Cubic graphs induced by bridge trisections

JEFFREY MEIER, ABIGAIL THOMPSON, AND ALEXANDER ZUPAN

Every embedded surface $\mathcal{K}$ in the 4-sphere admits a bridge trisection, a decomposition of $(S^4, \mathcal{K})$ into three simple pieces. In this case, the surface $\mathcal{K}$ is determined by an embedded 1-complex, called the *1-skeleton* of the bridge trisection. As an abstract graph, the 1-skeleton is a cubic graph Γ that inherits a natural Tait coloring, a 3-coloring of the edge set of Γ such that each vertex is incident to edges of all three colors. In this paper, we reverse this association: We prove that every Tait-colored cubic graph is isomorphic to the 1-skeleton of a bridge trisection corresponding to an unknotted surface. When the surface is nonorientable, we show that such an embedding exists for every possible normal Euler number. As a corollary, every tri-plane diagram for a knotted surface can be converted to a tri-plane diagram for an unknotted surface via crossing changes and interior Reidemeister moves. Tools used to prove the main theorem include two new operations on bridge trisections, crosscap summation and tubing, which may be of independent interest.

1. Introduction

A graph Γ is *cubic* if each of its vertices has valence three. A *Tait coloring* of a cubic graph is a function $\mathcal{C}$ from the edge set of Γ to the set {red, blue, green} such that each vertex is incident to one edge of each color. Bridge trisections of knotted surfaces in S^4 were defined by the first and third authors in [8] and extended to knotted surfaces in arbitrary 4-manifolds in [9]. A *bridge trisection* $\mathcal{T}$ of a knotted surface $\mathcal{K} \subset S^4$ is a decomposition

$$(S^4, \mathcal{K}) = (X_1, \mathcal{D}_1) \cup (X_2, \mathcal{D}_2) \cup (X_3, \mathcal{D}_3),$$

where, for each $i \in \{1, 2, 3\}$, $\mathcal{D}_i$ is a collection of trivial disks in the 4-ball X_i, and the pairwise intersection $\tau_{ij} = \mathcal{D}_i \cap \mathcal{D}_j$ is a trivial tangle in the 3-ball $B_{ij} = X_i \cap X_j$. It follows that the triple intersection $\mathcal{D}_1 \cap \mathcal{D}_2 \cap \mathcal{D}_3$ is a

collection $\mathbf{x}$ of *bridge points* in the *bridge sphere* $\Sigma = X_1 \cap X_2 \cap X_3$, where Σ is a 2-sphere.

The union $\Gamma = \tau_{12} \cup \tau_{23} \cup \tau_{31}$ along the points $\mathbf{x}$ is a 1-complex (that is, a graph), which we will call the *1-skeleton* associated to $\mathcal{T}$. Observe that Γ is cubic, and it has a natural Tait coloring $\mathcal{C}$ obtained by coloring the arcs of τ_{12} red, the arcs of τ_{23} blue, and the arcs of τ_{31} green. We say that the coloring $\mathcal{C}$ of Γ is *induced* by $\mathcal{T}$.

In this paper, we prove that the correspondence between bridge trisections and Tait-colored cubic graphs can be reversed: Given a cubic graph Γ with a Tait coloring $\mathcal{C}$, the subgraph induced by any pair of colors is a collection of disjoint cycles, and attaching 2-cells along every cycle for each of the three pairings gives rise to a surface S, which we call the surface *induced by $\mathcal{C}$.*

Theorem 1.1. *If Γ is a cubic graph with a Tait coloring $\mathcal{C}$, then there exists a bridge trisection $\mathcal{T}$ of an unknotted surface $\mathcal{U} \subset S^4$ such that the 1-skeleton $\mathcal{T}$ is graph isomorphic to Γ, with the coloring $\mathcal{C}$ induced by $\mathcal{T}$. Moreover, if the induced surface S is nonorientable, we may choose the embedding of $\mathcal{U}$ to have any possible normal Euler number.*

To prove the main theorem, we first prove that every Tait-colored cubic graph Γ admits a finite sequence of simplifications called *compressions* yielding the theta graph. The theta graph is the 1-skeleton of the 1-bridge trisection of the unknotted 2-sphere. Next, we show that there is a sequence of modifications to the 1-bridge trisection of the unknotted 2-sphere that undoes the sequence of compressions, eventually resulting in a bridge trisection of an unknotted surface inducing Γ. One of the moves, elementary perturbation, is well-known (see [8]), while the other two, crosscap summation and tubing, are new constructions that may be of independent interest.

To get a sense of the difficulty of the problem, the motivated reader is encouraged to attempt their own ad hoc construction of a bridge trisection with 1-skeleton isomorphic to either of the two examples shown in Figure 1. The graph in Figure 1a is the famous Heawood graph [3]; it induces an orientable surface, so we call it Γ_O. The graph in Figure 1b appears to be unnamed; it induces a nonorientable surface, so we call it Γ_N. We will refer back to these examples throughout the paper. In particular, the proof of the main theorem is constructive and is carried out for these two examples in Section 5.

A bridge trisection $\mathcal{T}$ induces additional structure: By choosing disks $E_{ij} \subset H_{ij}$ with a common boundary curve containing the bridge points $\mathbf{x}$,
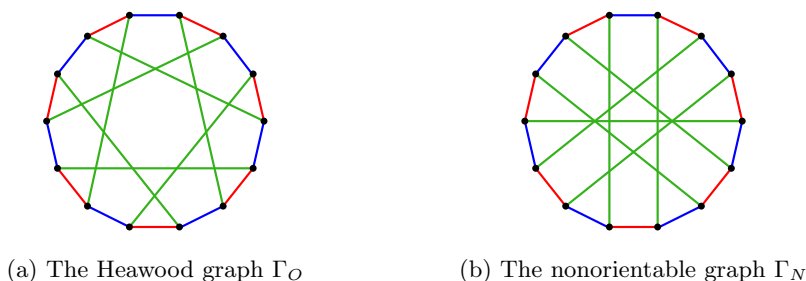
(a) The Heawood graph Γ_O (b) The nonorientable graph Γ_N

Figure 1. Two examples of Tait-colored cubic graphs.

we can project τ_{ij} onto E_{ij} to obtain a *tri-plane diagram* $\mathcal{P} = (\mathcal{P}_{12}, \mathcal{P}_{23}, \mathcal{P}_{31})$, a triple of planar diagrams of trivial tangles such that any pairwise union $\mathcal{P}_{ij} \cup \overline{\mathcal{P}}_{jk}$ yields a classical diagram for an unlink. Theorem 1.7 from [8] asserts that any two tri-plane diagrams corresponding to the same bridge trisection $\mathcal{T}$ are related by interior Reidemeister moves and mutual braid transpositions.

For tri-plane diagrams, we can leverage Theorem 1.1 to obtain the following corollary.

Corollary 1.2. *Every tri-plane diagram $\mathcal{P}$ of a knotted surface $\mathcal{K} \subset S^4$ can be converted to a tri-plane diagram $\mathcal{P}'$ for an unknotted surface $\mathcal{U}$ (homeomorphic to $\mathcal{K}$) by a sequence of interior Reidemeister moves and crossing changes.*

It is natural to wonder if Corollary 1.2 remains true when Reidemeister moves are disallowed when converting the given diagram $\mathcal{P}$ to the diagram $\mathcal{P}'$ corresponding to the unknotted surface; see Remark 4.3 following the proof of Corollary 1.2.

Question 1.3. *Does every tri-plane diagram $\mathcal{P}$ admit a sequence of crossing changes converting it to a tri-plane diagram $\mathcal{P}'$ for an unknotted surface?*

We proceed as follows: In Section 2, we set up some background material related to bridge trisections, knotted surfaces, and cubic Tait-colored graphs. In Section 3, we discuss connected summation, elementary perturbation, crosscap summation, and tubing of bridge trisections, and we define compression of Tait-colored cubic graphs. In Section 4, we prove Theorem 1.1 and Corollary 1.2. Finally, in Section 5, we carry out the process described in the proof of Theorem 1.1 for the examples in Figure 1.

Acknowledgements

2. Preliminaries

2.1. Bridge trisections

In [8], the first and third author proved that every knotted surface $(S^4, \mathcal{K})$ admits a bridge trisection $\mathcal{T}$, which can be encoded by a *shadow diagram*, a triple (a, b, c) such that each of a, b, and c is an embedded collection of pairwise disjoint arcs in Σ resulting from pushing the trivial arcs in each pairwise intersection (B_{ij}, τ_{ij}) into Σ. For two simple examples of shadow diagrams, see Figure 2. If two shadow diagrams are related by a sequence of shadow slides (replacing an arc with its band sum with boundary of a neighborhood of another shadow in the same set) and isotopies fixing the bridge points, then they correspond to the same bridge trisection $\mathcal{T}$. See Figure 16 for some examples of shadow slides.
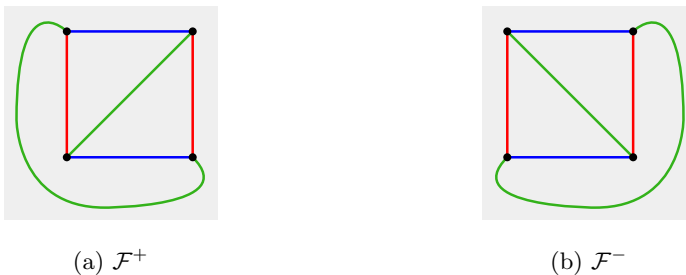

(a) $\mathcal{F}^+$ (b) $\mathcal{F}^-$

Figure 2. Examples of genus zero shadow diagrams.

2.2. Unknotted surfaces

We say that an orientable surface $\mathcal{K} \subset S^4$ is *unknotted* if $\mathcal{K}$ is the boundary of a smoothly embedded 3-dimensional handlebody. Equivalently, $\mathcal{K}$ is unknotted if and only if $\mathcal{K}$ is isotopic into $S^3 \subset S^4$ [4, Theorem 1.2]. (This is

analogous to the fact that a classical knot $K \subset S^3$ is the unknot if and only if K is isotopic into $S^2 \subset S^3$.) For a nonorientable surface $\mathcal{K}$, the situation is slightly more complicated: Each shadow diagram in Figure 2 corresponds to a bridge trisection of a embedding of $\mathbb{RP}^2$ into S^4. We call these two embeddings $\mathcal{F}^+$ and $\mathcal{F}^-$ as shown; they are the two unknotted embeddings of $\mathbb{RP}^2$, where the normal Euler number $e(\mathcal{F}^\pm)$ satisfies $e(\mathcal{F}^\pm) = \pm 2$. Following [4], we say that a nonorientable surface $\mathcal{K} \subset S^4$ is *unknotted* if $\mathcal{K}$ is isotopic to a connected sum of copies of $\mathcal{F}^+$ and $\mathcal{F}^-$. Since normal Euler number is additive under connected sum, for unknotted surfaces with nonorientable genus g, we have $e(\mathcal{K}) \in \{-2g, -2g+4, \ldots 2g-4, 2g\}$. The Whitney-Massey Theorem asserts that the normal Euler number of any embedded surface of nonorientable genus g also falls into this range [7].

2.3. Cubic graphs and surfaces

All cubic graphs in this paper are assumed to be connected. Throughout the paper, graphs drawn on a white background are abstract graphs. The graphs in Figures 3b and 5 are embedded in the torus and projective plane, respectively. All other graphs drawn on a gray, shaded background are shadow diagrams, with the grey shading representing the bridge sphere. We allow our cubic graphs to have parallel edges. If Γ is not the theta graph – i.e., the graph with two vertices and three parallel edges – then each edge is parallel to at most one other edge since Γ is cubic and connected.

Given a cubic graph Γ with a Tait coloring $\mathcal{C}$, recall that Γ and $\mathcal{C}$ give rise to the *induced surface* S obtained by attaching 2-cells to Γ along each bi-colored cycle determined by $\mathcal{C}$. See Figure 3 below for an example. Note that the induced surface is an abstract surface, abstractly trisected by Γ, not a bridge trisected surface in S^4.

The *patch numbers* (p_1, p_2, p_3) of a Tait-colored cubic graph Γ count the number of each type of bi-colored cycle. If $p_1 = p_2 = p_3 = p$, then we simply say that $\mathcal{T}$ is *p-patch*. Both examples Γ_O and Γ_N shown in Figure 1 are 1-patch. We also keep track of the orientability of the induced surface S, which can be verified by an easy condition, offered by the next lemma.

Lemma 2.1. *Given a Tait-colored cubic graph Γ, the induced surface S is orientable if and only if Γ is bipartite.*

Proof. First, suppose that S is oriented. The orientation of S induces an orientation of each 2-cell bounded by the bi-colored cycles of Γ with respect to $\mathcal{C}$. Orient the red edges to agree with the orientation of the 2-cells bounded
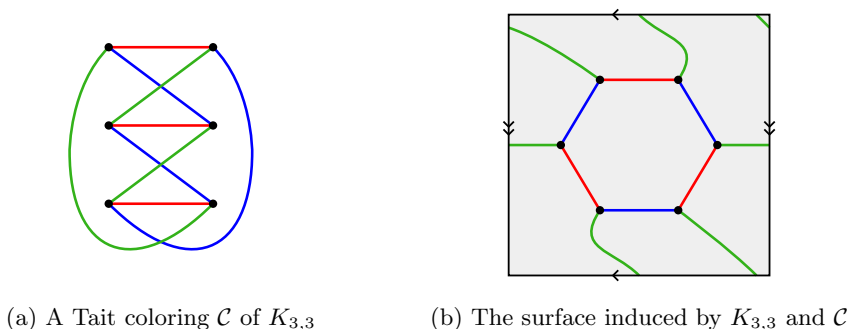
(a) A Tait coloring $\mathcal{C}$ of $K_{3,3}$ (b) The surface induced by $K_{3,3}$ and $\mathcal{C}$

Figure 3. An example of an induced surface.

by red-blue cycles, orient the blue edges to agree with the orientation of the blue-green 2-cells, and orient the green edges to agree with the orientation of the green-red 2-cells. Then the orientations of the blue edges, green edges, and red edges *disagree* with the orientations of the red-blue 2-cells, blue-green 2-cells, and green-red 2-cells, respectively. It follows that every vertex v in Γ is either at the head of an edge of each color, or v is at the tail of an edge of each color. Letting V^+ denote the vertices at the heads of a triple of edges and V^- the vertices at the tails, we have that Γ is bipartite.

Conversely, suppose that Γ is bipartite, with vertices partitioned into V^+ and V^-. Orient the edges of Γ so that each has a vertex in V^- at its tail and a vertex in V^+ at its head. Finally, orient the red-blue 2-cells so that they agree with the orientations of the red edges, orient the blue-green 2-cells to agree with the orientations of the blue edges, and orient the green-red 2-cells to agree with the orientations of the green edges. Then (as above) the orientations of blue edges, green edges, and red edges disagree with the orientations of the red-blue 2-cells, the blue-green 2-cells, and the green-red 2-cells, respectively. This implies that 2-cells are glued along oppositely oriented edges, so that the orientations of the 2-cells agree wherever they overlap. We conclude that their union, the surface S, is orientable. □

Following Lemma 2.1, we will say a Tait-colored cubic graph Γ is *orientable* if Γ is bipartite and *nonorientable* otherwise.

Remark 2.2. Given a Tait-colored cubic graph Γ, the induced surface S has a cell decomposition with Γ as its 1-skeleton and with $p_1 + p_2 + p_3$ 2-cells; hence $\chi(S) = \chi(\Gamma) + p_1 + p_2 + p_3$. The graph Γ_O in Figure 1a is orientable and induces a surface S_O, whereas the graph Γ_N in Figure 1b is nonorientable and induces S_N, satisfying $\chi(S_O) = \chi(S_N) = -4$.

Remark 2.3. A somewhat surprising consequence of Lemma 2.1 is that while the Euler characteristic of the induced surface S depends on a choice of Tait coloring, the orientability of S depends only on the underlying graph. In Figure 4, we depict two different Tait colorings $\mathcal{C}$ and $\mathcal{C}'$ of a graph Γ, inducing surfaces S and S', respectively, with $\chi(S) = 2$ and $\chi(S') = 0$.
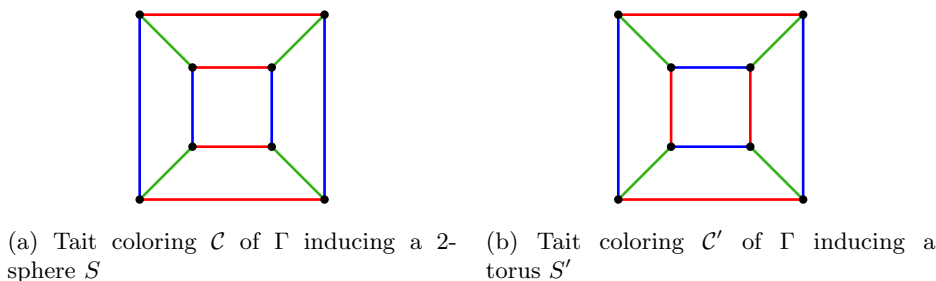

(a) Tait coloring $\mathcal{C}$ of Γ inducing a 2-sphere S

(b) Tait coloring $\mathcal{C}'$ of Γ inducing a torus S'

Figure 4. Distinct Tait colorings of Γ inducing different surfaces.

Remark 2.4. If Γ is a cubic graph embedded in a surface S such that $S \setminus \Gamma$ is a collection of disks, it does *not* necessarily imply that Γ has a Tait coloring inducing the surface S. Indeed, consider the embedding of $K_{3,3}$ in $\mathbb{RP}^2$ shown in Figure 5, where $\mathbb{RP}^2$ is obtained from the disk by identifying antipodal points on its boundary. Since $\mathbb{RP}^2$ is nonorientable and $K_{3,3}$ is bipartite, Lemma 2.1 implies that there does not exist a Tait coloring $\mathcal{C}$ of $K_{3,3}$ inducing $\mathbb{RP}^2$. We leave the following as an exercise for the reader: Let $\Gamma \subset S$ be an embedded cubic graph cutting S into disks. Then Γ admits a Tait coloring inducing S if and only if its graph dual Γ' admits a 3-coloring (of its vertex set). In the example shown in Figure 5, the dual Γ' contains a K_4 subgraph.

3. Operations on bridge trisections and cubic graphs

In this section, we describe two well-known operations (connected summation and elementary perturbation) and two novel operations (crosscap summation and tubing) that increase the number of bridge points in a given bridge trisection. We also discuss a way to simplify a cubic graph, called compression. These operations will be the basis for the construction in the proof of Theorem 1.1.
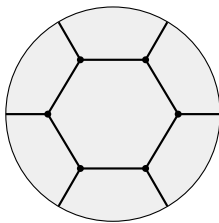
Figure 5. Embedding of $K_{3,3}$ in $\mathbb{RP}^2$, where antipodal points of the disk are identified. No Tait coloring of $K_{3,3}$ induces $\mathbb{RP}^2$.

3.1. Connected summation

Given two bridge trisections $\mathcal{T}_1$ of $(X_1, \mathcal{K}_1)$ and $\mathcal{T}_2$ of $(X_2, \mathcal{K}_2)$, with bridge spheres Σ_1 and Σ_2 and distinguished bridge points $x_1 \in \mathbf{x}_1$ and $x_2 \in \mathbf{x}_2$, the connected sum $\mathcal{T}_1 \# \mathcal{T}_2$ is the trisection for $(X_1 \# X_2, \mathcal{K}_1 \# K_2)$ obtained by removing a 4-ball neighborhood of each point x_i, which necessarily meets each component piece of $\mathcal{T}_i$ in a ball of the appropriate dimension, then identifying the component pieces of $\mathcal{T}_1$ with $\mathcal{T}_2$ along the resulting boundaries. On a diagrammatic level, a shadow diagram for $\mathcal{T}_1 \# \mathcal{T}_2$ is obtained by removing disk neighborhoods of the bridge points and gluing the two diagrams along the resulting boundaries. If Γ_1 and Γ_2 are the 1-skeleta of $\mathcal{T}_1$ and $\mathcal{T}_2$, then the 1-skeleton of $\mathcal{T}_1 \# \mathcal{T}_2$ is obtained by vertex summing Γ_1 and Γ_2 along the vertices v_1 and v_2 corresponding to the bridge points x_1 and x_2. An example is shown in Figure 8c.

3.2. Elementary perturbation

Another previously-known operation on bridge trisections is *elementary perturbation*. Every genus zero bridge trisection admits a shadow diagram (a, b, c) in which any one of the pairings, say (a, b), can be assumed to be *standard*, meaning that the union $a \cup b$ is a collection of embedded, polygonal curves that bound pairwise disjoint disks in Σ [8]. Choose a disk D bounded by one of the trivial curves in $a \cup b$, and let δ be an arc in D with one endpoint in the interior of an arc $a_i \in a$ and the other endpoint in the interior of an arc $b_j \in b$.

Now, replace $a_i \cup b_j \cup \delta$ via an IH-move, which changes a_i to two arcs a_i' and a_i'', converts b_j to two arcs b_j' and b_j'', and replaces δ with transverse arc d'. Let $a' = a \setminus \{a_i\} \cup \{a_i', a_i''\}$, $b' = b \setminus \{b_j\} \cup \{b_j', b_j''\}$, and $c' = c \cup \{d'\}$.

Then (a', b', c') is the shadow diagram for another bridge trisection $\mathcal{T}'$ of the same surface. This is called an *elementary perturbation* of $\mathcal{T}$. See Figure 6. Of course, our choice of pairing was arbitrary, and this construction will work for any of the three pairings.
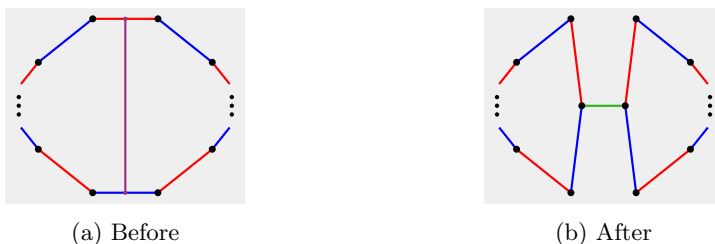

(a) Before   (b) After

Figure 6. Local pictures of a shadow diagram corresponding to an elementary perturbation.

The reverse operation is called *elementary deperturbation*: Suppose that a trisection $\mathcal{T}$ has a shadow diagram (a, b, c) such that the pairing (a, b) is standard, and there is an arc d in c and disjoint disks D_1 and D_2 bounded by arcs in $a \cup b$ such that d meets D_i in a single bridge point x_i. Then x_i is the endpoint of arcs a_i and b_i. Let a_1' be a push-off (to the outside of $D_1 \cup d \cup D_2$) of the arc $a_1 \cup d \cup a_2$, and let b_1' be a push-off (to the outside of $D_1 \cup d \cup D_2$) of the arc $b_1 \cup d \cup b_2$. Then, with $a' = a \setminus \{a_1, a_2\} \cup \{a_1'\}$, $b' = b \setminus \{b_1, b_2\} \cup \{b_1'\}$, and $c' = c \setminus \{d\}$, we have that (a', b', c') is a shadow diagram for another bridge trisection $\mathcal{T}'$ of the same surface. We say that $\mathcal{T}'$ is related to $\mathcal{T}$ by *elementary deperturbation*. By inspection, the net result of an elementary perturbation followed by an elementary deperturbation along the appropriate arc returns the original bridge trisection. For further details on these operations, see [8] and [9].

Remark 3.1. In defining elementary perturbation and deperturbation (and the moves that follow), we assume that the pairing $a \cup b$ is standard. This condition may be stronger than necessary in some instances. For example, to perform the perturbation shown in Figure 6a, it would suffice to assume that the red-blue circuit shown bounds an embedded disk in the bridge sphere that is disjoint from the other red arcs and blue arcs.

In addition, note that in the discussion above, we have standardized $a \cup b$; however, we can perform an elementary perturbation by standardizing any one of the three pairs of arc in (a, b, c).

3.3. Crosscap summation

Here we introduce a new type of local modification of a bridge trisection, which we call crosscap summation. As in the definition of elementary perturbation, suppose that a bridge trisection $\mathcal{T}$ admits a shadow diagram (a, b, c) in which one of the pairs, say (a, b), is standard. Choose a disk D bounded by one of the trivial curves in $a \cup b$, and let δ be an arc in D with one endpoint in the interior of an arc $a_i \in a$ and the other endpoint in the interior of an arc $b_j \in b$.

Next, replace $a_i \cup b_j \cup \delta$ via the following procedure: Introduce two new bridge points x_+ and x_- along a pushoff of δ, so that the subarc of the pushoff connecting x_+ to a_i does not contain x_-, and let d' be the subarc of the pushoff of δ connecting x_+ and x_-. Suppose the endpoints of a_i are x'_a and x''_a and the endpoints of b_j are x'_b and x''_b, labeled so that as we traverse ∂D in a counterclockwise direction, the order of these vertices is $\{x'_a, x''_a, x''_b, x'_b\}$. Let a'_i be an arc connecting x'_a to x_+, let a''_i be an arc connecting x''_a to x_-, let b'_j be an arc connecting x'_b to x_-, and let b''_j be an arc connecting x''_b to x_+, as shown in Figure 7b. With these choices of arcs, we have that $a'_i \cap b'_j = \emptyset$ and a''_i meets b''_j in a single point. Let $a' = a \setminus \{a_i\} \cup \{a'_i, a''_i\}$, $b' = b \setminus \{b_j\} \cup \{b'_j, b''_j\}$, and $c' = c \cup \{d'\}$. We say that the resulting triple (a', b', c') is obtained from (a, b, c) by *positive crosscap summation*.

If we repeat the construction but instead choose to label the endpoints of a_i and b_j as $\{x''_a, x'_a, x'_b, x''_b\}$ traversing ∂D in a counterclockwise direction, we obtain another triple of arcs, which is shown in Figure 7c and which we say is obtained from (a, b, c) by *negative crosscap summation*.

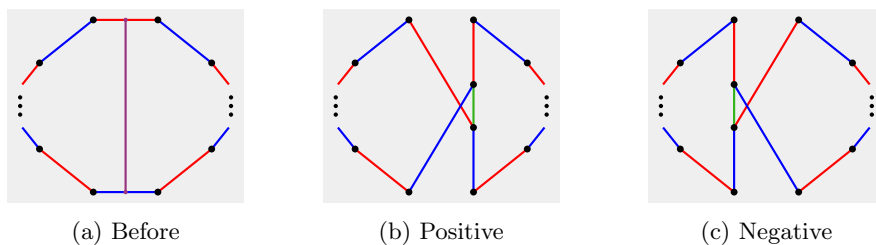| (a) Before | (b) Positive | (c) Negative |

Figure 7. Local pictures of a shadow diagram corresponding to positive and negative crosscap summations.

We have named these operations crosscap summation because the resulting surface is obtained by taking the connected sum with an unknotted projective plane, as demonstrated in the next lemma.

Lemma 3.2. *Suppose that $\mathcal{T}$ is a bridge trisection of $\mathcal{K}$ with shadow diagram (a, b, c), where the pairing (a, b) is standard. Then the result (a', b', c') of positive (resp. negative) crosscap summation is a shadow diagram for a bridge trisection $\mathcal{T}'$ of $\mathcal{K}\#\mathcal{F}^+$ (resp. $\mathcal{K}\#\mathcal{F}^-$).*

Proof. Suppose the triple (a', b', c') is obtained from (a, b, c) by a positive crosscap summation along an arc δ. Let the bridge trisection $\mathcal{T}''$ with shadow diagram (a'', b'', c'') be the result of an elementary perturbation of $\mathcal{T}$ along δ, where d'' is the newly created arc in c'', with endpoints v_1 and v_2. Let $\mathcal{T}^+$ denote the bridge trisection of $\mathcal{F}^+$ shown in Figure 2, and consider the bridge trisection $\mathcal{T}''\#\mathcal{T}^\pm$ of $\mathcal{K}\#\mathcal{F}^\pm$, where the connected sum is taken along a neighborhood of the bridge point v_1. Taking the connected sum of shadow diagrams for $\mathcal{T}''$ and $\mathcal{T}^\pm$, we can see that the $\mathcal{T}''\#\mathcal{T}^\pm$ admits an elementary deperturbation along the arc corresponding to d'' in the connected sum. The result of elementary deperturbation is a bridge trisection $\mathcal{T}'$ whose shadow diagram coincides with (a', b', c'). Since the three operations used in this proof result in bridge trisections, the statement of the lemma follows. See Figure 8. $\square$

(a) $\mathcal{T}$ (b) $\mathcal{T}''$ (c) $\mathcal{T}''\#\mathcal{T}^+$ (d) $\mathcal{T}'$
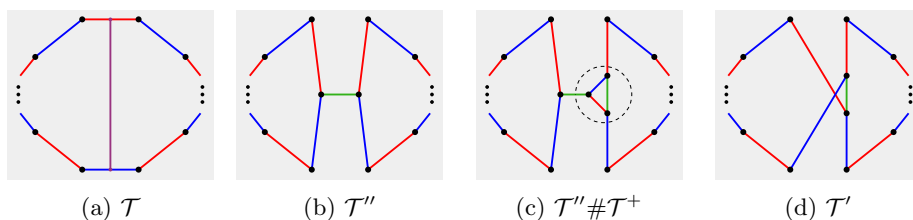
Figure 8. The sequence of bridge trisection moves described in the proof of Lemma 3.2. Choosing the opposite vertex of c'' and the shadow diagram for $\mathcal{F}^-$ in Figure 2b produces the negative crosscap summation shown in Figure 7c.

Remark 3.3. As in the case of elementary perturbation, we need not necessarily choose the standard pairing to be (a, b) in the construction of crosscap summation. Here, however, orientations are important, and so we note that the above constructions can be formulated for any cyclic permutation of (a, b, c).

3.4. Tubing

For this operation, suppose $\mathcal{T}$ is a bridge trisection with shadow diagram (a, b, c) such that one of the pairings, say (a, b) is standard, and let δ be an arc in Σ connecting an arc a_i in one component of $a \cup b$ to an arc b_j in another component, so that δ meets $a \cup b$ only at its endpoints. In a move locally identical to a perturbation along δ, we replace $a_i \cup b_j \cup \delta$ via an IH-move, which changes a_i to two arcs a_i' and a_i'', converts b_j to two arcs b_j', and b_j'', and replaces δ with transverse arc d'.

Let $a' = a \setminus \{a_i\} \cup \{a_i', a_i''\}$, $b' = b \setminus \{b_i\} \cup \{b_j', b_j''\}$, and $c' = c \cup \{d'\}$. We claim that (a', b', c') is a shadow diagram (for a bridge trisection). To see this, first note that (a', b') is embedded in Σ, since (a, b) was, and so (a', b') bounds a collection of disks in S^3. Next, note that after shadow slides of arcs in $c' \setminus \{d'\}$ over d' (if necessary), we find that $a' \cup c'$ is isotopic to $a \cup c$ (as immersed polygonal curves). Since (a, c) is a shadow diagram for a bridge splitting of an unlink, so is (a', c'). A similar arguments show that the same is true for (b', c').

It follows that (a', b', c') determines a trisection $\mathcal{T}'$ of some knotted surface $\mathcal{K}'$, although the relationship between $\mathcal{K}$ and $\mathcal{K}'$ is not immediately clear. We will show that $\mathcal{K}'$ is obtained from $\mathcal{K}$ by an operation called 1-handle addition, and to distinguish the operation on knotted surfaces from the operation on bridge trisections, we say that the bridge trisection $\mathcal{T}'$ is related to $\mathcal{T}$ by *tubing along* δ. See Figure 9 for a local picture of tubing. Note that the difference between tubing and elementary perturbation is global: in the case of tubing, the arc δ connects distinct components of $a \cup b$, whereas in an elementary perturbation, δ connects two arcs in the same component of $a \cup b$.
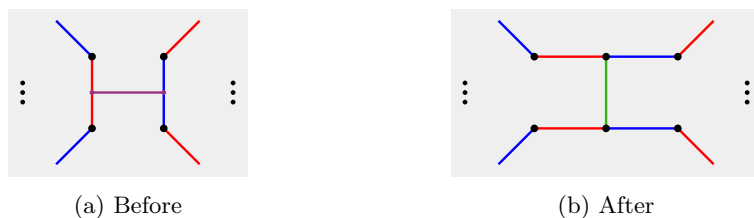

(a) Before   (b) After

Figure 9. A local picture of the tubing operation.

Suppose now that $\mathcal{K}$ is an embedded surface in S^4, and let $D^2 \times I$ be an embedded 1-handle for $\mathcal{K}$, so $(D^2 \times I) \cap \mathcal{K} = D^2 \times \{0, 1\}$. We can obtain a new surface $\mathcal{K}'$ by attaching the 1-handle to $\mathcal{K}$; that is, $\mathcal{K}' =$

$\mathcal{K} \setminus (D^2 \times \{0, 1\}) \cup (S^1 \times I)$. Letting $\delta = \{0\} \times I$, the core of the 1-handle, we say that $\mathcal{K}'$ is obtained from $\mathcal{K}$ by *1-handle addition along* δ [2] (or by *stabilization* [1]). Note that a 1-handle addition can be either orientable or nonorientable. This handle addition is a well-studied operation in knotted surface theory; it is known, for example, to be an unknotting operation [5]. In addition, the following lemma was established for orientable surfaces in [5] and nonorientable surfaces in [6] and [1]. We will require this lemma in the next section.

Lemma 3.4. *Suppose $\mathcal{K} \subset S^4$ is unknotted, and $\mathcal{K}'$ is obtained from $\mathcal{K}$ by 1-handle addition. Then $\mathcal{K}'$ is also unknotted.*

We now give the connection between tubing and 1-handle addition.

Lemma 3.5. *Suppose (a, b, c) is a shadow diagram for a bridge trisection $\mathcal{T}$ of a knotted surface $\mathcal{K} \subset S^4$, and let (a', b', c') be the shadow diagram for the bridge trisection $\mathcal{T}'$ of the knotted surface $\mathcal{K}'$ obtained by tubing $\mathcal{T}$ along an arc δ. The $\mathcal{K}'$ is related to $\mathcal{K}$ by a 1-handle addition along δ.*

Proof. As in the definition of tubing, suppose that the pairing (a, b) is standard and the arc $\delta \subset \Sigma$ connects arcs $a_i \in a$ and $b_j \in b$. By an isotopy of δ we can move it close to two bridge points x_1 and x_2 as shown in Figure 10a. Consider the shadow diagram for a bridge trisection $\mathcal{T}^*$ of a surface $\mathcal{K}^*$ shown in Figure 10b, which is contained in a neighborhood of $a_i \cup \delta \cup b_j$ in the surface Σ. We can see that $\mathcal{T}^*$ is the result of two elementary perturbations applied to the 1-bridge trisection of an unknotted 2-sphere, so that $\mathcal{K}^*$ is an unknotted 2-sphere bounding a 3-ball contained in a neighborhood of δ in S^4, which can be assumed to be disjoint from $\mathcal{K}$.

Now, we perform an operation similar to the connected sum at the bridge points x_1 and x_2 with the two closest bridge points of x_1^* and x_2^* of $\mathcal{T}^*$, as shown in Figure 10c. At the level of the embedded surfaces, this operation corresponds to removing disk neighborhoods of x_1 and x_2 in $\mathcal{K}$ and disk neighborhoods of x_1^* and x_2^* of $\mathcal{K}^*$ and identifying their respective boundaries to get a new surface $\mathcal{K}'$. Since $\mathcal{K}^*$ bounds a 3-ball, which is diffeomorphic to $I \times D^2$, we see that $\mathcal{K}'$ is obtained from $\mathcal{K}$ by a 1-handle attachment along δ. Finally, it follows from the discussion in the definition of tubing that pairs of arcs in (a', b', c') yield shadow diagrams for unlinks, so that (a', b', c') is a shadow diagram for a bridge trisection $\mathcal{T}'$ of $\mathcal{K}'$, and the diagram in Figure 10c is isotopic to Figure 9b.   □
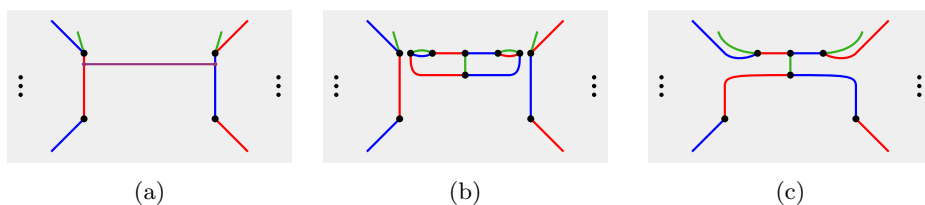
Figure 10. Shadow diagrams realizing the tubing operation as gluing an unknotted 2-sphere to $\mathcal{K}$ along two disks to get $\mathcal{K}'$, where $\mathcal{K}'$ is obtained from $\mathcal{K}$ by 1-handle addition.

Remark 3.6. As with elementary perturbation, we can perform a tubing operation by standardizing any one of the three pairs of arcs in (a, b, c).

3.5. Compression of cubic graphs

Surprisingly, at the level of the 1-skeleta, the inverses of each of the three moves described above correspond to a single abstract simplification move on a Tait-colored cubic graph Γ. Choose a distinguished edge in Γ, and suppose without loss of generality that the edge, call it e_g, is colored green. We also suppose further that e_g is not parallel to another edge of Γ. Let v^+ and v^- be the endpoints of e_g. Then there are red edges $e_r^\pm$ with one endpoint on $v^\pm$ and another endpoint on $v_r^\pm$, and there are blue edges $e_b^\pm$ with one endpoint on $v^\pm$ and another endpoint on $v_b^\pm$. Since e_g is not parallel to any other edge, it follows that the four edges $\{e_r^\pm, e_b^\pm\}$ are distinct, $v_r^+ \neq v_r^-$, $v_b^+ \neq v_b^-$, and no vertex in the set $\{v_r^\pm, v_b^\pm\}$ is v^+ or v^-.

We obtain a new Tait-colored cubic graph Γ' by removing the vertices $v^\pm$ and edge e_g, replacing the pair e_r^+ and e_r^- with a single red edge e_r between v_r^+ to v_r^-, and replacing the pair e_b^+ and e_b^- with a single blue edge between v_b^+ and v_b^-. We say that the new graph Γ' is obtained from Γ by *compression* along the edge e_g. By inspection, we can see that the graph in Figure 6a is obtained from the graph in Figure 6b by compression along the displayed green edge. Similarly, the graphs in Figures 7b and 7c are graph isomorphic, and the graph in Figure 7a is obtained from either of these graphs by compression along the displayed green edge. Finally, the graph in Figure 9a is obtained from the graph in Figure 9b by compression along the displayed green edge.

These three examples of compression are local identical but globally different, and so we further distinguish them. For that purpose, we define orientable and nonorientable edges. Suppose that Γ is Tait-colored cubic

graph, and e be an edge of Γ such that both endpoints of e are contained in the same bi-colored cycle C of the two colors opposite the color of e. Coherently orient C, so that the vertices have alternating $+$ and $-$ labels as in the proof of Lemma 2.1. If e connects vertices of opposite sign, we say e is *orientation-preserving*. Otherwise, e connects vertices of the same sign, and we say e is *orientation-reversing*. Equivalently, e is orientation-preserving if and only if it completes a path in C to a cycle of even length. Following the proof of Lemma 2.1, we also note that a one-patch graph Γ is orientable if and only if it does not contain an orientation-reversing edge with respect to some bicolored cycle.

By definition, an edge e with vertices in the same bi-colored cycle C of the two opposite colors is either orientation-preserving or orientation-reversing. If, on the the other hand, e connects distinct bi-colored cycles of the two opposite colors, we say that e is *connecting*. A compression performed along a connecting edge is called a *p-compression* (Figures 6a and 6b), a compression along an orientation-reversing edge is called a *c-compression* (Figures 7a, 7b, and 7c), and a compression along an orientation-preserving edge resulting in a connected graph is called a *t-compression* (Figures 9a and 9b). Note that p-compression, c-compression, and t-compression are operations that are inverses to the operations on the 1-skeleton of a bridge trisection induced by elementary perturbation, crosscap summation, and tubing, respectively. This implies that p-compression and c-compression don't disconnect the graph; by definition, neither does t-compression.

Remark 3.7. We observe that both p-compression and t-compression of an oriented graph result in orientable graphs. On the other hand, c-compression can only be applied to non-orientable graphs and may result in either an orientable graph or a non-orientable graph. Note that the above discussion pertains to abstract graphs, but refers to Figures 6–9, which show shadow diagrams. So, the reader should consider these figures as abstract graphs by forgetting the grey shading representing the bridge sphere.

4. Proof of the main theorem

The theta graph is the simplest Tait-colorable graph, and it is the 1-skeleton of the simplest bridge trisection, the 1-bridge trisection of the unknotted S^2. Before proving the main theorem, we require several technical results related to sequences of compressions reducing a given graph.

Proposition 4.1. *Suppose Γ is a one-patch nonorientable Tait-colored cubic graph. Then Γ admits a c-compression yielding a nonorientable graph or the theta graph.*

Proof. Note that up to isomorphism, K_4 has a unique Tait coloring, every edge is nonorientable, and any c-compression yields the theta graph. We will show that if every c-compression of Γ yields an orientable graph, then $\Gamma = K_4$, from which the statement of the proposition follows. Suppose every c-compression of Γ yields an orientable graph. We produce a convenient picture of Γ, in which the red and blue edges form a regular n-gon, and the green edges are drawn as chords of this n-gon (as in the examples in Figure 1). In this setting, every pair of green edges meets either once or not at all. Since Γ is nonorientable, it contains a nonorientable edge e; we suppose without loss of generality that e is colored green. Orient the vertices of the red-blue cycle C with $+$ and $-$. Since each orientation-preserving green edge connects two vertices of opposite sign, while each green orientation-reversing edge connects vertices of the same sign, and there are the same number of vertices labeled $+$ as there are labeled $-$, it follows that the number of green orientation-reversing edges is even.

Removing the red and blue edges incident to e separates the red-blue cycle C of Γ into paths p and p'. We assume further that the vertices incident to e are labeled $+$, so that both endpoints of the path p and both endpoints of the path p' are labeled $-$. Consider the graph Γ' obtained from c-compression of Γ along e. The red-blue cycle C' of Γ' is obtained by connecting the paths p and p' along new edges joining their endpoints, and thus, an orientation of C' can be obtained by preserving the orientation of p coming from C and reversing the orientation of p' coming from C. See Figure 11.

Figure 11. A nonorientable edge splits a red-blue cycle into paths p and p' (left), where the result of a c-compression preserves the orientation of p and reverses the orientation of p' (right).

It follows that if a green orientation-preserving edge of Γ crosses e, then that edge becomes orientation-reversing in Γ'. Similarly, if a green orientation-reversing edge of Γ avoids e, then that edge remains orientation-reversing in Γ'. By assumption, Γ' does not contain an orientation-reversing edge, and thus we see that every green orientation-preserving edge avoids e, while every green orientation-reversing edge crosses e. Moreover, this is true not just for e but for every green orientation-reversing edge. We conclude that every pair of green orientation-reversing edges in Γ meet in a single point, and no green orientation-preserving edge crosses a green orientation-reversing edge.

Suppose now that Γ contains a green orientation-preserving edge e', and let Γ^* be the graph obtained from Γ by deleting the green orientation-reversing edges along with any incident vertices and incident red or blue edges. Then Γ^* is a proper subgraph of Γ, and each component of Γ^* is spanned by a red-blue path p^*, with vertices $w_1, \ldots, w_m$ and edges $e_1, \ldots, e_{m-1}$ appearing in order. In Γ, each w_i is then the endpoint of a green orientation-preserving edge e^*, and since e^* crosses no green orientation-reversing edge, the other endpoint of e^* is contained in $\{w_1, \ldots, w_m\}$. It follows that m is even, so that the edges e_1 and e_{m-1} are the same color, say red. Note further that the valence of w_1 and w_m in Γ^* is two, while the valence of the other vertices is three. Thus, every vertex of Γ^* is the endpoint of both a red edge and a green edge, so Γ^* contains a red-green cycle. Since Γ^* is not all of Γ, we have that Γ contains more than one red-green cycle, contradicting the assumption that Γ is one-patch. See Figure 12 for an example.
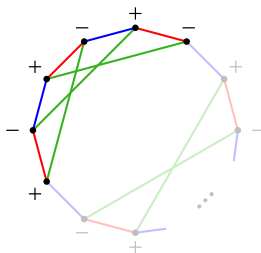

Figure 12. If every c-compression of Γ along a green edge yields an orientable graph, then a component of the subgraph Γ^* of Γ induced by the endpoints of green orientation-preserving edges contains a bi-colored cycle.

We are left with the case that every green edge of Γ is orientation-reversing. As noted above, there are an even number of such edges, so Γ has

$4k$ vertices for some integer k. Label the vertices of Γ in order, $v_1, \dots, v_{4k}$. Since every green edge must cross every other green edge, each green edge is a diameter of the red-blue cycle C, connecting v_i to v_{2k+i}; otherwise, there would exist disjoint green edges. Consider the green edges connecting v_1 to v_{2k+1} and v_2 to v_{2k+2}. Since the edges of C alternate colors, it follows that the edge between v_1 and v_2 is the same color as the edge between v_{2k+1} and v_{2k+2}. Thus, Γ contains a bi-colored cycle of length four. It follows that $\Gamma = K_4$, the unique one-patch graph with four vertices. $\square$

With this technical hurdle out the the way, we can swiftly prove the next lemma.

Lemma 4.2. *Suppose Γ is a Tait-colored cubic graph*

1) *If Γ is not 1-patch, then Γ can be reduced to a 1-patch graph by a finite number of p-compressions.*

2) *If Γ is an orientable 1-patch graph, then Γ can be reduced to the theta graph by a finite sequence alternating between t-compressions and p-compressions.*

3) *If Γ is a nonorientable 1-patch graph, then Γ can be reduced to the theta graph by a finite number of c-compressions.*

Proof. First, suppose that Γ is not 1-patch. We induct on the sum $p_1 + p_2 + p_3$ of the patch numbers of Γ. Suppose without loss of generality that Γ contains at least two red-blue cycles. Since Γ is connected, there is some green edge e_g connecting distinct red-blue cycles (and thus e_g is not parallel to another edge). Then a p-compression of Γ along e_g produces a new graph Γ' with one fewer red-blue cycle than Γ. Since no blue-green nor green-red cycles have been added, the claim holds by induction.

For the second part of the lemma, suppose Γ is orientable and 1-patch. Here we induct on the number of vertices of Γ. Suppose that the claim is true for all orientable 1-patch graphs with fewer vertices than Γ. By assumption, all edges of Γ are orientation-preserving. Choose one, and let Γ' be the result of a compression of Γ along it. Note that compression decreases the number of vertices of Γ by two but increases the sum of the patch numbers by one, so that Γ' is not 1-patch, although Γ' is spanned by a single bi-colored cycle, so it is still connected. Thus, the compression is a t-compression, since Γ is orientable. By the first step, Γ' admits a p-compression yielding a new one-patch graph Γ'' with four fewer vertices than Γ. Since t-compression and

p-compression of an orientable graph yield another orientable graph (see Remark 3.7), the claim holds by induction.

Finally, suppose Γ is a nonorientable 1-patch graph. Again, we induct on the number of vertices of Γ. Noting that c-compression of a 1-patch graph yields a 1-patch graph, we have by Proposition 4.1 that Γ has a nonorientable edge such that c-compression along e yields another nonorientable graph, and the third claim follows immediately. □

We remark that following the proof of Proposition 4.1, we can take the sequence of c-compressions guaranteed by claim (3) of Lemma 4.2 to occur along edges of the same color. As an example, seven c-compressions convert the nonorientable one-patch graph Γ_N from Figure 1b to the theta graph. This sequence of compressions is shown in Figure 13.
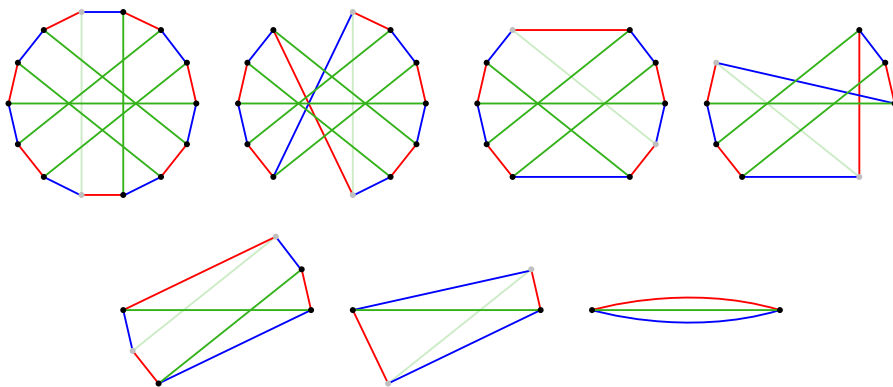

Figure 13. A sequence of c-compressions converting Γ_N to the theta graph.

We can now prove the main theorem, which we restate here for convenience.

Theorem 1.1. *If Γ is a cubic graph with a Tait coloring $\mathcal{C}$, then there exists a bridge trisection $\mathcal{T}$ of an unknotted surface $\mathcal{U} \subset S^4$ such that the 1-skeleton $\mathcal{T}$ is graph isomorphic to Γ, with the coloring $\mathcal{C}$ induced by $\mathcal{T}$. Moreover, if the induced surface S is nonorientable, we may choose the embedding of $\mathcal{U}$ to have any possible normal Euler number.*

Proof. Suppose first that Γ is an orientable 1-patch graph. By Lemma 4.2, there is a sequence of alternating t-compressions and p-compressions converting Γ to the theta graph. In the reverse direction, there is a sequence of alternating elementary perturbations and tubings performed on the 1-bridge

trisection of the unknotted 2-sphere that cancel the t-compressions and p-compressions, so that the 1-skeleton of the resulting bridge trisection $\mathcal{T}$ of a surface $\mathcal{K} \subset S^4$ is isomorphic to Γ. By Lemma 3.5, the surface $\mathcal{K}$ is obtained from the unknotted 2-sphere by a sequence of 1-handle additions, and thus by repeated applications of Lemma 3.4, we have that $\mathcal{K}$ is unknotted.

Next, suppose Γ is a nonorientable 1-patch graph. By Lemma 4.2, there is a sequence of n c-compressions converting Γ to the theta graph. In the reverse direction, a sequence of n crosscap summations on the 1-bridge trisection of the unknotted 2-sphere cancel the c-compressions, yielding a bridge trisection $\mathcal{T}$ with 1-skeleton isomorphic to Γ. By Lemma 3.2, the knotted surface $\mathcal{K}$ corresponding to $\mathcal{T}$ is the connected sum of n copies of $\mathcal{F}^+$ and $\mathcal{F}^-$, so $\mathcal{K}$ is unknotted. Moreover, without affecting the induced graph, we can choose the number of each type of summand by picking between positive and negative crosscap summations. Thus, with these choices we can construct $\mathcal{K}$ to have any possible normal Euler number between $-2n$ and $2n$, which are the only possible values by [7].

Finally, if Γ is not 1-patch, Lemma 4.2 asserts that Γ reduces to a 1-patch graph Γ' after a sequence of p-compressions. By the previous steps, there exists a bridge trisection $\mathcal{T}'$ of an unknotted surface $\mathcal{K}$ such that the 1-skeleton of $\mathcal{T}'$ is isomorphic to Γ', and if Γ' is nonorientable, $\mathcal{K}$ can be chosen with any possible normal Euler number. Then there is a sequence of elementary perturbations of $\mathcal{T}'$ canceling the sequence of p-compressions, yielding a bridge trisection $\mathcal{T}$ of the same unknotted surface $\mathcal{K}$, where the 1-skeleton of $\mathcal{T}$ is isomorphic to Γ, completing the proof. $\square$

Recall that for a bridge trisection $\mathcal{T}$, a choice of disks $E_{ij} \subset B_{ij}$ containing the bridge points, and a projection of the tangles τ_{ij} onto E_{ij} with crossing data determines a tri-plane diagram $\mathcal{P} = (\mathcal{P}_{12}, \mathcal{P}_{23}, \mathcal{P}_{31})$ representing $\mathcal{T}$, and any two tri-plane diagrams $\mathcal{P}$ and $\mathcal{P}'$ for $\mathcal{T}$ are related by interior Reidemeister moves and mutual braid transpositions. We now prove Corollary 1.2, which we restate for convenience.

Corollary 1.2. *Every tri-plane diagram $\mathcal{P}$ of a knotted surface $\mathcal{K} \subset S^4$ can be converted to a tri-plane diagram $\mathcal{P}'$ for an unknotted surface $\mathcal{U}$ by a sequence of interior Reidemeister moves and crossing changes.*

Proof. Suppose that $\mathcal{T}$ is a bridge trisection of an embedded surface $\mathcal{K}$ in S^4, where Γ is the 1-skeleton of $\mathcal{T}$ with induced Tait coloring $\mathcal{C}$. By Theorem 1.1, there exists a bridge trisection $\mathcal{T}'$ of an unknotted surface $\mathcal{K}'$ whose Tait-colored 1-skeleton is (graph) isomorphic to Γ and $\mathcal{C}$. Let $\mathcal{P}$ and $\mathcal{P}'$ be tri-plane diagrams corresponding to $\mathcal{T}$ and $\mathcal{T}'$, respectively. The graph isomorphism

of the 1-skeleta of $\mathcal{T}$ and $\mathcal{T}'$ induces a bijection from the bridge points of $\mathcal{P}$ and $\mathcal{P}'$. After performing some number of mutual braid transpositions on the tri-plane diagram $\mathcal{P}'$, we may assume that this bijection is the identity.

Viewing the tangles τ_{ij} and τ'_{ij} as being contained in the same 3-ball B, the graph bijection implies that τ_{ij} and τ'_{ij} are homotopic via a homotopy supported outside of a neighborhood of ∂B. Using the projection disk E_{ij}, the (generic) projection of this homotopy yields a sequence of interior Reidemeister moves and crossing changes taking $\mathcal{P}_{ij}$ to $\mathcal{P}'_{ij}$. Carrying out this process in each of the three sectors yields the corollary. $\qquad\square$

Remark 4.3. Although the end result of the the sequence of interior Reidemeister moves and crossing changes in the proof of Corollary 1.2 yields a tri-plane diagram $\mathcal{P}'$, there is no reason to expect that any of the intermediate diagrams is a tri-plane diagram, since changing a single crossing likely destroys the condition that tangles pair to give diagrams of unlinks.

5. The examples Σ_O and Σ_N

We conclude by working through the details of Theorem 1.1 with the two examples from Figure 1. First, the Tait-colored Heawood graph Σ_O is orientable and one-patch, and thus by Lemma 4.2, it admits an alternating sequence of t-compressions and p-compressions converting it to the theta graph. These compressions are shown in Figure 14.
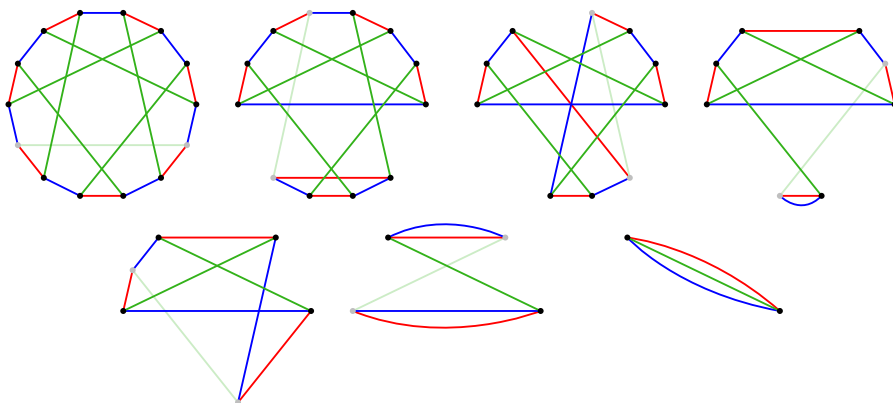

Figure 14. An sequence of t-compressions and p-compressions converting Γ_O to the theta graph.

At the level of bridge trisections, we can work our way backwards, starting with the 1-bridge splitting of the unknotted 2-sphere and performing an alternating sequence of elementary perturbations and tubings so that each subfigure of Figure 15 below is graph isomorphic to a corresponding graph in the sequence of compressions shown in Figure 14, in reverse order.
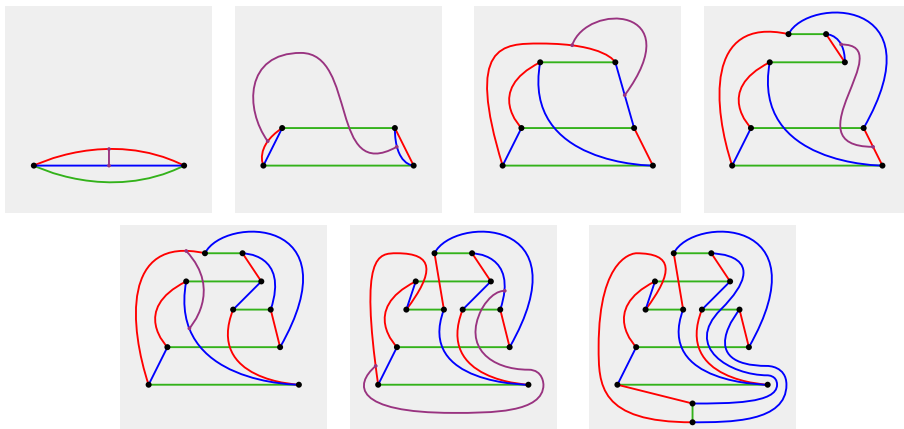

Figure 15. A sequence of elementary perturbations and tubings performed on the 1-bridge trisection of the unknotted 2-sphere, the end result of which is a bridge trisection of an unknotted genus three surface with 1-skeleton isomorphic to the Heawood graph Σ_O.

Turning to the other example, recall that for the nonorientable graph Σ_N, Theorem 1.1 also allows us to choose the normal Euler number of the resulting unknotted surface. In Figure 16, we have chosen to reverse the compressions shown in Figure 13 by performing three negative crosscap summations followed by three positive crosscap summations, so that the resulting surface $\mathcal{U}$ has normal Euler number zero. There is an added layer of complexity in this example, since a crosscap summation necessarily introduces shadow arcs that cross each other, and before we perform the next crosscap summation, we are required to standardize the red and blue pairing.

Finally, we convert the two diagrams with 1-skeleta Σ_O and Σ_N via diffeomorphisms to diagrams in which the red and blue arcs form a regular 14-gon. The final results are shown in Figure 17.

Remark 5.1. The operations of crosscap summation and tubing are new constructions appearing in this paper. They may very well be of independent interest, as they show some natural compatibility of the bridge trisection apparatus with existing well-understood constructions. We "discovered" the
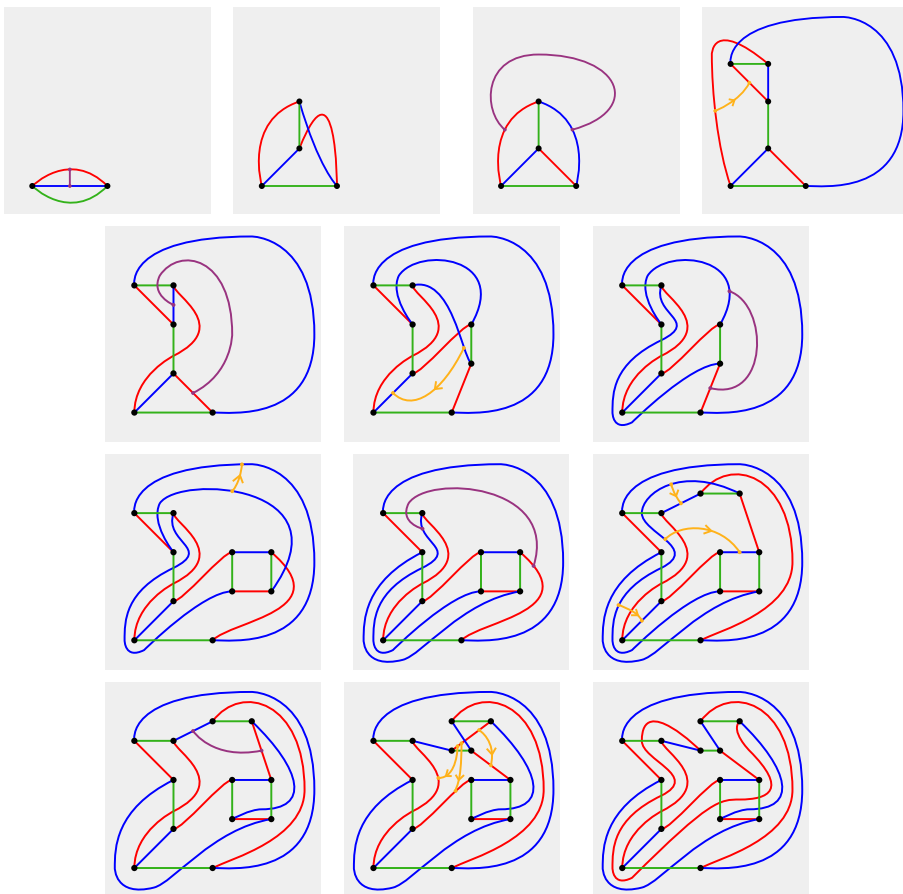
Figure 16. A sequence of moves, alternating between crosscap summations and shadow slides, starting with the 1-bridge trisection of the unknotted 2-sphere and ending with a bridge trisection of an unknotted nonorientable genus six surface with 1-skeleton isomorphic to Σ_N.

existence of crosscap summation and tubing by first understanding the different types of abstract graph compressions and investigating whether there could be corresponding topological operations. Thus, this work can be seen as a case of graph theory motivating topology.
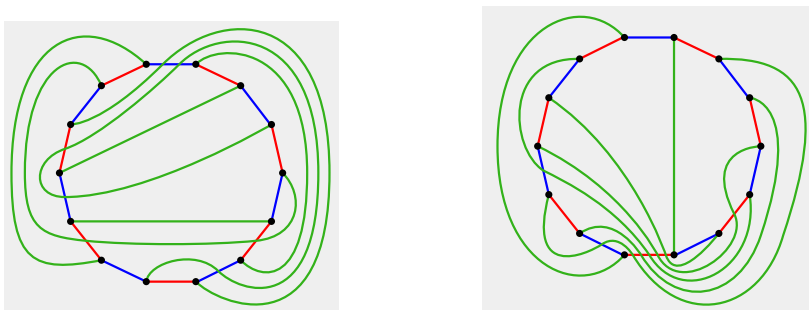
Figure 17. Versions of shadow diagrams of bridge trisections with 1-skeleta isomorphic to Σ_O (left) and Σ_N (right) in which the red-blue curve is drawn as a regular 14-gon.

References

[1] R. I. Baykur and N. Sunukjian, *Knotted surfaces in 4-manifolds and stabilizations*, J. Topol. **9** (2016), no. 1, 215–231.

[2] S. Carter, S. Kamada, and M. Saito, Surfaces in 4-space, Vol. 142 of *Encyclopaedia of Mathematical Sciences*, Springer-Verlag, Berlin (2004), ISBN 3-540-21040-7. Low-Dimensional Topology, III.

[3] P. J. Heawood, *Map-colour theorems*, Quart. J. Math. Oxford Ser. **24** (1890), 322–338.

[4] F. Hosokawa and A. Kawauchi, *Proposals for unknotted surfaces in four-spaces*, Osaka Math. J. **16** (1979), no. 1, 233–248.

[5] ———, *Proposals for unknotted surfaces in four-spaces*, Osaka Math. J. **16** (1979), no. 1, 233–248.

[6] S. Kamada, *Cords and 1-handles attached to surface-knots*, Bol. Soc. Mat. Mex. (3) **20** (2014), no. 2, 595–609.

[7] W. S. Massey, *Proof of a conjecture of Whitney*, Pacific J. Math. **31** (1969) 143–156.

[8] J. Meier and A. Zupan, *Bridge trisections of knotted surfaces in S^4*, Trans. Amer. Math. Soc. **369** (2017), no. 10, 7343–7386.

[9] ———, *Bridge trisections of knotted surfaces in 4-manifolds*, Proc. Natl. Acad. Sci. USA **115** (2018), no. 43, 10880–10886.

Department of Mathematics, Western Washington University
Bellingham, WA 98225, USA
E-mail address: jeffrey.meier@wwu.edu

Department of Mathematics, University of California, Davis
Davis, CA 95616, USA
E-mail address: thompson@math.ucdavis.edu

Department of Mathematics, University of Nebraska-Lincoln
Lincoln, NE 68588, USA
E-mail address: zupan@unl.edu