

SURFACE EMBEDDINGS IN $\mathbb{R}^2 \times \mathbb{R}$

WILLIAM W. MENASCO AND MARGARET NICHOLS

ABSTRACT. This is an investigation into a classification of embeddings of a surface in Euclidean 3-space. Specifically, we consider \mathbb{R}^3 as having the product structure $\mathbb{R}^2 \times \mathbb{R}$ and let $\pi : \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2$ be the natural projection map onto the Euclidean plane. Let $\varepsilon : S_g \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$ be a smooth embedding of a closed oriented genus g surface such that the set of critical points for the map $\pi \circ \varepsilon$ is a smooth (possibly multi-component) 1-manifold, $\mathcal{C} \subset S_g$. We say \mathcal{C} is the *crease set of ε* and two embeddings are in the same *isotopy class* if there exists an isotopy between them that has \mathcal{C} being an invariant set. The case where $\pi \circ \varepsilon|_{\mathcal{C}}$ restricts to an immersion is readily accessible, since the turning number function of a smooth curve in \mathbb{R}^2 supplies us with a natural map of components of \mathcal{C} into \mathbb{Z} . The Gauss-Bonnet Theorem beautifully governs the behavior of $\pi \circ \varepsilon(\mathcal{C})$, as it implies $\chi(S_g) = 2 \sum_{\gamma \in \mathcal{C}} t(\pi \circ \varepsilon(\gamma))$, where t is the turning number function. Focusing on when $S_g \cong S^2$, we give a necessary and sufficient condition for when a disjoint collection of curves $\mathcal{C} \subset S^2$ can be realized as the crease set of an embedding $\varepsilon : S^2 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$. From there, we give the classification of all isotopy classes of embeddings when $\mathcal{C} \subset S^2$ and $|\mathcal{C}| = 3$ —a simple yet enlightening case. As a teaser of future work, we give an application to knot projections and discuss directions for further investigation.

1. INTRODUCTION




FIGURE 1. The left illustration shows a generic neighborhood of a point in the crease set and its corresponding projection into \mathbb{R}^2 . The right illustration shows the neighborhood of a corner-point in the crease set and its corresponding projection into \mathbb{R}^2 .

1.1. Initial discussion of main results. We let $\varepsilon : S_g \hookrightarrow \mathbb{R}^2 \times \mathbb{R} (\cong \mathbb{R}^3)$ be a smooth embedding of a closed surface of genus g . Let $\pi : \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2$ be the natural projection coming from the product structure. Let $\mathcal{C} \subset S_g$ be the critical point set of $\pi \circ \varepsilon$ —the set of points, $x \in S_g$, where the differential, $d(\pi \circ \varepsilon)_x : T_x(S_g) \hookrightarrow T_{\pi \circ \varepsilon(x)}(\mathbb{R}^2)$ is not a surjective

Key words and phrases. surface embedding, branched surface, crease set, crossing balls.

map. We will refer to \mathcal{C} as the *crease set of ε* . By general position arguments we can assume $\varepsilon(\mathcal{C}) \subset \mathbb{R}^2 \times \mathbb{R}$ is a collection of smooth simple closed curves (s.c.c.'s) with a finite collection of marked *corners*—points where the immersion of the crease set into the plane, $(\pi \circ \varepsilon)(\mathcal{C}) \subset \mathbb{R}^2$, fails to be smooth, with a local picture as in Fig. 1. Moreover, we may assume that the only multi-point images of the projected crease set are transverse double points. We refer to such well-behaved embeddings as *regular*.




FIGURE 2. The crease curve, shown in green, projects to a piecewise smooth curve.

The simplest, yet still interesting, example might be an embedding of the 2-sphere, $S^2 (= S_0)$, that has its crease set a single curve with two corners. (See Fig. 2.) To metaphorically describe how this embedding might be obtained, one starts with the standard unit sphere in \mathbb{R}^3 centered at the origin. The reader could image forming a dimple in the “clay ball” the sphere bounds by taking a tilted rigid bar—say having a core line of $\{y + z = 1, x = 1.5\}$ —and pushing it into the ball via parallel translate from $x = 1.5$ to $x = 0$. Clay being clay, the bar deforms the round ball so that a dimple is created in the sphere. The crease curve—the equator of S^2 before this deformation—appears deformed by a type-I Reidemeister isotopy move. The new crease set projects to \mathbb{R}^2 with a single double-point crossing.

This note is an initial investigation into the classification of *isotopy classes of regular surface embeddings into $\mathbb{R}^2 \times \mathbb{R}$* . Specifically, we will consider two regular embeddings, $\varepsilon_0, \varepsilon_1 : S_g \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$, to be in the same isotopy class if there is a smooth isotopy $\varepsilon_t : S_g \rightarrow \mathbb{R}^2 \times \mathbb{R}$, $0 \leq t \leq 1$, such that the marked structure of $\mathcal{C} \subset S_g$, as the critical set of ε_t , is invariant for $t \in [0, 1]$ —we say such an isotopy is *regular*. The number of corners of a crease component is unchanged by such an isotopy, and in particular, the equatorial crease set of the standard unit sphere and the two-corner crease set of the dimpled sphere are in different isotopy classes.

For a given S_g and ε , and for a connected component $\gamma \subset \mathcal{C}$, let $N(\gamma)$ be a closed annular neighborhood of γ that is sufficiently small so that $N \cap \mathcal{C} = \gamma$. Let $\{\hat{\gamma}, \hat{\gamma}'\} = \partial N$. In §2.3 we give a scheme for assigning an orientation to each component of ∂N . For such oriented curves we can then define a “turning number” for their projections to \mathbb{R}^2 of their embeddings into $\mathbb{R}^2 \times \mathbb{R}$. We will then have a naturally associated 2-tuple, $(t(\hat{\gamma}), t(\hat{\gamma}')) \in \mathbb{Z}^2$ —the *turning numbers* associated with γ . (Please see §2 for precise definitions.) In the sphere-with-dimple example, this 2-tuple is $(1, 1)$ for the single component of \mathcal{C} .

When we restrict a regular isotopy to a smooth oriented curve in S_g , its turning number will be invariant. As such, we will see that our 2-tuple is invariant within the isotopy class of an embedding. The starting point for our investigation will be the establishment of relationships between our turning number 2-tuple and the Euler characteristic of the surface, $\chi(S_g)$. Specifically,

$$(1.1) \quad \chi(S_g) = \sum_{\gamma \in \mathcal{C}} t(\hat{\gamma}) + t(\hat{\gamma}').$$

This equality is restated in Theorem 4.3. It is properly seen as a Gauss-Bonnet type equation that governs the behavior of embeddings of S_g into $\mathbb{R}^2 \times \mathbb{R}$. Thus, we say that a 2-tuple weight assignment to components of \mathcal{C} which satisfies the equations of Theorem 4.3 corresponds to a *Gauss-Bonnet weighting* of \mathcal{C} .

When we restrict to embedding classes of S^2 into $\mathbb{R}^2 \times \mathbb{R}$, Equ. 1.1 becomes $\chi(S^2) = 2$. We then have the following natural questions.

1. Given an collection of disjoint s.c.c.'s with corners, $\mathcal{C} \subset S^2$, when is there a Gauss-Bonnet weighting of \mathcal{C} ?
2. Given a collection of s.c.c.'s with corners, $\mathcal{C} \subset S^2$, that has a Gauss-Bonnet weighting, is there an embedding, $\varepsilon : S^2 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$, that realizes \mathcal{C} as a crease set of ε ?

Focusing on question 1, it is easy to produce collections that do not have any Gauss-Bonnet weightings—three non-concentric circles in S^2 for example. However, when such a weighting exists it is unique.

Theorem 1.1. *A Gauss-Bonnet weighting on a collection of disjoint s.c.c.'s with corners in S^2 is uniquely determined.*

Let (S^2, \mathcal{C}) and (S^2, \mathcal{C}') be two pairs of 2-sphere/collection of disjoint smooth marked s.c.c.'s. Assume each pair has a Gauss-Bonnet weighting. We say the two pairs are *equivalent crease set configurations* if there exists a diffeomorphism of pairs, $h : (S^2, \mathcal{C}) \hookrightarrow (S^2, \mathcal{C}')$. With this equivalence in mind, Theorem 1.1 can be leveraged to prove the following finiteness result. For its statement, we denote the number of corners of curve $\gamma \in \mathcal{C}$ as $c(\gamma) (\in \mathbb{N})$.

Theorem 1.2. *Let $n_c, n_t \subset \mathbb{N}$ be two arbitrary integers. Then there are only finitely many possible crease set configurations in S^2 such that $|\mathcal{C}| \leq n_t$ and $\sum_{\gamma \in \mathcal{C}} c(\gamma) \leq n_c$.*

For our initial attempt at answering question 2, we will simplify to the special case where $\mathcal{C} \subset S^2$ is a collection of disjoint s.c.c.'s without any corners. In this case, for a component, $\gamma \in \mathcal{C}$, we will have $t(\hat{\gamma}) = t(\hat{\gamma}')$.

In this simplified case we have the following main result.

Theorem 1.3. *For any collection of disjoint smooth s.c.c.'s $\mathcal{C} \subset S^2$ admitting a Gauss-Bonnet weighting, there exists a regular embedding, $\varepsilon : S^2 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$, which realizes \mathcal{C} as the crease set with the corresponding Gauss-Bonnet weighting. That is, for each $\gamma \in \mathcal{C}$, $t(\gamma)$ is equal to its weight.*

We will establish Theorem 1.3 by construction. Our constructive argument will yield regular embeddings that are in an aesthetically nice form—the curvature function on each component of $\pi \circ \varepsilon(\mathcal{C})$ will never be zero.

Surprisingly (at least to the authors), this aesthetic feature is not always achievable. In particular, the relatively simple situation when $\mathcal{C} \subset S^2$ is just three concentric circles without corners there does exist a “non-intuitive” embedding where the curvature function on $\pi \circ \varepsilon(\mathcal{C})$ must have points of zero curvature. Regardless, it is still possible to perform a calculation that gives a complete classification of the isotopy classes for when $|\mathcal{C}| = 3$. We view this novel calculation as prescient of what is possible for when $|\mathcal{C}|$ and genus are higher.

Theorem 1.4. *Up to reflection, there are exactly three isotopy classes of regular S^2 embeddings into $\mathbb{R}^2 \times \mathbb{R}$ when \mathcal{C} is just three curves without corners.*

As a final remark to the section, the authors have conducted a wide literature search and have found no previous investigation into the critical set of $\pi \circ \varepsilon$ —the crease set—with the exception of the recent contribution of Joel Hass [H]. Since the Gauss-Bonnet theorem is a classical premier result of differential topology that has produced a significant body of applications and generalizations, this lack of previous interest in the crease set is puzzling. Regardless, as the arguments in this paper will illustrate the crease set is a source of significant control over the behavior and positioning of arbitrary surface embeddings in \mathbb{R}^3 .

1.2. A link projection application. Our original motivation for studying surface embeddings into $\mathbb{R}^2 \times \mathbb{R}$ comes from knot theory. Early work of the first author developed a “normal form” for representing essential surfaces in S^3 link complements with respect to a link projection [M1].

The key construction for this normal form involves positioning a link L to lie in \mathbb{R}^2 (coming from its projection) except at crossings, where the two crossing strands lie on the boundary of a small “crossing ball”. A surface $S \subset \mathbb{R}^3 \setminus L$ in normal form can be reconstructed entirely from its intersections with \mathbb{R}^2 and these crossing balls. To set notation, let \mathbf{R}_+^2 (respectively, \mathbf{R}_-^2) be \mathbb{R}^2 with each disk neighborhood of a crossing replaced by the upper (respectively, lower) hemisphere of a crossing ball. The surface intersects $\mathbf{R}_+^2 \cup \mathbf{R}_-^2$ in a graph, and when in normal form the crease set is a collection of disjoint circuits/curves of this graph.




FIGURE 3. The blue and green curves depict $S^2 \cap \mathbf{R}_+^2$ with the green portion indicating which part is contained in the crease set. There is an arbitrary 2-tangle contained inside the ball labeled T .

In Fig. 3 we depict the regular projection (with the crossing ball structure) of a link “template” realizing the dimpled S^2 of Fig. 2 as a surface in normal form. The crease set is a subset of the graph $\Gamma = S^2 \cap (\mathbf{R}_+^2 \cap \mathbf{R}_-^2)$ and corresponds to the green curve. Inside the region of the link template labeled T we can place any 2-strand tangle. Thus, there are infinitely many possible link projections that realize the dimpled S^2 as a surface in normal form. This example is illustrative of the general situation (Theorem 3.1) which we state here in a more colloquial manner.

Theorem 1.5. *Every isotopy class of regular surface embeddings into $\mathbb{R}^2 \times \mathbb{R}$ can be realized as in normal form with respect to infinitely many regular link projections.*

1.3. Outline of paper. In §2 we give the formal definitions of the crease set and turning number plus some additional concepts of our machinery. In §3 we give the proof of Theorem 1.5. In §4 we discuss the Gauss-Bonnet equations and results that are behind Equ. 1.1, which we then use to prove Theorems 1.1 and 1.2. Then in §5 we give a construction that establishes Theorem 1.3. In §6 we develop the additional machinery need to prove Theorem 1.4. In particular, in §6.1 we observe that for each isotopy class of a regular embedding of a surface into $\mathbb{R}^2 \times \mathbb{R}$ there is an associated naturally embedded *branched surface with boundary*. A salient feature of this natural branched surface is a correspondence between its boundary and branching locus curves, and the components of the crease set. (The reader may correctly suspect that the branched surface associated with the Fig. 2 embedding is one containing the Lorenz template [G-L].) Finally, in §9 we advance directions for further investigation.

ACKNOWLEDGEMENTS

The first author would like to thank Adam Sikora for the brief conversation pre-COVID-19 that was the spark for thinking about the crease set of an embedding. The second author is grateful to the Fields Institute for its support and hospitality. The authors are grateful to Joel Hass for alerting them to his work.

2. PRELIMINARIES

2.1. Definition of the crease set. Let S_g be a smooth closed oriented surface of genus g . Let $\varepsilon : S_g \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$ be a smooth embedding and $\pi : \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2$ be the natural projection onto the first factor in the product structure.

Definition 2.1. The *crease set* $\mathcal{C} \subset S_g$ of ε (with respect to π) is the critical point set of $\pi \circ \varepsilon$.

We would like to restrict to embeddings where \mathcal{C} is a nice subset of S_g —ideally, a smooth submanifold.

Fix an embedding $\varepsilon : S_g \rightarrow \mathbb{R}^3$. Note that for any matrix $A \in \text{SO}(3)$, post-composing ε by the map $x \mapsto Ax$ gives a new embedding, which we denote ε_A .

Let $N : S_g \rightarrow S^2$ denote the Gauss map corresponding to $\varepsilon(S_g)$, so $N(p)$ is the outward-pointing unit normal vector to $\varepsilon(p)$. Observe that $\mathcal{C} = N^{-1}(Z)$, where $Z \subset S^2$ denotes the equator of S^2 .

Define $f : S_g \times \text{SO}(3) \rightarrow S^2$ by $f(p, A) = A \cdot N(p)$. This map takes a point p to the outward normal vector at $\varepsilon(p)$, then rotates it by A . Alternatively, f maps p to its outward unit normal vector in ε_A ; seen this way, f parametrizes the Gauss maps associated to the different embeddings ε_A .

We claim that f is a submersion. To see this, observe that $df_{(p,A)}$ restricted to $T_{(p,A)}\text{SO}(3)$ already surjects: varying A infinitesimally directly corresponds to applying an infinitesimal rotation to S^2 , perturbing $f(p, A)$ in that direction.

As a submersion, f is consequently transverse to the equator $Z \subset S^2$. By parametric transversality, almost every $f(\cdot, A)$ is transverse to Z . For such $f(\cdot, A)$, the preimage of Z , which is the crease set \mathcal{C} corresponding to ε_A , is a smooth submanifold of S_g .

In this case, $\pi \circ \varepsilon$ restricts to a smooth map on \mathcal{C} . When $d(\pi \circ \varepsilon|_{\mathcal{C}})$ has rank 0, the tangent line to \mathcal{C} in $\varepsilon(S_g)$ is vertical (with respect to π). These are precisely the points where $\pi \circ \varepsilon(\mathcal{C})$ may fail to be smooth. Notice that unless the local picture of $\varepsilon(S_g)$ is like that of the right-side illustration in Fig. 1, ε can be perturbed by a generic small rotation to remove this rank-0 point. A rank-0 point cannot be eliminated precisely when it corresponds to a change in the *sign* of the crease set (see §2.2); we call such a point a *corner* of \mathcal{C} .

Definition 2.2. An smooth embedding $\varepsilon : S_g \rightarrow \mathbb{R}^2 \times \mathbb{R}$ is *regular* with respect to π if it satisfies the following conditions:

- (1) the Gauss map of $\varepsilon(S_g)$ is transverse to the equator of S^2 ;
- (2) the differential $d(\pi \circ \varepsilon|_{\mathcal{C}})$ has rank 1 except for a finite collection of points, which are all corners; and
- (3) the immersed, piecewise smooth $\pi \circ \varepsilon(\mathcal{C})$ self-intersects only in transverse double points.

The preceding discussion shows regular embeddings are generic. Unless otherwise stated, for the remainder of this paper all embeddings will be regular. We consider $\mathcal{C} \subset S_g$ to be a collection of smooth s.c.c.'s with marked points corresponding to the corners.

2.2. Folding sign of a crease curve. Since any embedding of an orientable closed surface into Euclidean 3-space is the boundary of a unique compact 3-submanifold, M_S , we have a well-defined orientation for $\varepsilon(S_g)$ coming from the outward pointing normal vector field with respect to this submanifold. For this natural orientation of $\varepsilon(S_g)$, in the neighborhood $N_x \subset S_g$ of a rank-1 point $x \in \mathcal{C}$ there are two possible embeddings. As illustrated in Fig. 4, either the outward pointing normal vectors are “pointing towards each other” or “away from each other”. For the former situation we say the *crease folds negatively* (or “−”) *along the segment*. (See left illustration of Fig. 4.) For the latter situation the *crease folds positively* (or “+”) *along the segment*. (See right illustration of Fig. 4.)




FIGURE 4. On the left, a positive fold, and on the right, a negative fold.

Alternatively, we can assume that $N_x \subset S_g$ has been chosen small enough neighborhood of $x \in S_g$ so that $\pi \circ \varepsilon(\mathcal{C}) \cap \pi \circ \varepsilon(N_x \setminus \mathcal{C}) = \emptyset$. Let $y \in \pi \circ \varepsilon(N_x \setminus \mathcal{C})$ and $\pi^{-1}(y) \cap \varepsilon(N_x) = \{x_1, x_2\} \subset \varepsilon(N_x \setminus \mathcal{C})$. Let $\overline{x_1 x_2} \subset \mathbb{R}^2 \times \mathbb{R}$ be the compact line segment for which $\pi(\overline{x_1 x_2}) = y$. Then $\varepsilon(N_x \cap \mathcal{C})$ is a negatively folding segment if $\overline{x_1 x_2} \cap M_S = \{x_1, x_2\}$. If $\overline{x_1 x_2} \cap M_S = \overline{x_1 x_2}$ then $\varepsilon(N_x \cap \mathcal{C})$ is a positively folding segment.

Note that along arcs of rank-1 points of \mathcal{C} , the folding direction is constant; it can only switch at rank-0 points. Any component of \mathcal{C} that is without corners is either a positively folding s.c.c. or a negatively folding s.c.c., and more generally a component passes through an even number of corners. We denote the collection of all positive arcs and components \mathcal{C}^+ and similarly all negative arcs and components \mathcal{C}^- .

The reader should observe that the folding sign of either a segment in \mathcal{C} between two corners, or a corner-free component of \mathcal{C} is an invariant of an embedding's regular isotopy class.

2.3. Partial turning numbers and external angles. We start by defining the “partial turning number” of an oriented smooth arc in \mathbb{R}^2 . Let $\alpha : [0, 1] \hookrightarrow \mathbb{R}^2$ be a smooth unit speed arc and define $\alpha'(0) := \lim_{s \rightarrow 0^+} \alpha'(s)$ and $\alpha'(1) := \lim_{s \rightarrow 1^-} \alpha'(s)$. Then $\alpha'(s) = (\cos(s(t)), \sin(s(t)))$ for some angle function $s(t)$. Then the *partial turning number* is

$$t_p(\alpha) := \frac{s(1) - s(0)}{2\pi} \in \mathbb{R}.$$

Note that in the case that α is a smooth closed curve— $\alpha(0) = \alpha(1)$ and $\alpha'(0) = \alpha'(1)$ —this definition agrees with the usual definition of the turning number:

$$t(\alpha) = t_p(\alpha) = \frac{s(1) - s(0)}{2\pi} \in \mathbb{Z}.$$

We observe that the above definitions are independent of parameterization, and are thus well defined.

Fix an orientation of $\mathbb{R}^2 \subset \mathbb{R}^2 \times \mathbb{R}$. Let $K \subset S_g \setminus \mathcal{C}$ be a connected component. We consider the closure of K , that is $\bar{K} = K \cup \partial K$ where $\partial K \subset \mathcal{C}$. (We will be utilizing this association between K and \bar{K} throughout this paper and \bar{K} will always be notation for a compact connected surface.) The orientation of \mathbb{R}^2 pulls back under $\pi \circ \varepsilon$ to give an orientation of \bar{K} ; this extends to an orientation on each curve $\gamma \subset \partial K$. The reader should observe that γ is necessarily a boundary component of two distinct components K, K' of $S_g \setminus \mathcal{C}$, and both will induce the same orientation on γ . Thus we can consider \mathcal{C} as a collection of oriented curves.

Next, we consider the “external angle” of a corner $x \in \gamma \subset \mathcal{C}$. Let $\alpha, \beta \subset \pi \circ \varepsilon(\gamma)$ be adjacent smooth arcs to $\pi \circ \varepsilon(x)$. For all such corners we require the technical assumption that the angle between α and β at $\pi \circ \varepsilon(x)$ be $\frac{\pi}{2}$ —the angle between their corresponding endpoint tangent vectors. The reader should observe that this assumption is not limiting since we can perform an isotopy (and, indeed, a regular isotopy) of $\varepsilon(S_g)$ in a neighborhood of each corner to obtain this technical assumption.

We consider the embedding of the two components of $N_x \setminus \mathcal{C}$ when x is a corner of \mathcal{C} . In Fig. 5 we have an illustration of the two embedded components of $\varepsilon(N_x \setminus \mathcal{C})$. We are interested in computing the *external angle with respect to K* , $\angle_K(x)$, at $\pi \circ \varepsilon(x)$, where K is the subsurface containing a given component of $N_x \setminus \mathcal{C}$. As shown in Fig. 5, this is the angular “swing” between the tangent vector coming into $\pi \circ \varepsilon(x)$ and the tangent vector coming out of $\pi \circ \varepsilon(x)$. As we have fixed the two arcs meeting at $\pi \circ \varepsilon(x)$ to meet at an angle of $\frac{\pi}{2}$, there are two possible values for $\angle_K(x)$: $+\frac{\pi}{2}$, as in Fig. 5a, and $-\frac{3\pi}{2}$, shown in Fig. 5b. (Here we are making the positive/negative sign assignments in the classical manner—counterclockwise is positive, clockwise is negative.)

An alternative approach to determining the external angle at a corner is to take a push-off of $\varepsilon(N_x \cap \mathcal{C})$ into the two pieces of $\varepsilon(N_x \setminus \mathcal{C})$ and compute the partial turning numbers of the projected arcs. In particular, let $\alpha^x = N_x \cap \mathcal{C}$ be a crease segment with corner x . We consider the two smooth push-offs, $\hat{\alpha}_{+1/2}^x$ and $\hat{\alpha}_{-3/2}^x$, into the two components of $N_x \setminus \alpha^x$. In both illustrations of Fig. 5, the oriented brown curve depicts such a push-off, with $\pi \circ \varepsilon(\hat{\alpha}_{+1/2}^x)$ on the left and $\pi \circ \varepsilon(\hat{\alpha}_{-3/2}^x)$ on the right. Then $t_p(\hat{\alpha}_{+1/2}^x) = +\frac{1}{4}$, since the tangent vector of




FIGURE 5. A top-down view of the two layers of $\varepsilon(N_x)$ at a corner x . The **green** segments depict the portion of N_x that is in \mathcal{C} . The **blue** vectors correspond to the associated external angles. The **dark green** arcs correspond to the push-offs of \mathcal{C} into N_x .

$\pi \circ e(\hat{\alpha}_{+1/2}^x)$ has a counterclockwise $\frac{1}{4}$ -turn; and, the partial turning number is $t_p(\hat{\alpha}_{-3/2}^x) = -\frac{3}{4}$ since the tangent vector of $\pi \circ \varepsilon(\hat{\alpha}_{+3/2}^x)$ has a clockwise $\frac{3}{4}$ -turn. The external angle of x with respect to a component of $N_x \cap \mathcal{C}$ is then 2π times the associated partial turning number.

2.4. Weighting of the crease set. Let $K \subset S_g \setminus \mathcal{C}$ be a connected component. Let $\gamma \subset \partial K$ be a boundary component with corners $\{x_1, \dots, x_n\} \subset \gamma$ connecting smooth subarcs $\{\alpha_1, \dots, \alpha_n\} \subset \gamma$. Then we define the *Gauss-Bonnet weight with respect to K* of γ to be

$$t_K(\gamma) = \sum_{1 \leq i \leq n} t_p(\pi \circ \varepsilon(\alpha_i)) + \frac{\angle_K(x_i)}{2\pi}.$$

When γ has no corners, this is the turning number of $\pi \circ \varepsilon(\gamma)$. The reader should observe that Gauss-Bonnet weights are well-defined, independent of our orientation choice for \mathbb{R}^2 .


An equivalent method of defining $t_K(\gamma)$ is to take a smooth push-off of γ into K , $\hat{\gamma} \subset K$, that inherits its orientation from γ ; then $t_K(\gamma) := t(\pi \circ \varepsilon(\hat{\gamma}))$. This is effectively applying the alternative definition of \angle_K to each corner along γ simultaneously. In particular, this shows that t_K always takes integral values and, moreover, does not depend on our choice to fix the corner angles of $\pi \circ \varepsilon(\mathcal{C})$ as $\frac{\pi}{2}$.

2.5. The decorated surface. We now consolidate the data that we have developed in previous sections into a *decorated surface*, $(S_g, \mathcal{C}, \mathbf{T})$, associated with an embedding, $\varepsilon : S_g \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$. Specifically, \mathbf{T} denotes the following function.

Turning number function: $\mathbf{T} : \mathcal{C} \hookrightarrow \mathbb{Z}^2$. Let $\gamma \subset \mathcal{C}$ and consider the distinct connected components $K, K' \subset S_g \setminus \mathcal{C}$ such that $\gamma \subset \partial \bar{K}$ and $\gamma \subset \partial \bar{K}'$. Then $\mathbf{T} : \gamma \mapsto (t_K(\gamma), t_{K'}(\gamma)) \in \mathbb{Z}^2$, the Gauss-Bonnet weights of γ with respect to K and K' . If γ has no corners then $t_K = t_{K'}$ and we simplify to a 1-tuple.

3. REALIZATION COMING FROM LINK COMPLEMENTS

Let $L : S^1 \sqcup \cdots \sqcup S^1 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$ be an embedding of a link such that $(\pi \circ L)(S^1 \sqcup \cdots \sqcup S^1) \subset \mathbb{R}^2$ is a regular projection of a link L . We now augment this link projection by replacing a small disk neighborhood of each crossing with a “crossing ball” as shown in Fig. 6a. Let \mathbf{R}_+^2 (respectively, \mathbf{R}_-^2) be \mathbb{R}^2 with each disk neighborhood of a crossing replaced by the upper (respectively, lower) hemisphere of a crossing ball. Let \mathbf{R}_+^3 (respectively, \mathbf{R}_-^3) be upper half-space (respectively, lower half-space) that has \mathbf{R}_+^2 (respectively, \mathbf{R}_-^2) as its boundary plane. The salient feature of this construction is that there is now an embedding, $\bar{L} : S^1 \sqcup \cdots \sqcup S^1 \hookrightarrow \mathbf{R}_+^2 \cup \mathbf{R}_-^2 \subset \mathbb{R}^2 \times \mathbb{R}$ such that $\pi \circ L = \pi \circ \bar{L}$. We will use \bar{L} to denote of the embedding of the link into $\mathbf{R}_+^2 \cup \mathbf{R}_-^2$ for the remainder of this discussion. The reader should observe that L and \bar{L} are isotopic to each other in $\mathbb{R}^2 \times \mathbb{R}$.

FIGURE 6. How an essential surface lies near a crossing of L .

3.1. Normal form. We consider an essential closed surface, $S \subset \mathbb{R}^3 \setminus \bar{L}$. By general position arguments we can assume that $S \cap \mathbf{R}_\pm^2$ is a collection of s.c.c.’s and that any nonempty intersection of S with a crossing ball is a collection of “saddle” disks. (See Fig. 6b.) Moreover, by requiring S to be an essential surface we can assume that $S \cap \mathbf{R}_+^3$ and $S \cap \mathbf{R}_-^3$ are collections of disks. The embedding of S is in *normal form* with respect to \mathbf{R}_\pm^2 if the following three conditions hold:

- (1) every s.c.c. of $S \cap \mathbf{R}_\pm^2$ intersects at least two distinct crossing balls;
- (2) for a crossing ball \mathbf{B} and s.c.c. $\gamma \in S \cap \mathbf{R}_\pm^2$, γ intersects \mathbf{B} in at most 2 arcs. If γ meets some \mathbf{B} in two arcs, they are on the boundary of a common saddle disk in \mathbf{B} .
- (3) $S \setminus (S \cap \mathbf{R}_\pm^2)$ is a collection of open disks in \mathbf{R}_\pm^3 and saddle disks in crossing balls. For any disk component $\Delta_\gamma \subset S \cap \mathbf{R}_\pm^3$, with boundary curve $\gamma \in S \cap \mathbf{R}_\pm^2$, we require π restricts to a diffeomorphism from Δ_γ onto its image, the disk that $\pi(\gamma) \subset \mathbb{R}^2$ bounds. (The reader should observe that requiring $S \setminus (S \cap \mathbf{R}_\pm^2)$ to be a collection of disks is natural due to the incompressibility of the essential surface S .)

Using general position arguments that are well known in the literature and a number of graduate texts (see for example [H-T-T1, H-T-T2, Li, M1, M2]), one can establish the following theorem. Thus, due to the repetitive nature of any proof we offer the following result without argument.

Theorem 3.1. *Let $L \subset \mathbb{R}^2 \times \mathbb{R}$ be a link and $\bar{L}, \mathbf{R}_\pm^2$ be the associated spaces as described above. Let $F \subset (\mathbb{R}^2 \times \mathbb{R}) \setminus L$ be an closed incompressible surface in the link’s exterior. Then there exists a closed incompressible surface $\bar{F} \subset (\mathbb{R}^2 \times \mathbb{R}) \setminus \bar{L}$ such that:*

- (a) \bar{F} is in normal form with respect to the \mathbf{R}_\pm^2 ;

- (b) $\chi(F) = \chi(\bar{F})$;
- (c) if L is a non-split link then (L, F) is pairwise isotopic to (\bar{L}, \bar{F}) in $\mathbb{R}^2 \times \mathbb{R}$;

We caution that not all surfaces in normal form with respect to a link projection are incompressible/essential, despite being “captured” by a link projection. For examples, projections of *hard unknots* ([G, K-L]) will have their non-essential peripheral torus in normal form with respect to their projection.

The salient feature of such a surface in normal form is that the curves of crease set live in the 3-valent graph, $\Gamma = S \cap (\mathbf{R}_+^2 \cup \mathbf{R}_-^2)$. This can be seen from condition (3) above. Since for each open disk $\Delta \subset S \setminus \Gamma$ we have $\pi|_\Delta$ is a diffeomorphism onto an open disk in \mathbb{R}^2 , we can place a flat structure on the disks of $S \setminus \Gamma$. Thus, the crease set of a surface in normal form will be exactly the set of edges of Γ where the a flat structure of a disk cannot be extended across an edge to an adjacent disk of $S \setminus \Gamma$.

The condition for deciding which edges of Γ live in \mathcal{C} is straight forward. Let $\alpha \subset S \cap (\mathbf{R}_+^2 \cap \mathbf{R}_-^2) \subset \Gamma$ be an arc which will necessarily lie away from crossing balls. Then α will be adjacent to two disks $\Delta_+ \subset S \cap \mathbf{R}_+^3$ and $\Delta_- \subset S \cap \mathbf{R}_-^3$. If $\pi(\text{int } \Delta_+) \cap \pi(\text{int } \Delta_-) \neq \emptyset$ near α then α is in the crease set of S . (Here, $\text{int } \Delta_\pm$ is the interior of the disks.) Similarly, consider an arc $\alpha \subset S \cap \partial \mathbf{B}$, where \mathbf{B} is a crossing ball. Then α is adjacent to a disk $\Delta \subset S \cap (\mathbf{R}_+^3 \cup \mathbf{R}_-^3)$ and a saddle disk $\Delta_{\mathbf{B}} \subset \mathbf{B}$. If $\pi(\text{int } \Delta) \cap \pi(\text{int } \Delta_{\mathbf{B}}) \neq \emptyset$ near α , then α is in the crease set of S .

3.2. Proof of Theorem 1.5. Since $\varepsilon(S_g) \subset \mathbb{R}^2 \times \mathbb{R}$ is an orientable surface, we choose an orientation and we let \mathbf{N} be the associate oriented normal bundle.

Next, for each component, $\gamma \in \mathcal{C}$, we take the two smooth normal push-offs, γ_+ in the N -direction and γ_- in the $-N$ -direction. After a small isotopy that places these push-offs in position so that we obtain a regular link projection under π , it is readily confirmed that $\varepsilon(S_g)$ is in normal form with respect to L and $\pi \circ \varepsilon(\mathcal{C}) \subset \pi(\Gamma) = \varepsilon(S_g) \cap (\mathbf{R}_+^2 \cup \mathbf{R}_-^2)$.

Once a single link projection $\pi(L)$ is obtained such that $\pi \circ \varepsilon(\mathcal{C}) \subset \Gamma$, an infinite family of links is obtained by letting strands that intersect a common region of $(\mathbf{R}_+^2 \cup \mathbf{R}_-^2) \setminus \Gamma$ of the link entangle each other. (See the tangle T in Fig. 3.)

4. GAUSS-BONNET EQUATIONS

In this section we will show how the classical Gauss-Bonnet Theorem from differential geometry governs the behavior of the turning number function of a decorated surface.

Specifically, let (F, \mathbf{g}) be a compact oriented Riemannian surface with boundary. (For a development of this topic please see [L].) We let \mathbf{K} denote Gaussian curvature and k_g denote geodesic curvature. We will be considering the case where the components of ∂F are piecewise smooth, so we let $\{v_1, \dots, v_m\} \subset \partial F$ be the collection of boundary vertices. We denote the external angle at a vertex by $\theta(v_i)$, $1 \leq i \leq m$. Then the Gauss-Bonnet Theorem states:

$$(4.1) \quad \int_F \mathbf{K} \, dA + \int_{\partial F} k_g \, ds + \sum_{v_i \in \partial F} \theta(v_i) = 2\pi \cdot \chi(F).$$

For a regular embedding $\varepsilon : S_g \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$, we recall the discussion of the subsurface $\bar{K} \subset S_g$ from §2.4. We now adapt Equ. 4.1 to our study of regular embeddings of surfaces into $\mathbb{R}^2 \times \mathbb{R}$.

Proposition 4.1. *Given a regular embedding $\varepsilon : S_g \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$ and an associated connected subsurface, $K \subset S_g \setminus \mathcal{C}$, for the restricted regular embedding $\varepsilon : \bar{K} \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$, we have*

$$(4.2) \quad \sum_{\gamma \in \partial \bar{K}} t_K(\gamma) = \chi(\bar{K}).$$

Proof. We start by giving \bar{K} the Riemannian metric that corresponds to pulling back the flat Riemannian metric of \mathbb{R}^2 to \bar{K} by the map $(\pi \circ \varepsilon)^{-1}$. Then the Gaussian curvature on \bar{K} is $\mathbf{K} \equiv 0$. Moreover, for any smooth arc, $\alpha \subset \partial \bar{K}$, the geodesic curvature, $k_g(x)$, $x \in \alpha$, corresponds to the signed curvature, $k_s(\pi \circ \varepsilon(x))$. That is, for the smooth arc $\pi \circ \varepsilon(\alpha) (\subset \mathbb{R}^2)$, the signed curvature $k_s|_{\pi \circ \varepsilon(x)}$ is equal to the classical Frenet-Serret curvature times ± 1 —it is $+1$ if the Frenet-Serret normal vector equals the inward pointing normal to $\pi \circ \varepsilon(\bar{K})$ at $\pi \circ \varepsilon(x)$, and -1 otherwise.

With this convention it is a result of classical differential geometry that

$$\int_{\alpha} k_g ds = \int_{\pi \circ \varepsilon(\alpha)} k_s ds = 2\pi \cdot t_p(\pi \circ \varepsilon(\alpha)).$$

A curve $\gamma \in \partial \bar{K}$ is piecewise smooth with this metric. It fails to be smooth at the corners of γ , and at such a corner x , the external angle is $\theta(\pi \circ \varepsilon(x)) = \angle_K(x)$. It follows for a boundary curve with corners, $\gamma \subset \partial \bar{K}$, with smooth arcs $\{\alpha_i\}_{i \in \mathcal{I}}$ and corners $\{x_i\}_{i \in \mathcal{I}}$,

$$(4.3) \quad \begin{aligned} \int_{\gamma} k_g ds + \sum_{i \in \mathcal{I}} \theta(x_i) &= \sum_{i \in \mathcal{I}} \left[\int_{\pi \circ \varepsilon(\alpha_i)} k_s ds + \theta(\pi \circ \varepsilon(x_i)) \right] \\ &= \sum_{i \in \mathcal{I}} [2\pi \cdot t_p(\pi \circ \varepsilon(\alpha_i)) + \angle_K(x_i)] \\ &= 2\pi \sum_{i \in \mathcal{I}} \left[t_p(\pi \circ \varepsilon(\alpha_i)) + \frac{\angle_K(x_i)}{2\pi} \right] \\ &= 2\pi \cdot t_K(\gamma). \end{aligned}$$

Then setting $\mathbf{K} \equiv 0$, it follows from Equ. 4.1 that

$$(4.4) \quad \begin{aligned} \int_{\partial \bar{K}} k_g ds + \sum_{i \in \mathcal{I}} \theta(x_i) &= 2\pi \sum_{\gamma \in \partial \bar{K}} t_K(\gamma) \\ &= 2\pi \cdot \chi(\bar{K}). \end{aligned}$$

The equality of our proposition then follows. \square

Remark 4.2. Following the alternative method of computing $t_K(\gamma)$ described in §2.4, can restate Equ. 4.2 as

$$(4.5) \quad \sum_{\gamma \in \partial \bar{K}} t(\pi \circ \varepsilon(\hat{\gamma})) = \chi(\bar{K}).$$

We are now situated to establish Equ. 1.1 mentioned in §1.

Theorem 4.3. *Let $\varepsilon : S_g \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$ be a regular embedding with crease set \mathcal{C} . Let $\{K_1, \dots, K_l\} \subset S_g$ be the collection of all the connected subsurface components of $S_g \setminus \mathcal{C}$. Then*

$$(4.6) \quad \sum_{1 \leq i \leq l} \sum_{\gamma \in \partial \bar{K}_i} t_{K_i}(\gamma) = \chi(S_g).$$

Moreover, when every component of \mathcal{C} is without corners then Equ. 4.6 implies the following equality.

$$(4.7) \quad 2 \sum_{\gamma \in \mathcal{C}} t(\gamma) = \chi(S_g).$$

In particular, when $g = 0$ and \mathcal{C} has no corners, twice the sum of all the turning numbers of s.c.c.'s of \mathcal{C} is equal to 2.

Proof. From classical topology it is a well known fact that the Euler characteristic is additive in the following sense. Let F_1 and F_2 be two orientable compact surfaces with boundary. Let $A \subset \partial F_1$ and $B \subset \partial F_2$ be two collections of boundary components such that $|A| = |B|$. For any homeomorphism, $f : A \hookrightarrow B$, consider the surface, $F_1 \cup_f F_2$, obtained by gluing F_1 to F_2 via f . Then $\chi(F_1 \cup_f F_2) = \chi(F_1) + \chi(F_2)$.

Applying this fact to S_g , we have the following sequence of equalities:

$$(4.8) \quad \begin{aligned} \chi(S_g) &= \chi(\bar{K}_1) + \cdots + \chi(\bar{K}_l) \\ &= \sum_{\gamma \in \partial \bar{K}_1} t_{K_1}(\gamma) + \cdots + \sum_{\gamma \in \partial \bar{K}_l} t_{K_l}(\gamma) \\ &= \sum_{1 \leq i \leq l} \sum_{\gamma \in \partial \bar{K}_i} t_{K_i}(\gamma) = \chi(S_g). \end{aligned}$$

Our first equality then follows.

The claimed equality for the situation where there are no corners follows from the observation that each $\gamma \in \mathcal{C}$ is a boundary curve of two K_i 's, and γ has the same Gauss-Bonnet weight for each. \square

We include the special case of the sphere in our statement since we will be restricting to that case alone for the remainder of the paper.

4.1. Proof of Theorem 1.1. The claim of the theorem can be restated as follows. Given a pair (S^2, \mathcal{C}) that can be geometrically realized by some regular embedding, there is a unique turning number function $\mathbf{T} : \mathcal{C} \hookrightarrow \mathbb{Z}^2$, that is, \mathbf{T} is independent of embedding.

We have three key observations that will be combined to establish this result. The first observation is that every s.c.c. on S^2 is a separating curve.

The second observation follows from the first, we can determine one of the Gauss-Bonnet weights of a component, $\gamma \in \mathcal{C}$, that is innermost on S^2 . That is, there will always exist a connected subsurface component with boundary, $\bar{K} \subset S^2$, that is homeomorphic to a 2-disk. With $\gamma = \partial \bar{K}$ we will have

$$+1 = \chi(\bar{K}) = t_K(\gamma),$$

and thus one of γ 's Gauss-Bonnet weights.

The third observation is that if we can determine one of the two Gauss-Bonnet weights for a s.c.c., $\gamma \in \mathcal{C}$, then we can determine the other. To see this, let \bar{K}, \bar{K}' be two connected subsurface components with boundary that share a common boundary curve. That is, we have a curve with corners $\bar{K} \cap \bar{K}' = \gamma \in \mathcal{C}$. We observe that if we know the value of t_K ,

we can determine the value of $t_{K'}$, since a $\frac{\pi}{2}$ corner in \bar{K} (respectively, $\frac{3\pi}{2}$ corner) becomes a $\frac{3\pi}{2}$ corner (respectively, $\frac{\pi}{2}$ corner) in \bar{K}' , while the partial turning numbers of smooth arcs remain the same.

Now combining these three observations, suppose there is a component of \mathcal{C} for which \mathbf{T} is undetermined. Again, let $\{\bar{K}_1, \dots, \bar{K}_l\}$ be the collection of all connected subsurface components coming from $S^2 \setminus \mathcal{C}$. Suppose there is a subsurface, \bar{K}_i , such that t_{K_i} is undetermined for one of its boundary components. If such a component exists then we can choose \bar{K}_i such that it has only one boundary component with t_{K_i} undetermined. But, using Equ. 4.2, we can solve for the “undetermined” Gauss-Bonnet weight associated to \bar{K}_i —we will know $\chi(\bar{K}_i)$ and all but one of the associated Gauss-Bonnet weights.

The theorem then follows.

4.2. Proof of Theorem 1.2. Once the number of components and corners of \mathcal{C} is bounded by n_t and n_c , respectively, the finiteness claim follows from topological bookkeeping arguments.

First, given a collection of disjoint s.c.c.’s, $\mathcal{C} \subset S^2$, we can naturally associate to it a graph. Each vertex of the graph will correspond to a connected component of $S^2 \setminus \mathcal{C}$ and the edges of the graph will correspond to the components of \mathcal{C} —two vertices associated to two subsurface components of $S^2 \setminus \mathcal{C}$ share a common edge if the curve of \mathcal{C} associated to that edge is a common boundary curve of the two subsurfaces. Since every s.c.c. in S^2 is separating the graph associated with a pair (S^2, \mathcal{C}) will be a tree. It is well known that the number of distinct tree-like graphs having at most n_t edges is a bounded set. (See [C].)

To take into account corners of \mathcal{C} , note that given a collection of disjoint s.c.c.’s, $\mathcal{C} \subset S^2$, there is only a finite number of ways to distribute no more than n_c marked points along the curves of \mathcal{C} .

5. CONSTRUCTING EMBEDDINGS OF S^2

In this section we revisit question (2) of §1: *given a collection of s.c.c.’s $\mathcal{C} \subset S^2$, when is \mathcal{C} realized as the crease set of a regular embedding into $\mathbb{R}^2 \times \mathbb{R}$?* We will answer this question in the simplified case where the crease set has no corners. From Theorem 1.1 we know that if such an embedding exists, the turning number function, \mathbf{T} , decorating the pair (S^2, \mathcal{C}) is uniquely determined. Since such a \mathbf{T} is totally determined by the topological information of the components of $S^2 \setminus \mathcal{C}$, one can readily decide whether such a function exists.

Specifically, let $(S^2, \mathcal{C}, \mathbf{T})$ be a triple satisfying the following $(*)$ conditions:

- (*-i) $\mathcal{C} \subset S^2$ is a collection of pairwise disjoint smooth s.c.c.’s.
- (*-ii) $\mathbf{T} : \mathcal{C} \hookrightarrow \mathbb{Z}$ such that for each connected surface $K \subset S^2 \setminus \mathcal{C}$, the associated surface with boundary, \bar{K} , satisfies the relationship

$$\sum_{\gamma \in \partial \bar{K}} \mathbf{T}(\gamma) = \chi(\bar{K}).$$

With the above statement of $(*)$ in mind, we now give an equivalent restatement of Theorem 1.3, which we will proceed to prove in the remainder of §5.

Theorem 5.1. *Let $(S^2, \mathcal{C}, \mathbf{T})$ be a triple satisfying $(*)$. Then there exists a regular embedding, $\varepsilon : S^2 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$, such that \mathcal{C} corresponds to the crease set of ε and \mathbf{T} corresponds to the turning number function associated with ε . In particular, the crease set has no corners.*

Our constructive argument has two steps. First, we develop an understanding of representing embeddings of flat planar surfaces in $\mathbb{R}^2 \times \mathbb{R}$. Second, we show how these representations can be “stacked” together to construct the needed embedding of S^2 .

5.1. Constructing planar surfaces \bar{K} . We start with the following definition.

Definition 5.2. An n -turning weight, W , is an $(n+1)$ -tuple of odd integers, $W_s = (t_0, \dots, t_n)$, such that

$$\sum_{i=0}^n t_i = 1 - n.$$

The *standard* n -turning weight set is

$$W_s = (1, \underbrace{-1, -1, \dots, -1}_n).$$

A *basis* for an n -turning weight, $W = \{t_0, \dots, t_n\}$, is a collection of integers, $\{k_0, \dots, k_n\}$, such that:

1. $t_0 = 1 + 2k_0$,
2. $t_i = -1 + 2k_i$, for $1 \leq i \leq n$.
3. Necessarily,

$$\sum_{i=0}^n k_i = 0.$$

Let \bar{K}_{1-n} be a smooth compact planar surface having Euler characteristic equal to $1 - n$. Let $W = (t_0, \dots, t_n)$ be an n -turning weight. We say an embedding $\varepsilon : \bar{K}_{1-n} \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$ is a *geometric realization of W* if:

- i. every point of $\text{int } \bar{K}_{1-n}$ is a regular point. Thus, the pull-back of $\pi \circ \varepsilon(\bar{K}_{1-n})$ gives a flat structure to \bar{K}_{1-n} .
- ii. There is an enumeration of the boundary components, $\partial \bar{K}_{1-n} = \{\gamma_0, \gamma_1, \dots, \gamma_n\}$, such that we have $t(\gamma_i) = t_i$ for $0 \leq i \leq n$. That is, the turning numbers of the components of $\partial \bar{K}_{1-n}$ realize W .

Consider a closed disk in $\mathbb{R}^2 \times \{\text{pt.}\}$. Delete from it n open subdisks with disjoint closures. The resulting “disk with n holes” will be an embedding of a \bar{K}_{1-n} that is a geometric realization of the standard n -turning weight, W_s . We will use the notation $\varepsilon_s : \bar{K}_{1-n} \hookrightarrow \mathbb{R}^2 \times \{\text{pt.}\}$ for this *standard embedding* of a planar surface of Euler characteristic $1 - n$. Observe that we can easily require that the Frenet-Serret curvature be nowhere zero on $\partial \varepsilon_s(\bar{K}_{1-n})$.

We now give a procedure for taking the standard embedding, $\varepsilon_s(\bar{K}_{1-n})$, and altering it so that we obtain a geometric realization of any chosen n -turning weight, W . Let $\{\gamma_0, \dots, \gamma_n\}$ be a fixed enumerated labeling of the boundary components $\partial \bar{K}$. Let $\{k_0, \dots, k_n\}$ be a basis for W . For each boundary component, γ_i , we chose $|k_i|$ distinct points, $p_i^1, \dots, p_i^{|k_i|} \subset \gamma_i$, and label them “+2” (respectively, “-2”) when $k_i > 0$ (respectively, when $k_i < 0$). (To keep down the clutter in our figures, we will use a dot, \bullet , for +2 label and a circle, \circ , for -2 label.) Note that all the labels on a particular boundary component are the same. See the initial illustration in the sequence in Fig. 7.

Next we choose a collection of properly embedded pairwise disjoint arcs, $\{a_j\} \subset \bar{K}_{1-n}$ such that each arc has one endpoint labeled -2 and the other endpoint labeled +2. Thus, each arc has a natural orientation from its negative to positive boundary endpoints. We will refer to such properly embedded oriented arcs as *twisting arcs of \bar{K} associated with W* .

Three observations are warranted. First, we must have

$$l = \frac{1}{2} \sum_{i=0}^n |k_i|$$

many arcs. Second, since the labels on a particular boundary component are homogeneous, the two endpoints of any twisting arc are on different boundary components. Finally, there is a great deal of choice in the embedding of the twisting arcs. Although it is not necessary, for convenience of description we will assume that these arcs are linear intervals. See the starting illustration in the sequence of Fig. 7.




FIGURE 7. The top sequence illustrates how the twisting operation in a neighborhood of the two oriented arcs takes a planar surface that realizes the standard 2-turning weight set, $(1, -1, -1)$, to a realization of the 2-turning weight set, $(-3, 1, 1)$. This top sequence continues with an isotopy of the resultant boundary so that $k_s \neq 0$ at every point on the **red boundary curves**. The bottom sequence illustrates our twisting operation in a neighborhood of an oriented arc.

Next, given a choice of twisting arcs, $\{a_j\}_{1 \leq j \leq l} \subset \bar{K}$, in a standard embedded planar surface we can perform the *twisting operation* illustrated in Fig. 7 in a neighborhood of each a_j . The resulting planar surface embedding will be a geometric realization of the associated n -turning weight set, W . The top sequence in Fig. 7 illustrates this procedure for the taking the standard embedding of a disk-minus-two-holes and obtaining an embedding associated with the 2-weight set, $(-3, 1, 1)$.

The reader should observe that the basis used for the $(-3, 1, 1)$ geometric realization is $(k_0, k_1, k_2) = (-4, 1, 1)$. However, one could also utilize the basis $(k_0, k_1, k_2) = (0, -1, 1)$. Utilizing this “change of basis” one will have a single properly embedded arc whose endpoints are on the two -1 boundaries of starting planar surface in the top-sequence of Fig. 7. The

2-turning weight then becomes $(1, -3, 1)$ or $(1, 1, -3)$, depending on the ordering of the boundary curves.

The above discussion is captured by the following proposition.

Proposition 5.3. *For any basis, $\{k_0, \dots, k_n\}$, associated with a n -turning weight set, W , there exists an embedding, $\varepsilon : \bar{K}_{1-n} \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$, that is a geometric realization of W in a manner that respects the basis. That is, the turning number of the γ_i boundary curve of \bar{K}_{1-n} is $1 + 2k_0$ when $i = 0$ and $-1 + 2k_i$ when $1 \leq i \leq n$.*

A number of remarks follow.

Remark 5.4. The reader should observe that the only possible way a boundary curve having $+1$ turning number can result from our twisting operation is if the associated basis value is either $k_0 = 0$ or $k_i = 1$, $1 \leq i \leq n$.

Remark 5.5. As illustrated in the initial planar surface of the top sequence of Fig. 7, we can readily assume that the geometric realization of a standard n -turning weighted set has its boundary curves in a “normal position”. That is, the curvature k_s is not zero at any point on the boundary. The reader should observe that the twisting operation preserves this normal position— $k_s \neq 0$ —on the γ_0 boundary curve when the associated basis value $k_0 \geq 0$.

Remark 5.6. Continuing an analysis of the behavior of the boundary curvature with respect to the twisting operation, for $2 \leq j \leq n$, if we have a basis value, $k_j \leq 0$, then $k_s \neq 0$ on the associated boundary curve after performing the twisting operation.




FIGURE 8. A depiction of the isotopy on a boundary curve of a planar component so it will be in normal form once the twist operations are performed. The isotopy fixes a neighborhood of the green oriented arcs where the twists will be done. In this case, the inner curve has $k_j = -3$, but this generalizes to any $k_j \leq -2$.

Remark 5.7. For some $1 \leq j \leq n$, if we have a basis value $k_j \leq -2$ then we can perform an isotopy that achieves $k_s \neq 0$ on the associated boundary curve after the twisting operation is applied. (See Fig. 8.)

Remark 5.8. To deal with final case for achieving a normal form for the resulting embedding of \bar{K} , Fig. 9 illustrates how the boundary curve initially having -1 turning number can achieve a $+1$ turning number after a single twist operation. This is the case where $k_j = -1$ for some $1 \leq j \leq n$. The embedding depicted in the middle illustration of the sequence has




FIGURE 9. The twist operation performed with a $k_j = -1$ curve, followed by an isotopy sliding the top portion of \bar{K} to the left, so that that $k_s \neq 0$ along the twisted blue boundary curve.

this $+1$ turning number boundary curve not in normal form since there will be two points where $k_s = 0$. The right-hand illustration shows this same boundary curve after an isotopy that places it in a position where $k_s > 0$.

5.2. Constructing embeddings and proof of Theorem 1.3.

5.2.1. *Labeled tree graphs.* We assume that we are given, (S^2, \mathcal{C}) , a 2-sphere/crease set pair. That is, $\mathcal{C} \subset S^2$ is a collection of pairwise disjoint smooth s.c.c.'s which have a Gauss-Bonnet weighting. We can associate to this pair a *labeled tree graph (LTG)*, \mathcal{G} , as follows:

1. For each connected planar component, $K \subset S^2 \setminus \mathcal{C}$, there is a corresponding vertex $v_K \in \mathcal{G}$, with label $\chi(\bar{K})$.
2. Each edge, $e_\gamma \in \mathcal{G}$, corresponds to a curve of $\gamma \in \mathcal{C}$ with the label $t(\gamma)$.
3. An edge, $e_\gamma \in \mathcal{G}$, is incident to vertices, $v_K, v_{K'} \in \mathcal{G}$, if $\gamma = \partial K \cap \partial K'$.

A LTG being a tree comes from the fact that all s.c.c.'s on a 2-sphere separate.

Example 5.9 (Part 1.). We offer a running example, $(S^2, \mathcal{C}) = (S^2, \{\gamma_1, \dots, \gamma_{13}\})$, with: $t(\gamma_i) = 1$, $i \in \{1, 2, 3, 4, 10, 11, 12, 13\}$; $t(\gamma_5) = t(\gamma_6) = t(\gamma_8) = t(\gamma_9) = -3$; and $t(\gamma_7) = +5$. Fig. 10 illustrates $\mathcal{C} \subset S^2$ and Fig. 11 depicts the associated LTG, \mathcal{G} .




FIGURE 10. A collection of curves on S^2 corresponding to \mathcal{C} . Each curve in \mathcal{C} is labeled γ_i for $1 \leq i \leq 13$. Additionally, each curve is labeled with its determined Gauss-Bonnet turning number weight.




FIGURE 11. The graph \mathcal{G} associated with Example 5.9. Labels in the vertices (circles) correspond to Euler characteristic of planar subsurface. Labels of edges correspond to the turning number of the adjoining curve of \mathcal{C} . The height axis between the two graph depictions gives the height/order for the stacking procedure that is illustrated in Fig. 14.

For a given (S^2, \mathcal{C}) configuration we can embed its associated LTG, $\mathcal{G} \hookrightarrow \mathbb{R} \times \mathbb{R} \subset \mathbb{R}^2 \times \mathbb{R}$. We use the coordinates $((x, y), z) \in \mathbb{R}^2 \times \mathbb{R}$ and $((x, y_0), z) \in \mathbb{R} \times \mathbb{R}$, where y_0 is an arbitrary, fixed value. We denote the z -height of a vertex $v \in \mathcal{G}$ by $z_v \in \mathbb{R}$, and the z -interval support of any edge $e_\gamma \in \mathcal{G}$, corresponding to $\gamma \in \mathcal{C}$, by $I_\gamma \subset \mathbb{R}$. As a tree, any LTG can be embedded in $\mathbb{R} \times \mathbb{R}$ so as to satisfy the following (\star) conditions:

- (\star) -0 For two distinct vertices, $v, v' \in \mathcal{G}$, we have $z_v \neq z_{v'}$.
- (\star) -1 Each edge is a straight line segment with slope being non-zero, i.e. $|\frac{dz}{dx}| > 0$.
- (\star) -2 Each vertex, $v \in \mathcal{G}$, having adjacent edges, $\{e_{\gamma_1}, \dots, e_{\gamma_p}\}$ for $\gamma_i \in \mathcal{C}$, $1 \leq i \leq p$, satisfies one of the following statements: (a) $\min I_{\gamma_i} \geq z_v$ iff $t(\gamma_i) > 0$; or, (b) $\max I_{\gamma_i} \leq z_v$ iff $t(\gamma_i) > 0$. (See Fig. 12.)

Regarding condition (\star) -2, the reader should observe that if $v, v' \in \mathcal{G}$ are two vertices adjacent to a common edge then one vertex satisfies (\star) -2a and the other satisfies (\star) -2b. Overall, it is an elementary inductive argument on the number of vertices that any LTG has an planar embedding satisfying condition (\star) . The reader can check that the LTG associated with our running example in (Fig. 11) satisfies (\star) .

5.2.2. *Consistent labeling of a LTG.* As in the embedding construction established in Proposition 5.3, we will initially need to designate a collection of twisting arcs in each \bar{K} , $\chi(\bar{K}) \leq 0$,




FIGURE 12. The two possible height configuration of edges at a vertex. The left (respectively, right) corresponds to (\star) -2a (respectively, (\star) -2b).

that is associated with a given (S^2, \mathcal{C}) that has a Gauss-Bonnet weighting. Such a designation of arcs is non-unique. As with the case of an individual compact connected planar surface, deciding the placement of twisting arcs in S^2 starts with deciding which curves of \mathcal{C} will correspond to the $+1$ turning number of a standard n -turning weight set. Specifically, for a (S^2, \mathcal{C}) configuration, let $C = \{\gamma_1, \dots, \gamma_m\} \subset \mathcal{C}$ be a subset collection such that for any \bar{K} , $\chi(\bar{K}) \leq 0$, associated with the configuration, exactly one curve of C lies in \bar{K} . For a subcollection C satisfying this condition we say C is a *consistent labeling* of \mathcal{C} . Reframing this in terms of the LTG, a subcollection of edges, $E(C) \subset \mathcal{G}$, is a consistent labeling if each vertex of \mathcal{G} that has a non-positive label is adjacent to exactly one edge in $E(C)$.

In our running example, the subcollection of \mathcal{C} with a \blacksquare corresponds to a consistent labeling of \mathcal{C} . See Fig. 13. Similarly, the corresponding edges of the LTG in our running example have a \blacksquare . See Fig. 11.

The salient feature of such a labeling is that if $\gamma \in \mathcal{C}$ is a crease curve that is a boundary curve of the two planar components, K, K' , then the ± 2 -labeling on γ when viewed as a boundary curve of K is the same as the ± 2 -labeling on γ when viewed as a boundary curve of K' .




FIGURE 13. The crease set, $\mathcal{C} \subset S^2$, of the running example with the enhancements of black-squares and green edge-paths of oriented arcs.

That label consistency is readily achievable follows from the *black-square tree decomposition* procedure that we now describe. Let $v \in \mathcal{G}$ be a vertex that is adjacent to a collection of edges, all but one of which carrying a “ $+1$ ” label. Note that the exceptional edge will have a negative integer as its label. We make an initial choice of edge for placing a black-square,

■: either a “+1” edge receives the square or the exceptional edge receives it. Next, we choose an “alternating” edge-path, $\mathbf{E}(v, v') \subset \mathcal{G}$, between v and another $v' \in \mathcal{G}$ that is also only adjacent to a collection of “+1” labeled edges and a single negative labeled edge. The edge-path is alternating in the sense that, referring to Fig. 12, each interior vertex of the edge-path is adjacent to one positive and one negative edge of $\mathbf{E}(v, v')$.

Now, if $\mathbf{E}(v, v')$ does not include our initial choice of the black-square edge, e , we extend it to include e . An extension would mean that the edge-path now goes between a +1 labeled vertex and v' , and we relabel this +1 vertex as our v . If our initial choice of a black-square edge is the exceptional negative labeled edge then there is nothing to do.

With this possible extension modification in place, we now traverse $\mathbf{E}(v, v')$ starting at v and traveling across the first edge which will have a black-square. We then place a \blacksquare on every other edge in the edge-path $\mathbf{E}(v, v')$. When we finally come to the ending vertex v' , if it black-square is not adjacent to an edge with a black-square then we add an extending edge that is adjacent to a +1 vertex and label this extending edge with a black-square. If this final extension is necessary, we again adjust the labeling of vertices so that v' is now a +1 labeled vertex.

Now let $e_1, \dots, e_k \subset \mathcal{G}$ be the collection of edges that are adjacent to $\mathbf{E}(v, v')$, and $T_1 \subset \mathcal{G} \setminus (\mathbf{E}(v, v') \cup \bigcup_i e_i)$ is a collection of trees—here $\mathbf{E}(v, v') \cup \bigcup_i e_i$ is the *link of $\mathbf{E}(v, v')$* . Any component of T_1 that is a single vertex will necessarily have a +1 label. For any component of T_1 that is not a single vertex we repeat the above procedure of choosing a black-square edge-path with the following proviso.

Observe that each vertex of T_1 inherits the positive/negative edge feature of Fig. 12. In our choice for v and v' we allow for the possibility that an edge path can end or begin at a vertex that has only negative/positive edge labels adjacent to it. With this in mind, we iterate the choice of a black-square edge-path in each component of T_1 . Then T_2 is obtained by removing the link of each black-square edge-path in each component of T_1 . We iterate this removal of the link of black-square edge-paths until what remains is a collection of vertices. We now reassemble \mathcal{G} to obtain a consistent labeling of \mathcal{G} .

Once we have a consistent labeling of $\mathcal{C} (\subset S^2)$ by black-squares we can label its curves with +2 green-dots and -2 green-circles. Specifically, for a curve $\gamma \in \mathcal{C}$ with a black-square having Gauss-Bonnet weight w_γ , if $w_\gamma < 0$ we place k_\circ green-circles along γ such that $w_\gamma = -1 - 2 \cdot k_\circ$. Whereas if $w_\gamma > 0$, we place k_\bullet green-dots along γ such that $w_\gamma = 1 + 2 \cdot k_\bullet$. Finally, in each \bar{K} component having both green-dots and green-circles we make a choice of disjoint twisting arcs connecting the dots to circles.

The reader should observe that the twisting arcs in S^2 define a collection of edge-paths in \mathcal{G} , and each edge-path has its endpoints on s.c.c.’s in \mathcal{C} that have Gauss-Bonnet weight +1. We return to our running example.

Example 5.10 (Part 2). Applying the above described procedure to the graph in Fig. 12, we can choose black-square designations for the subcollection, $\{\gamma_3, \gamma_5, \gamma_8, \gamma_{12}\}$, which is a consistent labeling of \mathcal{C} . Here the initial \mathbf{E} -edge-path could be the one that goes from γ_5 to γ_8 . Once the link of this edge-path is removed from the graph, we make ■-labeling choices of γ_3 and γ_{12} . We then transfer this labeling to the configuration in Fig. 13. Now referring to Fig. 13, to achieve a +1 Gauss-Bonnet weighting on curves $\gamma_1, \gamma_2, \gamma_4, \gamma_{10}, \gamma_{11}, \gamma_{12}$, we place a single +2 green-circle on each. There is no need for additional labeling of γ_3 and γ_{12} since their ■-label implies their Gauss-Bonnet weights are already +1. To achieve a -3 Gauss-Bonnet weight on γ_5 and γ_8 we label each of them with two -2 green-circles. Whereas, we

need only label γ_6 and γ_9 with a single -2 green-circle. Finally, to achieve the $+5$ Gauss-Bonnet weight on γ_7 we must label it with three $+2$ green-dots. Now connecting green-circle to green-dots in each planar component, we obtain three edge-paths.




FIGURE 14. A stacking of the planar surfaces from the running example which is consistent with the height/order of the associated vertices in Fig. 11.

5.2.3. *The height of vertices of a LTG.* For a LTG embedding, $\mathcal{G} \hookrightarrow \mathbb{R} \times \mathbb{R}$, satisfying conditions (\star) , there is still some ambiguity in establishing the heights of the vertices of \mathcal{G} . To address this issue we utilize our black-square tree decomposition to make the height assignments of the vertices of \mathcal{G} .

We first remark that a black-square decomposition results in having each edge of \mathcal{G} being in either a black-square edge-path of the decomposition or adjacent to two such edge-paths in the decomposition.

With that observation in mind, we consider the initial black-square edge-path $\mathbf{E}(v, v')$ of a decomposition. We can readily assign heights to its vertices that correspond to the order in which one encounters them as one traverses the path starting at v and ending at v' . This is due to the condition that the edge-path is alternating—height 0, height 1, and so on. Next, let $e_\gamma \in \mathcal{G}$ and $\mathbf{E}' \subset \mathcal{G}$ be an edge and black-square edge-path coming from the decomposition with $\mathbf{E}(v, v') \cap \mathbf{E}' = \emptyset$ such that e_γ is adjacent to both $\mathbf{E}(v, v')$ and \mathbf{E}' .

Let $\bar{v}, e_{\bar{\gamma}} \in \mathbf{E}(v, v')$ be a vertex and an edge that satisfy the following: \bar{v} is adjacent to both $e_{\bar{\gamma}}$ and $e_{\bar{\gamma}}$, and, $e_{\bar{\gamma}}$ and $e_{\bar{\gamma}}$ have the same \pm label. We then place the height assignment for the vertices of \mathbf{E}' within the z -interval support, $I_{\bar{\gamma}}$. Within $I_{\bar{\gamma}}$ we again assign heights to the vertices in \mathbf{E}' that correspond to the order in which they are encountered as \mathbf{E} is traversed.

In general, for two black-square edge-paths, \mathbf{E}_1 and \mathbf{E}_2 —their indices corresponding to the order in which they are chosen in the decomposition—and $e_{\bar{\gamma}}$ being the commonly adjacent edge, we will place the height assignments of the vertices of \mathbf{E}_2 within the z -interval support of the edge of \mathbf{E}_1 that shares a vertex with $e_{\bar{\gamma}}$ and has the same \pm label.

5.2.4. Stacking planar surfaces. We now utilize the embedding of the graph, \mathcal{G} , in $\mathbb{R} \times \mathbb{R}$ to guide a “stacking” of components of $S^2 \setminus \mathcal{C}$ in $\mathbb{R}^2 \times \mathbb{R}$. The setup is that we again have a configuration, (S^2, \mathcal{C}) , that has a Gauss-Bonnet weighting, and an associated embedding of its LTG, $\mathcal{G} \hookrightarrow \mathbb{R} \times \mathbb{R} (\subset \mathbb{R}^2 \times \mathbb{R})$, that satisfies conditions (\star) . Let $C \subset \mathcal{C}$ be a subcollection that corresponds to a consistent labeling of \mathcal{C} . Let \bar{K} be a compact connected planar surface coming from a connected component of $S^2 \setminus \mathcal{C}$ with $\chi(\bar{K}) \leq 0$. Finally, let $v_K \subset \mathcal{G}$ be the vertex associated with \bar{K} .

With the above setup we take $\varepsilon_s : \bar{K} \hookrightarrow \mathbb{R}^2 \times \{z_{v_K}\}$ to be a standard embedding such that for the curve, $\gamma = C \cap \partial\bar{K}$, we have $t(\varepsilon_s(\gamma)) = +1$. We have such a standard embedding for each planar component where $\chi(\bar{K}) \leq 0$. By condition (\star) -0 we know each plane $\mathbb{R}^2 \times \{\text{pt.}\}$ contains at most one planar surface.

Now suppose two distinct vertices, $v_K, v_{K'} \in \mathcal{G}$, are adjacent to an edge $e \in \mathcal{G}$. By the definition of \mathcal{G} , this implies \bar{K} and \bar{K}' share a common boundary curve, $\gamma_e \in \mathcal{C}$. We use ε_s (respectively, ε'_s) as notation for the standard embedding associated with \bar{K} (respectively, \bar{K}'). We now require that the \bar{K} -planar components associated with (S^2, \mathcal{C}) satisfy the following positioning conditions.

P1—For two planar components sharing some γ_e as a common boundary curve, we position two surfaces in their respective $\mathbb{R} \times \{\text{pt.}\}$ planes so that $\pi \circ \varepsilon_s(\gamma_e) = \pi \circ \varepsilon'_s(\gamma_e)$. Additionally, we position $\varepsilon_s(\gamma_e)$ and $\varepsilon'_s(\gamma_e)$ so that they are bijections on the set of green-circle/dots in γ_e .

P2—For two planar components sharing some γ_e as a common boundary curve, let $p \subset \gamma_e (= \partial\bar{K} \cap \partial\bar{K}')$ be a point that is either a green-circle or green-dot. Let $a \subset \bar{K}$ and $a' \subset \bar{K}'$ be two twisting arcs adjacent to p . We further position $\varepsilon_s(\gamma_e)$ and $\varepsilon'_s(\gamma_e)$ so that $\pi \circ \varepsilon_s(a) = \pi \circ \varepsilon'_s(a')$.

We claim that conditions P1 and P2 can be achieved simultaneously for all \bar{K} -planar components in their respective copies of $\mathbb{R}^2 \times \{\text{pt.}\}$. Indeed, if we allow our stacking of planar surfaces to mirror the black-square tree decomposition of \mathcal{G} , these two conditions are naturally met. That is, first stack only the \bar{K}' 's that correspond to vertices in the initial black-square edge-path $\mathbf{E}(v, v')$. For such a “linear” stacking of planar surfaces P1 and P2 are easily achieved. For the next edge-path, \mathbf{E}' , within the z -interval support of $I_{\bar{\gamma}}$, since it is also a linear stacking, again P1 and P2 are easily achieved. Iterating this mirroring of the decomposition, we obtain a stacking of planar surfaces satisfying P1 and P2.

Example 5.11 (Part 3). Returning to our running example, in Fig. 14 we show a stacking of \bar{K} , $\chi(\bar{K}) \leq 0$, components such that conditions P1 and P2 are satisfied. In particular, the double-headed arrows indicate how the projected image of the boundary curves along with the twisting arcs will coincide under the π projection map.

5.2.5. Gluing stacked curves. With the \bar{K} components ($\chi(\bar{K}) \leq 0$) associated with (S^2, \mathcal{C}) initially embedded in $\mathbb{R}^2 \times \mathbb{R}$ as describe in §5.2.4, we next glue together the two copies of




FIGURE 15. A gluing of two stacked planar components along a shared boundary curve. Note that this creates the stacking of green oriented arcs whose projection into the \mathbb{R}^2 -plane is a single oriented arc.

each embedded curve of \mathcal{C} . In particular, for any s.c.c. $\gamma \in \mathcal{C}$ with Gauss-Bonnet weight $t(\gamma) \neq +1$, there will be two planar components, \bar{K} and \bar{K}' for which $\gamma = \bar{K} \cap \bar{K}'$. As a result of our positioning conditions, **P1** and **P2**, we can perform an isotopy in $\mathbb{R}^2 \times [z_{v_{\bar{K}}}, z_{v_{\bar{K}'}}]$ that glues $\varepsilon_s(\gamma)$ and $\varepsilon'_s(\gamma)$ together. See Fig. 15.




FIGURE 16. The stack of oriented arcs projects onto a single oriented arc in the \mathbb{R}^2 -plane.

5.2.6. *Twisting the stacked planar surfaces.* Once a gluing of all $(\varepsilon_s(\gamma), \varepsilon'_s(\gamma))$ pairs of curves have been performed, we observe that condition (\star) -2 implies that there is an accordion-like-folding neighborhood of each edge-path of twisting arcs. See Fig. 16. Finally, we can perform our twisting operation to this entire edge-path neighborhood as depicted in Fig. 17. After performing these twisting operations the image of each s.c.c. of \mathcal{C} will realize its assigned Gauss-Bonnet weighting. In particular, all curves having +1 turning number can now be




FIGURE 17. Performing a twist on the stacked arcs.

placed in normal position—nowhere zero curvature—as depicted in Fig. 9 and then capped off with a disk. The interior of such disks will have empty intersection with the crease set.

6. EMBEDDINGS AND BRANCHED SURFACES

The main thrust of this section is to develop the machinery needed for the classification of isotopy classes for regular embeddings of S^2 into $\mathbb{R}^2 \times \mathbb{R}$ having a crease set of three curves without corners. Once our machinery has been developed we will carry out the classification calculation in §7.


We start by formalize our meaning of isotopy class with the following definition.

Definition 6.1. An isotopy $\varepsilon_t : S \rightarrow \mathbb{R}^2 \times \mathbb{R}$, $0 \leq t \leq 1$, is *regular* if the critical set of $\pi \circ \varepsilon_t$ is \mathcal{C} for all $t \in [0, 1]$.

Such an isotopy restricts to each subsurface \bar{K} as defined in §2.3. Note that a regular isotopy preserves the local data of the embedding, in that it cannot add or remove corners, nor change the sign of a crease curve, as defined in §2.2. However, a regular isotopy may pass through non-regular embeddings, as the assumption of transverse self-intersections of $\pi \circ \varepsilon(\mathcal{C})$ may be violated.

6.1. The natural branched surface. There is a natural branched surface with boundary, $\mathcal{B} \subset \mathbb{R}^2 \times \mathbb{R}$, associated with any regular embedding, $\varepsilon : S_g \rightarrow \mathbb{R}^2 \times \mathbb{R}$. This is due to the fact that the projection, $\pi : \mathbb{R}^2 \times \mathbb{R} \hookrightarrow \mathbb{R}^2$, induces an I -bundle fibration on M_S , the 3-manifold bound by $\varepsilon(S)$, where the fibers are the components of $\pi^{-1}(p) \cap M_S$, for any $p \in \mathbb{R}^2$. \mathcal{B} is then the quotient space, $M_S / \{I - \text{fibers}\}$. We will use $q_\varepsilon : \varepsilon(S) \rightarrow \mathcal{B}$ ($\subset \mathbb{R}^2 \times \mathbb{R}$) to denote the image of S under this quotient map.

Taking a brief aside for readers who are not familiar with branched surfaces, our \mathcal{B} will be a smooth surface with a codimension-1 branching locus singular set. Fig. 18 depicts the three neighborhood models for an interior point in \mathcal{B} . The boundary of \mathcal{B} is a compact 1-submanifold which may contain endpoints of the of the branching locus. For our discussion on the classification of embeddings of 2-sphere with three crease components, the branching locus and boundary of \mathcal{B} will be disjoint. The key point for our discussion is that at every interior point, $p \in \mathcal{B} \setminus \partial\mathcal{B}$, there is a well defined tangent plane $T_p\mathcal{B}$ for which the differential,

FIGURE 18. Neighborhoods of different points of \mathcal{B} .

$d\pi : T_p\mathcal{B} \rightarrow \mathbb{R}^2 (\subset \mathbb{R}^2 \times \mathbb{R})$, will be an isomorphism. This interior tangent bundle smoothly extends to a tangent bundle of the boundary for which $d\pi$ is again an isomorphism.

Although it is possible to connect our discussion of \mathcal{B} with the historical development of branched surfaces in the literature—specifying the horizontal and vertical boundary of the I -bundle structure of M_S and their interplay with \mathcal{C} —it is more straight forward to develop \mathcal{B} from our current machinery. To that end let, $\mathcal{C} = \mathcal{C}^+ \cup \mathcal{C}^-$, be the decomposition of the crease set into its the positive folding and negative folding components. Coming from the alternative way of defining the folding direction, we observe that there is a natural embedding and identification, $\mathcal{C}^+ \hookrightarrow \partial\mathcal{B} (\subset \mathcal{B})$, coming from the fact that a point, $p_0 \in \partial\mathcal{B}$, corresponds to a point component of $\pi^{-1}(p) \cap M_S$ for some $p \in \mathbb{R}^2$. Thus, we will abuse notion by having $\mathcal{C}^+ = \partial\mathcal{B}$.

Similarly, let \mathcal{L} denote the branching locus of \mathcal{B} . We now observe that there is a natural immersion, $\mathcal{C}^- \hookrightarrow \mathcal{L} (\subset \mathcal{B})$. This comes from the fact that a point, $p \in \mathcal{L}$, corresponds to a fiber component, $I \subset \pi^{-1}(\pi(p)) \cap M_S$, for which $\text{int } I \cap \varepsilon(\mathcal{C}^-) \neq \emptyset$. (Here the $\text{int } I$ is the interior of the fiber.) Since $\pi \circ \varepsilon(\mathcal{C})$ is a 4-valent graph, $\text{int } I \cap \varepsilon(\mathcal{C}^-)$ is either one or two points. If it is just one point then p lies in a neighborhood modeled on Fig. 18b. If there are two points then the neighborhood model corresponds to Fig. 18c, a double point in \mathcal{L} . Again, we may abuse notion by writing $\mathcal{C}^- = \mathcal{L} \subset \mathcal{B}$. We refer readers interested in a fuller development of branched surfaces to [Le].

A primary use of branched surfaces is their “book-keeping” function for carrying embedded surfaces in a fibered neighborhood—surfaces that correspond to assigning non-negative weights to the components of $\mathcal{B} \setminus \mathcal{L}$. (Again, see [Le] for a detailed treatment.)

For \mathcal{B} in our setting, there is a natural way to see the embedding of the components, $\{\bar{K}_1, \dots, \bar{K}_l\} \subset S_g$, “carried” by \mathcal{B} if we allow both $\partial\mathcal{B}$ and \mathcal{L} to serve as boundary. To start, it is natural to decompose the \bar{K}_i ’s into two sets. Recall that each \bar{K}_i inherits a normal vector field coming from the outward pointing normal of M_S . As such, for a given open K_i , its normal vectors project onto $\lambda\mathbf{j}$, where \mathbf{j} is the positively oriented unit associated with the z -axis—the \mathbb{R} factor of $\mathbb{R}^2 \times \mathbb{R}$. Let \mathbf{K}^+ (respectively, \mathbf{K}^-) be the sub-collection of \bar{K}_i ’s for which $\lambda > 0$ (respectively, $\lambda < 0$).

With the above in place the reader should observe that a weight assignment of +1 to each component of $\mathcal{B} \setminus \mathcal{L}$ corresponds to \mathcal{B} carrying either \mathbf{K}^+ or \mathbf{K}^- . Specifically, in Fig. 19 we locally depict how \mathbf{K}^+ or \mathbf{K}^- are configured near $\partial\mathcal{B}$ and \mathcal{L} . The reader should observe that we treat \mathcal{L} as lying in the boundary of the resulting surface, so that the weighting satisfies the branching equations.

6.2. Regular isotopies and the branched surface. Given a regular isotopy, $\varepsilon_t : S \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$, $0 \leq t \leq 1$, there is an associated family of branched surfaces $\bar{\varepsilon}_t(S) = q_{\varepsilon_t} \circ \varepsilon_t(S)$,




FIGURE 19. The branched surface carries S_g by splitting appropriately at \mathcal{L} .

$0 \leq t \leq 1$. When all self-intersections of $q_{\varepsilon_t} \circ \varepsilon_t(\mathcal{C})$ are transverse, the topological type of $\bar{\varepsilon}_t(S)$ is fixed, so $\bar{\varepsilon}_t$ is an isotopy of the associated branched surface \mathcal{B} such that all points of $\pi(\mathcal{B} \setminus (\partial\mathcal{B} \sqcup \mathcal{L}))$ are regular values; in this case we abuse notation and write simply $\bar{\varepsilon}_t(\mathcal{B})$. To account for those isolated values of $t \in [0, 1]$ where $q_{\varepsilon_t} \circ \varepsilon_t(\mathcal{C})$ has non-transverse intersections, we expand the meaning of a regular isotopy $\bar{\varepsilon}_t$ of \mathcal{B} to include two classical operations on branched surfaces, *pinching* and its inverse *splitting*. Fig. 20 depicts these local operations along a neighborhood of the branching locus, $\mathcal{L} \subset \mathcal{B}$. Both operations can result in the introduction or removal of double points.




FIGURE 20. A regular isotopy of $\varepsilon(S)$ which descends to a pinching move on \mathcal{B} .

6.2.1. *Boundary based isotopies.* One primary source of regular isotopies comes from moving $\partial\mathcal{B}$ along a portion of $\text{int } \mathcal{B}$. Specifically, we say $R \subset \mathcal{B} \setminus \mathcal{L}$ is a *boundary annular region* if R is an annulus such that $\partial R = a \cup b$ with $a \subset \partial\mathcal{B}$ and $b \subset \text{int } \mathcal{B} \setminus \mathcal{L}$. We say $R \subset \mathcal{B} \setminus \mathcal{L}$ is a *boundary half-disk* if $\partial R = a \cup b$ where $a \subset \partial\mathcal{B}$ is an arc and $b \subset \mathcal{B} \setminus \mathcal{L}$ is a properly embedded arc— $\partial b = \partial a \subset \partial\mathcal{B}$ —and $(\partial\mathcal{B} \setminus \text{int } a) \cup b$ is a smooth 1-manifold. We then have the following lemma.

Lemma 6.2. *Let $R \subset \mathcal{B} \setminus \mathcal{L}$ be either a boundary annular region or a boundary half-disk region. Then there is a regular isotopy $\bar{\varepsilon}_t(\mathcal{B})$, $0 \leq t \leq 1$, such that $R \cap \bar{\varepsilon}_1(\mathcal{B}) = b$ and $\bar{\varepsilon}_t(\mathcal{B} \setminus R) = \mathbf{id}$, $0 \leq t \leq 1$.*

Proof. Let ϕ_t , $0 \leq t \leq 1$, be an smooth flow coming from extending the inward pointing normal vector field along a ($\subset \partial R$) to all of R such that along b the vector field is an outward (with respect to R) pointing normal. Viewing $R \cong a \times [0, 1]$, with $a \times \{1\}$ identified with a , we take $\bar{\varepsilon}_t(\mathcal{B})$ to be $\phi_t(a, s)s$ on R and $\bar{\varepsilon}_t(\mathcal{B} \setminus R) \cong \mathbf{id}$, $0 \leq t \leq 1$. \square

6.2.2. *Pinching and double-cusp disks.* Next, we develop the tools for pinching a branched surface \mathcal{B} . This pinching will correspond to a regular isotopy of the embedding ε , thus preserving \mathcal{C}^- (which corresponds to \mathcal{L}) and \mathcal{C}^+ (which corresponds to $\partial\mathcal{B}$). Let $d^+, d^- \subset \mathcal{B}$ be two simple arcs such that:

- $\text{int } d^\pm \subset \mathbf{K}^\pm$;
- $\partial d^+ = \partial d^- = d^+ \cap d^- \subset \mathcal{L}$; and,
- $\pi(d^+) = \pi(d^-)$.

Then $d^+ \cup d^-$ bounds a *double cusp disk*, \mathcal{D} , for which $\pi(\mathcal{D}) = \pi(d^\pm)$. Necessarily, $\partial\mathcal{D} \cap \mathcal{L}$ will contain the two *cusp points* of $d^+ \cap d^-$ (along with all points of $\text{int } d^\pm \cap \mathcal{L}$).

The double cusp disks that will be of interest to us are ones that are associated with a *pinching 3-ball*. Specifically, suppose $d^\pm \subset \mathbf{K}^\pm$ split off disks $\delta^\pm \subset \mathbf{K}^\pm$ such that

$$\Sigma = \mathcal{D} \cup \delta^+ \cup \delta^-$$

is a 2-sphere bounding a 3-ball, B_Σ . To parse out the gluing: \mathcal{D} and δ^+ are glued along d^+ ; \mathcal{D} and δ^+ are glued along d^+ ; and, δ^+ and δ^- are glued along a segment $\ell \subset \mathcal{L}$ for which $\partial\ell = \partial d^+ = \partial d^-$. We observe that $\pi(B_\Sigma) = \pi(\delta^+) = \pi(\delta^-)$. As such, for any point $p \in \pi(B_\Sigma)$ the segment components of its pre-image give us an I -bundle structure for B_Σ . Specifically,

$$B_\Sigma = \bigsqcup_{t \in [-1, 1]} D_t,$$

where

- $D_{-1} = \delta^+$ and $D_1 = \delta^-$;
- for all $t \in [-1, 1]$, $\pi(D_t) = \pi(\delta^\pm)$;
- for all $t \in [-1, 1]$ and any $p \in \text{int } D_t$, $d\pi_p : T_p D_t \rightarrow \mathbb{R}^2$ is an isomorphism;
- for all $t \in [-1, 1]$, $\partial D_t = \ell \cup d_t$ where $d_t \subset \mathcal{D}$; and
- $\mathcal{D} = \bigsqcup_{t \in [-1, 1]} d_t$ with $d_{-1} = d^+$ and $d_1 = d^-$.

Observe that when $\text{int } \mathcal{D} \cap \mathcal{B} = \emptyset$, then $\text{int } B_\Sigma \cap \mathcal{B} = \emptyset$. We can then use the I -bundle structure of B_Σ to pinch \mathcal{B} along ℓ , with the isotopy moving ℓ through the disk $\delta^+ (\subset B_\Sigma)$ to the arc $d^+ (\subset \mathcal{D})$.

When we utilize the above pinching procedure in §7, we will be concerned with, first, identifying such double-cusp disks and, second, ensuring that an associated $\text{int } B_\Sigma$ has empty intersection with \mathcal{B} . We return to this second point in §6.2.4.

6.2.3. *Pinching and ∂ -close branching locus segments.* For a given, $\varepsilon : S_g \rightarrow \mathbb{R}^2 \times \mathbb{R}$, and associated branched surface, \mathcal{B} , let $\ell \subset \mathcal{L}$ be a closed segment with $\partial\ell = \{p_1, p_1\}$. Let $R \subset \mathbf{K}^+$ or \mathbf{K}^- be a rectangular region (for simplicity, we take $R \subset \mathbf{K}^+$), with $\partial R = \ell \cup r_1 \cup b \cup r_2$ where:

- r_i , $i \in \{0, 1\}$, are arcs in \mathbf{K}^+ ;

- $b = R \cap \partial\mathcal{B} \subset \mathbf{K}^+$ is a segment with $\partial b = \{q_1, q_2\}$;
- r_i is glued to ℓ at p_i , $i \in \{0, 1\}$; and
- r_i is glued to b at q_i , $i \in \{0, 1\}$.

We say $\ell (\subset \mathcal{L})$ is ∂ -close with respect to \mathbf{K}^+ , or just ∂ -close, if R exists such that $\text{int } R \cap \mathcal{L} = \emptyset$.

Similarly, let $\ell \subset \mathcal{L}$ be a circle component. Let $\ell \subset \partial\bar{K}$ where $\bar{K} \subset \mathbf{K}^+$ or \mathbf{K}^- is an annulus. If $\text{int } \bar{K} \cap \mathcal{L} = \emptyset$ then we say ℓ is ∂ -close with respect to \mathbf{K}^+ , or just ∂ -close.

The salient feature of ∂ -close is, in the annular case, we have a regular isotopy (by Lemma 6.2) that positions $\bar{K} \cap \partial\mathcal{B}$ arbitrarily close to ℓ . In the case where R is a rectangular region we can take a half-disk, $R' \subset R$, for which $\partial R' \setminus b$ is arbitrarily close to $r_1 \cup \ell \cup r_2$. Then an application of Lemma 6.2 to R' positions $\partial\mathcal{B}$ arbitrarily close to ℓ .




FIGURE 21. Pinching along a portion of R to remove two double points along the green ∂ -close arc ℓ .

Now, let $R \subset \bar{K} (\subset \mathbf{K}^\pm)$ be either a disk or an annulus. Suppose R is an annulus with $\partial R = \ell \cup b$, where $\ell \subset \mathcal{L}$ is a circle component and $b \subset \text{int } \bar{K}$; we say R is an \mathcal{L} -annulus. For R a disk, suppose we may write $\partial R = \ell \cup b$ where $\ell \subset \partial\mathcal{L}$ is an segment and $b \subset \bar{K}$ is a properly embedded arc— $\partial b = \partial\ell \subset \partial\mathcal{L}$ —and moreover, $(\mathcal{L} \setminus \text{int } \ell) \cup b$ is a smooth 1-manifold. Then we say R is an \mathcal{L} -disk. We then have the following lemma.

Lemma 6.3. *Let $R \subset \bar{K} (\subset \mathbf{K}^\pm)$ be either an \mathcal{L} -annulus or \mathcal{L} -disk. Let $\ell \subset \partial R$ be its associated boundary circle or segment in \mathcal{L} . If ℓ is ∂ -close with respect to \mathbf{K}^\mp then we can pinch \mathcal{B} along ℓ so as to replace ℓ in \mathcal{L} with b . (See Fig. 21.)*

Proof. For simplicity of notation, assume $\ell \subset \bar{K} \subset \mathbf{K}^+$. Then ℓ is ∂ -close with respect to \mathbf{K}^- and we can perform a regular isotopy of \mathcal{B} such that ℓ is arbitrarily close to either a component of $\partial\mathcal{B}$ (when R is an annulus) or a segment of $\partial\mathcal{B}$ (when R is a disk). We can then pinch \mathcal{L} through R and extend the isotopy to the portion of \mathbf{K}^- that is close to ℓ ; importantly, since this is an arbitrarily small strip of \mathbf{K}^- , we do not need to worry about interactions with the remainder of \mathcal{B} . Fig. 21 illustrates the pinching operation. \square

6.2.4. *Emptying out a B_Σ ball.* The final operation we now discuss will allow us to assume that the interior of the B_Σ 3-balls introduced in §6.2.2 have empty intersection with \mathcal{B} . In that discussion we were given a double cusp disk, \mathcal{D} , and we assumed that there was an associated pinching 3-ball. If we restrict to embeddings of 2-spheres, then the bounded component of $\mathbb{R}^2 \times \mathbb{R} \setminus (\varepsilon(S^2) \cup \mathcal{D})$ is always a 3-ball. When we consider the situation where \mathcal{C} has only three components, as we will in §6, the condition that $\pi(B_\Sigma) = \pi(\delta^+) = \pi(\delta^-)$ will be almost immediate. As such, we have the following lemma.

Lemma 6.4. *Let B_Σ be a pinching 3-ball as described in §6.2.2. Let \mathcal{D} be the associated double cusp disk and Σ be the associated 2-sphere. Then by a regular isotopy of \mathcal{B} we can assume that $\text{int } B_\Sigma \cap \mathcal{B} = \emptyset$.*

Proof. As a technical point we take a small collar neighborhood of $N(\mathcal{D}) = \mathcal{D} \times (-\varepsilon, \varepsilon)$ where $N(\mathcal{D}) \cap B_\Sigma$ is supported only over the $[0, \varepsilon)$ interval. We then extend the D_t -disk bundle structure of B_Σ to $B' = B_\Sigma \cup N(\mathcal{D})$. We will use the notation

$$B' = \bigsqcup_{t \in [-1, 1]} D'_t$$

for the resulting bundle structure. Similarly extending the notation of §6.2.2, we have $\partial D'_t = \ell' \cup d'_t$ where $\ell \subset \ell' \subset \mathcal{L}$ —an extension of ℓ —and $d'_t \subset \mathcal{D} \times \{-1\}$.

Then we choose a non-zero vector field flow, ϕ_s , $s \in (-1, 1)$, on $\text{int } B'$ such that each restricted flow $\phi_s|_{\text{int } D'_t}$, $t \in [-1, 1]$ has an extension to a flow $\phi_{s,t}$ such that on ℓ' , $\phi_{s,t}$ is an inward pointing normal vector field and on d'_t , it is an outward pointing normal vector field.

With this setup, we use ϕ_s to perform a smooth isotopy moving $\mathcal{B} \cap \text{int } B_\Sigma$ to lie in the $(-\varepsilon, 0)$ interval support of $N(\mathcal{D})$. Such a isotopy will necessarily preserve the regular points of $\pi : \mathcal{B} \rightarrow \mathbb{R}^2$. \square

7. CLASSIFICATION OF S^2 EMBEDDINGS WHEN $|\mathcal{C}| = 3$

In this section we give the classification of isotopy classes of regular embeddings of S^2 into $\mathbb{R}^2 \times \mathbb{R}$ having $\mathcal{C} = \{\gamma_i, \gamma_m, \gamma_o\}$ —three smooth curves without corners. Seemingly a simple case, we will see that the subtleties of what can occur are instructive for what may happen in the general setting.

To begin, it is readily seen that there are topologically two possible configurations of three disjoint curves on S^2 : one where the curves split S^2 into three disks plus a pair-of-pants; and, one where the curves split S^2 into two disks and two annuli. The former case is not allowable since, as previously observed, there is no Gauss-Bonnet weighting of \mathcal{C} . For the latter case, we let γ_i and γ_o (“i” for “inner” and “o” for “outer”) be the two curves bounding disks Δ_i and Δ_o , respectively. Since both curves bound disks, by Proposition 4.1, their Gauss-Bonnet weightings must both be $+1$. Additionally, we let γ_m (“m” for “middle”) be the curve that co-bounds an annulus, $A_{i,m}$, with γ_i and co-bounds an annulus $A_{o,m}$, with γ_o . Again, applying Proposition 4.1 we have -1 for the Gauss-Bonnet weight for γ_m .

Theorem 7.1. *Let $\varepsilon : S^2 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$ be a regular embedding such that the crease set \mathcal{C} has three components with no corners. Then $\varepsilon(S^2)$ is regularly isotopic to one of the following, up to the obvious (reflection) symmetries: the “saucer” (Fig. 22a), the “mushroom” (Fig. 22b), or the “toric” sphere (Fig. 22c).*

To be precise about the meaning of obvious symmetry, we allow pre-composition by a homeomorphism of pairs $h : (S^2, \mathcal{C}) \rightarrow (S^2, \mathcal{C})$ and post-composition by an isometry $i : \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2 \times \mathbb{R}$ satisfying $\pi \circ i = i|_{\mathbb{R}^2 \times \{0\}} \circ \pi$. For instance, the saucer and its reflection through the $\mathbb{R}^2 \times \{0\}$ -plane are not regularly isotopic, but are related by an obvious symmetry.

7.1. Proof approach. As remarked earlier, the folding data of the crease set are a regular isotopy invariant. Theorem 7.1 may be understood as saying that in the case $|\mathcal{C}| = 3$, this is a complete invariant—at most one isotopy class occurs per folding direction assignment—and only three folding direction assignments actually occur.

Initially the number of possible assignments for the folding signs is $2^{|\mathcal{C}|} = 2^3$. However, since $\partial \mathcal{B} \neq \emptyset$ and $\partial \mathcal{B}$ corresponds to \mathcal{C}^+ , some component of \mathcal{C} must be a $+$ crease. Similarly, if there is no $-$ crease, then $\mathcal{L} = \emptyset$ and \mathcal{B} is in fact an embedded planar surface




FIGURE 22. Model representatives of the three regular isotopy classes of embeddings of S^2 with three corner-less crease curves. The toric sphere is realized as a plumbing of two saucer spheres.

with boundary. Since \mathcal{B} must carry $\varepsilon(S^2)$, \mathcal{B} is necessarily a disk. But, then $|\mathcal{C}| = 1$, not 3, a contradiction. Thus, both \mathcal{C}^+ and \mathcal{C}^- are nonempty.

Next we observe that we can interchange of roles of γ_i and γ_o , an allowed symmetry. This leaves us with four bijective maps of ordered sets that correspond to a folding assignment.

- i. $f_1 : [\gamma_i, \gamma_m, \gamma_o] \longleftrightarrow [+ , + , -]$.
- ii. $f_2 : [\gamma_i, \gamma_m, \gamma_o] \longleftrightarrow [+ , - , +]$.
- iii. $f_3 : [\gamma_i, \gamma_m, \gamma_o] \longleftrightarrow [- , + , -]$.
- iv. $f_4 : [\gamma_i, \gamma_m, \gamma_o] \longleftrightarrow [+ , - , -]$.

Lemma 7.2. *The folding assignment f_4 cannot occur.*

Proof. Suppose $\varepsilon : S^2 \rightarrow \mathbb{R}^2 \times \mathbb{R}$ realizes this folding assignment. Observe that γ_i , as the single crease curve with positive folding, bounds an embedded disk Δ in \mathcal{B} , and moreover $\partial\mathcal{B} = \partial\Delta$. Then $\mathcal{B} \setminus \Delta$ is a branched subsurface of \mathcal{B} . The boundary of $\mathcal{B} \setminus \Delta$ is precisely $\mathcal{L} \cap \Delta$. In particular, since Δ is a disk, every loop $\ell \subset \partial(\mathcal{B} \setminus \Delta)$ can be realized as the image of the boundary of some map $f : D^2 \rightarrow \Delta$.

Suppose ℓ is homotopically nontrivial in $\partial(\mathcal{B} \setminus \Delta)$. If ℓ is nullhomotopic in $\mathcal{B} \setminus \Delta$, then there is a second map $f' : D^2 \rightarrow \overline{\mathcal{B} \setminus \Delta}$ mapping the boundary to ℓ . Together these two maps define $f \cup f' : S^2 \rightarrow \mathcal{B}$, which is evidently non-nullhomotopic.

If ℓ is essential in $\mathcal{B} \setminus \Delta$, it is homotopic to another loop $\ell' \subset \partial(\mathcal{B} \setminus \Delta)$. Let $g : S^1 \times [0, 1] \rightarrow \overline{\mathcal{B} \setminus \Delta}$ realize this homotopy. We also have $f' : D^2 \rightarrow \Delta$ mapping ∂D^2 to $\ell' \subset \Delta$. Then $f \cup g \cup f' : S^2 \rightarrow \mathcal{B}$ represents a nontrivial class in $\pi_2(\mathcal{B})$.

Since \mathcal{B} is the quotient of a 3-ball, \mathcal{B} is contractible. However, we have shown $\pi_2(\mathcal{B}) \neq 1$, a contradiction. \square

Our approach to proving Theorem 7.1 is to show, in each of the remaining three cases of folding assignments, the branched surface associated to any embedding with given folding data may be taken by a regular isotopy to one of the following three models.

Saucer-folding assignment $f_1 : [\gamma_i, \gamma_m, \gamma_o] \longleftrightarrow [+ , + , -]$. The branched surface model associated with the saucer embedding has $\pi \circ \varepsilon(\gamma_i \sqcup \gamma_m \sqcup \gamma_o)$ ($= \pi(\partial\mathcal{B} \sqcup \mathcal{L})$) as three concentric circles in \mathbb{R}^2 with $\pi \circ \varepsilon(\gamma_i)$ the largest radius and $\pi \circ \varepsilon(\gamma_m)$ the smallest radius.

We take \mathbf{K}^+ to be the images of $\{A_{i,m}, \Delta_o\}$ and for convenience of notation relabel their images $\{A^+, \Delta^+\} \subset \mathcal{B}$, respectively. Similarly, we take \mathbf{K}^- to be the images of $\{A_{m,i}, \Delta_i\}$ and relabel their images $\{A^-, \Delta^-\} \subset \mathcal{B}$, respectively. To satisfy the folding assignment, we take the z -support of A^+ and Δ^\pm in $\mathbb{R}^2 \times \mathbb{R}$ to be $\{0\}$, and the z -support of $\text{int } A^-$ to be $(0, \frac{1}{2})$.

Mushroom-folding assignment $f_2 : [\gamma_i, \gamma_m, \gamma_o] \longleftrightarrow [+ , - , +]$. The branched surface model associated with the mushroom embedding will have $\pi \circ \varepsilon(\gamma_i \sqcup \gamma_m \sqcup \gamma_o) = \pi(\partial\mathcal{B} \sqcup \mathcal{L}) \subset \mathbb{R}^2$ again be three concentric circles with $\pi \circ \varepsilon(\gamma_i)$ the largest radius and $\pi \circ \varepsilon(\gamma_m)$ the smallest radius.

We take \mathbf{K}^+ to be the images of $\{A_{i,m}, \Delta_o\}$ and relabel their images $\{A^+, \Delta^+\} \subset \mathcal{B}$, respectively. Similarly, we take \mathbf{K}^- to be the images of $\{A_{m,i}, \Delta_i\}$ and relabel their images $\{A^-, \Delta^-\} \subset \mathcal{B}$, respectively. Again, in order to satisfy the folding assignment, we take the z -support of A^+ and Δ^- to be $\{0\}$, and the z -support of $\text{int } A^-$ to be $(0, \frac{1}{2})$.

Toric-folding assignment $f_3 : [\gamma_i, \gamma_m, \gamma_o] \longleftrightarrow [- , + , -]$. To construct the branched surface model associated with this toric embedding, we start with $S^1 \vee S^1 \subset \mathbb{R}^2 \times \mathbb{R}$ such that $\pi(S^1 \vee S^1) \subset \mathbb{R}^2$ looks like the standard two crossing projection of the Hopf link except that one crossing is now the common point of the wedge sum; this choice, as well as the choice of over-strand for the intact crossing, is a matter of symmetry. Next, let $T \subset \mathbb{R}^2 \times \mathbb{R}$ be a torus-minus-open-disk which contains and deformation retracts onto our $S^1 \vee S^1$. Further, we require that $\pi(\text{int } T) \subset \mathbb{R}^2$ be an immersion. Then, adapting Proposition 4.1, observe the turning number of ∂T must be -1 . As such we let $\partial T = \varepsilon(\gamma_m)$. Finally, to construct our \mathcal{B} , we attach to T two disks, Δ^+ and Δ^- , one along each of the S^1 's in our wedge sum, so that $\pi(\Delta^\pm) \subset \mathbb{R}^2$ is an embedding. To specify, we identify $\partial\Delta^+$ (respectively, $\partial\Delta^-$) with the under-crossing (respectively, over-crossing) S^1 -factor of the remaining crossing in the Hopf link projection. From this construction we have that $\partial\Delta^+$ and $\partial\Delta^-$ are the images of γ_i and γ_o in \mathcal{B} . The boundary and branch locus of \mathcal{B} are shown in Fig. 23.




FIGURE 23. The crease set of a toric sphere as identified in the branched surface. The green \mathcal{C}^- curves have been identified at a single point in \mathcal{L} .

Summary of the 3 models. We summarize the “geography” of each of our three models. Referring back to Fig. 19, a choice of splitting into \mathbf{K}^+ or \mathbf{K}^- produces a 2-disk and an annulus, independent of model. A salient feature of any such splitting will be the “membership” of the boundary components of the resulting disk/annulus pair—a boundary component will either come from a component of $\partial\mathcal{B}$ or a component of \mathcal{L} . Table 1 categorizes the membership of each possible K^\pm surface. Superscripts of \pm correspond to the superscripts associated with the splitting and superscripts, where used, correspond to the membership of the boundary component.

TABLE 1. Geography of the three models.

\mathbf{K}^\pm	f_1 (saucer)	f_2 (mushroom)	f_3 (toric)
Δ^+	$\partial\Delta^+ \subset \mathcal{L}$	$\partial\Delta^+ \subset \partial\mathcal{B}$	$\partial\Delta^+ \subset \mathcal{L}$
A^-	$\partial A^- = \partial_{\mathcal{B}}^- \cup \partial_{\mathcal{L}}^-$	$\partial A^- = \partial_{\mathcal{B}}^- \cup \partial_{\mathcal{L}}^-$	$\partial A^- = \partial_{\mathcal{B}}^- \cup \partial_{\mathcal{L}}^-$
A^+	$\partial A^+ \subset \partial\mathcal{B}$	$\partial A^+ = \partial_{\mathcal{B}}^+ \cup \partial_{\mathcal{L}}^+$	$\partial A^+ = \partial_{\mathcal{B}}^+ \cup \partial_{\mathcal{L}}^+$
Δ^-	$\partial\Delta^- \subset \partial\mathcal{B}$	$\partial\Delta^- \subset \partial\mathcal{B}$	$\partial\Delta^- \subset \mathcal{L}$

7.2. Proof of Theorem 1.4. We will establish that any regular embedding $\varepsilon(S^2)$ with $|\mathcal{C}| = 3$ is regular isotopic to either the saucer, mushroom, or toric embedding (up to symmetry) by showing its associated branched surface, $\mathcal{B} \subset \mathbb{R}^2 \times \mathbb{R}$, is equivalent to one of our three model branched surfaces through a sequence of boundary-based regular isotopies (see §6.2.1), ∂ -close pinching (see §6.2.3), and pinching double-cusp disks (see §6.2.2). Applying appropriate symmetries, we make the following assumptions on a given embedding to match our models: in all cases, Δ_i will always correspond to Δ^- ; any f_1 -folding embedding will have $\partial\Delta^+ \subset \mathcal{L}$; and, in any f_3 -folding embedding, the signed intersection number of the image in \mathcal{B} of the ordered pair (γ_i, γ_o) is -1 .

To give justification for why, in an f_3 -folding, \mathcal{L} must always contain at least one double point, we consider the surface with boundary $T = \mathcal{B} \setminus \{\text{int } \Delta^+ \cup \text{int } \Delta^-\}$. Since $\partial T = \partial\mathcal{B}$, the turning number of its boundary is -1 . Moreover, since every point in $\text{int } T$ is regular with respect to $\pi : T \rightarrow \mathbb{R}^2$, the crease set of T is empty. Thus, T is a torus minus a disk. But, the two components of \mathcal{L} are naturally contained in $\text{int } T$ as homotopically non-trivial curves, with each component bounding an embedded disk—the disks Δ^\pm . Such curves will have algebraic intersection ± 1 .

To power our argument we first define a complexity measure. The setup is as follows. When $\varepsilon(S^2)$ is associated with the f_1 folding assignment, we let $\mathbf{X} = \{A^-, \Delta^+\}$. For folding assignment f_2 , we let $\mathbf{X} = \{A^+, A^-\}$. And, for f_3 we let $\mathbf{X} = \{A^+, A^-, \Delta^+, \Delta^-\}$. In each case, \mathbf{X} is comprised of those components with some boundary component in \mathcal{L} .

With this in place, we consider the graph $G = \bigsqcup_{X \in \mathbf{X}} X \cap \mathcal{L}$, where $X \in \mathbf{X}$ is a single component. (We allow for the possibility that G has S^1 -“edges”, which occur as isolated components of G ; these occur when some component of \mathcal{L} has no double points.) Let $D(\mathcal{L})$ denote the set of double points of \mathcal{L} and V_G the vertex set of \mathcal{L} . Each $v \in V_G$ corresponds to a point of $D(\mathcal{L})$, and by our choices of \mathbf{X} , any double points of \mathcal{L} will be captured as vertices in G . Referring back to Fig. 18 and 19, observe that each component of $\mathcal{L} \setminus D(\mathcal{L})$ appears as an edge at most four times in G while each double point of \mathcal{L} appears as a vertex at most six times in G . Both bounds are realized when considering all four components of K^\pm .

For a component $X \in \mathbf{X}$, we say a vertex $v \in V_G \cap \partial X$ is ∂ -close in X if there exists a segment neighborhood, $\ell \subset \partial X \cap \mathcal{L}$, such that $v \in \ell$ and ℓ is ∂ -close as defined in §6.2.3. Let N_2 be the number of vertices of V_G that are ∂ -close in some component of \mathbf{X} and N_1 the number of vertices of V_G which are not. Then our complexity measure for a branched surface, $\mathcal{B} \subset \mathbb{R}^2 \times \mathbb{R}$, is the lexicographically ordered 2-tuple, $\chi(\mathcal{B}) = (N_1, N_2)$.

Since $D(\mathcal{L}) = \emptyset$ for our saucer and mushroom standard branched surface models, the minimal possible complexity for folding assignments f_1 and f_2 is $(0, 0)$ with $G = \partial\mathbf{X} \cap \mathcal{L}$.

The complexity for the toric model is $(6, 0)$, which we calculate for each component of \mathbf{X} as follows. $\Delta^\pm \cap \mathcal{L}$ is just the $\partial\Delta^\pm$ with each containing a single vertex of G and, since $\Delta^\pm \cap \partial\mathcal{B} = \emptyset$, these single vertices will not be ∂ -close in D^\pm . $A^\pm \cap \mathcal{L}$ will be subgraphs having three edges, $e_1^\pm, e_2^\pm, e_3^\pm$ and two vertices, v_1^\pm, v_2^\pm , with: $e_1^\pm \cup v_1^\pm = \partial A^\pm \cap \mathcal{L}$, $e_2^\pm \cup v_2^\pm$ a loop bounding a disk in $\text{int } A^\pm$, and e_3^\pm having endpoints v_1^\pm and v_2^\pm . As such, v_1^\pm and v_2^\pm are not ∂ -close in A^\pm .

This is indeed the minimum possible complexity for f_3 , as there must be at least one double point in \mathcal{L} , which is necessarily not ∂ -close anywhere it appears in G .

We note that the subgraphs in A^\pm in the toric model are informative. Having an edge loop in G implies that the number of components of \mathcal{L} is greater than 1. Similarly, having e_3^\pm as ‘‘cut’’ edges implies that $|\mathcal{L}| \geq 1 - e_3^\pm$ share a vertex with e_1^\pm and e_2^\pm so in the quotient they cannot be identified with either edge.

Case 1: $\chi(\mathcal{B})$ is minimal. For the all three folding assignments, we can perform a boundary based isotopy that positions $\partial\mathcal{B}$ arbitrarily close to \mathcal{L} .

For the saucer and mushroom cases, we pick a point p in the interior of the disk component of $\mathcal{B} \setminus \mathcal{L}$. We can then choose a small enough neighborhood, $N(p)$, such that $C_r = \pi(\partial N(p)) \subset \mathbb{R}^2 \times \{0\}$ is circle of some fixed radius.

Now, pick a vector flow on the closure of disk component of $\mathcal{B} \setminus \mathcal{L}$ that flows inward from \mathcal{L} to C_r . We then use this flow to isotopy the neighborhood of \mathcal{L} that contains the annular components of $\mathcal{B} \setminus \mathcal{L}$ to a neighborhood of C_r . After this isotopy we will have the projection of the resultant branched surface, $\pi(\mathcal{B}')$, be one that corresponds the that of our saucer or mushroom models. That is, $\pi(\partial\mathcal{B}' \sqcup \mathcal{L})$ will correspond to three concentric circles. We can then employ a linear isotopy that positions \mathcal{B}' so as to satisfy the z -support conditions of our saucer and mushroom models. If Δ^- projects to $\pi(\Delta^-)$ with $\partial\pi(\Delta^-)$ being the outermost concentric circle, point-wise we could use $H_t(p) = (1-t)[p] + t[\pi(p)]$, $p \in \Delta^-$, $0 \leq t \leq 1$, and extend to the rest of \mathcal{B}' . Being linear, it is evident that $dH_t : T_p\mathcal{B}' \rightarrow \mathbb{R}^2$, $0 \leq t \leq 1$ is an isomorphism on $p \in \mathcal{B} \setminus \partial\mathcal{B}$.

For the toric case we pick two arbitrarily small circles $S^+ \subset \Delta^+$ and $S^- \subset \Delta^-$ such that $S^+ \cap S^- = D(\mathcal{L})$ and $\pi(S_1 \cup S_2)$ corresponds to a ‘‘nice Hopf link’’ projection with a wedge point and a single crossing point.

Next, we choose vectors flows on Δ^\pm that flow inward from $\partial\Delta^\pm (\subset \mathcal{L})$ to S^\pm . We use these two flows to perform an isotopy moving \mathcal{L} to $S^+ \cup S^-$. This isotopy extends to one moving the neighborhood of \mathcal{L} containing the annular component of $\mathcal{B} \setminus \mathcal{L}$ to a small neighborhood of $S^+ \cup S^-$. This describes a regular isotopy of \mathcal{B} to \mathcal{B}' such that the associated branching locus, \mathcal{L}' , has a nice Hopf link type projection. Finally, let $T' = \mathcal{B}' \setminus (\text{int } \Delta'^+ \cup \text{int } \Delta'^-)$. (Here we are extending the ‘‘prime’’ notation to all of the component pieces of \mathcal{B}' .) As in the saucer and mushroom cases, we can use a linear isotopy as a regular isotopy to take T' to the T of the toric model, extending to take \mathcal{B}' to the toric model \mathcal{B} .

Case 2: $\chi(\mathcal{B})$ is not minimal. Consider an arc $\ell \subset \mathcal{C}^-$ whose image connects two double points of \mathcal{L} , or ℓ an entire component of \mathcal{C}^- . By our definition of \mathbf{X} there will be two components, $X^+, X^- \in \mathbf{X}$ such that $\ell \subset \partial X^\pm \cap \mathcal{L}$. (By way of example, in the mushroom case we necessarily have $\ell \subset \partial A^+ \cap \mathcal{L}$ and $\ell \subset \partial A^- \cap \mathcal{L}$.) Let $v_1^\pm, \dots, v_n^\pm \in \partial X^\pm \cap V_G$ be those vertices which lie along ℓ in X^\pm . If one or both of X^\pm is annular, then there is the possibility that ℓ is ∂ -close in that component—though by the assumption that ℓ has double points, it cannot be ∂ -close in both. We break into two cases: when ℓ is ∂ -close in one of

X^\pm , and when ℓ is not ∂ -close in either. These correspond exactly to when we can use each of the pinching moves discussed in §6.2.

Pinching ∂ -close as in §6.2.3. Suppose that ℓ is ∂ -close in, say, X^+ , but not in X^- —there exists a connected subgraph, $J \subset G$, such that $J \cap \partial X^- = \{v_1^-, \dots, v_n^-\}$ and either $X^- = \Delta^-$ or $J \cap \text{int } X^-$ is nonempty. We further suppose $J \cap \text{int } X^-$ is connected, possibly by passing to a subarc of ℓ , and that ℓ is maximal in the sense that either $G \cap X^- \setminus \{v_1^-, v_n^-\}$ is disconnected or $\ell \cong S^1$. Then a neighborhood $R(J) \subset X^-$ defines either an \mathcal{L} -disk or \mathcal{L} -annulus. We then perform the pinching procedure of §6.2.3, replacing \mathcal{B} with the resulting \mathcal{B}' . This eliminates the $2n$ points v_1^\pm, \dots, v_n^\pm from $D(\mathcal{L})$, reducing N_1 and N_2 each by n . Other points of $D(\mathcal{L})$ may become ∂ -close (decreasing N_1 while increasing N_2). However, the conditions on ℓ guarantee no new double points are introduced, and any ∂ -close double point is either eliminated or stays ∂ -close. As a result, $\chi(\mathcal{B}') \leq \chi(\mathcal{B})$. This same operation may be performed with X^+ and X^- interchanged, so that ℓ is ∂ -close in X^- and not in X^+ .

Pinching double-cusp disks as in §6.2.2. Now we assume that ℓ is not ∂ -close in either X^+ or X^- . Similar to before, we assume ℓ is maximal in the sense that either $G \cap X^\pm \setminus \{v_1^\pm, v_n^\pm\}$ are both disconnected or $\ell \cong S^1$, and that no subarc of ℓ is maximal.

The reader should observe that if there is no such ℓ , either every double point of \mathcal{L} is ∂ -close and we can always apply the pinching procedure of §6.2.3 to eliminate points in $D(\mathcal{L})$, or there is a single point $p \in D(\mathcal{L})$ which is not ∂ -close in either X^\pm —the toric case. In either case, if \mathcal{L} has at least two double points, some pair must be connected by a ∂ -close arc.

Suppose $\ell \subset \mathcal{L}$ is such a maximal arc. Let $J^\pm \subset G$ be connected subgraphs such that $J^\pm \setminus \ell \subset \text{int } X^\pm$ and $J^\pm \cap \partial X^\pm = \ell$. At least one of J^\pm must contain an edge outside of ℓ .

Next, using the product structure of our space, $\mathbb{R}^2 \times \mathbb{R}$, we project J^+ (respectively, J^-) onto X^- (respectively, X^+). Let $\pi_{X^\pm}(J^\mp) \cup J^\pm \subset X^\pm$ denote the resulting graphs, taking the (transverse) intersections of $\pi_{X^\pm}(J^\mp) \cup J^\pm$ as additional vertices. Observe that

$$\pi(\pi_{X^+}(J^-) \cup J^+) = \pi(\pi_{X^-}(J^+) \cup J^-).$$

Now, let $E^+ \subset \pi_{X^+}(J^-) \cup J^+ \subset X^+$ be a simple edge-path that starts and ends in $\ell \subset \partial X^\pm$ with $\text{int } E^+ \subset \text{int } X^+$. From the common projection, we can lift a dual edge-path $E^- \subset \pi_{X^-}(J^+) \cup J^- \subset X^-$ such that $\pi(E^+) = \pi(E^-)$. Thus, $E^+ \cup E^-$ is the boundary of a double-cusp disk, \mathcal{D} . We then take δ^\pm to be the half-disks that E^\pm split off in X^\pm . Thus, we have the setup for Lemma 6.4 and can now assume $\text{int } \mathcal{D} \cap \mathcal{B} = \emptyset$. With this assumption in place we can performing the pinching procedure described in §6.2.2.

In the resulting branched surface, \mathcal{B}' , we have reduced N_1 , so $\chi(\mathcal{B}') < \chi(\mathcal{B})$. We caution the reader that a decrease in N_1 may result in an increase in N_2 . However, the combined procedures of §6.2.3 and §6.2.2 will always terminate, at which point either $|D(\mathcal{L})| = 0$ —the saucer and mushroom—or $D(\mathcal{L}) = 1$ —the toric case. We are then in the case where $\chi(\mathcal{B})$ is minimal.

Example 7.3 (Mushroom with five double points). It is helpful to see how the argument of §7.2 can be carried out in actual examples. In Fig. 24a, we offer an example of a branched surface that is equivalent to that of a mushroom S^2 embedding. There are five double points in the branching locus. We draw the reader's attention to the green curve, the unique component of \mathcal{L} . We have marked a point with “ \times ”. Starting at this point one can produce

a “Gaussian notation” sequence by labeling the double points as they are encountered when traversing the \mathcal{L} . Specifically, we have the numeric sequence

$$\times, 1, 2, 3, 3, 4, 2, 5, 5, 4, 1, \times.$$

We also draw the reader’s attention to a single crossing in the projection of \mathcal{L} which does not make a contribution to this numeric sequence.

To finish out the narrative for Fig. 24a, $\partial\mathcal{B}$ consists of the two blue curves. The light-blue “train-tracks” illustrate how \mathcal{B} branches at \mathcal{L} . And the regions with the light-blue swirls correspond to portions of Δ^+ . Here Δ^- is not visible, as it is the underside of this depiction of \mathcal{B} .

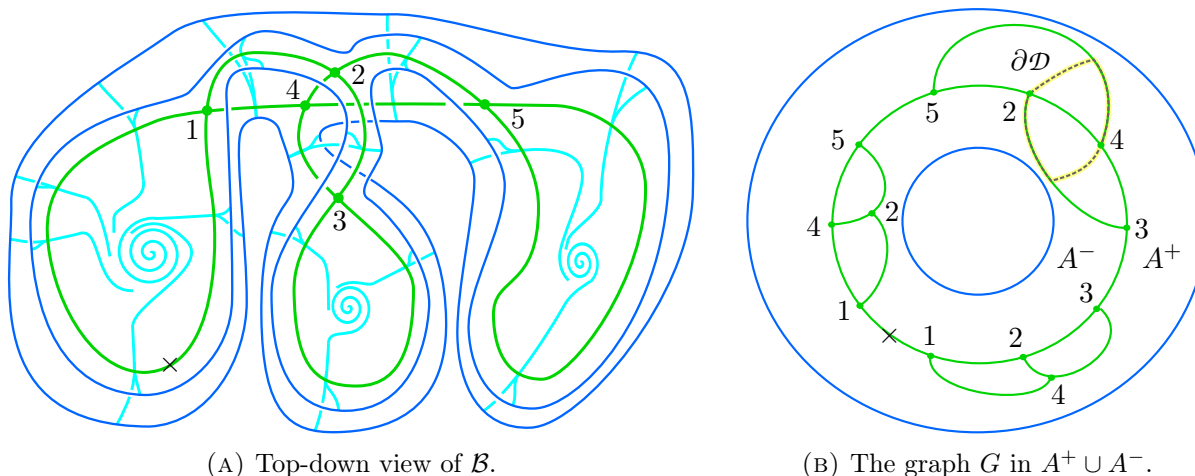


FIGURE 24. A mushroom-type branched surface \mathcal{B} and its associated graph G . The green curve corresponds to \mathcal{L} and the blue curves correspond to the two components of $\partial\mathcal{B}$. The three “swirling” regions of \mathcal{B} lie in Δ^+ . The dotted loop on the right is the boundary of a double-cusp disk.

Fig. 24b depicts A^+ and A^- , which are X^+ and X^- in the mushroom case. For convenience, we adjoin them together along their common branching boundary component—the core green circle. Note that the vertices along this curve appear twice in G —once in the A^+ component and once in the A^- component. In $\text{int } A^+$ and $\text{int } A^-$, we have the subgraphs of G arising from $A^+ \cap \mathcal{L}$ and $A^- \cap \mathcal{L}$, respectively.

Initially, $\chi(\mathcal{B}) = (14, 8)$. To see this with respect to A^+ , when we traverse counterclockwise the core green circle— $\partial A^+ \cap \mathcal{L}$ —starting at the \times -point, the first occurrences of 1, 2, 3, 4, and 5 are not ∂ -close. Additionally, the second occurrence of 2 and the interior vertex 4 are not ∂ -close. However, the second occurrences of 3, 5, 4, and 1 are all ∂ -close. Thus, $A^+ \cap G$ contributes a count of 7 to N_1 and 4 to N_2 .

Replicating this count with respect to A^- , the second occurrences of 3, 2, 5, 4, and 1 along with the first occurrence of 4 and the interior 2-vertex are not ∂ -close. But, the first occurrences of 1, 2, 3, and 5 are all ∂ -close. Thus, $A^- \cap G$ also contributes 7 to N_1 and 4 to N_2 .

We have a number of ways to use subgraphs to eliminate double points of \mathcal{L} , thereby reducing $\chi(\mathcal{B})$. We refer the reader to Fig. 24b. To perform a ∂ -close pinching, we consider the subgraph J^+ that consists the vertex labeled 4 in $\text{int } A^+$, its three adjacent vertices on the green core circle, that in our numeric sequence correspond to $\times, 1, 2, 3$, and their adjoining

edges. With respect to A^- , the arc through the vertices in the same partial numeric sequence, $\times, 1, 2, 3$, is ∂ -close. Then we take a neighborhood $N(J^+) \subset A^+$ as the needed \mathcal{L} -disk to perform a ∂ -close pinching. There is a similar possible ∂ -close pinching coming from the subgraph $J^- \subset A^-$ containing the 2-labeled vertex in $\text{int } A^-$, its three adjacent vertices in ∂A^- , and their adjoining edges.

Performing the ∂ -close pinching associated with the above J^+ eliminates the double points 1, 2, and 3, leaving two double points, 4 and 5. The resulting numeric sequence for the new branched surface, \mathcal{B}' , is

$$\times, 4, 5, 5, 4, \times.$$

Specifically, the first occurrences of 4 and 5 are ∂ -close with respect to A^- , but the second occurrences of 5 and 4 are not. This pattern flips with respect to A^+ —the first occurrences of 4 and 5 are not ∂ -close with respect to A^+ , but the second occurrences of 5 and 4 will be. Thus, the new complexity measure is $\chi(\mathcal{B}') = (4, 4)$.

Alternatively, the arc $\ell \subset \mathcal{L}$ containing the subsequence 3, 4, 2, 5 is not ∂ -close in either A^+ or A^- , as it cannot be entirely “seen” from $A^+ \cap \partial \mathcal{B}$ or $A^- \cap \partial \mathcal{B}$. So one can employ pinching. To do so, we first identify the boundary of a double-cusp disk.

To obtain J^- , we add to $\ell \subset A^-$ the edge in $\text{int } A^-$ with endpoints 4 and 5, and to obtain J^+ , we take $\ell \subset A^+$ and the edge in $\text{int } A^+$ with endpoints 3 and 2. We then project J^- into A^+ and J^+ into A^- . The dotted loop $\partial \mathcal{D}$, shown in Fig. 24b, is a boundary of the required double cusp disk. The “rectangular” region $\partial \mathcal{D}$ bounds corresponds to $d^+ \cup d^-$. Performing the associated pinching will increase $|D(\mathcal{L})|$ from 5 to 6.

To see this move in Fig. 24a, the reader should imagine sliding the 2-labeled double point along \mathcal{L} towards-and-past the double point labeled 4. This will create two new double points, one in the $\overline{14}$ segment of \mathcal{L} —label it 6—and another in the $\overline{43}$ segment—label it 7. In Fig. 24b, this slide corresponds to shrinking all three edges having point labels (2, 4). In the resulting branched surface, we drop the use of 2 as a double point label, yielding the numeric sequence,

$$\times, 1, 6, 7, 3, 3, 7, 4, 5, 5, 4, 6, 1, \times.$$

To focus on the complexity change, N_1 drops from 14 to 12 while N_2 increases from 8 to 12. Thus, the complexity will be reduced since lexicographically, $(12, 12) < (14, 8)$. Moreover, every subgraph of $\text{int } A^\pm \cap G$ will be associated with a ∂ -close pinching. Specifically, the reader can check that in $\text{int } A^+ \cap \mathcal{L}$, there will be three edges with endpoint pairs (1, 6), (7, 3), and (4, 5); and in $\text{int } A^- \cap \mathcal{L}$, there will be three edges with endpoint pairs (3, 7), (5, 4) and (6, 1)—all outermost edges in their respective annuli.

8. POSSIBLE ISOTOPY CLASSES OF $\varepsilon(S^2)$ FOR $|\mathcal{C}| = 5$.

We expect such a classification to become much more difficult as the number of crease curves increases. As an illustration of the variety of behavior that might arise, we examine one of the configurations consisting of five curves. In this configuration, shown in Fig. 25, four of the curves bound disks in S^2 , while the fifth curve has turning number -3 .

We begin by considering which crease sign assignments are realizable. Drawing a parallel to the three-curve case, we might hope each choice of sign assignment for the four disk-bounding curves determines that of the fifth, -3 -turning number curve, and in turn a unique regular isotopy class of embedding.

To explore this, we introduce a method for constructing embeddings.




FIGURE 25. The five-crease configuration under discussion.

Definition 8.1. Let $\varepsilon, \varepsilon' : S^2 \rightarrow \mathbb{R}^2 \times \mathbb{R}$ be normal embeddings with smooth crease sets \mathcal{C} and \mathcal{C}' , and fix curves $\gamma \in \mathcal{C}, \gamma' \in \mathcal{C}'$ such that at least one of γ, γ' has positive crease sign. Let $\Delta, \Delta' \subset S^2$ be disks such that $\Delta \cap \mathcal{C} \subset \gamma$ and $\Delta' \cap \mathcal{C}' \subset \gamma'$. Then the *connect sum* $\varepsilon \#_{\gamma \# \gamma'} \varepsilon'$ of ε and ε' along γ and γ' is the embedding formed by taking a smooth connect sum of $\varepsilon(S^2)$ and $\varepsilon'(S^2)$ along $\varepsilon(\Delta)$ and $\varepsilon'(\Delta')$.

It is straightforward to see that, up to regular isotopy, the resulting embedding is well-defined and determined by the choice of curves γ and γ' . The crease sign condition on γ and γ' is necessary for the result to be an embedding: if, say, γ has negative folding, in order to create a smooth connect sum, the ε' factor must lie inside the ball bound by ε . The crease signs of the resulting embedding depend on the crease signs of γ and γ' as well, in that if, say, γ has negative folding, $\gamma \# \gamma'$ will have negative sign, the other curves of \mathcal{C} will remain unchanged, but the signs of all other curves of \mathcal{C}' will flip.

Returning to the five-curve configuration, this connect sum operation yields examples for every choice of crease sign for the four disk-bounding curves, shown in part in Table 2. The unlisted assignments may be obtained through a symmetry of one of the connect sum factors or the resulting embedding.

TABLE 2. Realizations of embeddings based on the crease signs of disk-bounding curves, enumerated in the first four columns. The first factor in each sum corresponds to γ_1 and γ_2 ; the second to γ_3 and γ_4 . The last column indicates the crease sign of γ_5 in the given realization.

γ_1	γ_2	γ_3	γ_4	Realization	γ_5
+	+	+	+	$\varepsilon_M \# \varepsilon_E$	-
-	-	-	-	$\varepsilon_E \# \varepsilon_E$	+
+	+	\pm	\mp	$\varepsilon_M \# \varepsilon_S$	-
-	-	\pm	\mp	$\varepsilon_S \# \varepsilon_E$	+
\pm	\mp	\pm	\mp	$\varepsilon_S \# \varepsilon_S$	-

Note that this construction leaves a gap: we only give an example for one of the two possibilities for the crease sign of γ in each case. However, γ_5 's sign is determined by those of $\gamma_1, \dots, \gamma_4$ in rows 1-2 (since we must have at least one + and one - crease) and in row 4 (by the same proof as of Lemma 7.2). We do not know whether examples with the opposite sign for γ_5 exist for rows 3 and 5.

Moreover, it seems plausible that non-regularly isotopic embeddings may have the same crease sign data, even accounting for further symmetry. Let $\bar{\varepsilon}_E$ be the mirror image of ε_E

through the \mathbb{R}^2 -plane. Then $\varepsilon_E \# \varepsilon_E$ and $\varepsilon_E \# \bar{\varepsilon}_E$ have the same crease sign data, yet do not seem to be regularly isotopic.

9. OPEN QUESTIONS

Although there are large number of directions an investigation into the crease set could take—for example, surfaces with boundary in $\mathbb{R}^2 \times \mathbb{R}$ —we supply only a few that immediately come to mind for closed surfaces. The motivation of these problems is to advance the further development of the machinery of §5 and §6.

Problem 9.1 (Analyze the S^2 crease set with corners). As previously mentioned, from §5 onward our analysis deals with understanding the embeddings of S^2 into $\mathbb{R}^2 \times \mathbb{R}$ in the situation where the crease set is without corners. An obvious direction of investigation is to carry out a similar analysis in the situation where corners occurred. There should be some interaction between this line of investigation and the one described in the previous sections. In particular, referring back to Fig. 2 and the accompanying discussion, one sees that a pair of “canceling corners” can be introduced by the creation of a dimple. Fig. 26 illustrates of another method of adding/canceling corners to the crease set. The proposed problem is to not only give a constructive procedure for the geometric realization of a crease set with corners that satisfies a Gauss-Bonnet weighting criteria, but give a finite list of operations for moving between embeddings.




FIGURE 26. The green curves represent \mathcal{C} . The dashed arc between the corners in the upper diagram gives a strip to perform a corner-canceling isotopy. The arc in the lower diagram, at first flat, crosses the new crease set twice.

Problem 9.2 (Estimate the growth rate of embedding equivalence classes for S^2). The construction procedure for the geometric realization of a crease configuration given in §5 clearly involves making a lot of choices. In §7, we see there are exactly three non-isotopic embeddings with the same crease configuration of three curves; in §8 we produce five with a common crease configuration. It is reasonable to suspect this is the case in general, and that a given crease configuration will have numerous geometric realizations. The proposed problem is to calculate the grow rate of set of equivalency classes of embeddings for S^2 or any other closed orientable surface into $\mathbb{R}^2 \times \mathbb{R}$, as measured in comparison to the number of crease curves and corners.

Problem 9.3 (Analyze the crease set for embeddings of $S^1 \times S^1 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$). We restrict our analysis here to embeddings of $S^2 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$. It is natural to consider which phenomena persist or collapse in positive genus. Moving just one step along this path, an investigation of the crease set for embeddings $\varepsilon : S^1 \times S^1 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}$ might proceed as follows. Given a configuration pair, $(S^1 \times S^1, \mathcal{C})$, whose closed components of $S^1 \times S^1 \setminus \mathcal{C}$ satisfy Equ. 4.2, any component of \mathcal{C} is a s.c.c. of one of three types: (1) homologically trivial in $S^1 \times S^1$; (2) a curve that is the boundary of an embedded disk contained in the solid torus that $\varepsilon(S^1 \times S^1)$ bounds; or, (3) a longitudinal curve on $\varepsilon(S^1 \times S^1)$. If $\varepsilon(S^1 \times S^1)$ is unknotted then this last type of curve *may* also bound an embedded disk. If type-(2) occurs then type-(3) cannot occur and vice-versa. With this in mind, the proposed problem is to give a construction of the embedding for geometrically realizing the crease set $\mathcal{C} \subset S^1 \times S^1$.

Problem 9.4 (Analyze the crease set for embeddings of $S^2 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}^2$). It is very tempting to push an analysis of the crease set into dimension 4. Since \mathbb{R}^4 has a natural product structure of $\mathbb{R}^2 \times \mathbb{R}^2$, one would in fact have two crease configurations to work with—one for each factor projection map. The obvious first interesting examples one might consider are that of knotted 2-spheres.

REFERENCES

- [C] A. Cayley, *A theorem on trees*. Quart. J. Pure Appl. Math. 23 (1889), 376–378.
- [G] L. Goeritz, *Bemerkungen zur Knotentheorie*, Abh. Math. Sem. Univ. Hamburg 18 (1997) 201–210.
- [G-L] E. Ghys and J. Leys, *Lorenz and modular flows: a visual introduction*, www.ams.org/publicoutreach/feature-column/fcarc-lorenz, 2006
- [H] Joel Hass, *Surfaces and their Profile Curves*. eprint, May 2022, arXiv:2205.01737.
- [H-T-T1] Joel Hass, Abigail Thomspson, and Anastasiia Tsvietkova, *Tangled decompositions of alternating link complements*. **Illinois J. Math.** 65(3): 533–545 (September 2021).
- [H-T-T2] Joel Hass, Abigail Thomspson, and Anastasiia Tsvietkova, *The number of surfaces of fixed genus in an alternating link diagram*, **Int. Math. Res. Not.** (2017), no. 6, 1611–1622.
- [K-L] L. Kauffman, S. Lambropoulou, *Hard unknots and collapsing tangles*, in *Introductory Lectures on Knot Theory—Selected Lectures presented at the Advanced School and Conference on Knot Theory and its Applications to Physics and Biology ICTP*, Trieste, Italy, 11–29 May 2009, Edited by L. Kauffman, S. Lambropoulou, S. Jablan, and J. Przytycki, World Scientific, Singapore, 2011.
- [L] John M. Lee, *Riemannian manifolds. An introduction to curvature*. Graduate Texts in Mathematics, 176. Springer–Verlag, New York, 1997.
- [Le] Tao Lee, *Laminar Branched Surfaces in 3-manifolds*. Geometry & Topology Volume 6 (2002) 153–194.
- [Li] W. B. Raymond Lickorish, *An introduction to knot theory*. Graduate Texts in Mathematics, 175. Springer–Verlag, New York, 1997.
- [M1] William W. Menasco, *Closed incompressible surfaces in alternating knot and link complements*. **Topology** 23 (1984), no. 1, 37–44.
- [M2] William W. Menasco, *Determining incompressibility of surfaces in alternating knot and link complements*, **Pacific J. Math.** 117 (1985), no. 2, 353–370.

DEPARTMENT OF MATHEMATICS, UNIVERSITY AT BUFFALO
Email address: menasco@buffalo.edu

FIELDS INSTITUTE FOR RESEARCH IN MATHEMATICAL SCIENCES
Email address: mnichols@fields.utoronto.ca