyperbolic knot with a prime, twist reduced, alternating diagram, suppose $K_{B,i}$ is a $(2, 2)$ -torus link. Then $L_{B,i}$ cannot have a crossing circle encircling the two crossings of $K_{B,i}$, for $i = 0, 2$.*

Proof. If $L_{B,i}$ has a crossing circle encircling the two crossings of a $(2, 2)$ -torus link, then both crossings are associated with that crossing circle. When $i = 0$, this contradicts the definition of $L_{B,0}$: it has at most one crossing associated with any crossing circle.

When $i = 2$, the diagram of $L_{B,2}$ has at least one crossing associated with any crossing circle. Since the diagram of $K_{B,2}$ has just two crossings, and both are associated with the given crossing circle, there can be no other crossing circles in $L_{B,2}$. Hence when we put in the crossings to go from the diagram of $L_{B,2}$ to that of K , we only add crossings to the given twist region. Thus K is a $(2, 2q)$ -torus link. This contradicts the assumption that K is a hyperbolic alternating link. \square

4. SURFACE PROPERTIES

In this section, as well as the next two, we will give restrictions on the graphs Γ_B coming from Lemmas 2.1 and 2.2, as well as a similar graph obtained in Lemma 7.2. To analyze these graphs, we will actually be considering three surfaces in $S^3 \setminus L_{B,i}$. The first surface is the blue surface B_i , which becomes $S_{B,i}$ in $S^3 \setminus K$ by attaching annuli or Möbius bands. The second surface is the red surface R_i , which we have noted is a checkerboard surface for $K_{B,i}$, and is embedded in $S^3 \setminus K$, $S^3 \setminus L_{B,i}$, and $S^3 \setminus K_{B,i}$. The third surface we color *green*. It consists of all crossing disks bounded by the crossing circles of $L_{B,i}$. The green surface is embedded in $S^3 \setminus L_{B,i}$, since each crossing disk is embedded and disjoint from the others.

4.1. Graphs on a disk. Lemmas 2.1 and 2.2 give us a graph Γ_B on a disk D . In fact, the results in this section, as well as in Sections 5 and 6, apply to any graph Γ_B on a disk D

coming from the pull back of $S_{B,i}$ under a map $\phi: D \rightarrow S^3 \setminus \text{int}(N(K))$, as in Lemmas 2.1 and 2.2, and in Lemma 7.2 in Section 7.

For all these graphs Γ_B , the subscript B stands for *blue* because Γ_B is the set of points in D that map to the blue surface $S_{B,i}$. Note that Γ_B has a finite collection of isolated vertices on D and ∂D . Hence removing these points, we obtain an embedding of a punctured disk $\phi': D' \rightarrow S^3 \setminus L_{B,i}$. The map ϕ' is transverse to the blue, red, and green surfaces in $L_{B,i}$. Pulling back the intersections of these surfaces to D' , we obtain a graph Γ_{BRG} with blue, red, and green edges on a punctured disk. So, Γ_B is a subgraph of Γ_{BRG} . There is also an intermediate subgraph Γ_{BR} consisting of only the blue and red edges.

Each edge of these graphs maps to an arc in the diagram $L_{B,i}$. The arrangement of these arcs in this diagram will play a central role in this paper. Note that we are focusing on the diagram $L_{B,i}$ and hence every crossing circle punctures only blue regions.

We record now how corresponding edges may meet. The proof follows immediately from the definitions, but we put the information into a lemma for future reference.

Lemma 4.1. *Suppose $\phi: D \rightarrow S^3 \setminus K$ is a map of a disk meeting the blue, red, and green surfaces transversely, such that the pull back of these surfaces under ϕ gives an embedded planar graph Γ_{BRG} with blue, red, and green edges, and the restriction $\phi': D' \rightarrow S^3 \setminus L_{B,i}$ is a map of a punctured disk, with punctures mapping to crossing circles.*

- (1) *Blue edges meet blue edges only at vertices (punctures of D'). In the diagram of $L_{B,i}$, these map to arcs that lie in blue regions of the complement of $K_{B,i}$ with endpoints on crossing circles. Adjacent blue edges at a vertex map to arcs that meet the corresponding crossing circle in the two distinct regions of Lemma 3.2(3). See Figure 6, (a) and (b).*
- (2) *Because red and green surfaces are embedded, red edges do not meet red edges, and green edges do not meet green.*
- (3) *Green edges meet red on a crossing disk, along the arc in the plane of projection that runs between the two punctures of the twice-punctured disk.*
- (4) *Green edges meet blue either in the interior of blue edges, which correspond to intersections in the interior of crossing disks, or at a high valence vertex where many blue edges come together, which correspond to green edges meeting the boundary of the crossing disk.*
- (5) *Blue edges meet red edges at a crossing of the diagram. If the blue and red edges together bound a region of D' that is mapped to lie above the plane of projection, then the blue and red edges are mapped to meet the crossing as in Figure 6(c). If they bound a region mapped below the plane of projection, then they are mapped as in Figure 6(d). These are the only possibilities.*
- (6) *The projection plane of the diagram of $K_{B,i}$ is made up of red and blue surfaces; we also refer to this as the projection plane of $L_{B,i}$. Thus each region of D' is mapped above or below the plane of projection of $L_{B,i}$, with regions switching from above to below or vice-versa across red or blue edges.*

□

Note that it is, in principle, possible for the inverse image of one of the red, green or blue surfaces to contain a simple closed curve component that is disjoint from the surfaces




FIGURE 6. Adjacent blue edges of Γ_{BRG} meeting at a vertex as in (a) will lie in distinct regions as in (b). In (c), blue and red edges meeting above the projection plane. In (d), meeting below.

with other colours. However, any such component bounds a disk in one of the red, blue or green surfaces that is disjoint from the remaining surfaces, and so can easily be removed by a homotopy. We therefore assume that no such closed curves arise. Hence, as explained in the above lemma, the inverse image of the red, green and blue surfaces forms a graph Γ_{BRG} in D .

We want to arrange this graph to be as simple as possible, in a certain suitable sense, so that various trivial arrangements can be ruled out.

Definition 4.2. Define the *complexity* of $\phi': D' \rightarrow S^3 \setminus L_{B,i}$ to be the ordered set

$$C(\phi') = (\#\text{vertices}(\Gamma_B), \#\text{vertices}(\Gamma_{BR}), \#\text{vertices}(\Gamma_{BRG}), \#\text{edges}(\Gamma_{BRG}))$$

Order complexity lexicographically. We will assume that ϕ' has been chosen so that the complexity is as small as possible.

4.2. Monogons, bigons, and triangles. In this subsection and the next, we give technical results with combinatorial proofs to show that certain configurations of the graph Γ_{BRG} cannot hold. Together, these results will give the Lemmas used in the proof of Theorems 1.5 and 1.6.

We will be considering complementary regions of the graph Γ_{BRG} of red, blue, and green edges on D' . The complementary regions in D' need not be embedded in $S^3 \setminus L_{B,i}$, but they are disjoint from the red, blue, and green surfaces in $S^3 \setminus L_{B,i}$. We refer to those complementary regions that are disks by the number of edges they have, namely monogons meet one edge, bigons meet two, and triangles three. Sometimes, we will also consider subgraphs of Γ_{BRG} , for example, the subgraph Γ_{BR} consisting only of the red and blue edges. We will also refer to complementary regions of these subgraphs as monogons, bigons, triangles, and so on.

Lemma 4.3. *The graph Γ_{BRG} has no bigons with one red side and one green, disjoint from blue. More generally, there are no green edges disjoint from blue that have both endpoints on red.*

Proof. The first statement follows immediately from the second, and so we prove the second. The red surface runs through the interior of a crossing disk in $S^3 \setminus L_{B,i}$. Hence if a green edge on D' has both endpoints on red, it has endpoints meeting the same

embedded interval on the green 2-punctured disk. Because it is disjoint from blue, the arc of intersection of green and red, along with the green edge, bounds a disk on the green surface. Use this disk to homotope away the intersections with the red surface. That is, use the disk to push the green arc of intersection to the other side of the projection plane. In a neighborhood of the crossing disk, this will remove two red points of intersection with that crossing disk. It will not affect any intersections of D with crossing circles (vertices of Γ_B) or intersections of blue and red (vertices of Γ_{BR}). Hence, the graph Γ_{BRG} in D has been simplified: its number of vertices has been reduced. This contradicts our minimality assumption on complexity. \square

Lemma 4.4. *The graph Γ_{BRG} has no bigons with one green side and one blue, disjoint from red, and with at least one endpoint not being a vertex of Γ_B . More generally, Γ_{BRG} has no green edge disjoint from red with both endpoints on the same blue edge of Γ_{BR} and with at least one endpoint not being a vertex of Γ_B .*

Proof. Such a green edge must lie on a single crossing disk. The blue edge will either have both endpoints in the same region adjacent to the corresponding crossing circle, or it will have endpoints in the two distinct regions meeting that crossing circle. In the latter case, a single blue edge that meets no red must belong to a single region, contradicting Lemma 3.2(3). So the former must happen, that is, the blue edge has both of its endpoints in the same region adjacent to the crossing circle. Then the green arc must bound a disk E in the portion of the crossing disk on one side of the projection plane. We may use this disk to homotope the disk D . A homotopy of D along E removes the intersection of the green and blue in the interior of the blue edge, without affecting the number of vertices of Γ_B or Γ_{BR} , yet decreasing the number of vertices of Γ_{BRG} . This contradicts the assumption that the complexity is minimal. \square

There is another option for a bigon with green and blue sides, namely when the green edge has its endpoints on two high valence vertices of Γ_B . In this case, the blue edge must be trivial (as in Definition 4.8 below) and we will deal with trivial blue edges and trivial bigons separately (as in Section 2).

Lemma 4.5. *The graph Γ_{BR} has no bigons with one blue side and one red, whether or not the bigon meets the green surface.*

Proof. Suppose there is a bigon with one red side and one blue. Then it is mapped either completely above or completely below the projection plane. Without loss of generality, say it is mapped above. Then by Lemma 4.1(5), the edges of the bigon meet the diagram of $K_{B,i}$ as shown in Figure 7, left. But then the union of these two edges forms a simple closed curve γ meeting the diagram of $K_{B,i}$ twice with crossings on either side. If $i = 2$, or if $i = 0$ and the red edge meets no green disks, this contradicts Lemma 3.1.

So suppose $i = 0$ and the red edge of γ meets green disks. Then Lemmas 4.3 and 4.4 imply that green edges run from the blue edge to the red. Thus the green disks meet the diagram in one of the two ways shown in Figure 7, right, i.e. either with blue above or blue below.

Follow γ along the red edge. Consider the first green disk G that the red edge intersects. Replace γ with two new closed curves γ_1 and γ_2 by drawing an arc along G from a red edge




FIGURE 7. Left: A red–blue bigon that does not meet green. Right: case that it meets green

of γ to a blue edge, and then splitting γ along this arc. Both closed curves γ_1 and γ_2 meet the diagram exactly twice. One of them, say γ_1 , bounds crossings on either side. But note that γ_1 has one fewer points of intersection with the green disk than γ . Thus, by induction on the number of intersections of the red edge with green, we obtain a contradiction to Lemma 3.1. \square

Lemma 4.6. *Consider the blue graph Γ_B alone. There are no simple closed curves of intersection of the blue surface. That is, $\phi(D)$ does not meet the blue surface in any component disjoint from crossing circles.*

Proof. Suppose there is a simple closed curve of intersection of the blue surface. By Lemmas 4.4 and 4.5, it cannot meet either the green or red surface, for an outermost arc of intersection would give an illegal bigon. So, it is a blue simple closed curve disjoint from the red and green surfaces. But as explained above, we have arranged that D contains no such curves. \square

Lemma 4.7. *In the blue graph Γ_B , there are no monogons. That is, no edges of intersection run from one vertex (corresponding to a crossing circle) back to that same vertex forming a monogon.*

Proof. Lemmas 4.4 and 4.5 imply that any blue edge forming a monogon in Γ_B cannot meet red or green edges of Γ_{BRG} in its interior. Hence the edge of a monogon runs from one side of a crossing circle through a region of the projection plane to the other side. This implies that there is a single complementary region of the diagram on both sides of a crossing circle, contradicting Lemma 3.2(3). \square

The above two lemmas imply there are no simple closed curves or monogons in Γ_B . There will be bigons, and these can be trivial or non-trivial, as mentioned in Section 2. We are now ready to define trivial edges and trivial bigons. Since Lemma 2.5 implies not all bigons can be trivial, we will then move to studying non-trivial bigons.

Definition 4.8. Define an edge of Γ_{BRG} to be *trivial* if it is a blue arc that is disjoint from the red edges, and its endpoints are distinct vertices of Γ_B in D , but correspond to the same crossing circle in $L_{B,i}$.

Let α be a trivial edge. Since it is a blue edge disjoint from the red edges of Γ_{BRG} , it corresponds to an arc in a blue region of the diagram $L_{B,i}$. By assumption, this arc has both its endpoints on the same crossing circle, and the two blue regions at the punctures




FIGURE 8. A blue arc starting and ending on the same crossing circle

of this crossing circle are distinct. Hence, α must have both its endpoints on the same puncture. It therefore forms a closed loop in $L_{B,i}$, bounding a disk E in the diagram with interior which is disjoint from the red, blue and green surfaces and which can only meet green in single edges with both endpoints on high valence blue vertices. See Figure 8.

There are four types of trivial arcs: those with neither endpoint on ∂D , those with exactly one endpoint on ∂D , those with both endpoints on ∂D , and those lying entirely in ∂D . We will deal with these types in different ways.

Lemma 4.9. *In the graph Γ_{BRG} , there are no trivial arcs with both endpoints in the interior of D .*


Proof. As explained above, such an arc α bounds a disk E in the diagram $L_{B,i}$ with interior that is disjoint from the red, blue and green surfaces. We may perform a homotopy to D , sliding α along E , so that it ends up in $\partial N(C)$, where C is the crossing circle at the endpoints of α . (See Figure 9.) The two vertices at the endpoints of α lie in meridian disks of $N(C)$ that are subsets of D . The union of these two meridian disks with a regular neighbourhood of α is a disk D' in the interior of D , that maps to $N(C)$. Now $\partial D'$ is a curve on the boundary of the solid torus $N(C)$ that is homotopically trivial. Hence, we may homotope $\partial D'$ so that it is a multiple of a meridian curve for $N(C)$. If this multiple is zero, then we may homotope D' to lie in $\partial N(C)$. It thereby misses C , and so this reduces the number of vertices of Γ_B . This contradicts our assumption that complexity is minimal. On the other hand, if $\partial D'$ represents a non-zero multiple of a meridian, then we may homotope D' so that it maps to the crossing circle at a single point, where it forms a branch point. In this way, we end up with a single vertex of Γ_B , and again we have reduced its complexity. Again we reach a contradiction. \square

Note that the procedure described in the above proof may create vertices with valence that is a non-zero multiple of n_j . It was for this reason that vertices of this form are permitted in Lemmas 2.1 and 2.2.

Lemma 4.10. *In the graph Γ_{BRG} , there are no trivial arcs that are subsets of ∂D .*

Proof. We may perform a similar homotopy to the one in Lemma 4.9, but we may take the disk E to be a subset of $S_{B,i}$. In this way, ∂D remains in $S_{B,i}$, but the complexity of D is reduced. Again, this is a contradiction. \square

Finally, suppose that there is a trivial arc in Γ_{BRG} properly embedded in D with one or two endpoints on the boundary of D . There is no obvious way of eliminating such an arc, and so we deal with them another way in the proof of Lemma 2.5.

FIGURE 9. Modifying D by a homotopy

Lemma 4.11. *If all but one of the edges of a region in Γ_B are trivial, then the remaining edge is also trivial.*

Proof. Consider such a region in Γ_B . The trivial edges in the boundary of the region are disjoint from the red and green surfaces. So, if any red arc enters the region, it must intersect the remaining blue side twice. Hence, an outermost such arc gives a bigon in Γ_{BR} with one red side and one blue side, contradicting Lemma 4.5. So, there are no red arcs in the region. If there are any green arcs intersecting the remaining blue edge in its interior, then we similarly deduce that there is a bigon in Γ_{BRG} with one blue side and one green side, and which is disjoint from the red edges, contradicting Lemma 4.4. We therefore deduce that all the edges of the region are disjoint from the red and green surfaces. The trivial edges start and end at the same crossing circle. Hence, the remaining edge does also. It is therefore trivial. \square

Lemma 4.12. *In the graph Γ_{BRG} , there are no non-trivial bigons with two blue sides, disjoint from red and green edges.*

Proof. For such a bigon, the vertices would correspond to crossing circles of $L_{B,i}$. Since the bigon is non-trivial, the vertices correspond to distinct crossing circles. Put in the crossings at the two crossing circles, and elsewhere, to form the diagram of K . The blue edges are disjoint from crossing disks, and so they remain on the projection plane in the complement of the diagram of K . They connect across (former) crossing disks to give a simple closed curve in the diagram of K that meets the diagram at exactly two crossings. Because K is twist reduced, these crossings correspond to the same twist region. But then they would have corresponded to the same crossing circle in $L_{B,i}$. This is a contradiction. \square

Lemma 4.13. *In Γ_{BRG} , there are no pairs of red–green–blue triangles (RGB triangles) adjacent across a green edge.*

Proof. Such a pair of triangles would give a red–blue bigon in Γ_{BR} , contradicting Lemma 4.5. \square




FIGURE 10. Two RGB triangles adjacent across a blue edge lead to the configurations shown.

Lemma 4.14. *In Γ_{BRG} , there are no pairs of RGB triangles adjacent across a blue edge. More generally, no pair of green edges can be added to the graph Γ_{BR} in such a way that the result is a pair of triangles adjacent across a blue edge.*

Proof. If a pair of green edges exists as in the statement of the second part of the lemma, then there exists an innermost such pair. That is, red, blue, and green edges of Γ_{BRG} bound two triangular regions between them with red, green, and blue sides, identified along the blue side. In addition, there may be other green edges in Γ_{BRG} that intersect these triangular regions, but the fact that it is innermost implies none of the green edges cut off a pair of such triangular regions adjacent across the blue side.

Because there are no additional red or blue edges inside these triangular regions, one of the triangular regions is mapped above the plane of projection, one below. Because there is only one vertex where red meets blue in the two triangles, there would be just one crossing in the diagram where the triangles meet, and the edges must run from the crossing as shown in Figure 10, middle. The green arc lies on a single crossing circle, although one edge lies above the plane of projection and one below. Hence the two red edges running from the crossing in the figure must run to the same crossing disk, and in fact to the same arc on the projection plane between the two punctures of the crossing disk.

If the red edges have endpoints on opposite sides of the crossing disk, then one of the triangles links the crossing circle nontrivially, contradicting the fact that it bounds a disk disjoint from the crossing disk. Hence the red edges have endpoints on the same side of the crossing disk. But then we can form a simple closed curve meeting the diagram of $K_{B,i}$ twice with crossings on either side, by taking the union of the red edge labeled b , the blue edge, and a portion of the crossing disk in Figure 10, right.

By Lemma 3.1, we must have $i = 0$ and this red edge must meet another crossing disk, i.e. another component of the green surface. This gives a new green edge in the graph, which runs from a red edge in an RGB triangle. It can't run to the green side of the triangle, since green edges don't meet green. Thus it runs to blue. But then the green arc continues into the other RGB triangle, and must exit through red, and we have another pair of RGB triangles adjacent across the blue, contained in the original pair. This contradicts the fact that the pair was innermost. \square

Lemma 4.15. *No green edge of Γ_{BRG} can have both its endpoints on the same vertex of Γ_B .*

Proof. Such a green edge bounds a disk E on a crossing disk. Use this disk to homotope the disk D past the crossing disk, removing a green edge of Γ_{BRG} , contradicting the fact that the graph had minimal complexity. \square

Lemma 4.16. *The graph Γ_{BRG} has no blue–blue–green triangles. More generally, no green edge can be added to Γ_{BR} to cut off a triangular region with two blue sides and one green.*

Proof. Suppose some green edge can be added to Γ_{BR} to cut off such a triangular region. Then there must be an innermost such green edge, so that by Lemmas 4.15 and 4.4, the triangular region it bounds is disjoint from all other green edges. Hence it suffices to prove the first claim.

The two blue edges of such a triangle meet at a vertex corresponding to a crossing circle, bounding a crossing disk. The green edge lies on a crossing disk. If these crossing disks agree, then each of the blue edges runs from a crossing circle back to the same crossing disk. We may form two closed curves in the projection plane, one on either side of the crossing disk, by connecting the endpoints of the blue edge together in the projection plane, and we may do so in such a way that the closed curves bound disks in the projection plane disjoint from the diagram of $L_{B,i}$. We can use one of these two disks to homotope the triangle through the crossing circle. In the graph Γ_{BRG} , this has the effect of sliding the endpoint of the green edge so that it ends at the blue vertex. This reduces the number of vertices of the graph Γ_{BRG} without affecting the number of vertices of Γ_B and Γ_{BR} , which contradicts our minimality assumption.

So suppose the crossing disk of the green edge has boundary on a different crossing circle than that corresponding to the vertex. Then as we argued in Lemma 4.12, the blue edges connected to the crossing circles can be adjusted to give a closed curve in the diagram of K meeting the diagram at exactly two crossings, each in a distinct twist region of K . This gives a contradiction to the fact that the diagram of K is twist reduced. \square

4.3. Adjacent blue–blue–red triangles. In the following lemmas, we will consider adjacent blue–blue bigons meeting red edges, which will lead to considering blue–blue–red triangles by Lemma 4.5. In this subsection, we will give restrictions on adjacent blue–blue–red triangles.

Lemma 4.17. *Suppose the graph Γ_{BR} contains three adjacent blue–blue–red triangles, adjacent at the vertex meeting blue edges, with edges labelled as in Figure 11. Then when we include again the green edges, no green edge meets edge b or c . Moreover, if a green edge meets the edge j , then it must run from the common vertex of the triangles to j . In particular, it comes from an intersection of the disk D with the crossing disk associated to the vertex.*

Proof. If green meets b , then a green edge runs through the triangles abi and bcj in Figure 11. By Lemma 4.16 (no blue–blue–green triangles), a green edge meeting b cannot meet either a or c . By Lemma 4.4 (no blue–green bigons), it cannot run from b back to b , or even to the vertex meeting the blue edges. Hence it must run to i and to j . But this contradicts Lemma 4.14 (no RGB triangles adjacent along blue). So green cannot meet b . The argument for c is symmetric.




FIGURE 11. Labels on three adjacent blue–blue–red triangles, adjacent at the blue vertex.

If green meets j , then a green edge runs through the triangle with edges b , c , and j . Since green cannot meet b or c , and cannot form a green–red bigon, it must run directly to the vertex. \square

Note that in a single blue–blue–red triangle, there are two points where blue edges meet red, each of which corresponds to a crossing of the diagram, by Lemma 4.1(5).

Lemma 4.18. *Restrict to the graph Γ_{BR} . In the case $i = 2$, a single blue–blue–red triangle meets two distinct crossings.*

Proof. If not, then the red edge of the triangle runs from one crossing back to the same crossing. By Lemma 4.1(1), the blue edges run to opposite sides of that crossing. Then since the triangle is either mapped entirely above or entirely below the plane of projection, Lemma 4.1(5) implies that the endpoints of the red edge are on opposite sides of the crossing as well. But this contradicts Lemma 3.2(2), for the diagram $L_{B,2}$. \square

In the case $i = 0$, the two crossings at endpoints of the red edge in a red–blue–blue triangle may actually not be distinct. However, when we have three adjacent triangles, we are able to rule out too much overlap of crossings.

Lemma 4.19. *Restrict to the graph Γ_{BR} . If three blue–blue–red triangles are adjacent, then each triangle must meet two distinct crossings.*

Proof. In the case $i = 2$, this is immediate from Lemma 4.18. So we will only consider the case $i = 0$, and the diagram $L_{B,0}$.

Label the triangles as in Figure 11 again. Without loss of generality, we may assume that triangles abi and cdk map above the plane of projection, and triangle bcj maps below.

First, we prove that triangle bcj cannot meet exactly one crossing. Suppose by way of contradiction that it does. Then edges b and c of the triangle must meet opposite sides of this crossing, and endpoints of j must also meet opposite sides of the crossing as shown in Figure 12, left.

It follows that j runs through the crossing disk G corresponding to the crossing circle of the vertex, as in Figure 12. By Lemma 4.17, no other green edge can meet j . But now put the crossings of K back into the diagram, to obtain the diagram as in Figure 12, right. Note that the union of j and b form a closed curve meeting the diagram of K twice. This contradicts the fact that the diagram of K is prime and hyperbolic.

So if one of the triangles meets exactly one crossing, it cannot be the middle one. It must be triangle abi or triangle cdk . The two cases are symmetric, and so we show that




FIGURE 12. If triangle bcj meets just one crossing, edges b , c , and j meet the crossing as shown left. Right: this gives a simple closed curve in the diagram of K meeting the diagram twice.

triangle abi meets two crossings. If not, then as in the previous argument, edges a and b meet opposite sides of the same crossing, call it x , and edge i must run through the crossing disk G to meet x on either side. Again if i meets no other green crossing disks, then by putting in twists corresponding to the twist region shown, we would obtain a contradiction to the fact that the diagram of K is prime and hyperbolic, just as in the above paragraph. Thus we must conclude that i meets some other green disk, a possibility which was ruled out for j by Lemma 4.17, but which is not impossible for i .

Thus there must be some green edge g' , corresponding to a green disk G' distinct from the crossing disk G at the vertex of a , b , and c , such that g' lies in the triangle abi and meets i . Note that g' cannot have its other endpoint at the vertex where a and b meet, since G' is distinct from G . By Lemmas 4.17 and 4.3, the other endpoint of g' meets a , forming a triangular region of Γ_{BRG} .

Consider such a triangle nearest the vertex where a and i meet. That is, take g' such that the triangle cut off by g' and portions of a and i in Γ_{BRG} contains no other green edges disjoint from G . Notice that with this choice of g' , the triangular region either contains no other green edges, or it contains a green edge g coming from the crossing disk G , cutting off an interior triangle.

If g' cuts off a triangle in Γ_{BRG} that contains no other green edges, then the boundary of that triangle maps to give a curve in the projection plane meeting the diagram of $K_{B,0}$ twice, disjoint from all crossing disks. If instead g' cuts off a triangle meeting a green edge g , then the boundary of the quadrilateral with sides g' , g , and portions of a and i , maps to give a closed curve in $K_{B,0}$ meeting the diagram exactly twice. In either case, by Lemma 3.1, the curve bounds no crossings. This is shown on the left of Figure 13 for the triangle case. Draw an arc in the diagram running over the strand of the knot between a and i , and label this arc ℓ .

Now put triangle bcj into the figure. One endpoint of j meets the crossing x on the side of b determined by Lemma 4.1(5). By Lemma 4.17, j does not meet the crossing disk G' . This is shown on the right of Figure 13. Because i and j are in the same region, we may connect an endpoint of the arc ℓ to the edge j without meeting the diagram of $K_{B,0}$. Similarly, because a and c are in the same region, we may connect the other endpoint of the arc ℓ to the edge c without meeting the diagram. The arc ℓ is shown by the thin dashed line in Figure 13.




FIGURE 13. Possible configuration of diagram when triangle abi meets just one crossing. The thin dashed line is the arc ℓ in the proof.




FIGURE 14. Labeling for adjacent blue–blue–red triples

Then the union of the arc ℓ , the portion of j running from ℓ to the crossing where it meets c , and the portion of c running back to ℓ , forms a closed curve γ in the diagram meeting the diagram of $K_{B,0}$ exactly twice. Observe that γ is disjoint from the crossing disk G . Lemma 3.1 implies that the portion of γ in the red face must run through a crossing disk, which implies j must run through a crossing disk, distinct from G . But Lemma 4.17 implies that j can only run through G . Together, these give a contradiction. \square

Lemma 4.20. *The graph Γ_{BR} cannot contain two pairs of three adjacent blue–blue–red triangles adjacent across the red edges, as in Figure 14.*

Proof. Suppose not.

We first argue that the two crossing circles corresponding to the two vertices where blue edges meet cannot agree. For if they do, then a , c , and f are in the same region, and b , e , and g are in the same region, and so the crossings at endpoints of a and e , at endpoints of b and f , and at endpoints of c and g are all associated with the crossing circle (Lemma 3.2(4)). By Lemma 4.19, at least two of those crossings are distinct. This immediately gives a contradiction for the diagram of $L_{B,0}$, by Definition 1.3. However, if two of those crossings agree, then we do not immediately have a contradiction in the case of $L_{B,2}$. However, in $L_{B,2}$, the following unions of edges will give simple closed curves in the diagram of $K_{B,2}$ meeting the diagram twice: $c \cup i \cup f$, $e \cup i \cup b$, $f \cup a \cup j$, $b \cup j \cup g$. Because the diagram of $K_{B,2}$ is prime (Lemma 3.1), these curves encircle portions of the diagram meeting no crossings, and the diagram of $K_{B,2}$ is that of a $(2, 2)$ -torus link, with a crossing circle of $L_{B,2}$ encircling the two crossings. This contradicts Lemma 3.5. Thus the two vertices where blue edges meet correspond to two distinct crossing circles.




FIGURE 15. Left: configuration of three adjacent blue–blue–red triangles when the triangles meet three distinct crossings. Right: diagram of $L_{B,2}$ in this case.

It follows that neither of the crossing disks corresponding to the vertices can intersect the edge j , for by Lemma 4.17, if the green meets j , then a green edge must run from j to the vertex where b and c meet, and a green edge must run from j to the vertex where g and f meet. This would imply that those two vertices correspond to the same crossing disk, which we showed cannot happen.

Additionally, we claim that we cannot have both crossing disks meeting the edge i . For if the crossing disk corresponding to the vertex of e , f , and g meets i , then there is a corresponding green edge in the triangle abi with an endpoint on i . By Lemma 4.17, its other edge meets a . Then a is in the same region as f . Similarly, if the crossing disk corresponding to the vertex of a , b , and c meets i , then e is in the same region as b . Lemma 3.2(4) implies that the crossing at the endpoints of a and e , as well as the crossing at the endpoints of b and f , must both be associated to both crossing circles. But any single crossing is associated to at most one crossing circle. This contradiction implies that we cannot have both crossing disks meeting i . Without loss of generality, assume that the crossing disk corresponding to the vertex at a , b , and c does not meet i .

Now we claim that the crossing x at the endpoints of a and e and the crossing y at the endpoints of c and g must be distinct. For if not, then $i \cup j$ maps to a closed curve in the diagram, and the two crossing circles must link $i \cup j$. But then the crossing disk corresponding to the vertex where a , b , and c meet would have to intersect either i or j , and this contradicts the above arguments.

Thus we have established: (1) The crossings x and y are distinct. (2) The crossing disk corresponding to the vertex meeting a , b , and c does not intersect i or j . By Lemma 4.19, the crossing at the endpoints of b and f is also distinct from x and y . Hence we may sketch images of the triangles abi and bcj , and endpoints of edges e , f , and g as in Figure 15, left.

Note that g starts within the region of the diagram bounded by $b \cup j \cup c \cup (\text{crossing disk})$ and e starts outside it. Now e and g share a vertex that is mapped to the same side of the same crossing circle. Hence, e and g lie in the same blue region of the diagram. Neither e nor g can meet the region of c , by Lemma 3.2(1). They cannot meet j because they are the wrong color. Hence one must meet b or the crossing disk in the region of b , and so both are in the same region as b . Then Lemma 3.2(4) implies that the crossing x at the endpoints of a and e and the crossing y at the endpoints of c and g both are associated with the




FIGURE 16. Left: A triangle–square pair. Right: Three adjacent triangle–square pairs will be labeled as shown.

crossing circle corresponding to the vertex of a, b , and c . This contradicts Definition 1.3 in the case of $L_{B,0}$: each crossing circle is associated with at most one crossing.

So we continue for the diagram $L_{B,2}$ only. In Figure 15 right, the diagram of $L_{B,2}$ is shown, with crossings associated with the crossing circle attached to that crossing circle, and with the crossing circle cutting through the bigon formed by the two crossings, as required in Definition 1.3. Now note that the edge f is enclosed by a region bounded by strands of the diagram, as well as red edges i and j . The edge f cannot cross any of these. Similarly, e and g cannot cross these. Thus the crossing circle corresponding to the vertex of e, f , and g , which is distinct from the crossing circle corresponding to the vertex of a, b , and c , must straddle either i or j . It cannot straddle j , by the above work, so it straddles i . Then a green edge with endpoint on i meets triangle abi . By Lemma 4.17, its other endpoint is on a . Then f is in the region of a and c . By Lemma 3.2(4), all three crossings shown are associated with the crossing circle of a, b , and c . This contradicts the definition of $L_{B,2}$, Definition 1.3. \square

5. TRIANGLES AND SQUARES

The main result of this section is that the graph Γ_B of Lemma 2.1 cannot contain three and five adjacent bigons, for $L_{B,0}$ and $L_{B,2}$, respectively. To prove this, we examine adjacent triangles and squares.

Suppose a blue–blue–red triangle in Γ_{BR} is adjacent to a blue and red square, adjacent across the red edge, as illustrated on the left of Figure 16. We will call this a *triangle–square pair*. In this section, we examine adjacent triangle–square pairs, adjacent at the vertex of the triangle. We show that in $L_{B,0}$ there cannot be three adjacent triangle–square pairs, and in $L_{B,2}$ there cannot be five adjacent triangle–square pairs.

Label three adjacent triangles and squares as in Figure 16. By Lemma 4.19, each triangle shown meets distinct crossings. Thus the diagram may have two, three, or four crossings coming from the intersections of blue and red edges of triangles. Up to symmetry, there are exactly four possibilities, shown in Figure 17.

Note the diagrams shown are only sketches. In options I and II, the curve $i \cup j$ must run through the crossing disk, and in option III, the curve $i \cup j \cup k$ must run through the crossing disk. We have chosen to show j running through the disk in I and II, and k in III. Additionally, we have chosen to show none of i, j , or k running through the crossing disk in option IV. However, a priori, these are arbitrary choices. Hence we must be careful to ensure that our arguments below do not depend on these choices.




FIGURE 17. Possibilities for three adjacent triangles.

For all four possibilities, we have the following lemma.

Lemma 5.1. *Suppose the edge labeled f in Figure 16 is not mapped into the region of the diagram of $K_{B,i}$ containing images of edges labeled a and c . Then f cannot have both of its endpoints mapped to the same crossing of $K_{B,i}$. Similarly, if g is not mapped into the same region as edges labeled b and d , then g cannot have both of its endpoints mapped to the same crossing.*


Proof. We will prove the lemma for f . The argument for g is symmetric. Again we will refer to an edge in Figure 16 and its image in the diagram of $L_{B,i}$ by the same name.

Suppose f is not in the region of a and c , but f has both endpoints on the same crossing. Then f forms a loop in a blue region. By our assumption that complexity is minimal, f cannot be homotoped away from this blue region. Thus it meets a crossing disk G . Note that since f does not meet the region of a and c by assumption, and f cannot meet the region of b by Lemma 3.2(1), G is disjoint from the crossing disk corresponding to the vertex at a, b, c and d .

First, we claim that G meets the edge i . Because f intersects G , there are green edges γ_1 and γ_2 , each with one endpoint on f , running through squares $eifl$ and $fjgm$ of Figure 16, respectively, corresponding to intersections with G . If γ_1 meets i , then the claim is proved. So suppose γ_1 meets e or l . By Lemma 4.17, γ_2 cannot meet j , so γ_2 meets either m or g . Finally, Lemma 4.14 implies that we cannot simultaneously have γ_1 meeting l and γ_2 meeting m . Thus either γ_1 meets e or γ_2 meets g . After mapping to the diagram, endpoints of e and g are separated from those of f by the closed curve $a \cup c \cup i \cup j$. Note that Lemma 3.2(1) implies e cannot meet the region of a (which is also the region of c), and g cannot meet the region of c (also the region of a). Hence the only way for γ_1 to meet e or γ_2 to meet g is if G runs over i or j . It does not meet j by Lemma 4.17. Hence G meets i , as claimed.

Now, G intersects i , so a green edge lies in the triangle abi . By Lemma 4.17, it does not run to b . Thus it runs from i to a . Since f and a are in distinct regions, by assumption, a and f must meet this green disk on opposite sides, as in Figure 18.

Now, edges e and l meet at a crossing. The edge e cannot be in the region of a by Lemma 3.2(1). The edge e cannot lie in the region containing f because these are separated by $a \cup i \cup j \cup c$. So e cannot meet G . Therefore, the only way for l to meet the endpoint of e is for l to run through G . Hence a green edge runs through the square $eifl$ with one


 FIGURE 18. Edges a , f , and i intersect some green crossing disk

endpoint on ℓ . That edge cannot meet i , by Lemma 4.3. It cannot meet e because e is disjoint from G . So it meets f and continues into a green edge in the square $fmgj$. This green edge cannot meet j , by Lemma 4.17. It cannot meet m by Lemma 4.14. It cannot meet g , or a and g , therefore c and g , are in the same region, contradicting Lemma 3.2(1). This gives a contradiction. \square

Lemma 5.2. *If the edge labeled f in Figure 16 is mapped to the same region as edges labeled a and c , then the crossing at the endpoints of images of b and f is associated to the crossing circle corresponding to the common vertex of the triangles of Figure 16, and in addition, at least one of the crossings at the endpoints of images of h and d or of g and c is associated to the same crossing circle.*

Notice that Lemma 5.2 implies something symmetric for g by flipping the adjacent triangle–square pairs upside down and relabeling: If g is mapped to the region of images of b and d , then the crossing at the endpoints of images of c and g is associated to the crossing circle of the vertex, in addition to at least one of the crossings at the endpoints of images of e and a or b and f .

Proof. Again we refer to edges and their images by the same name. Suppose that f is in the region of a and c . Then Lemma 3.2(4) implies that the crossing meeting endpoints of f and b is associated with the crossing circle C corresponding to the vertex of the graph.

Suppose that the crossing at the endpoints of d and h is not associated to C . This implies that h cannot be in the region of a and c , by Lemma 3.2(4). In cases I and III of Figure 17, h is automatically either in the region of f or in the region of a and c , so these cases are impossible. Thus we are either in case II or IV.

Notice that there are two versions of case II, where the crossings at the ends of a and c are the same, and where the crossings at the ends of b and d are the same. In the latter case, we deduce that the crossing at the endpoints of d and h must be associated with C . So we may assume in case II that the crossings at the ends of a and c are the same, as shown in Figure 17.

Now suppose that the crossing at the endpoints of c and g is not associated to the crossing circle. Then g cannot meet the region of b and d by Lemma 3.2(4). By Lemma 5.1 it cannot have both endpoints on the same crossing. But the endpoint of g is contained in the region bounded by $b \cup j \cup k \cup d$. Thus the other endpoint of g must also be contained in that region, and must lie on a new crossing disjoint from the existing crossings in the diagram, as in Figure 19.




FIGURE 19. Edge g meets a new crossing within the region bounded by $b \cup j \cup k \cup d$

Edges n and h have endpoints on the same crossing, but their other endpoints are bounded away from each other by $b \cup j \cup k \cup d$. Because h does not meet the region of b or d , the edge n must either meet k or j .

Suppose first that n meets j . Then n crosses into a region bounded away from the endpoint of h by $a \cup c \cup i \cup j$. We may assume, by applying a small homotopy which does not change the graphs Γ_B , Γ_{BR} , and Γ_{BRG} , that n intersects j at most once. Since h does not meet the region of a or c , the edge n must then run through i . But then f is separated from the region of a and c by red edges, contradicting assumption.

So we conclude that n meets the edge k . Now, the edge g along with a portion of n and a portion of k form a closed curve γ meeting the diagram of $K_{B,i}$ twice with crossings on either side. Lemma 3.1 implies that we are in the diagram of $K_{B,0}$, and one of the red arcs, either n or k , must meet a green disk. Moreover, the closed curve γ intersects this green disk transversely, exactly once.

Suppose first that k meets a green disk G . Then there is a green edge inside the triangle cdk , which must run from k to d by Lemma 4.17, hence d also meets G . Since the portions of $k \cup n$ making up γ intersect G only once, the diagram must be as in Figure 20, left. (Note that d cannot cross g since they are in distinct regions.) Now consider the portion of d shown in that figure, along with a portion of k and an arc on G . This forms a new closed curve γ' meeting the diagram exactly twice, with crossings on either side. If the portion of γ' on the red meets another green disk, again the disk must meet d , so we may replace γ' by a smaller closed curve consisting of an arc on the new green disk, and remnants of γ' , as in the proof of Lemma 4.5. As in the proof of that lemma, by induction, we obtain a simple closed curve meeting the diagram of $K_{B,0}$ twice without meeting any crossing disks, and with crossings on both sides, contradicting Lemma 3.1.

Hence n must meet a crossing disk G between the endpoint of n at g and the point where n meets k . Thus there is a green edge with endpoint on n in the square $gnhk$. This edge cannot have its other endpoint on k , by Lemma 4.3. We claim it cannot have an endpoint on h . This is because h has both endpoints in the region bounded by $b \cup j \cup k \cup d$. If the green edge happened to end on h , then the green disk must cross $b \cup j \cup k \cup d$. It cannot cross j , by Lemma 4.17. If it crosses b or d , then either g or h must be in the region of b and d , which is impossible by assumption, and by Lemma 3.2(1). If G meets




FIGURE 20. Left: Crossing disk meeting k and d . Right: Crossing disk meeting n and g .

k , then there is a green arc in the triangle cdk , which again implies G meets d , and so g or h is in the region of b and d , which is impossible.

Thus the green edge in square $gnhk$ with one endpoint on n has its other endpoint on g . Because the portion of $n \cup k$ making up γ meets the crossing disk exactly once, the disk must be as shown in Figure 20, right. Again portions of n , g , and the crossing disk form a simple closed curve γ' meeting the diagram of $K_{B,0}$ exactly twice. Again, as in the proof of Lemma 4.5, replacing this closed curve if necessary, we obtain a contradiction to Lemma 3.1. \square


Lemma 5.3. *If there are three adjacent triangle-square pairs, as in Figure 16, then for one of the following pairs of crossings, both crossings are associated to the crossing circle corresponding to the vertex:*

- (1) crossings at the endpoints of f and b and of d and h , or
- (2) crossings at the endpoints of f and b and of g and c , or
- (3) crossings at the endpoints of g and c and of a and e .

Proof. We prove the lemma by contradiction. By Lemmas 5.2 and 5.1, if there are three adjacent triangles and squares but the result does not hold, then edges f and g cannot be in the regions of edges a and c , and b and d , respectively, and neither f nor g can have endpoints on the same crossing. Thus f ends in a new crossing in the region $a \cup i \cup j \cup c$, and g ends in a new crossing in the region $b \cup d \cup j \cup k$.

We may add these new crossings to Figure 17. However, first note that in options I and II from that figure, we have sketched the diagram so that j runs through the crossing disk corresponding to the vertex of the graph. In fact, if j runs through that crossing disk, then a green edge corresponding to this crossing disk runs through the square $fmgj$, with one endpoint on j . By Lemma 4.3, the other endpoint must be either on f or g . This implies that either f meets the region of a and c , or g meets the region of b and d , both of which we have ruled out. Hence we need to adjust the diagrams of I and II so that i meets the crossing disk corresponding to the vertex instead of j . In III and IV, i or k might meet that crossing disk as well, but the same argument shows that j cannot. Updated sketches of the four possibilities are shown in Figure 21.

Note that the two endpoints of m are shown in each figure. In order for these endpoints to join, either m must intersect j , or the crossing disk corresponding to the vertex. This is because the endpoints of m are separated by the closed curve $b \cup j \cup c \cup (\text{crossing disk})$.

FIGURE 21. Four possibilities for endpoints of f and g

Suppose first that m intersects j . Then f , along with a portion of j and a portion of m , form a closed curve meeting the diagram of $K_{B,i}$ twice with crossings on either side. Lemma 3.1 immediately implies that the diagram is that of $K_{B,0}$, not $K_{B,2}$. Lemmas 3.1 and 4.17 imply that the portion of m meets some green crossing disk. Then there is a green edge in the square $fmgj$ meeting either the edge f or g . If f , then we may choose such a crossing disk closest to the crossing shared by m and f , and argue as in the proof of the previous lemma that we obtain a contradiction to Lemma 3.1. So the green edge runs from m to g . But this green crossing disk must be contained in the region bounded by $i \cup j \cup a \cup c$, or have endpoints in the region of a and c . Since g cannot meet the region of a and c , we obtain a contradiction.

The only remaining possibility is that m intersects the crossing disk corresponding to the vertex. Then there is a corresponding green edge running through the square $fmgj$, with one endpoint on m . Lemma 4.3 implies its other endpoint is not on j . Hence its endpoint is on f or g . But that is possible only if f is in the region of a and c , or g is in the region of b and d . \square

Lemma 5.4. *If there are three adjacent triangle-square pairs as in Figure 16, then it cannot be the case that the crossings at the endpoints of images of b and f and of c and g are both simultaneously associated to the crossing circle corresponding to the vertex.*

Proof. By Lemma 4.19, the crossing at the endpoints of f and b and the crossing at the endpoints of g and c must be distinct. Suppose both are associated to the crossing circle corresponding to the vertex, call it C . Then there are two distinct crossings associated to C , so the diagram cannot be that of $L_{B,0}$ by definition. Hence we may assume that the diagram is that of $L_{B,2}$ and no additional crossings are associated to C .

Consider each of the options from Figure 17. We will show none of these can occur.

First consider option I. The diagram has two crossings shown, and by assumption both of those crossings are associated to the crossing circle shown. This means the region of a and c is the region of f and h , and the region of b and d is the region of e and g . Then we can draw arcs in these regions from one crossing to another. These arcs, along with the edges i and j , will give closed curves in the diagram of $K_{B,2}$ meeting the diagram twice. Because the diagram of $K_{B,2}$ is prime, each bounds a strand of the knot and no

crossings. Thus the diagram of $K_{B,2}$ is that of a $(2,2)$ -torus link, with a crossing circle of $L_{B,2}$ encircling the two crossings. This contradicts Lemma 3.5.

Now consider option II. By assumption, the edge f is in the region of a and c , the edge g is in the region of b and d . Draw an arc α inside the region of b and d from the crossing at the endpoint of e (and g) to the crossing at the endpoint of b . Now $\alpha \cup i$ meets the diagram of $K_{B,2}$ twice, with crossings on either side. This contradicts Lemma 3.1.

For option III, note the crossing at endpoints of a and d must be associated to the crossing circle shown, by Lemma 3.2(4). By assumption, the other two crossings are associated to that crossing circle as well. Thus three distinct crossings are associated to the crossing circle, contradicting the definition of $L_{B,2}$.

Finally, consider option IV. Since the crossings at endpoints of b and f and of c and g are associated with the crossing circle shown, the crossing disk it bounds intersects the edge labeled j , by our convention on the diagram of Definition 1.3. Now consider the green edges of intersection in the graph. One green edge runs from the vertex to j . From there, a green edge must run to f or g , but not m by Lemma 4.3. The cases f and g are symmetric, so suppose it runs to f . Now there must be a green edge in the square $eifl$. It cannot run from f to i , by Lemma 4.14. It cannot run from f to e , or e would be in the region of b and d , and Lemma 3.2(4) would imply that a third crossing is associated to the crossing circle shown, which is a contradiction. Thus the green edge runs from f to ℓ , and ℓ must be parallel to j in the red bigon region containing j . But then e meets one of the crossings at the endpoints of j , either at the endpoint of b or at the endpoint of c . If e meets b , then e belongs to the region of b , and again we have too many crossings associated to the crossing circle. If e meets c , we get a contradiction to Lemma 3.2(3). \square

Lemma 5.5. *In three adjacent triangle-square pairs as in Figure 16, if the crossings at the endpoints of f and b and of h and d are both associated to the crossing circle of the vertex, then so is the crossing at the endpoint of a and e .*

Proof. Suppose first that the crossings at the endpoints of f and b and at the endpoints of h and d are actually the same crossing. If that crossing is associated to the crossing circle of the vertex, call it C , then Lemma 5.4 implies that the crossing at the endpoints of g and c is not associated to C . Assume by way of contradiction that the crossing at the endpoints of a and e is also not associated C . Then Lemma 3.2(4) implies that neither g nor e can meet a region meeting C , hence both g and e are disjoint from the crossing disk bounded by C .

Since C is associated to the crossing at endpoints of f and b , we know that the corresponding crossing disk intersects either i or j , by our convention on diagrams (Definition 1.3). The arguments for both are nearly identical (in fact, symmetric in the squares $fjgm$ and $iefl$). We will walk through the argument for j and leave the case of i to the reader. So suppose there is a green edge with one endpoint on j running through the square $fjgm$. The other endpoint cannot lie on m , by Lemma 4.3. It cannot lie on g by our above observation that g is disjoint from this crossing disk. Hence the green edge runs from j to f . See Figure 22.

But then there is a green edge in the square $eifl$ with one endpoint on f . It cannot have its other endpoint on i by Lemma 4.14. It cannot have its other endpoint on e since e is disjoint from the region of C . Thus its final endpoint is on ℓ , and so ℓ runs through




FIGURE 22. Left: Form of diagram if crossing at endpoints of b and d agree, and is associated to the crossing circle of the vertex.

the crossing disk. But then ℓ has an endpoint on the crossing at endpoints of b , j , and f . Because the square $eif\ell$ and the triangle bcj are both mapped either above or below the projection plane (in Figure 22, they are shown mapped below), Lemma 4.1(5) implies that j and b meet in the same regions of the diagram as ℓ and a blue edge meeting ℓ in a vertex. So either e or f is the same region as b . The edge f cannot be in the region of b by Lemma 3.2(1). But the edge e cannot be in the region of b either, because e is not in a region meeting C .

This contradiction implies that if both the crossing at the endpoints of f and b and the crossing at the endpoints of h and d are associated with C , then they must be distinct crossings. So, the diagram cannot be that of $L_{B,0}$.

But now notice that if two crossings are associated to the same crossing circle, then they come from the same twist region, so there is a (red) bigon between them. Then either j lies in that bigon region or i does. If j lies in the bigon, then its two endpoints must lie on the two crossings of the bigon. Hence c has an endpoint on one of those two crossings, and it is associated with the crossing circle of the vertex. This contradicts Lemma 5.4. Thus i lies in the bigon region between the two crossings. It follows that a has its endpoint on the crossing meeting d . Thus the crossing at the endpoints of a and e is the same as the crossing at the endpoint of d and h , and so it must be associated to the crossing circle corresponding to the vertex. \square

Lemma 5.6. *There cannot be three adjacent triangle–square pairs for the diagram of $L_{B,0}$.*

Proof. Lemma 5.3 implies that images of edges from the triangle–square pairs must meet at two crossings associated to the crossing circle C corresponding to the vertex. In $L_{B,0}$, at most one crossing can be associated to any crossing circle. So the two crossings from Lemma 5.3 must actually be the same crossing of the diagram. Lemma 4.19 implies that endpoints of b and of c map to distinct crossings, so those two cannot map to crossings associated to C . Lemma 5.5 implies that if endpoints of b and d map to a crossing associated to C , then so does a , and again Lemma 4.19 implies we have two crossings associated to C , contradicting the definition of $L_{B,0}$. The only remaining possibility is that endpoints of c and of a map to a crossing associated to C . But then by relabeling, we may again apply Lemma 5.5 to obtain a distinct crossing associated to C . In all cases, we have a contradiction. \square

Lemma 5.7. *There cannot be five adjacent triangle–square pairs for the diagram of $L_{B,2}$.*




FIGURE 23. Green circles correspond to points which map to crossings associated with the crossing circle of the vertex. Shown are three, four, and five adjacent triangle-square pairs. Note that in the case of five, the squares shown in gray contradict Lemma 5.4.

Proof. In three adjacent triangle-square pairs, there are four points where red edges meet blue edges to form triangles. Lemmas 5.3, 5.4, and 5.5 imply that three of these four points must map to crossings associated to the crossing circle corresponding to the blue vertex, and moreover, two of those four are adjacent to the top (or bottom by symmetry), and the last is adjacent to the bottom (resp. top). Start with three such triangle-square pairs, and without loss of generality assume that the top two points are associated to the crossing circle. This is shown on the left of Figure 23.

Now attempt to add a fourth adjacent triangle-square pair. This new pair will form another set of three adjacent triangle-square pairs. It cannot be added to the top of the three, else the new set of three adjacent triangle-square pairs on the top will contradict Lemma 5.4. Hence it must be added to the bottom, and the new point on the new triangle where blue meets red must be associated to the crossing circle of the blue vertex. This is shown in the center of Figure 23.

Finally, attempt to add a fifth adjacent triangle-square pair. This cannot be added to top or bottom, or we contradict Lemma 5.4. Thus there cannot be five adjacent triangle-square pairs. \square

Proposition 5.8. *When $i = 0$, the graph Γ_B cannot contain three or more adjacent non-trivial bigons.*

Proof. Suppose the graph Γ_B contains three adjacent non-trivial bigons. By Lemma 4.12, these bigons cannot be disjoint from the red and green surfaces. The bigons cannot be disjoint from the red surface, for the following reason. The bigons would have to contain at least one green edge, by Lemma 4.12 (no non-trivial blue bigons). If there is a green edge which intersects the interior of a blue edge, this contradicts Lemma 4.4 (no blue-green bigons) or 4.16 (no blue-blue-green triangles). On the other hand, if both endpoints of the green edge are the same vertex of Γ_B , this contradicts Lemma 4.15 (no green monogons). If the endpoints of the green edge are distinct vertices of Γ_B , this implies that the edges of the blue bigons are trivial, which is contrary to hypothesis. So, the blue bigons intersect the red surface. Lemma 4.5 (no red-blue bigons) implies any intersection with the red surface must run straight through all three adjacent bigons. An outermost such intersection cuts off three adjacent blue-blue-red triangles.

By Lemma 4.20, there cannot be a single such intersection of red, cutting off two pairs of three adjacent blue–blue–red triangles adjacent across the red edges. Hence there are at least two such intersections of red, and the outermost two cut off three adjacent red triangles and three adjacent red squares. But now Lemma 5.6 implies that this is impossible when $i = 0$. \square

Proposition 5.9. *When $i = 2$, the graph Γ_B cannot contain five or more adjacent non-trivial bigons.*

Proof. Suppose the graph contains five adjacent non-trivial bigons. As in the proof of Proposition 5.8, we argue that the bigons cannot be disjoint green and red, by Lemma 4.12, cannot be disjoint red, by Lemmas 4.4, 4.15 and 4.16, and red intersections run through each bigon, splitting off triangles and squares by Lemma 4.5. Again Lemma 4.20 implies there is more than one intersection of red, and the two outermost intersections cut off five adjacent triangles adjacent five adjacent squares across red edges. This contradicts Lemma 5.7. \square

This completes the proof of Theorem 1.5.

6. BOUNDARY- π_1 -INJECTIVITY

In this section, we finish proofs of lemmas needed to show twisted surfaces are boundary- π_1 -injective. To do so, we analyze further the graph Γ_B from Lemma 2.2. In the graph from that lemma, blue edges may have at least one endpoint on the part of ∂D which maps to $\partial N(K)$. Where two blue edges leave the same high valence vertex and both end on $\partial N(K)$, we obtain a triangle with two blue sides and one side on $\partial N(K)$. We show in this section that we cannot have many adjacent triangles of this form, by restricting the graph Γ_B .

Lemma 6.1. *In $L_{B,0}$, there cannot be two adjacent triangles of Γ_B each with two blue sides and one side on $\partial N(K)$, unless the triangles meet the red surface.*

Similarly, in $L_{B,2}$, there cannot be three adjacent triangles of Γ_B with two blue sides and one on $\partial N(K)$, unless the triangles meet the red surface.

Proof. Suppose two adjacent triangles in D do not meet the red surface. Label the three blue edges of the triangle as in Figure 24, left. Without loss of generality, we may assume that the triangle with edges labeled a and b is mapped above the plane of projection, and that the one with edges labeled b and c is mapped below. Since edges a and b start on opposite sides of the same crossing circle, they must end in different regions of the diagram, by Lemma 3.2(3). Because the third side of the triangle they form lies on the link, that third side must run over a single over-crossing of the diagram. This is shown in Figure 24, second from left. Similarly, the triangle with sides b and c must have third side running over a single under-crossing of the diagram. Putting these together, the result must be as in Figure 24, third from left. Note each of these crossings is associated to the crossing circle corresponding to the blue vertex of the triangles, by Lemma 3.2(4). Hence the crossing circle has at least two crossings which are associated to it. This contradicts our construction of $L_{B,0}$: at most one crossing belongs to any crossing circle.




FIGURE 24. Left to right: Labels on edges of three adjacent triangles. A triangle above the plane of projection. Two adjacent triangles. Three adjacent triangles meeting only two crossings.




FIGURE 25. Labels on triangles.

If a third triangle is adjacent, then either a third crossing will be associated to the crossing circle, which is impossible for $L_{B,2}$, or the crossing straddled by endpoints of edges in the third triangle will agree with one of the existing crossings, as on the right of Figure 24. However, repeated applications of primality, Lemma 3.1, implies that in this case the diagram is that of a $(2,2)$ -torus link encircled by a crossing circle. This contradicts Lemma 3.5. \square

Lemma 6.2. *In the graph Γ_B , three adjacent triangles, each with two blue sides and one side on $\partial N(K)$, must meet the red surface more than once.*

Proof. Suppose not. By Lemma 6.1, any three adjacent triangles must meet the red surface at least once. Suppose three adjacent triangles meet the red surface only once. Because there are no red–blue bigons (Lemma 4.5), the red must run across the three triangles, meeting each triangle in both of its blue edges. Label the edges of adjacent triangles meeting the red once as in Figure 25. By Lemma 4.19, each of the blue–blue–red triangles meets two distinct crossings. As before, the three triangles can meet two, three, or four crossings, and the four possibilities are sketched in Figure 17. However, now e and f must straddle an undercrossing, f and g straddle an overcrossing, and g and h straddle an undercrossing.

We step through the cases in Figure 17 one by one, ruling out each case.

Case I. The three blue–blue–red triangles meet just two crossings. In this case, $i \cup j$ separates endpoints of e and f . Since e and f straddle a crossing, but cannot meet i or j , they must straddle one of the crossings between i and j . Similarly, f and g straddle the other crossing between i and j . Then both crossings shown in case I of Figure 17 belong to the crossing circle shown, and a primeness argument shows that the diagram is a $(2,2)$ -torus link encircled by a crossing circle. This violates Lemma 3.5.

Case II. Again $i \cup j$ separates endpoints of e and f and of f and g , so again e and f straddle one of the crossings between i and j , and f and g straddle the other, and both crossings are associated to the crossing circle shown. So the diagram is that of $L_{B,2}$. Then one of e or g runs from the crossing at the endpoint of a to the crossing at the endpoint of b . Taking this edge, along with the edge i , gives a closed curve meeting the diagram of $K_{B,2}$ twice with crossings on either side. This contradicts Lemma 3.1.

Case III. This time, $i \cup j \cup k$ separates endpoints of e and f , and of f and g , and of g and h . Hence these pairs of edges straddle crossings between i , j , and k . There are three such pairs and three such crossings. In all cases, it can be shown that the pairs of edges straddle distinct crossings, and all three crossings are associated to the crossing circle shown. This contradicts the definition of $L_{B,0}$ and $L_{B,2}$: at most 1 or 2 crossings, respectively, can be associated to a given crossing circle.

Case IV. In this case, f and e are not required to straddle one of the four crossings shown. However, if they do not, then since endpoints of f and e are separated by $a \cup i \cup j \cup c$, and since e cannot meet any of those curves, f must intersect a or c . Thus the crossing straddled by f and e meets the same blue regions on either side as the crossing at the endpoint of a . By the fact that the diagram of $K_{B,i}$ is blue twist reduced, Lemma 3.4, these two crossings must bound a sequence of (at least one) bigon regions of the diagram between them. Then f and g straddle the next crossing in the bigon sequence, hence e and g are in the same region of the diagram. But endpoints of e and g are bounded away from each other by $b \cup j \cup k \cup d$, hence e and g are both in the region of b and d . That implies that crossings at the endpoints of a and c belong to the crossing circle shown. Since f is in the region of a , the crossing at the endpoint of b also belongs to that crossing circle. This contradicts the definition of $L_{B,i}$.

Finally, it remains to show that e and f cannot straddle any of the crossings shown in the diagram on the left of Figure 17. Arguments similar to those above imply that for each of these crossings, if e and f straddle the crossing then three of the four crossings shown must be associated to the crossing circle shown. We leave the details to the reader. \square

Lemma 6.3. *In the graph Γ_B , when $i = 0$, there cannot be three adjacent triangles each with one side on $\partial N(K)$. When $i = 2$, there cannot be five such triangles.*

Proof. Lemma 6.1 implies the triangles with two blue edges and one edge on $\partial N(K)$ must meet the red surface. Lemma 6.2 implies that the triangles must meet the red surface at least twice. But then the two intersections closest to the vertex of the triangles cut the triangles into blue–blue–red triangles adjacent to red and blue squares. Lemma 5.6 gives a contradiction in case $i = 0$: no disk can be mapped into $L_{B,0}$ such that three adjacent blue–blue–red triangles meet adjacent red and blue squares. In the case $i = 2$, Lemma 5.7 gives a contradiction. \square

This completes the proof of Theorem 1.6.

7. PROPERTIES OF TWISTED SURFACES

In this section, we investigate homotopy classes of arcs in twisted checkerboard surfaces. This requires machinery of the previous sections, and has applications in [7].




FIGURE 26. The subsurface of a checkerboard surface associated with a twist region.

Consider two arcs that are distinct and essential in the surface $S_{B,i}$, for $i = 0, 2$, but homotopic when mapped into $S^3 \setminus K$. We determine when this can happen.

First, we introduce terminology. A small regular neighborhood of a twist region in S^3 is a 3-ball which intersects both checkerboard surfaces of K . We say that the intersection of the ball and a checkerboard surface is the subsurface *associated* with the twist region. If the twist region has c crossings, then the intersection with one checkerboard surface is a disk, and the other has Euler characteristic $2 - c$. See Figure 26.

In the case of twisted checkerboard surfaces, we can make a similar definition. Consider a twist region of K_i . The checkerboard surfaces for K_i have subsurfaces associated with this twist region. If the twist region is encircled by a crossing circle of L_i , we choose the subsurface in R_i or B_i so that it is punctured twice by this crossing circle. However, we arrange that the subsurfaces are disjoint from all other crossing circles. Since there are inclusions $R_i \subset S_{R,i}$ and $B_i \subset S_{B,i}$, we obtain subsurfaces of $S_{R,i}$ and $S_{B,i}$ which are *the subsurfaces associated with the twist region of K_i* .

Theorem 7.1. *Suppose K is a link with prime, twist reduced, alternating diagram (which we also call K). Let K_i , $i = 0, 2$, be the diagram obtained from that of K by removing pairs of crossings from each twist region of K with more than N_{tw} crossings, where $N_{\text{tw}} \geq 26$ if $i = 0$, and $N_{\text{tw}} \geq 39$ if $i = 2$, so that the diagram of K_i has one or i remaining crossings in any such twist region. Finally, suppose that two distinct essential arcs in the surface $S_{B,i}$ have homotopic images in $S^3 \setminus K$, but are not homotopic in $S_{B,i}$. Then the two arcs are homotopic in $S_{B,i}$ into the same subsurface associated with some twist region of K_i .*

We will prove Theorem 7.1 in a sequence of lemmas. The first is an analogue of Lemmas 2.1 and 2.2.

Lemma 7.2. *Suppose homotopically distinct essential arcs a_1 and a_2 in $S_{B,i}$ map by $f: S_{B,i} \rightarrow S^3 \setminus K$ to homotopic arcs e_1 and e_2 in $S^3 \setminus K$. Then there is a map of a disk $\phi: D \rightarrow S^3 \setminus \text{int}(N(K))$ with ∂D expressed as four arcs, with opposite arcs mapping by ϕ to e_1 and e_2 , and the other two arcs mapping to $\partial N(K)$.*

Moreover, $\Gamma_B = \phi^{-1}(f(S_{B,i}))$ is a collection of embedded closed curves and arcs and an embedded graph on D whose edges have endpoints either at vertices where $\phi(D)$ meets a crossing circle, or on $\phi^{-1}(\partial N(K))$ on ∂D . Each vertex in the interior of D has valence a non-zero multiple of $2n_j$, where $2n_j$ is the number of crossings removed from the twist region at the relevant crossing circle. Each vertex in the interior of an arc in ∂D that

maps to $S_{B,i}$ has valence $n_j + 1$. Each vertex on an arc in ∂D that maps to $\partial N(K)$ has valence one.

Proof. The homotopy between e_1 and e_2 gives a map of a disk $\phi: D \rightarrow S^3 \setminus \text{int}(N(K))$, with ∂D mapped to the four arcs required by the lemma. As in Lemmas 2.1 and 2.2, we may modify ϕ in two stages, first near ∂D and then in the interior, to obtain the map with the desired properties.

The arcs e_1 and e_2 both lift to the orientable double cover of $S_{B,i}$, which is transversely orientable, and by pushing e_1 and e_2 in this transverse direction, we obtain the map ϕ in a neighborhood of these arcs. Similarly, using the fact that $\partial N(K)$ is transversely orientable, we can extend the definition of ϕ over a collar neighborhood of ∂D . Now extend ϕ over all of D , and make it transverse to all crossing circles, and transverse to $f(S_{B,i})$.

Let $\Gamma_B = \phi^{-1}(f(S_{B,i}))$ on D . Because $S_{B,i}$ is embedded in $S^3 \setminus K$ except at crossing circles, Γ_B consists of embedded closed curves, embedded arcs (edges) with endpoints corresponding to points of intersection of crossing circles (vertices), or with endpoints on $\partial N(K)$.

As in the proof of Lemmas 2.1 and 2.2, a vertex in the interior of D corresponds to a transverse intersection of $\phi(D)$ with a crossing circle in $S^3 \setminus K$. Hence the vertex has valence $2n_j$. For a vertex on an arc in ∂D that maps to $S_{B,i}$, a neighborhood of the vertex maps to half a meridian disk for a crossing circle, and so the vertex has valence $n_j + 1$. At a vertex on an arc in ∂D that maps to $\partial N(K)$, the arc in ∂D is transverse to $S_{B,i}$, and so this vertex of Γ_B has valence one. \square

We view the disk D of the previous lemma as a square with west side mapping to e_1 , east side mapping to e_2 , and north and south sides mapping to $\partial N(K)$. As before, we also have graphs Γ_{BR} and Γ_{BRG} , and complexity as in Definition 4.2, ordered lexicographically. We will take our graph to make the complexity as small as possible. All the results of Sections 4, 5, and 6 will apply to these graphs.

The following is an analogue of Lemma 2.4.

Lemma 7.3. *Let Γ be a graph in the disk $I \times I$. Suppose that Γ includes $\partial I \times I$. Suppose also that Γ contains no bigons and no monogons, and that its intersection with $I \times \partial I$ is a collection of valence one vertices. Suppose also that there are no triangular regions, with one edge on $I \times \partial I$. Finally, suppose that Γ contains at least one vertex that does not lie in $I \times \partial I$. Then either there is some vertex in the interior of $I \times I$ with valence at most 6, or there is a vertex on $\partial I \times (I - \partial I)$ with valence at most 4.*

Proof. Double the disk $I \times I$ along the two arcs $I \times \partial I$ to obtain an annulus. Then double the annulus to obtain a torus. At both stages, double the graph, and thereby obtain a graph Γ^+ in the torus. The 1-valent vertices of Γ on $I \times \partial I$ become midpoints of edges of Γ^+ . This graph Γ^+ has no monogons and no bigons, by our assumptions about Γ . It is well known that a graph in the torus with no monogons and no bigons contains a vertex with valence at most 6. The proof is analogous to the Euler characteristic calculation in Lemma 2.3. This vertex restricts to the required vertex in Γ . \square

The next lemma is analogous to Lemmas 2.5 and 2.6.

Lemma 7.4. *Let Γ_B be the graph in D provided by Lemma 7.2. Suppose that Γ_B has no monogons, and there are no trivial edges of Γ_B in ∂D . Finally, suppose that Γ_B contains at least one blue vertex, i.e. a vertex mapping to a crossing circle in $S^3 \setminus K$. Then either Γ_B has at least $(R_{\text{tw}}/6) - 2$ adjacent non-trivial bigons, or there are at least $(R_{\text{tw}}/6) - 2$ adjacent triangles, each with one arc on $\phi^{-1}(\partial N(K))$.*

Proof. We argue as in the proof of Lemmas 2.5 and 2.6.

We again need to deal with trivial arcs. Recall from Definition 4.8 that a *trivial arc* is a blue arc of Γ_B that is disjoint from the red and green edges, that has endpoints on the same crossing circle, but that does not form a loop in Γ_B . According to Lemma 4.9, each trivial arc must have at least one endpoint on ∂D . Note also that, by definition, trivial arcs must end at crossing circles, and so their endpoints do not lie on the part of ∂D that maps to $\partial N(K)$.

By Lemma 4.11, if all but one of the edges of a region of Γ_B are trivial, then the remaining edge is also trivial. Hence, if one edge of a bigon of Γ_B is trivial, then so is the other. We call this a *trivial bigon*. Any bigon that shares an edge with a trivial bigon is also trivial, and therefore trivial bigons patch together to form discs called *trivial bigon families*. If more than one trivial bigon family is incident to an interior vertex, then we consider all the trivial bigon families incident to this vertex, and call it a *trivial star*.

As in Lemmas 2.5 and 2.6, we consider the collection of trivial stars and trivial arcs with both endpoints on ∂D . The complement of these is a collection of disks. If one of these disks intersects ∂D in a single point, we pass to this subdisk, and the argument proceeds exactly as in Lemma 2.5. Otherwise, we consider a disk that is outermost in D . We can arrange that such a subdisk intersects at most one of the arcs that map to $\partial N(K)$. If the disk is disjoint from both these arcs, then the argument proceeds again as in Lemma 2.5.

So, suppose that we pass to a subdisk that contains one of the arcs that maps to $\partial N(K)$. This subdisk may have exceptional vertices, just as in the proof of Lemma 2.5. We double this disk along the arc that maps to $\partial N(K)$, to obtain a new disk. The valence one vertices on $\phi^{-1}(\partial N(K))$ become midpoints of edges. Hence, every vertex in the interior of the new disk has valence at least R_{tw} , and every unexceptional vertex on the boundary has valence at least $(R_{\text{tw}}/2) + 1$. There are two collections of exceptional vertices on the boundary, and we collapse these to two exceptional vertices. The proof of Lemma 2.5 and 2.6 still applies, since when Lemma 2.4 is used, it produces a vertex in the interior of the disk with valence at most 5 or at least three vertices on the boundary with valence at most 3, and in the latter case, we therefore deduce that at least one of these vertices is unexceptional.

So suppose now that D contains no trivial arcs with both endpoints on ∂D , and no trivial stars. Then we do not pass to a subdisk of D . Instead, we apply the procedure given in the proof of Lemma 2.5, where trivial bigon families are removed, and replaced by vertices on ∂D , producing a graph Γ' . Every interior vertex of Γ' has valence at least R_{tw} and every boundary vertex disjoint from $\phi^{-1}(\partial N(K))$ has valence at least $(R_{\text{tw}}/2) + 1$.

We now suppose that the lemma is not true, and that Γ has no collection of $(R_{\text{tw}}/6) - 2$ adjacent non-trivial bigons and no collection of $(R_{\text{tw}}/6) - 2$ adjacent triangles. Then, as in the proof of Lemma 2.5, Γ' has no collection of $R_{\text{tw}}/6$ adjacent bigons and no collection of $R_{\text{tw}}/6$ adjacent triangles. Hence, every interior vertex of Γ' valence more than 6. Also,

every boundary vertex of Γ' disjoint from $\phi^{-1}(\partial N(K))$ has valence more than 4. This contradicts Lemma 7.3. \square

Lemma 7.5. *If $i = 0$ and $N_{\text{tw}} \geq 26$, then the graph Γ_B of Lemma 7.2 contains no blue vertices. Similarly, if $i = 2$ and $N_{\text{tw}} \geq 39$, then the graph Γ_B contains no blue vertices. That is, $\phi(D)$ meets no crossing circles.*

Proof. Lemma 4.7 implies that Γ_B contains no monogons. Lemma 4.10 implies that it contains no trivial arc that is a subset of ∂D . So if Γ_B contains a blue vertex, then Lemma 7.4 applies. If $N_{\text{tw}} \geq 26$ and $i = 0$, then by equation (1), $R_{\text{tw}} \geq 26$, so Γ_B must contain more than two adjacent non-trivial bigons or adjacent triangles. Similarly, if $i = 2$ and $N_{\text{tw}} \geq 39$, then $R_{\text{tw}} \geq 38$, and Γ_B must contain more than four adjacent non-trivial bigons or adjacent triangles. This contradicts either Proposition 5.8 or 5.9, or Lemma 6.3. Hence there can be no blue vertices on Γ_B . \square

As before, consider the red surface R_i embedded in $S^3 \setminus L_{B,i}$, and the graph Γ_{BR} , with Γ_B coming from Lemma 7.2.

Lemma 7.6. *The graph Γ_B consists only of arcs whose two endpoints are on north and south sides of ∂D . The red edges of the graph Γ_{BR} consist of arcs with endpoints on distinct sides of D (north, south, east, west).*

Proof. By Lemma 7.5, we may assume that Γ_{BR} consists of red and blue arcs and simple closed curves, but no blue vertices. By Lemma 4.6, we may assume there are no blue simple closed curves. Because there are no blue vertices, no blue edges can have an endpoint on the east or west sides (mapping to e_1 and e_2). By Theorem 1.6, any blue arc with both endpoints on $\partial N(K)$ on the north (or both on the south) is trivial in the blue surface, and so we may replace D' with a disk that does not meet that blue arc. This will reduce the number of vertices of Γ_B , contradicting our assumption that the graph has minimum complexity. Thus blue arcs run from north to south.

As for the red, a red simple closed curve disjoint from blue bounds a disk on red, so D' could be replaced by a disk that does not meet this red curve. A red simple closed curve that is not disjoint from the blue would imply the existence of a red–blue bigon, contradicting Lemma 4.5. Hence we may assume there are no red simple closed curves, whether or not they meet blue. Because there are no red–blue bigons by Lemma 4.5, no red edge of Γ can have both endpoints on the east or west side of D . Finally, the red surface R_i is a checkerboard surface for K_i , hence is boundary incompressible. Thus any red arc with both endpoints on $\partial N(K)$ on the north (or both on the south) is trivial in the red surface, and so we may replace D' with a disk that does not meet that red arc, without increasing the number of vertices of Γ_B . This move contradicts the assumption that complexity is as small as possible. The result follows. \square

Lemma 7.7. *The graph Γ_{BR} consists of red and blue arcs with endpoints on opposite sides of D (north–south, or east–west).*

Proof. By the preceding lemma, blue arcs run north to south, as desired. Hence we need to show there are no red arcs running from north to east, north to west, south to east, or south to west. Any such red arc would have an endpoint on the east or west side mapping



FIGURE 27. A triangle with one edge on red, one on blue, and one on the knot

to the arc e_1 or e_2 . It may meet other blue arcs, running north to south, but in any case, the arc cuts off a triangle with one side on blue, one side on red, and one side on $\partial N(K)$. By considering the region of D where the arc meets the north or south side, we may take such a triangle to have interior disjoint from all other red and blue edges.

Because the triangle is disjoint from all crossing circles, we may sketch its image into the diagram of $L_{B,i}$, assuming without loss of generality that the triangle maps into the region above the plane of projection. The blue arc has one endpoint on a strand of the link and one endpoint at a crossing. The red has one endpoint on the same strand of the link (i.e. the portion of K_i running between two adjacent under-crossings), and the other on the same crossing. There are two ways that a red and a blue endpoint can meet the same strand of the link between under-crossings. One way is if they are on the same side of that strand, but with that strand running over a crossing between them, as in Figure 27, left. The other is if they are on opposite sides of a strand that does not run over a crossing, as in Figure 27, right.

Consider the left of Figure 27. By connecting the endpoints of arcs on the link, and pushing the triangle into the plane of projection at that point, we obtain a red–blue bigon. Arguing as in the proof of Lemma 4.5, we conclude there is a contradiction to primality. Thus if there is a triangle, it must lie in the diagram as on the right of Figure 27.

As for the figure on the right, we may connect the endpoints of the arcs to form a closed curve γ meeting the diagram twice. We will use Lemma 3.1 to show there are no crossings in the interior of γ . If there are such crossings, then the red arc of γ must run through a crossing circle, and since the triangular region bounds a disk disjoint from the crossing circle, the blue arc of γ must also run through the crossing disk. Then that crossing disk splits γ into two closed curves, which can be pushed to meet the diagram twice, meeting one fewer crossing circle. By induction on the number of crossing circles met by such a curve, we conclude there are no crossings of $K_{B,i}$ in the interior of γ .

Thus the red and blue arcs in that figure are in fact parallel to a single strand of the link. Hence they are both homotopic to a portion of the arc running from the top of the crossing shown in that figure to its base. Use this homotopy to slide the image of D to this crossing arc, removing the intersection of the red and blue arcs and removing the triangle. The result has one fewer vertex of Γ_{BR} , and does not affect the number of vertices of Γ_B , contradicting our assumption that complexity is as small as possible. \square

Since red arcs cannot intersect red arcs, Lemma 7.7 implies that either all red arcs run from north to south, or all red arcs run from east to west.

Suppose all blue arcs and all red arcs run from north to south. Because there are no red–blue bigons by Lemma 4.5, all such edges are disjoint. Then either all edges are blue,




FIGURE 28. Possible images of red–blue rectangles

and there are no red edges of intersection at all, or there is a rectangle with one blue side and one red side, with north and south edges on K , and with interior disjoint from the red and blue surfaces. The next lemma deals with the latter case.

Lemma 7.8. *If the graph Γ_{BR} cuts D into a subrectangle with two opposite sides mapped to $\partial N(K)$, one side on blue, one side on red, and interior disjoint from blue and red, then the blue and red sides of that rectangle are homotopic to the same crossing arc, and the homotopies can be taken to lie entirely in the blue and red surfaces, $S_{B,i}$ and R_i , respectively.*

Recall that a *crossing arc* is an arc in the link complement that runs straight from the top of a crossing to the bottom in the diagram.

Proof. Consider such a rectangle. Abuse notation slightly and give the blue arc on the west side the label e_1 , and the red arc on the east the label e_2 , and call the rectangle D . Because the interior of D does not meet blue or red surfaces, it can be mapped into the complement of $N(K_{B,i})$, missing checkerboard surfaces, hence it must be mapped completely above or below the plane of projection of $K_{B,i}$. Without loss of generality, assume it is mapped above. Then the arcs of ∂D on $\partial N(K_{B,i})$ each lie on a single strand of the diagram, i.e. on a strand running between two undercrossings. Endpoints of e_1 and e_2 either straddle an overcrossing or lie on either side of a strand.

There are three cases to consider: first, endpoints of e_1 and e_2 straddle overcrossings at both ends; second, one set of endpoints straddles an overcrossing and one set lies on opposite sides of a strand; and third, both sets of endpoints lie on opposite sides of strands. The cases are shown in Figure 28.

In the first case, we may connect red and blue arcs to form a closed curve meeting the diagram twice at crossings. As in the proof of Lemma 4.5, this gives a contradiction.

In the second and third case, we connect red and blue arcs to form a closed curve γ meeting the diagram twice. As in the proof of the previous lemma, we may use Lemma 3.1 to argue that there are no crossings of $K_{B,i}$ on one side of γ . In the third case, blue and red edges are parallel and not essential, and we can modify D by homotoping off this region of the diagram, removing these intersections with red and blue surfaces, and reducing complexity, contradicting our minimality assumption. In the second case, e_1 and e_2 are homotopic to the same crossing arc, as desired. Notice that the homotopies taking these arcs to the crossing arc lie entirely in the blue or red surface, respectively. \square

Lemma 7.9. *If the graph Γ_{BR} consists of disjoint red and blue arcs on D , all running north to south, then e_1 and e_2 are each homotopic in the blue surface to arcs in the same subsurface associated with a twist region of K_i .*

Proof. We will show that in this case, each arc of Γ_{BR} is homotopic, in the surface $S_{B,i}$ or R_i containing it, into a subsurface corresponding to a twist region of K_i . We will also show that successive arcs lie in subsurfaces corresponding to the same twist region. Now, two arcs, both lying in $S_{B,i}$ or both lying in R_i , and lying in subsurfaces corresponding to distinct twist regions, cannot be homotopic in those surfaces. Hence, we will deduce that e_1 and e_2 are homotopic in $S_{B,i}$ to the subsurface corresponding to the same twist region.

In the case where successive arcs are red and blue, Lemma 7.8, applied to the part of D lying between these arcs, gives this claim. Thus we need to show that when there are two adjacent blue arcs of Γ_{BR} , or two adjacent red arcs, the result still holds.

In the blue case, we have a rectangle with two sides on blue, two sides on $\partial N(K_{B,i})$, and interior disjoint from red and blue. It must be mapped by ϕ entirely to one side of the projection plane. Its edges on $N(K_{B,i})$ must run over crossings, else we could homotope D to remove both intersections, contradicting our requirement that Γ_{BR} be minimal. Hence it defines a simple closed curve γ in the diagram meeting the knot in exactly two crossings. Because $K_{B,i}$ is blue twist reduced by Lemma 3.4, the curve γ encircles a collection of red bigons. Hence we may isotope the blue arcs on $S_{B,i}$, relative to their endpoints on $\partial N(K_{B,i})$, to lie in a neighborhood of the twist region of $K_{B,i}$ containing those red bigons. Note this is a twist region of K_i as well, since the blue surfaces lie on the outside of the twist region, so the result follows in this case.

In case that there are two adjacent red arcs in D , the argument is similar. Again the rectangle must be mapped entirely to one side of the projection plane of $K_{B,i}$, and it defines a simple closed curve γ in the diagram meeting the knot in exactly two crossings, with arcs γ_1 and γ_2 in the red surface running between these crossings. The diagram of $K_{B,i}$ may not be red twist reduced, which means γ is not required to bound blue bigons on one side, a priori. However, by induction we may assume that one of the red arcs, say γ_1 , is homotopic to a crossing arc in a twist region of K .

The arc γ_1 has endpoints on overstrands of distinct crossings. Since γ_1 is homotopic in R_i to a crossing arc of a single crossing, one of the overstrands must run directly to the understrand of the other crossing, and the arc must be homotopic to that strand of the knot. Then we may slide γ in the diagram of $K_{B,i}$ to the opposite side of this strand, forming a closed curve α consisting of γ_2 and an arc parallel to the knot strand parallel to γ_1 , and meeting the diagram twice. This closed curve α has zero linking number with each of the crossing circles of $L_{B,i}$, because other γ would link with a crossing circle, and D would contain a blue vertex corresponding to this crossing circle, which is a contradiction.

Moreover, α must be disjoint from all the crossing disks, for if it were to intersect a crossing disk, it would do so twice, once in the blue surface and once in the red. Then γ would also intersect this crossing disk twice, once in γ_1 , and once in γ_2 . We may then form a simple closed curve in the diagram for $L_{B,i}$ that starts at a crossing where γ_1 and γ_2 meet, runs along γ_1 as far as the crossing disk, runs along the crossing disk to γ_2 and then back along γ_2 to the original crossing. By choosing the crossing disk appropriately, we may ensure that this curve intersects no other crossing disks. After a small isotopy, it can then be made disjoint from the crossing disks, and so it then specifies a simple closed curve in the diagram for K which is disjoint from K except at a single crossing. This implies that K was not prime, which is a contradiction.

Hence, α corresponds to a simple closed curve in the diagram of K that hits K twice. By the primality of K , it bounds a region of the diagram with no crossings. Therefore, γ_2 is homotopic in R_i to a crossing arc, as required. \square

It follows from Lemma 7.9 that if D meets the red surface in vertical arcs, then Theorem 7.1 holds. So we assume that D does not meet the red surface in vertical arcs. Then the blue surface cuts D into rectangles with north and south sides on K and east and west sides on $S_{B,i}$, and interiors disjoint from blue. Because each rectangle is disjoint from all vertices (crossing circles), it can be embedded in $S^3 \setminus K_{B,i}$. The embedding will put east and west sides of the boundary of the rectangle on the blue checkerboard surface of $K_{B,i}$, north and south sides on the link $\partial N(K_{B,i})$, and will map the interior into the complement of the blue checkerboard surface in $S^3 \setminus \text{int}(N(K_{B,i}))$. A rectangle embedded in $S^3 \setminus \text{int}(N(K_{B,i}))$ in this way is a well-known object: it is an *essential product disk*.

Definition 7.10. An *essential product disk* for the blue checkerboard surface B of a knot K is an essential disk properly embedded in $S^3 \setminus \text{int}(N(B))$, whose boundary is a rectangle with two opposite sides on $N(B)$ and two opposite sides on $\partial N(K)$.

Essential product disks have been studied in many other contexts (for example, to identify the guts of a manifold [5, 6, 4]).

To conclude the proof of Theorem 7.1, we will consider essential product disks for $K_{B,i}$, and show the boundary of such a disk gives two arcs in the neighborhood of a twist region in the diagram of K_i .

Lemma 7.11. *Let e_1 and e_2 denote the boundary arcs on the blue surface in an essential product disk for the blue checkerboard surface of $K_{B,i}$. Then there is a subsurface associated with a twist region of the diagram of K_i , and arcs a_1 and a_2 in that subsurface, such that e_1 is homotopic in $S_{B,i}$ to a_1 , and e_2 is homotopic in $S_{B,i}$ to a_2 .*

Proof. Let E denote the essential product disk. If E is disjoint from the red checkerboard surface for $K_{B,i}$, then an argument as in the proof of Lemma 7.9, using the fact that the diagram of $K_{B,i}$ is blue twist reduced, implies the blue edges of E are homotopic to arcs in a subsurface associated to a single twist region of K_i .

So suppose the essential product disk E meets the red surface. By Lemma 7.7, we may assume intersections with the red surface do not run from a blue edge to $N(K_{B,i})$. By Lemma 7.8, if intersections with red have both endpoints on $N(K_{B,i})$, then the blue edges of E are both homotopic to the same crossing arc, hence can be homotoped to lie in the same twist region of the diagram of K_i . Hence we assume the red surface meets E in a sequence of horizontal arcs, cutting it into rectangles.

It is well known that the checkerboard surfaces of a connected alternating link diagram cut the complement into two identical 4-valent ideal polyhedra. These two polyhedra have edges corresponding to edges of the diagram of the knot, and ideal vertices corresponding to vertices of the knot. The knot complement is obtained by gluing the same red faces of the two polyhedra by a single clockwise rotation. Blue faces are glued by a counter clockwise rotation. See, for example, [10] or [6].

Our sequence of rectangles making up E has boundary components which lie in the checkerboard surfaces of $K_{B,i}$. The sequence alternates lying in one polyhedron and then

the other, but their boundaries can be sketched into the diagram graph of the knot $K_{B,i}$, using the identification of edges and vertices of the polyhedron with edges and vertices of the diagram.

Consider the rectangle at the north of E , denote it by E_1 , and the rectangle glued just under it, E_2 . The rectangle E_1 at the north has one side running along $N(K_{B,i})$, which means it has a side cutting through an ideal vertex of a polyhedron. Push this off the ideal vertex slightly, to cut off a single vertex of a red face. Now consider its side of E_1 in the other red face. This is glued by a clockwise turn to a side of E_2 . Impose the image of E_2 under this clockwise turn on the polyhedron containing E_1 , and denote it by \bar{E}_2 . By [6, Lemma 7] (see also [4, Lemma 4.9]), if these sides of E_1 and \bar{E}_2 intersect in this red face, then they must intersect in two red faces. But the side of E_1 cuts off a single vertex in its other red face, so it can be homotoped such that it does not intersect another side of \bar{E}_2 . Hence E_1 and \bar{E}_2 , and thus E_2 , must each cut off a single vertex in the red face in which they are glued. The same argument shows that E_2 and E_3 also both cut off a single vertex in the red face in which they are glued. By induction, each rectangle making up E has sides in red faces cutting off a single ideal vertex.

These rectangles map to regions of the diagram meeting the diagram exactly four times, adjacent to two crossings. We can push the sides in the red faces onto these crossings. Because the diagram is blue twist reduced by Lemma 3.4, each such rectangle bounds a (possibly empty) string of red bigons. Hence their boundaries all lie in a neighborhood of the same twist region of the diagram of $K_{B,i}$. Note in this case, the twist region must also be a twist region of K_i , since the blue surface is outside the twist region. \square

We can put this together to finish proof of Theorem 7.1.

Proof of Theorem 7.1. Lemma 7.2 implies that two edges e_1 and e_2 that are homotopic in $S^3 \setminus K$ give rise to a mapping of a disk $\phi: D \rightarrow S^3 \setminus K$ and a graph Γ_B . Provided $N_{\text{tw}} \geq 26$ if $i = 0$, and $N_{\text{tw}} \geq 39$ if $i = 2$, the graph Γ_B contains no blue vertices, by Lemma 7.5. By Lemmas 7.6 and 7.7, blue edges of Γ_{BR} run north to south, and red edges either all run north to south or all run east to west. If red run north to south, then Lemma 7.9 implies that e_1 and e_2 are homotopic to arcs in the same subsurface associated with a twist region. If red run east to west, then Lemma 7.11 implies they are homotopic to arcs in the same subsurface associated with a twist region of K_i . \square

We finish with a result for regular checkerboard surfaces that follows almost immediately from the previous work.

Proposition 7.12. *Let S denote the disjoint union of the two checkerboard surfaces of a prime, twist reduced alternating diagram of a hyperbolic knot K . Suppose a_1 and a_2 are disjoint essential embedded arcs in S that are not homotopic in S , but are homotopic in $S^3 \setminus K$ after including S into $S^3 \setminus K$. Then either a_1 and a_2 are isotopic in S to crossing arcs in the same twist region of the diagram, or they both lie on the same checkerboard surface and both are isotopic in that checkerboard surface to arcs in the same subsurface associated with some twist region of K .*

Proof. Given a prime, twist reduced alternating diagram K , let C denote the maximal number of crossings in any twist region of K . Let N_{tw} be the maximum of 39 and C . Then

the diagram K_2 obtained from that of K by removing pairs of crossings from each twist region of K with more than N_{tw} crossings is identical to the diagram of K . Moreover, the diagrams of K_i and $K_{B,i}$ are also identical, and the surfaces R_i and $S_{B,i}$ are identical to the checkerboard surfaces of K . We therefore apply the above results to these knots and these surfaces. In particular, Theorem 7.1 implies that if e_1 and e_2 lie on the same surface, either red or blue, then the two arcs are homotopic in that surface into the same subsurface associated with a twist region. Thus we only need to finish the case that e_1 and e_2 lie on different surfaces, say e_1 on blue and e_2 on red.

First, modify Lemma 7.2 in a straightforward way to allow these two arcs to lie on distinct surfaces. We obtain a graph Γ_{BR} , which we assume, as usual, has minimal complexity. There will be no vertices at all in the graph Γ_{BR} , since there are no crossing circles added to adjust the diagram of K , so Lemma 7.5 trivially holds. We may argue just as in Lemma 7.6 that no red or blue arc has both endpoints on the same side of D (north, south, east, or west). By considering the same triangles of Lemma 7.7, or those with red and blue switched, we may argue that red and blue arcs either run north to south, or east to west, just as in the conclusion of that lemma. However, because we have two distinct colors on east and west, and because the red and blue surfaces are embedded, no arc may run east to west. Thus all arcs run north to south, and they cut D into rectangles with two sides on $\partial N(K)$, and two sides on red or blue, and interior disjoint from these surfaces.

If there are only the two original edges, and Γ_{BR} is disjoint from the red and blue surfaces otherwise, then we are done by Lemma 7.8. Similarly, if the edges of Γ_{BR} alternate red and blue, then each subrectangle gives a homotopy to the same crossing arc. Since e_1 and e_2 are on different surfaces, there must be at least one subrectangle with sides on different surfaces, so Lemma 7.8 implies that the two arcs of this subrectangle are homotopic to a crossing arc, with the homotopy taken within the respective surface.

Now suppose there is a subrectangle with both sides on the same (red or blue) surface. By induction, we may assume that one of the arcs is homotopic to a crossing arc in its surface. Arguing as in Lemma 7.9, the two arcs on either side of the rectangle will be homotopic to arcs which together encircle a twist region. Since one of those arcs is homotopic to a crossing arc, that arc runs from one side of a single crossing of that twist region to the other. Hence both arcs encircle a trivial twist region, consisting of one crossing, and the other arc is also homotopic to the same crossing arc. Putting this all together implies the result. \square

REFERENCES

1. Colin C. Adams, *Augmented alternating link complements are hyperbolic*, Low-dimensional topology and Kleinian groups (Coventry/Durham, 1984), London Math. Soc. Lecture Note Ser., vol. 112, Cambridge Univ. Press, Cambridge, 1986, pp. 115–130. MR 903861 (89f:57003)
2. Iain R. Aitchison, E. Lumsden, and J. Hyam Rubinstein, *Cusp structures of alternating links*, Invent. Math. **109** (1992), no. 3, 473–494. MR MR1176199 (93h:57007)
3. Oliver T. Dasbach and Xiao-Song Lin, *A volumish theorem for the Jones polynomial of alternating knots*, Pacific J. Math. **231** (2007), no. 2, 279–291. MR 2346497 (2008h:57009)
4. David Futer, Efstratia Kalfagianni, and Jessica S. Purcell, *Guts of surfaces and the colored Jones polynomial*, Lecture Notes in Mathematics, vol. 2069, Springer, 2013.
5. David Gabai and William H. Kazez, *Group negative curvature for 3-manifolds with genuine laminations*, Geom. Topol. **2** (1998), 65–77 (electronic). MR 1619168 (99e:57023)

6. Marc Lackenby, *The volume of hyperbolic alternating link complements*, Proc. London Math. Soc. (3) **88** (2004), no. 1, 204–224, With an appendix by Ian Agol and Dylan Thurston. MR MR2018964 (2004i:57008)
7. Marc Lackenby and Jessica S Purcell, *Cusp volumes of alternating knots*, In preparation.
8. W. Menasco, *Closed incompressible surfaces in alternating knot and link complements*, Topology **23** (1984), no. 1, 37–44. MR 721450 (86b:57004)
9. William Menasco and Morwen Thistlethwaite, *The classification of alternating links*, Ann. of Math. (2) **138** (1993), no. 1, 113–171. MR 1230928 (95g:57015)
10. William W. Menasco, *Polyhedra representation of link complements*, Low-dimensional topology (San Francisco, Calif., 1981), Contemp. Math., vol. 20, Amer. Math. Soc., Providence, RI, 1983, pp. 305–325. MR MR718149 (85e:57006)

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, OXFORD, OX2 6GG, UK

DEPARTMENT OF MATHEMATICS, BRIGHAM YOUNG UNIVERSITY, PROVO, UT 84602, USA