

Coplanar constant mean curvature surfaces

KARSTEN GROSSE-BRAUCKMANN, ROBERT B. KUSNER
AND JOHN M. SULLIVAN

We consider constant mean curvature surfaces with finite topology, properly embedded in three-space in the sense of Alexandrov. Such surfaces with three ends and genus zero were constructed and completely classified by the authors. Here we extend the arguments to the case of an arbitrary number of ends, under the assumption that the asymptotic axes of the ends lie in a common plane: we construct and classify the entire family of these genus-zero, coplanar constant mean curvature surfaces.

*Dedicated to Hermann Karcher
on the occasion of his 65th birthday.*

1. Introduction

The search for examples of surfaces with constant mean curvature (CMC) dates back to the nineteenth century. Unlike minimal surfaces, which have a Weierstrass representation, the case of nonzero constant mean curvature does not readily yield explicit examples; it thus appeared that constant mean curvature surfaces were scarce. In recent decades, however, several construction methods have evolved. Wente's discovery of an immersed torus of constant mean curvature initiated the integrable-systems approach explored by Abresch, Bobenko, Pinkall, Sterling and others; this in turn led to a generalized Weierstrass representation by Dorfmeister, Pedit and Wu. The ground-breaking analytic perturbation (or gluing) techniques of Kapouleas have been extended by Mazzeo, Pacard, Pollack, Ratzkin and others. Conjugate surface techniques, introduced by Lawson, were developed further by Karcher and — combined with new analytic methods — by the present authors. An abundance of known CMC surfaces has emerged, inviting their systematic study and classification.

With the conjugate surface method, we can study CMC surfaces with sufficient symmetries. This method is powerful in that it enables us to study entire families of surfaces, that is, components of suitably defined

moduli spaces. In contrast, the gluing techniques mentioned above, while more widely applicable, tend to produce examples only near a boundary of the relevant moduli space. The generalized Weierstrass representation also has the potential to deal with moduli spaces, even without any symmetry assumptions, but it has difficulty detecting key geometric features such as embeddedness.

The *Delaunay unduloids* [5] are the noncompact embedded CMC surfaces of revolution. These singly periodic surfaces decompose along their *necks* (minimizing closed geodesics) into *bubbles*. All known embedded CMC surfaces display similar bubbles. In light of experience with doubly and triply periodic surfaces [7, 15, 22], the discovery that bubbles can be created and deleted continuously on surfaces of finite topology [10] was surprising. This insight is what led us to expect our moduli spaces to be connected, and thus to look for the existence proof presented here, based on the continuity method we used in [11].

Let us now make the setting of our paper precise. A surface has *finite topology* when it is homeomorphic to a compact surface Σ of genus g from which a finite number k of points has been removed; a neighborhood of any of these punctures is called an *end*. If the surface has nonzero constant mean curvature, then by a scaling we may assume that the mean curvature is identically 1; from now on CMC will denote this condition. As in [11], we consider a class of surfaces more general than the embedded ones, namely those bounding an immersed solid. To be precise, we define a finite topology CMC surface M to be *Alexandrov-embedded* if M is properly immersed, if each end of M is embedded and if there exists a compact 3-manifold W with boundary $\partial W = \Sigma$, and a proper immersion $F: W \setminus \{q_1, \dots, q_k\} \rightarrow \mathbb{R}^3$ whose boundary restriction $f: \Sigma \setminus \{q_1, \dots, q_k\} \rightarrow \mathbb{R}^3$ parametrizes M . Moreover, we assume that the (inward) mean-curvature normal ν of M points into W .

A Delaunay unduloid is an embedded CMC surface with two ends and genus zero; it is determined by its *necksize* $n \in (0, \pi]$, the circumference of any neck. More generally, we define a *k-unduloid* to be an Alexandrov-embedded CMC surface with finite genus and a finite number $k \in \mathbb{N}$ of ends. Our terminology is justified by the asymptotics result of Korevaar, Kusner and Solomon.

Theorem 1.1 [19]. *Each end of a k-unduloid is exponentially asymptotic to a Delaunay unduloid, and thus has an asymptotic axis and necksize $n \in (0, \pi]$.*

For $k \leq 2$, all k -unduloids have genus zero, and their classification is well-known: Alexandrov [1, 2] showed that the round sphere is the only 0-unduloid; Meeks [24] showed there are no 1-unduloids; and the Delaunay unduloids themselves are the only 2-unduloids [19].

A CMC surface is called *coplanar* if it is contained in a half-space of \mathbb{R}^3 . For $k \leq 3$, all k -unduloids are coplanar [24]. The Alexandrov reflection method extends to coplanar k -unduloids [19], showing that any such surface M is *Alexandrov-symmetric* in the following sense: it has a reflection plane P such that $M \setminus P$ has two connected components M^\pm whose Gauss images $\nu(M^\pm)$ are contained in disjoint open hemispheres; moreover $\nu(M \cap P)$ is contained in the equator.

Theorem 1.2 [19]. *A coplanar k -unduloid M is contained in a slab and is Alexandrov-symmetric across a plane P within that slab. The plane P contains the asymptotic axes of the ends of M .*

We want to study moduli spaces of coplanar k -unduloids with genus zero. Thus, we identify two such k -unduloids when they differ by a rigid motion in space and we label the ends to avoid orbifold points at surfaces with extra symmetry. It is convenient to represent any coplanar k -unduloid with the horizontal xy -plane as its fixed mirror plane P , such that the labels of the ends occur in cyclic order:

Definition 1.3. Let us fix $k \geq 0$ points q_1, \dots, q_k in cyclic order along the equator of \mathbb{S}^2 . We consider proper Alexandrov-embedded CMC immersions of $\mathbb{S}^2 \setminus \{q_1, \dots, q_k\}$ into \mathbb{R}^3 , which are equivariant under reflection across the xy -plane. The *moduli space* \mathcal{M}_k of coplanar k -unduloids of genus zero consists of all such immersions, modulo diffeomorphisms of the domain fixing the points q_i and modulo isometries of \mathbb{R}^3 . Uniform convergence in Hausdorff distance on compact subsets of \mathbb{R}^3 defines the topology on \mathcal{M}_k .

Our main theorem classifies coplanar k -unduloids of genus zero in \mathcal{M}_k . This generalizes our classification of *triunduloids* (with $k = 3$) by triples of points in \mathbb{S}^2 modulo rotation [11], which shows \mathcal{M}_3 is homeomorphic to the open three-ball. For $k \geq 4$, there is still a natural map from coplanar k -unduloids of genus zero to k -tuples on \mathbb{S}^2 , but this no longer suffices to classify the surfaces. Instead, we need to employ a space \mathcal{D}_k of spherical metrics, whose precise definition will be given below. Our classification says that \mathcal{M}_k is homeomorphic to \mathcal{D}_k .

Main Theorem. *For each $k \geq 3$, there is a homeomorphism Φ from the moduli space \mathcal{M}_k of coplanar k -unduloids of genus 0 to the connected $(2k - 3)$ -manifold \mathcal{D}_k of k -point metrics. The necksizes of a k -unduloid $M \in \mathcal{M}_k$ are the spherical distances of the k consecutive pairs of completion points in $\Phi(M)$.*

As a consequence, we have the following constraints on the necksizes of a coplanar k -unduloid of genus 0, as illustrated in figure 1. (These do not hold for higher genus, as shown by the coplanar k -unduloids of genus 1 with all ends cylindrical [13].)

Corollary 1.4. *Let $k \geq 3$. The necksizes $n_1, \dots, n_k \in (0, \pi]$ of any coplanar k -unduloid of genus 0 satisfy*

$$(1.1) \quad n_1 + \dots + n_k \leq (k - 1)\pi \text{ for } k \text{ odd,}$$

$$(1.2) \quad n_1 + \dots + n_k < k\pi \text{ for } k \geq 4 \text{ even,}$$

and these inequalities are sharp. At most $k - 2$ ends of the surface can be cylindrical, and examples of coplanar k -unduloids with $k - 2$ cylindrical ends exist for each $k \geq 3$.

Our theorem should be compared with the classification of coplanar minimal surfaces by Cosin and Ros [4]. The classifying space for their minimal k -noids consists of planar polygons, whose combinatorial description is well-known. Another simplification arises because they are able to prove directly




Figure 1: Our main theorem characterizes which sets of k asymptotic necksizes a genus-zero coplanar k -unduloid can attain. In addition, for $k = 3$ the necksizes uniquely determine the angles between the ends, as follows from force balancing. For $k \geq 4$, however, there is no such dependence: the 4-unduloids shown above all have all necksizes equal to $\pi/2$, but the angles between the ends vary.

that all their surfaces are nondegenerate. We are not aware of any other classification results for moduli spaces of minimal or CMC surfaces with more than three ends.

Very recently, in joint work [9] with Korevaar and Ratzkin (following up on [18]), we have augmented the results of this paper as follows. All coplanar k -unduloids of genus zero are nondegenerate, so that \mathcal{M}_k is a real-analytic manifold, and a tubular neighborhood of it in the space of all k -unduloids is a $(3k - 6)$ -manifold. Furthermore, our classifying map Φ is a (real-analytic) diffeomorphism.

We conclude this introduction with an overview of our paper. In Section 1, we define the classifying map Φ as a composition of the following maps. By coplanarity and Theorem 1.2, we first associate to each $M \in \mathcal{M}_k$ its upper half M^+ . Since M has genus zero, M^+ is an open disk. It thus has a conjugate cousin, an isometric minimal immersion to the 3-sphere. Hopf projection of this cousin disk to the 2-sphere produces a metric on M^+ locally isometric to \mathbb{S}^2 , which we call the spherical k -point metric $\Phi(M)$.

Central for our paper is the analysis (in Sections 2 and 3) of the space \mathcal{D}_k of spherical k -point metrics. In Section 2, we start with a combinatorial decomposition of each metric $D \in \mathcal{D}_k$. In view of the classifying map, analyzing metrics amounts to analyzing coplanar k -unduloids, yielding a decomposition of the surface into bubbles. On cylindrical parts of a k -unduloid, such a decomposition is necessarily ambiguous. For this reason, our decomposition cannot be unique, making the analysis somewhat involved.

In Section 3, we apply our combinatorial decomposition to show that \mathcal{D}_k is a connected manifold. We analyze the example \mathcal{D}_4 completely to find that $\mathcal{D}_4 \cong \mathcal{M}_4$ is a 5-ball, generalizing the result of [11] that $\mathcal{D}_3 \cong \mathcal{M}_3$ is a 3-ball. Our combinatorial description is not explicit enough to decide whether \mathcal{D}_k is always a $(2k - 3)$ -ball. (In ongoing joint work with Korevaar, using a Möbius-invariant decomposition of k -point metrics, we expect to be able to show that indeed $\mathcal{D}_k \cong \mathbb{R}^{2k-3}$ for all $k \geq 3$.)

In Section 4, we prove injectivity and properness of our classifying map. Injectivity follows by applying a maximum principle to the cousin surfaces (just as in the case $k = 3$ of [11]). Properness involves *a priori* estimates, corresponding to the closedness part of a traditional continuity method. The proof here is somewhat different than in the case $k = 3$ because of the possibility of small necks occurring away from the ends. Still, our characterization of compact subsets of \mathcal{D}_k allows us to prove the needed area and curvature estimates.

In Section 5, we establish the surjectivity of our classifying map. As in [11], the crucial point is to apply a structure theorem for real-analytic

varieties. Thus, we must show that our spaces \mathcal{M}_k , consisting of coplanar surfaces with a mirror symmetry, are real-analytic. For this, we adapt the methods of [21], which dealt with the space of all k -unduloids. More generally, we show that for any fixed finite symmetry group, the moduli space of symmetric k -unduloids is real-analytic, a result which has applications beyond this paper [6]. Then we can apply our version of the continuity method, as introduced in [11], to prove the main theorem.

We gratefully acknowledge partial support from the Deutsche Forschungsgemeinschaft and the National Science Foundation. Our figures of CMC surfaces were produced with the Graphics Programming Environment *Grape*; our sketches of spherical polygons, with *Spherical Easel*.

2. The classifying map

In this section, we define the space \mathcal{D}_k of k -point spherical metrics and the classifying map $\Phi: \mathcal{M}_k \rightarrow \mathcal{D}_k$. We assume $k \geq 2$ to avoid trivial cases.

2.1. Conjugate cousins and necksizes

Consider a coplanar k -unduloid $M \in \mathcal{M}_k$; it is Alexandrov symmetric across the xy -plane P . Take a component M^+ of $M \setminus P$ and orient M^+ such that the boundary labels occur in increasing cyclic order: this means that M^+ lies to the left of ∂M^+ . By a rotation of M in \mathbb{R}^3 , we may assume that $M^+ := M \cap \{z > 0\}$; we call M^+ the *upper half*.

When M is embedded the ends occur in counterclockwise order in P . In other words, if at each end of M we truncate $P \cap M$ and join in a locally embedded way, then we obtain an embedded loop in P with turning number one.

However, in general ends of M may cross (as in figure 2) and thus no longer appear in the same order in P . We note that for large k , most coplanar k -unduloids have such crossed ends and thus fail to be embedded. By Alexandrov symmetry, the projection of M^+ to P gives an immersed disk, so that the turning number of the corresponding boundary loop is still one.

We now review the conjugate cousin relationship [22], in the first-order form presented in [11, Theorem 1.1]; see also [8]. Given any simply connected CMC surface, such as M^+ , there is a cousin minimal surface \tilde{M}^+ in the 3-sphere \mathbb{S}^3 . We regard \mathbb{S}^3 as the Lie group of unit quaternions and identify \mathbb{R}^3 with its Lie algebra $T_1\mathbb{S}^3 = \text{Im } \mathbb{H}$, as in [11]. The cousin is determined uniquely up to left translation in \mathbb{S}^3 .




Figure 2: A coplanar 4-unduloid and its intersection with the symmetry plane P . Typical coplanar k -unduloids with $k > 3$ have such crossed ends and are not embedded.

The upper half $M^+ \subset \mathbb{R}^3$ is a disk isometric to its minimal cousin $\tilde{M}^+ \subset \mathbb{S}^3$. This isometry has the following effect on normals and tangent vectors.

If $\nu \in \mathbb{S}^2$ is the normal at $p \in M^+$, then the normal $\tilde{\nu}$ at the corresponding point $\tilde{p} \in \tilde{M}^+$ is simply the left translation of ν , that is, $\tilde{\nu} = \tilde{p}\nu$. Left translation preserves the inner product, and so for the unit quaternion $\mathbf{k} \in \mathbb{S}^2 \subset \text{Im } \mathbb{H}$ we find

$$\langle \tilde{\nu}, \tilde{p}\mathbf{k} \rangle_{\mathbb{S}^3} = \langle \nu, \mathbf{k} \rangle_{\mathbb{R}^3}.$$

Since \mathbf{k} is vertical, Theorem 1.2 gives $\langle \nu, \mathbf{k} \rangle < 0$ on the upper half M^+ , and thus the cousin \tilde{M}^+ is transverse to the \mathbf{k} -Hopf circles $t \mapsto \tilde{p}(\cos t + \mathbf{k} \sin t)$.

If X is any tangent vector to M^+ at p , then its isometric image $\tilde{X} = \tilde{p}JX$ at \tilde{p} is obtained through a rotation J by $\pi/2$ within T_pM , followed by left translation to \tilde{p} . When X is the (horizontal) unit tangent to ∂M^+ at p , the tangent plane at p is vertical and so JX agrees with the constant conormal \mathbf{k} to ∂M^+ . Thus $\tilde{X} = \tilde{p}\mathbf{k}$, implying that the component of $\partial \tilde{M}^+$ through \tilde{p} covers precisely the \mathbf{k} -Hopf circle $t \mapsto \tilde{p}(\cos t + \mathbf{k} \sin t)$.

The \mathbf{k} -Hopf circles are the fibers of the \mathbf{k} -Hopf projection $\Pi_{\mathbf{k}}: \mathbb{S}^3 \rightarrow \mathbb{S}^2$, defined by $\Pi_{\mathbf{k}}(q) := q\mathbf{k}\bar{q}$ in terms of quaternion multiplication. Thus each component of $\partial \tilde{M}^+$, corresponding to the arc from some end i to end $i + 1$, projects to a single point p_i . Collecting these points in \mathbb{S}^2 , we obtain an ordered k -tuple

$$\Psi(M) := \Pi_{\mathbf{k}}(\partial \tilde{M}^+) = (p_1, \dots, p_k).$$

As in [11, Section 2] it follows that $\Psi(M)$ is well-defined up to rotation in $\text{SO}(3)$.

Theorem 1.1 says that the ends of a k -unduloid M are asymptotic to certain unduloids, defining *asymptotic necksizes* n_1, \dots, n_k of M . These are determined by $\Psi(M)$.

Theorem 2.1. *Let M be a coplanar k -unduloid of genus zero. Its asymptotic necksizes n_1, \dots, n_k are the distances between consecutive points of the spherical k -tuple $\Psi(M) = (p_1, \dots, p_k)$:*

$$\text{dist}_{\mathbb{S}^2}(p_{i-1}, p_i) = n_i.$$

Proof. As in [11] the result is immediate for $k = 2$: the cousin of the upper half of an unduloid of necksize n is a spherical helicoid bounded by two k -Hopf circles; these circles Hopf-project to two points in \mathbb{S}^2 at distance n . Then, the proof given in [11, Theorem 2.1] for the case $k = 3$ carries over to the general case: the asymptotics of Theorem 1.1 imply for any k that each end of the cousin \tilde{M}^+ has the same bounding great circle rays as its asymptotic unduloid cousin. □

Corollary 2.2. *The set of k necksizes of a coplanar k -unduloid obeys the inequalities satisfied by the edgelengths of a spherical k -gon: for each subset $P \subset \{1, \dots, k\}$ with an odd number $|P|$ of elements*

$$(2.1) \quad \sum_{i \in P} n_i \leq \sum_{i \notin P} n_i + (|P| - 1)\pi.$$

Proof. Note that the case of $|P| = k$ odd in (2.1) is (1.1). We prove this case only; for the general case, see [3], where it is also shown that these inequalities are sharp.

We proceed by induction, noting that for $k = 3$ we recognize (1.1) as the perimeter inequality for spherical triangles [11].

For $k \geq 5$, we draw minimizing geodesics of length $a \geq 0$ and $b \geq 0$ from p_k to p_{k-3} and p_{k-2} , as in figure 3. These divide the k -gon into a $(k - 2)$ -gon and two triangles, which degenerate when a or b vanishes. By induction the $(k - 2)$ -gon satisfies

$$n_1 + \dots + n_{k-3} + a \leq (k - 3)\pi.$$

The perimeter and triangle inequalities for the two triangles give

$$n_{k-1} + n_k + b \leq 2\pi \quad \text{and} \quad n_{k-2} \leq a + b.$$

Adding these three inequalities gives (1.1) as desired. Since all necksizes are positive at most $k - 2$ necksizes can be π . □




Figure 3: Edges (dashed) of length a and b from p_k to p_{k-3} and p_{k-2} divide a k -gon into two triangles and a $(k - 2)$ -gon; the pieces overlap in this example.

Remark. This corollary extends to a broader class of immersed (and not Alexandrov embedded) CMC surfaces, including the surfaces constructed in [16]. If M has genus 0 with k -ends, each of which is asymptotically unduloidal, and M has a mirror plane, then it can be shown that the disk M^+ is still bounded by k curves of planar reflection. Thus, the map Ψ is still defined, so that (1.1) and (2.1) hold. However, inequality (1.2) depends on Proposition 4.1, and need not extend to the broader class.

2.2. Spherical metrics and the classifying map

For $k \leq 3$, the spherical k -tuple $\Psi(M)$ determines the coplanar k -unduloid M uniquely. Since for $k > 3$ this is no longer the case, we are led to introduce a family of metrics to serve as our classifying space.

A metric is *spherical* if it is locally isometric to \mathbb{S}^2 . We will consider spherical metrics D on the oriented open two disk. We denote the completion of such a metric by \hat{D} and refer to $\partial D := \hat{D} \setminus D$ as the *completion boundary* of D . To the spherical metric D , we can associate an oriented *developing map* $\varphi: D \rightarrow \mathbb{S}^2$, unique up to rotation, which extends to \hat{D} .

Definition 2.3. Consider a spherical metric on the disk D whose completion boundary consists of fixed points $\{q_1, \dots, q_k\}$, called *vertices*, in increasing

cyclic order. Two such metrics are equivalent if they are related by an isometry extending to fix each vertex q_1, \dots, q_k . A k -point metric is an equivalence class of such metrics, and we let \mathcal{D}_k denote the space of all k -point metrics. The developing map defines a topology on \mathcal{D}_k : a sequence of metrics in \mathcal{D}_k is said to converge if there are developing maps (from the disk to \mathbb{S}^2), which converge in the compact-open topology.

As we noted above, if M is a coplanar k -unduloid, \tilde{M}^+ is transverse to the \mathbf{k} -Hopf circles. Hence $\Pi_{\mathbf{k}}: \tilde{M}^+ \rightarrow \mathbb{S}^2$ is an immersion well-defined up to rotation. We use $\Pi_{\mathbf{k}}(\tilde{M}^+)$ to denote the spherical metric on M^+ pulled back from \mathbb{S}^2 . As in [11, Lemma. 3.4], the exponential asymptotics of the ends of M to unduloids shows that this spherical metric is a k -point metric. The k curves in ∂M^+ correspond to the vertices, which develop to the k -labelled points $(p_1, \dots, p_k) = \Psi(M)$. Thus, we have shown.

Theorem 2.4. *For each $k \geq 2$, there is a map $\Phi: \mathcal{M}_k \rightarrow \mathcal{D}_k$ defined by*

$$\Phi(M) := \Pi_{\mathbf{k}}(\tilde{M}^+).$$

We call this map Φ the *classifying map*. Our main theorem asserts that Φ is a homeomorphism.

3. Decompositions of a k -point metric

We analyze a k -point metric $D \in \mathcal{D}_k$ by decomposing it into standard pieces. We will use this combinatorial analysis later to prove that the space \mathcal{D}_k is connected. We assume $k \geq 2$ throughout this section.

3.1. Polygonal metrics

It will simplify the exposition to consider a further class of spherical metrics on the disk, whose completions are compact.

Definition 3.1. A *polygonal metric* or (*spherical*) k -gon is an isometry class of spherical metrics on the open disk whose completion boundary is topologically a circle consisting of k vertices and k edges each developing to a great-circle arc in \mathbb{S}^2 of length at most π . We denote the space of all k -gons by \mathcal{P}_k .

At each vertex v of a polygonal metric, there is an interior *angle* $\alpha(v)$, well-defined as a positive real number. We introduce some further terminology.




Figure 4: Left: the complement of the great circle arc is a slit sphere. Right: two geodesics of length π , connecting a pair of antipodal points, cut the sphere into two complementary lunes, one convex and one concave.

A *diagonal* in a spherical metric D is a geodesic of length at most π in D , whose completion in \hat{D} joins two vertices (which may be consecutive). A diagonal develops onto a minimizing great circle arc.

Two kinds of 2-gons (see figure 4) will be important throughout. A *slit sphere* is the complement of a minimizing great circle arc in \mathbb{S}^2 . Each of its vertices has angle $\alpha = 2\pi$; its *slit length* is the length of each edge. A *lune* of angle α is a 2-gon whose vertices develop to antipodal points of \mathbb{S}^2 ; it has a continuous family of diagonals.

A *ray of spheres* is a spherical metric D whose completion boundary is a single arc, developing to a minimizing geodesic γ . The preimage of γ decomposes D into an infinite union of slit spheres. Let us indicate an argument for this fact, which is more direct than [11, Lemma 3.2]. Note that each point of a slit sphere is the endpoint of a geodesic emanating from one of the endpoints of the slit; each of these geodesics is minimizing and disjoint from the bounding slit. Thus, the completeness of $D \cup \gamma$ gives that D contains a slit sphere attached across the completion boundary γ . Continuing, by induction, we see D contains a ray R of slit spheres attached across γ . But $R \cup \gamma$ is complete, and thus, R is relatively open and closed in D . This shows $D = R$.

The analysis of k -point metrics can be reduced to k -gon metrics.

Proposition 3.2. *Given any k -gon, we can attach rays of spheres along its edges to give a k -point metric. Conversely, any k -point metric D can be decomposed into a k -gon D' together with k rays of spheres; we call D' a truncation of D .*

Proof. Given a k -gon D , we can attach to each edge e a ray of spheres with slit length equal to the length of e . The resulting metric has no completion boundary except for the k -vertices of D , so it is a k -point metric.

Conversely, given a k -point metric D , note that its vertices are cyclicly ordered. If we fix one vertex v , the others inherit a linear ordering; let w be the last vertex in this ordering. We claim that there is a diagonal from v to w .

First, consider some reference geodesic arc c_0 from v into D . By completeness, for each angle $\vartheta \in \mathbb{R}$, there is a geodesic arc c_ϑ making a signed angle ϑ with c_0 at v ; we continue each arc c_ϑ only until either it hits ∂D (at another vertex) or it reaches length π staying inside D . Whenever the first case happens, we have found a diagonal from v .

Suppose for all ϑ in some interval (α, β) the second case occurs. All arcs c_ϑ then end inside D at the same point v' (developing antipodally to v). Since the angle around v' is only 2π , we must have $\beta - \alpha \leq 2\pi$, meaning there is never an angular gap of more than 2π between diagonals. Thus, we find a sequence $(d_i)_{i \in \mathbb{Z}}$ of diagonals from v , indexed in increasing order of ϑ . (If some diagonal has length π , it is part of a continuous family of diagonals for some range of ϑ ; in this case, we include only a discrete subset of these in our sequence.)

Let v_i be the endpoint of d_i . Since each d_i cuts D into two subdiscs, the sequence (v_i) of vertices is increasing with respect to their linear order. Thus for $i \geq I$, the point v_i is a constant vertex. For each $i > I$, the loop $d_i \cup d_I$ bounds a 2-gon D_i . The union $\bigcup D_i =: R$ has no boundary other than d_I , and so is a ray of slit spheres. Since such a ray has no further completion points, v_I must be w , the neighbor of v . This proves the claim: each vertex has a diagonal to the consecutive vertex.

Consider the set of k diagonals in D , connecting the pairs of consecutive vertices. Since these are minimizing geodesics between consecutive vertices, they cannot intersect one another. Together, they bound a k -gon whose complement consists of k rays of spheres. \square

Example. We exhibit polygonal metrics with $k - 2$ edges of length π , for each $k \geq 3$. By Proposition 3.2, these examples can be extended with k -rays of spheres to form k -point metrics with $k - 2$ edgelengths π . By the Main Theorem, these metrics establish the existence of k -unduloids with $k - 2$ cylindrical ends.

Let D_n be a 3-gon developing to a hemisphere, with edgelengths π , n , $\pi - n$ where $0 < n < \pi$. Reflecting across the edge of length $\pi - n$, we obtain a 4-gon with edgelengths n , π , π and n , as in figure 5.




Figure 5: Left: the solid arcs have lengths π and n , while the dashed arc has length $\pi - n$. The 3-gon D_n develops to the upper hemisphere, while the 4-gon described in the text (with edgelengths n, π, π, n) develops to the complement of the solid arcs. Right: the classifying map takes this 4-unduloid (with necksizes n, π, π, n) to the 4-point metric one of whose truncations is the 4-gon shown on the left.

We can reflect a copy of D_n across an edge of length n of the 4-gon to obtain a 5-gon with edgelengths n, π, π, π and $\pi - n$. Continuing the same way, we obtain k -gons with $k - 2$ edgelengths π ; the remaining edgelengths are n, n when k is even, and $n, \pi - n$ when k is odd. Note that the k vertices of these gluings always just develop to the three original vertices of D_n .

3.2. Polygonal metrics without diagonal

We will use diagonals to decompose a k -point metric in \mathcal{D}_k into embedded pieces. Let us introduce the particular k -gons, which arise in the decomposition, beginning with triangles (3-gons).

Definition 3.3. A triangle is *small* if its completion develops into an open hemisphere. Equivalently, its developing image has area less than 2π , or it has angles less than π . Its edge lengths are also less than π .

In the case of a flat metric with piecewise geodesic boundary, there is a set of diagonals, which decomposes the metric into flat triangles. This fact was used by Cosín and Ros to analyze coplanar minimal k -noids [4]. For spherical metrics, the situation is more involved. Besides slit spheres and small triangles, our decomposition will include pieces of the following types.




Figure 6: Left: a lune with $k - 2$ marked points on one of its edges forms an embedded k -gon called a marked lune. Right: the complement of the small convex pentagon shown on the front of the sphere is an embedded concave 5-gon D . It has no diagonals. The great circle arc pq shown, for instance, is longer than π and so by definition it is not a diagonal. (A minimum-length path from p to q exists in \hat{D} , but it stays within ∂D .)

Definition 3.4. A k -gon in \mathcal{P}_k is *concave* if all its angles are at least π and no edge has length π . It is *embedded* if its developing map is an embedding into \mathbb{S}^2 (its completion need not be embedded). A *marked lune* is a lune of angle at most 2π with $k - 2$ additional vertices marked on one of its edges, as in figure 6 (left); when the angle is less than π we call it *convex*.

Slit spheres of slit length less than π are concave 2-gons. A triangle without diagonals is embedded and is either concave or a marked lune or small. For larger k , figure 6 (right) shows an example of an embedded concave k -gon, with no diagonals. Another concave example is a metric isometric to a hemisphere of \mathbb{S}^2 with k vertices on its bounding equator; it admits a diagonal if and only if some pair of vertices is antipodal. For slit length less than π , slit spheres or chains of slit spheres are always concave.

Our goal (in Proposition 3.7 below) is to show that a concave k -gon with no diagonals is embedded; for the special case of triangles this was shown in [11]. To formulate a lemma, let us fix a vertex v of D and consider geodesics in D emanating from v , as in the proof of Proposition 3.2. For each angle $\vartheta \in (0, \alpha(v))$, there is a unique geodesic arc c_ϑ making an angle ϑ with the edge after v ; again we continue c_ϑ only until either it hits ∂D , or it reaches length π staying inside D . Thus, the only points of ∂D contained in an arc c_ϑ are its initial point v and perhaps its endpoint.

Lemma 3.5. *Let D be a k -gon, which is not isometric to a hemisphere, and suppose that D has no diagonals. Then all edges of D have length less than π . Moreover, at each vertex v the following alternative holds. Either*

- (i) all arcs c_ϑ can be continued to length π , where they converge at a point $v' \in D$, which develops antipodally to v , or
- (ii) all arcs c_ϑ have length at most π and end on the interior of the same edge e (not incident to v), where they meet e transversely.

The second case can only arise when $\alpha(v) < \pi$.

Proof. If some edge e has length π , then there is a sufficiently thin lune contained in D whose boundary includes e . But such a lune contains a continuous family of diagonals, joining the endpoints of e . This contradiction shows the edges have length less than π .

Let us now prove the alternative. Since D has no diagonals, no arc c_ϑ can hit ∂D at a vertex. For any fixed ϑ , either c_ϑ reaches length π without encountering ∂D , or c_ϑ first hits ∂D in an edge e transversely. Since the edges have length less than π , the edge e in the second case is one of the $k - 2$ edges not incident to v . We claim that each of these $k - 1$ possibilities occurs for an open set of angles ϑ . Since $(0, \alpha(v))$ is connected, it then follows that we must have the same endpoint behavior for all angles ϑ .

To prove the claim, consider c_ϑ . Let X be ∂D with v and its two incident edges removed; if c_ϑ ends on an edge e we also remove e . Then the two disjoint closed subsets c_ϑ and X of \hat{D} have a positive distance δ . Within the δ -neighborhood of c_ϑ , we can construct the required nearby geodesics starting from v and with the same endpoint behavior.

It remains to show that for $\alpha(v) \geq \pi$, case (ii) cannot occur. If $\alpha > \pi$, then we have two opposite arcs from v hitting the same boundary edge e transversely. The two intersection points of this great circle with e must be antipodal, contradicting the fact that e has length less than π . To rule out $\alpha = \pi$, we similarly use two limit arcs in \hat{D} , which hit the same edge e . If e is transverse to the limit arcs, we get the same contradiction. Otherwise e lies along the limit arcs; however, by assumption D is not a hemisphere, so in this case one of the limit arcs must include a diagonal of D , a contradiction. \square

Proposition 3.6. *If a k -gon $D \in \mathcal{P}_k$ has no diagonal, then either D is concave or it is a small triangle.*

Proof. For $k = 2$, the 2-gons without diagonals are slit spheres, and hence concave. So we assume $k \geq 3$. We first consider the case that all angles of D equal π . The developing image of the completion boundary is then contained in a single great circle and, by the Gauss–Bonnet formula, the metric must have area 2π . Therefore, D is isometric to a hemisphere. It




Figure 7: The rays c_ϑ from v initially hit \overline{wx} so by Lemma 3.5 they all do, and we deduce that, unlike in this picture, $\alpha(v) < \pi$.

cannot have an edge of length π for then it would have diagonals near this edge; thus it is concave and we are done.

We may thus assume that D is not a hemisphere, meaning the hypotheses of Lemma 3.5 hold. We suppose D is not concave and claim D must be a small triangle. In the following, we use \overline{vw} to denote the edge between consecutive vertices v and w , not including its endpoints.

There is at least one vertex w with $\alpha(w) < \pi$. Consider the consecutive vertices t, u, v, w, x (not necessarily distinct for small k). We apply the lemma to the arcs c_ϑ at v . Because the length of \overline{vw} and the angle at w are both less than π , as in figure 7, arcs c_ϑ near \overline{vw} end on edge \overline{wx} . The lemma then says all c_ϑ end on \overline{wx} ; this implies that we also have $\alpha(v) < \pi$.

Repeating this argument from v , we find next that $\alpha(u) < \pi$. Now consider again the arcs c_ϑ at v . They have length less than π and end on the edge \overline{wx} when ϑ is small, while they end on the edge \overline{tu} for ϑ close to $\alpha(v)$. By the lemma, we must have $\overline{wx} = \overline{tu}$, implying that $k = 3$; moreover all edges and angles are less than π . That is, D is a small triangle, as claimed. □

We can characterize concave k -gons without diagonals further.

Proposition 3.7. *A concave k -gon with no diagonals is embedded.*

Proof. It follows as in the previous proof that if all vertex angles equal π then equivalently the metric is a hemisphere and hence embedded. Otherwise Lemma 3.5 is applicable.

We first claim that $\alpha(w) \in [\pi, 2\pi]$ for each vertex w . Since D is concave, $\alpha(w) \geq \pi$, so if we apply Lemma 3.5 we are in case (i). That is, all arcs c_{ij} from w continue to length π , where they end at a common point $w' \in D$. We call their union the *vertex lune* of D at w . Since w' is an interior point of D , and D immerses into \mathbb{S}^2 under the developing map, as in the proof of Proposition 3.2, the angle $\alpha(w)$ of this lune cannot exceed 2π . This proves the claim.

We then use induction on k . For $k = 2$, we know a 2-gon without diagonals is an embedded slit sphere. For the induction step, we will show that given a k -gon D , we can join a triangle to it to obtain a $(k - 1)$ -gon D' that still has no diagonals. By the induction hypothesis, D' is embedded, and so the same must hold for its subset D .

Since the k -gon is not a hemisphere, we can choose a vertex v with $\alpha(v) \in (\pi, 2\pi]$. Let us assume for now that $\alpha(v) < 2\pi$. Then, the three consecutive vertices u, v, w of D develop to the three vertices of a small triangle Δ , which is locally exterior to the image of D . We define a polygonal metric $D' \in \mathcal{D}_{k-1}$, as in figure 8, isometric to the union of D with Δ , glued across the two edges incident to v . More precisely, we glue the closure of Δ to D and then remove \overline{uv} . The completion boundary of D' agrees with that of D , except that \overline{uv} and \overline{vw} are replaced by \overline{uw} .

The angles at u and w have only increased under this gluing, so D' is still concave. We claim that D' does not have a diagonal. To see this, consider the open hemisphere $H \subset D'$, bounded by \overline{uw} , and contained in the union of Δ with the vertex lune at $v \in D$. No vertices of D' are in H , so a diagonal of D' , being a minimizing geodesic with endpoints outside H , stays outside H . Hence any diagonal of D' stays outside Δ and




Figure 8: By induction, the $(k - 1)$ -gon D' , obtained by gluing the triangle Δ to the concave k -gon D , is embedded.

is thus also a diagonal of D . However, by assumption there are no such diagonals. This completes the inductive proof under the assumption that $\alpha(v) < 2\pi$.

When $\alpha(v) = 2\pi$, essentially the same ideas apply, except that Δ degenerates to a segment. When u and w develop to distinct points, we obtain exactly as above a metric $D' \in \mathcal{P}_{k-1}$ with no diagonals. When u and w develop to the same point, note that we must have $k \geq 4$. Here we get $D' \in \mathcal{P}_{k-2}$ by gluing the edge \overline{vw} (being the degenerate triangle Δ) to D , and removing the vertex $u = w$. Again by the same argument, the metric D' cannot have any diagonals. \square

Note that an embedded concave polygon has diagonals only if it has two antipodal vertices; in this case, it is a lune with angle in $[\pi, 2\pi]$, possibly with extra vertices marked along its edges.

3.3. Decomposition of polygonal and k -point metrics

We will now use the previous proposition to decompose any polygonal metric into a finite union of embedded pieces. Each piece is either a small triangle, or an embedded concave polygon with no diagonals, or a marked lune of angle less than π . It will follow immediately that any k -point metric can be decomposed into the same pieces.

Theorem 3.8. *Any polygonal metric $D \in \mathcal{P}_k$ can be decomposed by diagonals into a finite union of slit spheres, small triangles, embedded concave polygons without diagonals, and marked lunes.*

Proof. By Propositions 3.6 and 3.7, if D has no diagonals, it is already one of the allowed pieces. Otherwise, we work by induction on k , using diagonals to simplify D . Suppose that D has a diagonal d connecting two nonconsecutive vertices. In this case, the induction step is immediate since d splits D into two pieces each with fewer edges.

The remaining case is when D has a diagonal connecting consecutive vertices, but no diagonals connecting nonconsecutive ones. This includes the base case ($k = 2, 3$) of the induction on k . Here, we will establish the decomposition directly, without using the induction hypothesis. Instead, we work by induction on the area of D (in multiples of 4π).

So we have a metric D , of area less than $4\pi n$ for some $n \in \mathbb{N}$, where all diagonals connect consecutive vertices. Choose a shortest diagonal d , and

call its endpoints v and w . Suppose first d has length less than π . Then the 2-gon bounded by d and the edge \overline{vw} of D is a union of m slit spheres, where $1 \leq m \leq n$. The k -gon on the other side of d has area $4\pi m$ less than that of D . Thus, for the base case $n = 1$, this is impossible, meaning that no such diagonal can exist so that D itself is embedded. On the other hand, for $n > 1$, the metric decomposes into the allowed embedded pieces by induction.

The more difficult case is when d has length π . In this case, we consider the lune L of maximal angle, which is contained in D and incident to the edge \overline{vw} . Then, the other edge d' bounding L must include some point of ∂D . Since there are no diagonals shorter than π , the arc d' must lie entirely in ∂D . Thus D is a marked lune, though with angle α possibly greater than 2π . It has a continuous family of diagonals, so we can use these to decompose D into n slit spheres and one marked lune of angle at most 2π (adjacent to d'). Thus, also in this case, we have arrived at a decomposition into the desired embedded pieces. \square

Given a k -point metric D we can remove its rays as in Proposition 3.2 and then apply Theorem 3.8 to obtain the following decomposition.

Corollary 3.9. *Any k -point metric $D \in \mathcal{D}_k$ can be decomposed into k -rays of slit spheres together with a finite union of slit spheres, small triangles, embedded concave polygons with no diagonals and marked lunes.*

When applied to the Hopf projection $D = \Phi(M)$ of a coplanar k -unduloid $M \in \mathcal{M}_k$ this decomposition can be regarded as a decomposition of M into bubbles.

There is an equivalent reformulation of the decomposition in terms of graphs dual to the tessellating pieces. This graph is a tree, topologically equivalent to the skeletal graphs considered earlier by Schoen [26], Kusner [20] and Kapouleas [14].

Note that the decomposition of a k -point metric in general is not unique. Clearly, the location of diagonals of length π is not determined; in addition, many metrics admit combinatorially different choices of diagonals, with the simplest example being a convex square.

In Proposition 3.2, we removed rays (of slit spheres) from a k -point metric. In view of the decomposition established by the corollary above, we can make this more precise.




Figure 9: This 6-point metric D on the open disk is decomposed by the six thick arcs into its small truncation D^\square and six rays of slit spheres. In turn, the polygon metric D^\square is, as in Corollary 3.9, decomposed into two triangles, three slit spheres, and one 4-gon (necessarily embedded, concave, and without diagonals). A 6-point metric with such a decomposition will correspond (under our classifying map Φ) to a 6-unduloid with two midsegments, one having two unduloid bubbles and the other having just one.

Definition 3.10. Given a k -point metric $D \in \mathcal{D}_k$, a *small truncation* is a metric $D^\square \in \mathcal{P}_k$ obtained from D by removing (closed) rays so that no slit sphere remains incident to any bounding edge.

When D has no pair of consecutive vertices at distance π , the small truncation is uniquely defined. In the general case, the ambiguity is by less than one slit sphere at each of the k ends.

Figure 9 illustrates the results of this section.

4. The space of k -point metrics

We now show that \mathcal{D}_k is a connected manifold. As for $k = 3$, we consider the projection Π from \mathcal{D}_k to k -tuples of points on \mathbb{S}^2 , given by restricting the developing map of a k -point metric to its vertices. For $k = 3$, the mapping Π is a diffeomorphism from \mathcal{D}_3 to the open 3-ball of triples in \mathbb{S}^2 , as in [11]. For larger k , this is only true locally. From now on we assume $k \geq 3$.

4.1. \mathcal{D}_k is a manifold

Let us define \mathcal{S}_k to be the set of k -tuples of sequentially distinct points in \mathbb{S}^2 .

$$\mathcal{S}_k := \{(p_1, p_2, \dots, p_k) \in (\mathbb{S}^2)^k, p_i \neq p_{i+1} \text{ for } i = 1, \dots, k\}.$$

Define $\mathcal{A}_k \subset \mathcal{S}_k$ to be the set $\{(p, q, \dots, p, q) : p \neq q \in \mathbb{S}^2\}$ for k even and the empty set for k odd. Note that \mathcal{A}_k contains all k -tuples in \mathcal{S}_k consisting only of a pair of antipodal points (namely, when $q = -p$). Thus, the action of $\text{SO}(3)$ is free on the manifold $\mathcal{S}_k \setminus \mathcal{A}_k$, and the quotient

$$\mathcal{T}_k := (\mathcal{S}_k \setminus \mathcal{A}_k) / \text{SO}(3),$$

is a manifold.

Proposition 4.1. *For $k \geq 3$, the completion boundary of any k -point metric develops to a k -tuple in \mathcal{T}_k . This gives a surjective map $\Pi: \mathcal{D}_k \rightarrow \mathcal{T}_k$.*

Note, however, that for $k = 2$ the image of \mathcal{D}_2 includes $(p, -p) \in \mathcal{A}_2 = \mathcal{S}_2$. This fact explains why $\mathcal{D}_2 \cong (0, \pi]$ fails to be a manifold but instead has a boundary point π . Of course, π corresponds to the cylinder in $\mathcal{M}_2 \cong \mathcal{D}_2$, where rotations also fail to act freely.

Proof. For $D \in \mathcal{D}_k$, Proposition 3.2 gives a k -gon $D^\square \in \mathcal{P}_k$, obtained by removing rays from D . The spheres of each ray have nonzero slit length, and so consecutive vertices must develop to distinct points. Thus, any developing image of the completion boundary ∂D lies in \mathcal{S}_k .

To show the image is not in \mathcal{A}_k , suppose to the contrary that the completion points of $D \in \mathcal{D}_k$ develop onto just one pair $p \neq q$ of points in \mathbb{S}^2 . Then, the decomposition of Corollary 3.9 can only contain slit spheres of slit length $\text{dist}(p, q)$. Such slit spheres cannot be joined together to produce a k -gon for $k > 2$.

Finally, we show Π is surjective. For $k = 3$ this is obvious, while for $k > 3$ this follows by induction. There exists a nonconsecutive pair p_i, p_j of vertices with distinct images. By joining two metrics along a diagonal developing to $\overline{p_i p_j}$, we obtain the desired k -point metric. \square

Corollary 4.2. *For k even, at least two ends of a k -unduloid are non-cylindrical, and so $n_1 + \dots + n_k < k\pi$ holds.*

Proof. If $k - 1$ ends were cylindrical, then the k -tuple $\Psi(M)$ would alternate between a pair of antipodal points (making also the k th necksize cylindrical). But this is impossible by Proposition 4.1. \square

The map Π relates the k -tuple map Ψ and classifying map Φ introduced in Section 1. For any k -unduloid $M \in \mathcal{M}_k$, we have $\Phi(M) \in \mathcal{D}_k$ and $\Psi(M) = \Pi(\Phi(M)) \in \mathcal{T}_k$.

Let $\mathcal{T}_k^* \subset \mathcal{T}_k$ denote the open set of all pairwise distinct k -tuples in $\mathbb{S}^2 \pmod{\text{SO}(3)}$, and let $\mathcal{D}_k^* := \Pi^{-1}(\mathcal{T}_k^*) \subset \mathcal{D}_k$. For $k = 3$, we have $\mathcal{T}_3^* = \mathcal{T}_3$ and so $\mathcal{D}_3^* = \mathcal{D}_3$.

Theorem 4.3. *For any $k \geq 3$, the map $\Pi: \mathcal{D}_k \rightarrow \mathcal{T}_k$ is a covering map over \mathcal{T}_k^* . It is a local diffeomorphism everywhere, implying that \mathcal{D}_k is a real-analytic manifold of dimension $2k - 3$.*

Proof. Consider a k -point metric $D \in \mathcal{D}_k$, which maps to some k -tuple $T \in \mathcal{T}_k$. Pick any representative $(v_1, \dots, v_k) \in (\mathbb{S}^2)^k$ for T ; the v_i are the images of the completion boundary points q_i of D under some developing map φ .

By introducing a finite set of additional vertices, we can find a triangulation Δ of \mathbb{S}^2 by small triangles whose vertices include the v_i . Passing to a subdivision if necessary, we may assume that the v_i are not adjacent in Δ , in other words that their star neighborhoods $S_i \subset \mathbb{S}^2$ are disjoint when $v_i \neq v_j$. (Note that we may have $v_i = v_j$ for $i \neq j$ if $T \notin \mathcal{T}_k^*$.) We claim that Π is a local diffeomorphism from a neighborhood of D onto $U := (S_1 \times \dots \times S_k)/\text{SO}(3)$. It follows that \mathcal{D}_k is a manifold of dimension $2k - 3$; we endow it with the real-analytic structure coming from these charts.

To prove the claim, note that the pullback $\varphi^{-1}(\Delta)$ gives a triangulation of the metric space \hat{D} . For any $(w_1, \dots, w_k) \in S_1 \times \dots \times S_k$, we construct a completed k -point metric with that developing image as follows. Each triangle of $\varphi^{-1}(\Delta)$ incident to the completion boundary ∂D has exactly one vertex $q_i \in \partial D$. We continuously change the metric on any such triangle according to the motion of v_i to w_i , but we fix the metric on the remaining triangles (those not incident to ∂D). Under this new k -point metric, q_i develops to w_i ; this confirms that we have constructed a continuous local inverse for Π .

To show that Π is a covering over \mathcal{T}_k^* , note that \mathcal{T}_k^* is exactly the set where the v_i are distinct. In this case, the neighborhood U as above is trivially covered by $\Pi^{-1}(U)$ because we can use the same triangulation Δ for any $D \in \Pi^{-1}(U)$. \square




Figure 10: This family of 4-unduloids with fixed necksizes $\pi/2$ exhibits a curvature blow-up on the three necks attached to the two central bubbles. Its singular limit is a union of two unduloids and 2-spheres. Under Ψ , the surfaces Hopf project to the two poles of \mathbb{S}^2 as well as two points on the equator, which approach one another; thus the vertex distance has limit 0 and Lemma 5.3 does not apply.

4.2. \mathcal{D}_k is connected

What initially led us to believe that \mathcal{M}_k is connected was that we can imagine a path from any given k -unduloid $M \in \mathcal{M}_k$ to a standard one, say a dihedrally symmetric one. In particular, we could deform M such that one of its midsegments becomes cylindrical (as in figure 10 viewed from right to left), and then (as in figure 1) telescope inward along that midsegment. Repeating this process until all midsegments have been retracted, we could finally adjust the necksizes of the ends until they are equal.

This intuitive idea will guide our proof of connectedness of \mathcal{D}_k . Our first lemma singles out a class of k -point metrics for which the decomposition of Corollary 3.9 is unique. Let us introduce the subset $\mathcal{E}_k \subset \mathcal{T}_k$ of equatorial k -tuples subject to the condition that all points occur in counterclockwise order on the equator, and consecutive pairs have distances in $(0, \pi)$.

Lemma 4.4. *Suppose that $D \in \mathcal{D}_k$ is a metric such that $\Pi(D) \in \mathcal{E}_k$ and D has no diagonal of length π . Then the decomposition of D asserted in Corollary 3.9 is unique. More specifically, the small truncation D^\square is a union of some number $m(D) \geq 0$ of pieces isometric to slit spheres and exactly one piece isometric to a hemisphere.*

Proof. The decomposition pieces can only be slit spheres or embedded concave k -gons without diagonals. Let us consider the developing image of D . The diagonals of D can only develop to the equator, and so the concave k -gons are isometric either to hemispheres or to slit spheres. Moreover, all

rays of slit spheres are joined from the same side to the equator, meaning that the small truncation D^\square has area an odd multiple of 2π .

Thus, there is at least one decomposition piece isometric to a hemisphere. The orientation of D means that any such hemisphere H has its bounding vertices in increasing cyclic order; however, the developing image of ∂H covers the equator once, so H must develop to the upper hemisphere. If there were two or more hemispheres in the decomposition, we could find a pair connected by a chain of pieces isometric to slit spheres; this pair would develop to complementary hemispheres, contradicting the observation above.

Let us now show the decomposition is unique. We use preimages of the equator to subdivide D into hemispherical tiles, developing alternately to the upper and lower hemisphere. This canonical decomposition is the same as the one above, except that each slit sphere has been divided (along an internal arc of length greater than π) into two hemispheres. Thus, the hemisphere piece in the original decomposition is distinguished as the unique tile, which has all edge lengths less than π . □

We can now prove the following.

Theorem 4.5. *The manifold \mathcal{D}_k is connected.*

Proof. By Theorem 4.3 the complement of \mathcal{D}_k^* has codimension 2 in \mathcal{D}_k ; for $k = 3$ it is empty. Thus, proving \mathcal{D}_k is connected is equivalent to showing that \mathcal{D}_k^* is connected. To prove the latter, it is sufficient to pick some $T \in \mathcal{T}_k^*$ and show that the deck transformation group of the covering $\Pi: \mathcal{D}_k^* \rightarrow \mathcal{T}_k^*$ acts transitively on the particular fiber $F := \Pi^{-1}(T) \subset \mathcal{D}_k^*$.

To define F , we pick a k -tuple $T \in \mathcal{E}_k$ with no pairs of points antipodal. For instance, if k is odd, then T could be k points with maximal dihedral symmetry, while for even k , a slight perturbation of this configuration within \mathcal{E}_k would work. Since a metric $D \in F$ cannot have diagonals of length π , Lemma 4.4 is applicable to D and gives a well-defined number $m(D)$ of pieces isometric to slit spheres in the unique small truncation D^\square . Note there is a unique metric in $D_0 \in F$ with decomposition number $m(D_0) = 0$, isometric to a hemisphere with k -rays of spheres attached.

For $k = 3$, any $D \in F$ has $m(D) = 0$ since the small truncation D^\square is the unique triangle determined by $\Pi(D) \in \mathcal{T}_3$. This shows $\mathcal{D}_3^* = \mathcal{T}_3^*$; but $\mathcal{T}_3^* = \mathcal{T}_3$ is an open 3-ball (see [11]). So let us from now on assume $k \geq 4$.

For $D \in F$, we will construct a loop in \mathcal{T}_k^* based at T whose lift in \mathcal{D}_k^* runs from D to D_0 . We can assume D has decomposition number $m(D) > 0$; by induction it is sufficient to construct a loop γ whose lift $\tilde{\gamma}$ connects D

to some $D' \in F$ for which m has decreased by one. We will choose γ as the concatenation $\alpha * \lambda * \alpha^{-1}$ of three subpaths, where λ is a loop in \mathcal{T}_k^* and α is a path in \mathcal{E}_k . On the level of CMC surfaces, the lift of λ telescopes a midsegment, while α adjusts necksizes.

Since $m(D) > 0$, the metric D has at least one diagonal d between two nonconsecutive vertices v_i, v_j . There is a path α in \mathcal{E}_k from $T = \Pi(D)$ to some k -tuple $S \in \mathcal{E}_k$ in which v_i and v_j are antipodal. Among pairs of vertices joined by a diagonal, we can assume that no pair becomes antipodal along our path α , since otherwise we instead take the first such diagonal to be d , and replace α by the shortened path.

The path α has a unique lift $\tilde{\alpha}$ in \mathcal{D}_k^* through D . We let $E \in \Pi^{-1}(S)$ be its terminal metric. By Lemma 4.4, the decomposition of each metric along $\tilde{\alpha}$ is unique. Hence up until the terminal time, the diagonal d as well as the entire decomposition continues uniquely along the lift $\tilde{\alpha}$. At E , the diagonal and decomposition are no longer unique, but each can be defined by taking the limit along $\tilde{\alpha}$. The limiting diagonal e in E then develops to a semicircle contained in the equator, bounded by the antipodal vertices v_i and v_j of S .

We now define λ so that it lifts to a path in \mathcal{D}_k^* , which removes a slit sphere whose boundary contains the diagonal e . Of the pieces adjacent to e at least one — call it P — is isometric to a slit sphere, because there is only one piece isometric to a hemisphere in the decomposition. Now consider the vertices developing into the image of e . Note that by the definition of \mathcal{E}_k , these vertices are either all the vertices with index between i and j or all the complementary ones. But e decomposes the metric into two connected halves, and these are exactly the vertices that lie to one side of e . In order for the lift of λ to remove P , we revolve the vertices, which develop into e by an angle of 2π about the fixed endpoints v_i, v_j ; we must revolve in the direction of P . By a slight perturbation of α and S , if necessary, we can assume that the orbits of the revolving vertices stay disjoint from the fixed vertices. This ensures that λ lies in \mathcal{T}_k^* and so has a unique lift $\tilde{\lambda}$ in \mathcal{D}_k^* .

The complement in E of the slit sphere P has two components. We can explicitly define the path $\tilde{\lambda}$ by attaching these components to a path of lunes whose angles decrease from 2π to 0. By leaving the rest of the decomposition unchanged, we obtain a natural decomposition of the metrics along $\tilde{\lambda}$. The decomposition of the terminal metric E' of $\tilde{\lambda}$ then agrees with that of E , except that the one slit sphere is removed.

There is a unique lift $\tilde{\alpha}^{-1}$ starting from E' ; it terminates at some $D' \in F$. As with $\tilde{\alpha}$, the metrics along $\tilde{\alpha}^{-1}$ (aside from E') have a unique

decomposition. This path of decompositions agrees with that along the reversed path $\tilde{\alpha}$ except one slit sphere has been removed; the limiting decomposition of this path therefore agrees with the decomposition at E' . The small truncation of D' has one less slit sphere than the small truncation of D , so $m(D') = m(D) - 1$. \square

4.3. The space \mathcal{D}_4

We now proceed to analyze in detail the structure of the space \mathcal{D}_4 of 4-point metrics, proving that it is diffeomorphic to a 5-ball. Recall that \mathcal{T}_4^* is the 5-manifold of all pairwise distinct spherical 4-tuples (p_1, p_2, p_3, p_4) modulo rotation; the 5-manifold \mathcal{T}_4 additionally includes the 4-tuples where either $p_1 = p_3$ or $p_2 = p_4$.

Lemma 4.6. *The space \mathcal{T}_4^* is a \mathbb{C} -bundle whose base is the open 3-ball with two proper arcs removed; it thus has the homotopy type of $\mathbb{S}^1 \vee \mathbb{S}^1$, a figure-8. The space \mathcal{T}_4 is also a \mathbb{C} -bundle; the base is an open 3-ball with a proper arc and a point removed, so the homotopy type is $\mathbb{S}^1 \vee \mathbb{S}^2$.*

Proof. Let us identify \mathbb{S}^2 with $\hat{\mathbb{C}} := \mathbb{C} \cup \infty$. We use a rotation to move p_4 to $\infty \in \hat{\mathbb{C}}$. Then

$$\mathcal{T}_4 \cong (\{(p_1, p_2, p_3) : p_1, p_3 \in \mathbb{C}, p_2 \in \hat{\mathbb{C}}, p_1 \neq p_2 \neq p_3\} \setminus \{(p, \infty, p)\}) / \mathbb{S}^1,$$

where \mathbb{S}^1 acts diagonally by multiplication by $e^{i\theta}$.

We define $\mathcal{X}_4 \subset \mathcal{T}_4$ to be the subset where $p_3 = 0$. Then

$$\pi : \mathcal{T}_4 \longrightarrow \mathcal{X}_4, \quad (p_1, p_2, p_3) \mapsto (p_1 - p_3, p_2 - p_3, 0)$$

is well-defined and indeed a deformation retract. In fact, π is the desired bundle with fiber \mathbb{C} : the \mathbb{C} -valued expression $p_3/(p_2 - p_3)$ is \mathbb{S}^1 -invariant, hence well-defined on \mathcal{T}_4 , and it gives a trivialization away from $p_2 = \infty$; by symmetry a similar trivialization exists away from $p_1 = p_3$.

In our picture, $\mathcal{T}_4^* \subset \mathcal{T}_4$ is the subspace where $p_1 \neq p_3$ and $p_2 \neq \infty$. So if we define

$$\mathcal{X}_4^* := \{(p_1, p_2, 0) : p_1, p_2 \in \mathbb{C} \setminus \{0\}, p_1 \neq p_2\} / \mathbb{S}^1,$$

the bundle $\mathcal{T}_4 \rightarrow \mathcal{X}_4$ restricts to a \mathbb{C} -bundle $\mathcal{T}_4^* \rightarrow \mathcal{X}_4^*$.

Now, consider the smooth map $\varphi: \mathcal{X}_4 \rightarrow \mathbb{C} \times \mathbb{R}$ defined by $\varphi(p_1, p_2, 0) := (z, t)$, where

$$z := \frac{p_1}{p_2}, \quad t := \log(1 + |p_1|) - \log\left(1 + \frac{1}{|p_2|}\right) = \log\left(\frac{1 + |p_1|}{1 + |z|/|p_1|}\right).$$

The map φ is well-defined, since it is unchanged when multiplying p_1 and p_2 by the same factor $e^{i\theta}$. Note that z is a cross-ratio of the 4-tuple

$$(p_1, p_2, p_3, p_4) = (p_1, p_2, 0, \infty).$$

We now examine the behavior of φ on \mathcal{X}_4^* ; here z takes every value except 0 and 1. For any fixed $z = p_1/p_2$, the function

$$r \mapsto \log\left(\frac{1 + r}{1 + |z|/r}\right)$$

is strictly increasing from \mathbb{R}^+ onto \mathbb{R} . Thus \mathcal{X}_4^* is diffeomorphic to its image $\varphi(\mathcal{X}_4^*) = (\mathbb{C} \setminus \{0, 1\}) \times \mathbb{R}$, and so has the homotopy type of $\mathbb{S}^1 \vee \mathbb{S}^1$. This completes the proof for \mathcal{T}_4^* .

The space \mathcal{X}_4 includes two open arcs in addition to \mathcal{X}_4^* , namely the arcs γ_{13} , where $p_1 = p_3$, and γ_{24} , where $p_2 = p_4$. Under φ , both arcs map to the line $z = 0$; the coordinate t maps $(0, p_2, 0) \in \gamma_{13}$ to $-\log(1 + 1/|p_2|)$, and $(p_1, \infty, 0) \in \gamma_{24}$ to $\log(1 + |p_1|)$. Hence this coordinate maps \mathbb{R}^+ diffeomorphically onto \mathbb{R}^- and \mathbb{R}^+ , respectively. Thus, $\varphi(\mathcal{X}_4)$ is the union of $\varphi(\mathcal{X}_4^*)$ with $0 \times (\mathbb{R} \setminus 0)$. We conclude φ is a diffeomorphism from \mathcal{X}_4 to its image, which is all of $\mathbb{C} \times \mathbb{R}$ except for the origin and the line $\{1\} \times \mathbb{R}$. This is homotopy-equivalent to $\mathbb{S}^2 \vee \mathbb{S}^1$, completing the proof for \mathcal{T}_4 . \square

Remembering that \mathcal{D}_4^* is a covering space of \mathcal{T}_4^* , we know it will be a \mathbb{C} -bundle over the corresponding cover \mathcal{Y}_4^* of \mathcal{X}_4^* . Letting $\mathbb{Z}_{\text{odd}} := 2\mathbb{Z} + 1$ denote the odd integers, we have

Lemma 4.7. *The space \mathcal{D}_4^* is a \mathbb{C} -bundle over $\mathcal{Y}_4^* \cong (\mathbb{C} \setminus \mathbb{Z}_{\text{odd}}) \times \mathbb{R}$.*

Proof. The fundamental group of \mathcal{T}_4^* is $\mathbb{Z} * \mathbb{Z}$. The first generator β is a loop represented by p_1 circling counterclockwise around p_2 ; in terms of $\varphi(\mathcal{X}_4^*) = (\mathbb{C} \setminus \{0, 1\}) \times \mathbb{R}$, it corresponds to a loop around the deleted line at 1. The second generator γ is represented by p_1 circling clockwise around p_3 (or equivalently, counterclockwise around both p_2 and p_4); in $\varphi(\mathcal{X}_4^*)$, it is a loop around the deleted line at 0.

We claim that every lift of γ to \mathcal{D}_4^* is a loop, while the lifts of β are not. This means that the covering $\mathcal{D}_4^* \rightarrow \mathcal{T}_4^*$ is normal, corresponding to the normal subgroup generated by γ . The covering map $\mathcal{Y}_4^* \rightarrow \mathcal{X}_4^*$ can be then explicitly given as $(w, t) \mapsto (1 + e^{\pi i w}, t)$ with $\mathcal{Y}_4^* \cong (\mathbb{C} \setminus \mathbb{Z}_{\text{odd}}) \times \mathbb{R}$, and we can represent \mathcal{D}_4^* as the pullback \mathbb{C} -bundle over \mathcal{Y}_4^* .

To prove the claim, we will work over the deformation retract \mathcal{X}_4^* . Because $p_3 = 0$ and $p_4 = \infty$ are antipodal, truncations of metrics projecting to \mathcal{X}_4^* are not unique. To overcome this difficulty, we adopt the convention that we always truncate that end along an arc developing to the positive reals. This corresponds to viewing the point $(p_1, p_2, 0)$ as a limit of (p_1, p_2, ε) for ε positive real.

We fix a basepoint $P := (1, 2, 0) \in \mathcal{X}_4^* \subset \mathcal{T}_4^*$. We can use our decomposition results to analyze any metric in the fiber over P . Because of the order of the four points in the circle, and our convention for drawing diagonals p_3p_4 , each decomposition piece is isometric to a slit sphere. Furthermore, exactly two pieces are 3-gons, with vertices either $p_1p_2p_3$ and $p_1p_3p_4$, or $p_1p_2p_4$ and $p_2p_3p_4$. Metrics in the fiber over P are completely characterized by the choice of one of these two types, together with the nonnegative number of slit spheres between the two triangles.

More explicitly, suppose we truncate a 4-point metric D in the fiber over P , using again a diagonal p_3p_4 which projects to the positive reals. The 4-gon D' obtained has four interior angles $\alpha_i := \alpha(p_i)$; each of these is an integer multiple of 2π . The alternating sum

$$\alpha_1 - \alpha_2 + \alpha_3 - \alpha_4 = 2\pi n$$

is unchanged if we add 2π to two consecutive angles, and hence is independent of the truncation D' we have chosen. Note that the area of D' and the number $2(\alpha_2 + \alpha_4)$ are each multiples of 4π . Thus, the Gauss–Bonnet Theorem shows that n must be odd; the odd integer n uniquely characterizes a metric $D = D_n$ in the fiber over P . This also follows from our combinatorial picture: indeed, D_n has a diagonal p_1p_3 when $n > 0$ and a diagonal p_2p_4 when $n < 0$; the number of slit spheres along this diagonal is $(|n| - 1)/2$.

Now, choose a truncation D' of D_n with at least one slit sphere added at each end. The loops β and γ in \mathcal{X}_4^* can be lifted to paths in \mathcal{D}_4 with initial point D_n . These paths also have a unique lift starting at D' in the space \mathcal{P}_4 of 4-gons, provided no angle α_i becomes zero. Indeed, this is easy to see if the paths are chosen to avoid antipodal pairs in the 4-tuple: then the unique geodesics between them lead to a unique continuation of the truncation.

Let us now observe the effect of these lifts of β and γ on the angles α_i . Since only p_1 moves and the edge p_3p_4 stays fixed (by our convention), the angle α_3 remains constant. For γ , as p_1 moves around p_3 (which is antipodal to p_4) the edge $\overline{p_1p_4}$ rotates once around the sphere. This means that α_1 and α_4 each decrease by 2π . Since the angle α_2 stays nearly constant all along γ , it must return to its initial value. We conclude that n is fixed, and so γ lifts to a loop in \mathcal{D}_4^* . For β , on the other hand, α_4 is nearly constant, α_2 increases by 2π and α_1 decreases by 2π ; thus n decreases by 2. Thus β changes the metric D_n to D_{n-2} . This completes the proof of the claim. \square

Proposition 4.8. *The space \mathcal{D}_4 is diffeomorphic to an open 5-ball.*

Proof. Remember that $\mathcal{X}_4 \setminus \mathcal{X}_4^*$ consists of two arcs γ_{13} and γ_{24} . The decision of where to fill in copies of these arcs into the covering space \mathcal{Y}_4^* is purely combinatorial, based on the decompositions of the metrics in $\mathcal{D}_4 \setminus \mathcal{D}_4^*$.

We can connect our basepoint $P = (1, 2, 0)$ to the fibers over these arcs by simple paths: the first path moves only p_1 (along the reals to 0); the second moves only p_2 (along the reals to ∞). We use these paths to extend the odd labels n from the fiber over P . Now, we note that $p_i = p_j$ is allowed in \mathcal{D}_4 exactly when there is no diagonal $p_i p_j$. Thus, the arc γ_{13} gets filled in exactly at the levels with negative label $n < 0$. Similarly, the arc γ_{24} gets filled in exactly at the levels with $n > 0$.

This means that at each level, exactly one of the two arcs gets filled in. Again, the decision is purely combinatorial: whenever we fill in a single point, we fill in an entire arc into \mathcal{Y}_4^* , and \mathcal{D}_4 contains the entire \mathbb{C} -bundle over this arc. In the picture $\mathcal{Y}_4^* \cong (\mathbb{C} \setminus \mathbb{Z}_{\text{odd}}) \times \mathbb{R}$, we fill in the lower half ($t < 0$) of the missing lines where $z = n < 0$, and the upper half ($t > 0$) of the others where $z = n > 0$; this leaves us with a space \mathcal{Y}_4 diffeomorphic to a 3-ball. Finally, \mathcal{D}_4 is a \mathbb{C} -bundle over this ball, which must be trivial, meaning that \mathcal{D}_4 is a five-ball. \square

5. Injectivity and Properness

To show the classifying map Φ is continuous, injective and proper, we follow the approach we used for the case $k = 3$ of triunduloids. The proof of properness is where substantial modifications are needed. In particular, since the classifying space is \mathcal{D}_k , not \mathcal{T}_k , we need to characterize the compact subsets of \mathcal{D}_k for general k .

5.1. Injectivity

As for triunduloids we have for coplanar k -unduloids of genus zero.

Theorem 5.1. *The classifying map Φ is injective.*

Proof. The proof is essentially the same as for triunduloids. We sketch the main steps and refer to [11] for details.

Suppose two k -unduloids $M, M' \in \mathcal{M}_k$ map to the same k -point metric $D := \Phi(M) = \Phi(M') \in \mathcal{D}_k$. We need to show $M = M'$. We pull back the \mathbf{k} -Hopf bundle $\mathbb{S}^3 \rightarrow \mathbb{S}^2$ to the disk D by the developing map. Taking the universal cover of the fibers, we obtain a line bundle $E \rightarrow D$. It carries a Riemannian structure, locally isometric to \mathbb{S}^3 . The \mathbf{k} -Hopf flow acts isometrically on E , moving points vertically along the fibers.

The cousin disks \widetilde{M}^+ and \widetilde{M}'^+ can be considered sections of E . Over each of the k -rays of the decomposition of D given by Theorem 3.8, the \mathbf{k} -Hopf flow foliates the bundle E with unduloid cousins. Over each ray, the cousin disks are asymptotic to one of the leaves. But across each end of M and M' the *periods* (see [11, Section 3.3]) vanish, which by [11, Lemma 3.7] implies for the cousins that the \mathbf{k} -Hopf flow can move \widetilde{M}^+ to become asymptotic to \widetilde{M}'^+ over all rays simultaneously. This allows us to deal with \widetilde{M}^+ and \widetilde{M}'^+ as if they were compact; as in [11, Theorem 3.8] the maximum principle then shows $\widetilde{M}^+ = \widetilde{M}'^+$. Therefore $M = M'$, up to translation. \square

5.2. Properness

To show properness, we want to characterize compact subsets of \mathcal{D}_k . We need the following notion. The *vertex distance* $d(D)$ of $D \in \mathcal{D}_k$ is the minimum distance between any pair of vertices in the completion \hat{D} .

Proposition 5.2. *For any compact $K \subset \mathcal{D}_k$, there exist $m, \varepsilon > 0$ such that for all $D \in K$*

- (i) *the vertex distance $d(D)$ is at least ε , and*
- (ii) *there is a small truncation D^\square with at most m decomposition pieces.*

Proof. (i) The distance between any two given vertices depends continuously on $D \in \mathcal{D}_k$ and is always positive; thus, it assumes a positive minimum over the compact set K . Minimizing over the finitely many pairs of distinct vertices still yields a positive minimum for $d(D)$. (It is not hard to see that,

in fact, $d(D)$ is realized by the spherical distance of some pair of distinct points of $\Pi(D) \in \mathcal{T}_k$.)

(ii) Consider a metric $D \in \mathcal{D}_k$ which admits a small truncation with n decomposition pieces. If no piece is a lune or a concave polygon with angle π then D has a neighborhood of metrics with the same number of pieces. Otherwise, at most k additional pieces are needed.

Now let $V_j \subset \mathcal{D}_k$ be the set of metrics, which admit a small truncation with at most $j - k$ decomposition pieces. The above discussion shows V_j has an open neighborhood $U_j \subset \mathcal{D}_k$, such that any metric in U_j admits a truncation with no more than j decomposition pieces. The $\{U_j\}$ define an open exhaustion of \mathcal{D}_k , so that the compact set K is contained in some U_m . □

Lemma 5.3. *Suppose K is a subset of \mathcal{D}_k such that the vertex distance of each $D \in K$ is at least $\varepsilon > 0$. Then there is $C = C(\varepsilon) > 0$ such that the curvature estimate $\sup_M |A| < C(\varepsilon)$ holds for each $M \in \Phi^{-1}(K)$.*

Proof. We modify the argument given in [11, Section 4]. Suppose on the contrary that $\sup_{M_j} |A| \rightarrow \infty$ with $\Phi(M_j) \in K$. Then by [11, Lemma 4.2] we find a sequence of rescaled surfaces M'_j , converging to some minimal surface N , which is Alexandrov-symmetric across the plane P . Note that, as in [11, Proposition 4.4], the M_j have genus zero, meaning that $M_j \cap P$ has no closed loops. Since the convergence is with multiplicity one, each component of $N \cap P$ is likewise an unbounded curve.

By [11, Lemma 4.3], the minimal surface N contains a simple closed curve γ with nonzero force, symmetric with respect to P and not contained in P . So the surfaces M'_j contain closed curves γ'_j , whose lengths and forces converge to those of γ . Now, consider the closed curves γ_j , on the unrescaled surfaces M_j . Note that the length of γ_j converges to 0. For sufficiently large j , however, the force of γ'_j is nonzero, and so γ'_j as well as γ_j must be homologically nontrivial. But since the M_j have genus zero this means that the curve γ_j must connect two different components of $M_j \cap P$. Therefore, the conjugate curve $\tilde{\gamma}_j^+$ to its upper half γ_j^+ Hopf projects to a curve η_j in $\Phi(M_j) \in \mathcal{D}_k$, which connects two different vertices.

By definition, the length of any such curve is at least the vertex distance $d(D) \geq \varepsilon$. Since, the Hopf projection $\Pi_{\mathbf{k}}$ has Lipschitz constant 2, it follows that the length of $\tilde{\gamma}_j^+$, and thus also the length of γ_j , is bounded below. This contradicts the fact that γ_j has length converging to 0. □

Lemma 5.4. *Given a compact subset $K \subset \mathcal{D}_k$ there is $R > 0$ with the following property. For each $M \in \Phi^{-1}(K)$ there exists a ball $B_R \in \mathbb{R}^3$ of*

radius R such that the noncompact components of $M \setminus B_R$ are exactly k annular ends, each of which is contained in a solid cylinder of radius 3.

Proof. The enclosure theorem, proved in [17] for embedded surfaces, carries over without change to say that any Alexandrov-embedded CMC surface M is contained in the union of k solid cylindrical rays as well as a finite number of balls of radius 11 and solid cylindrical segments of radius 3; the numbers of these depend only on k and the genus. Collecting the balls and finite cylindrical segments into one larger ball, we see that for each $M \in \mathcal{M}_k$ there is an $R = R(M)$ with the property desired. Note that \mathcal{M}_k contains sequences of k -unduloids where the length of some midsegment, and hence R , diverges. However, we claim that, over K , there is a uniform choice of R .

Suppose, by contradiction, that for some sequence $M_j \in \Phi^{-1}(K)$ we have $R(M_j) \rightarrow \infty$. Consequently, the enclosing sets for the M_j have solid cylindrical segments with length $l(j)$ tending to infinity. These segments contain annuli $A_j \subset M_j$. By the uniform bound on the vertex distance from Proposition 5.2(i), the force carried by these annuli are uniformly bounded away from 0. Thus, as $l(j) \rightarrow \infty$, it follows as in [19, Theorem 5.1] that after rigid motion a subsequence of the A_j converge to some (complete) unduloid.

Since this unduloid has a cousin which Hopf-projects to a line of slit spheres, we find that the spherical metrics for A_j must contain increasing numbers of slit spheres. This contradicts Proposition 5.2(ii). \square

Theorem 5.5. *The classifying map $\Phi: \mathcal{M}_k \rightarrow \mathcal{D}_k$ is continuous and proper.*

Proof. Given a compact $K \subset \mathcal{D}_k$, Lemma 5.3 (combined with Proposition 5.2(ii)) gives a uniform curvature estimate for k -unduloids in $\Phi^{-1}(K)$. The continuity of Φ on this set then follows just as in [11, Theorem 4.5].

To show that $\Phi^{-1}(K)$ is compact, we now prove that any sequence $M_j \in \Phi^{-1}(K)$ subconverges with respect to the topology on \mathcal{M}_k . The M_j are given only up to horizontal translation; to get convergence, we choose representatives such that the enclosing balls B_R from Lemma 5.4 are centered at the origin. Note also that by [17], there is a uniform area bound, for any coplanar genus-zero k -unduloid M , of the form $\text{area}(M_j \cap B_r) \leq Cr^2$. This area bound, together with the curvature estimate of Lemma 5.3, now gives a subsequence of the M_j that converges (on each compact subset) to some CMC surface M .

By smooth convergence, it is clear that M is Alexandrov-embedded. Because a sequence of contractible loops has a contractible limit, M has

genus zero and at most k -ends. On the other hand, by Lemma 5.4, on each M_j there are k closed curves in $B_{R+1}(0) \setminus B_R(0)$ which bound the k -ends. These k sequences of curves subconverge to k closed curves, which are disjoint in the limit surface M , and thus bound k different ends. Thus, M is a coplanar k -unduloid, as desired. \square

6. Surjectivity of the classifying map

6.1. Moduli spaces of coplanar k -unduloids

Given a minimal or CMC surface M , a solution u to the linearized mean-curvature equation $L(u) := \Delta_M u + |A|^2 u = 0$ is called a *Jacobi field*. The surface M is called *nondegenerate* if all square-integrable Jacobi fields $u \in L^2(M)$ vanish.

Let $\mathcal{M}_{g,k}$ denote the moduli space of all k -unduloids of genus g . Kusner, Mazzeo and Pollack [21] showed that $\mathcal{M}_{g,k}$ is a real-analytic variety (for $k \geq 3$), and that near any nondegenerate surface it is a $(3k - 6)$ -dimensional manifold. We now prove a similar result for coplanar k -unduloids:

Theorem 6.1. *The moduli space \mathcal{M}_k of coplanar k -unduloids of genus zero is locally a real-analytic variety of finite dimension. In a neighborhood of a nondegenerate coplanar k -unduloid, it is a manifold of dimension $2k - 3$.*

In their proof, [21] consider the *premoduli space* $\tilde{\mathcal{M}}_{g,k}$, which is the space of k -unduloids of genus g *before* dividing by euclidean motions. At a nondegenerate surface M , they show the premoduli space $\tilde{\mathcal{M}}_{g,k}$ is locally a $3k$ -dimensional Lagrangian submanifold of a symplectic vector space W . Here, each end of M has a six-dimensional space of geometric Jacobi fields, induced by translations, rotations, and change of asymptotic necksize. The direct sum of these, over the k ends of M , forms the $6k$ -dimensional space W . The tangent space to $\tilde{\mathcal{M}}_{g,k}$ at M is the Lagrangian subspace consisting of all globally defined Jacobi fields on M .

Proof. In the premoduli space $\tilde{\mathcal{M}}_{0,k}$ of all genus-zero k -unduloids, let $\tilde{\mathcal{M}}_k$ denote the subset consisting of coplanar k -unduloids with Alexandrov symmetry across the xy -plane P . The reflection τ across P acts on $\tilde{\mathcal{M}}_{0,k}$, fixing $\tilde{\mathcal{M}}_k$.

Near any surface M , the variety $\tilde{\mathcal{M}}_{0,k}$ is defined by real-analytic equations on a finite-dimensional vector space $W' \supset W$. Our premoduli space $\tilde{\mathcal{M}}_k$ is obtained locally by adding one more set of linear equations, $u \circ \tau = u$, and

is thus still real-analytic. Dividing by euclidean motions of the xy -plane P , we find the same is true for \mathcal{M}_k .

To compute the dimension of \mathcal{M}_k , we show that $\tilde{\mathcal{M}}_k$ has dimension $2k$ in the neighborhood of a nondegenerate surface M . We use the following transversality argument. The reflection τ induces an automorphism τ^* on the $6k$ -dimensional vector space W , which encodes the asymptotic behavior of Jacobi fields on M . This automorphism preserves the natural symplectic form on W . The global Jacobi fields on M form a Lagrangian subspace $L \subset W$, which corresponds to the tangent space to $\tilde{\mathcal{M}}_{0,k}$ at M . Since τ takes one global Jacobi field on M to another, τ^* carries L into itself. The fixed point set of τ^* is a $4k$ -dimensional vector space V , consisting of those asymptotic changes that keep the ends in the plane P . The following lemma then says that $L \cap V$, corresponding to the tangent space of $\tilde{\mathcal{M}}_k$ at M , is $2k$ -dimensional. \square

Lemma 6.2. *Suppose L is a Lagrangian (maximal isotropic) subspace of a symplectic vector space W , and suppose G is a compact group of symplectic automorphisms of W preserving L . Then the fixed-point set W^G of G is a symplectic subspace $W^G \subset W$, and $L \cap W^G$ is Lagrangian in W^G .*

Proof. By averaging over the group G , we can choose a compatible almost complex structure J on W that is also invariant under G . That means $g \circ J = J \circ g$ for all $g \in G$. Let $W^g := \ker(g - I)$, which is preserved by J . Then the fixed set W^G is the intersection $W^G = \bigcap_{g \in G} W^g$. In particular, $J(W^G) = W^G$, showing this is symplectic.

Because L is Lagrangian, we have $W = L \oplus JL$, so any vector in W can be represented as $v + Jw$, where $v, w \in L$. If, moreover, $v + Jw$ is in W^g then $g(v + Jw) = v + Jw$ and so $-(g - I)v = (g - I)Jw = J(g - I)w$. But since g (and therefore $g - I$) preserves L , this equates a vector in L with one in JL . Thus $(g - I)v = 0 = (g - I)w$, so that in fact $v, w \in L \cap W^g$. Since the preceding argument holds for all $g \in G$ we obtain a decomposition $W^G = (L \cap W^G) \oplus J(L \cap W^G)$, meaning that $L \cap W^G$ is Lagrangian in W^G . \square

6.2. Existence of nondegenerate coplanar k -unduloids

A *minimal k -noid* is a finite-total-curvature minimal surface of finite topology with genus zero which has only catenoidal ends. Jorge and Meeks determined the k -noids which are coplanar and dihedrally symmetric. They have only two umbilic points, so that the result of Montiel and Ros [25] applies (in the form mentioned in [11]) to give

Proposition 6.3. *The Jorge–Meeks k -noids are nondegenerate.*

As in the case of triunduloids, we want to use these minimal surfaces in the gluing technique of Mazzeo and Pacard [23] to produce nondegenerate coplanar CMC surfaces. Their theorem generalizes to an equivariant setting as follows:

Theorem 6.4. *Let $M_0 \subset \mathbb{R}^3$ be a nondegenerate minimal k -noid that can be oriented so that each end has normal pointing inwards towards its asymptotic axis. Assume M_0 has symmetry group G in \mathbb{R}^3 , consisting of isometries preserving the normal. Then for some ε_0 , there is a family M_ε , $0 < \varepsilon < \varepsilon_0$, of nondegenerate CMC surfaces with k embedded ends, also with symmetry group G , such that on any compact set in \mathbb{R}^3 , the dilated surfaces $\frac{1}{\varepsilon}M_\varepsilon$ converge smoothly and uniformly to M_0 .*

Sketch of proof. The proof of this theorem follows exactly the argument of Mazzeo and Pacard [23], so we merely outline how to adapt their argument to our G -invariant setting. They first glue exactly unduloidal ends to M_0 , truncated along k circles, and then consider normal perturbations. They slightly perturb the bounding circles, and derive solution operators for the respective Dirichlet problems for the mean curvature equation, both over M_0 and over the ends. Unique solutions to these Dirichlet problems are established as fixed points of a contraction mapping [23, Corollaries 5, 8]; thus G -invariant Dirichlet data lead to G -invariant solutions. The surface M_ε is then obtained by a Cauchy matching of a pair of Dirichlet solutions. This matching is accomplished [23, Proposition 32] by applying the Leray–Schauder fixed-point theorem to a nonlinear operator F_ε on the Dirichlet data. The operator F_ε is designed to have a fixed point exactly when the first-order Cauchy data match.

Given a G -invariant initial surface M_0 we note that G acts by composition on all relevant function spaces. Also, the operators used in the various fixed-point arguments, being geometrically defined, all commute with G . Therefore, we can simply replace the function spaces by their closed linear subspaces consisting of G -invariant functions. Clearly, all estimates in [23] continue to hold, and so the contraction-mapping and Leray–Schauder fixed-point theorems still apply, giving us a G -invariant CMC surface M_ε .

The arguments of [23, Section 12] show that for sufficiently small ε any function in $L^2(M)$ (whether G -invariant or not) satisfying the linearized mean curvature equation vanishes. Thus, M_ε is nondegenerate. \square

Now let M_0 be the Jorge–Meeks k -noid of Proposition 6.3, and G be its symmetry group. Applying Theorem 6.4 gives

Corollary 6.5. *For each $k \geq 3$ there exists a nondegenerate coplanar genus-zero k -unduloid.*

6.3. Proof of main theorem and corollary

In [11, Section 5.3] we show

Theorem 6.6. *If $f: X \rightarrow Y$ is a continuous, proper, injective map from a d -dimensional real-analytic variety X to a connected d -manifold Y , then f is surjective and thus a homeomorphism.*

We now have all the ingredients needed to prove

Theorem 6.7. *The classifying map $\Phi: \mathcal{M}_k \rightarrow \mathcal{D}_k$ is a homeomorphism between connected $(2k - 3)$ -manifolds.*

Proof. To apply Theorem 6.6 to Φ we need to verify its assumptions. The space \mathcal{D}_k is a connected $(2k - 3)$ -manifold by Theorems 4.3 and 4.5. The map Φ is proper and injective by Theorems 5.5 and 5.1.

From Theorem 6.1, the moduli space \mathcal{M}_k of coplanar k -unduloids is locally a real-analytic variety; injectivity of Φ shows it has dimension at most $2k - 3$. But by Corollary 6.5 there is a nondegenerate coplanar k -unduloid, so then by Theorem 6.1, \mathcal{M}_k is a $(2k - 3)$ -manifold in some neighborhood of this k -unduloid.

Thus \mathcal{M}_k has dimension exactly $2k - 3$, so Theorem 6.6 applies to conclude that Φ is a homeomorphism. \square

Together with Theorem 2.1, this proves our Main Theorem.

The corollary on necksizes stated in the introduction collects the statements of Corollary 2.2 and Corollary 4.2. The existence statement follows from the example of Section 3.1; in particular the example shows that (1.1) and (2.1) are sharp.

References

- [1] A.D. Alexandrov, *Uniqueness theorems for surfaces in the large V*, Vestnik Leningrad Univ. **19**(13) (1958), 5–8; Amer. Math. Soc. Transl. (Ser. 2) **21** (1962), 412–416.

- [2] A.D. Alexandrov, *A characteristic property of spheres*, Ann. Math. Pure Appl. **58** (1962), 303–315.
- [3] R. Buckman, and N. Schmitt, *Spherical polygons and unitarization*, Preprint, 2002; www.gang.umass.edu/reu/2002/polygon.html
- [4] C. Cosín, and A. Ros, *A Plateau problem at infinity for properly immersed minimal surfaces with finite total curvature*, Indiana Univ. Math. J. **50** (2001), 847–879.
- [5] C. Delaunay, *Sur la surface de révolution, dont la courbure moyenne est constante*, Journal de mathématiques **6** (1841), 309–320.
- [6] P. Groisman, *A new family of symmetric constant-mean-curvature surfaces*, PhD thesis, University of Illinois, 2004.
- [7] K. Große-Brauckmann, *New surfaces of constant mean curvature*, Math. Z. **214** (1993), 527–565.
- [8] K. Große-Brauckmann, *Cousins of constant mean curvature surfaces*, Global theory of minimal surfaces, Clay Mathematics Proceedings. Amer. Math. Soc., Providence, RI, 2005, pp. 747–767.
- [9] K. Große-Brauckmann, N. Korevaar, R. Kusner, J. Ratzkin and J.M. Sullivan, *Coplanar k -unduloids are nondegenerate*, Preprint, 2007.
- [10] K. Große-Brauckmann, and R. Kusner, *Embedded constant mean curvature surfaces with special symmetry*, Man. Math. **99** (1999), 135–150.
- [11] K. Große-Brauckmann, R. Kusner, and J.M. Sullivan, *Triunduloids: Embedded constant mean curvature surfaces with three ends and genus zero*, J. Reine Angew. Math. **564** (2003), 35–61. arXiv:math.DG/0102183
- [12] K. Große-Brauckmann, R. Kusner, and J.M. Sullivan, *Constant mean curvature surfaces with three ends*, Proc. Natl. Acad. Sci. USA **97** (2000), 14067–14068. arXiv:math.DG/9903101
- [13] K. Große-Brauckmann, R. Kusner, and J.M. Sullivan, *Constant mean curvature surfaces with cylindrical ends*, Mathematical visualization, Springer, Berlin, 1998, pp. 106–117.
- [14] N. Kapouleas, *Complete constant mean curvature surfaces in Euclidean three-space*, Ann. of Math. **131** (1990), 239–330.

- [15] H. Karcher, *The triply periodic minimal surfaces of A. Schoen and their constant mean curvature companions*, Man. Math. **64** (1989), 291–357.
- [16] M. Kilian, N. Schmitt, and I. Sterling, *Dressing CMC n -noids*, Math. Z. **246** (2004), 501–519.
- [17] N. Korevaar, and R. Kusner, *The global structure of constant mean curvature surfaces*, Invent. Math. **114** (1993), 311–332.
- [18] N. Korevaar, R. Kusner, and J. Ratzkin, *On the nondegeneracy of constant mean curvature surfaces*, Geom. Funct. Anal. **16** (2006), 891–923; arXiv:math.DG/0407153
- [19] N. Korevaar, R. Kusner, and B. Solomon, *The structure of complete embedded surfaces with constant mean curvature*, J. Diff. Geom. **30** (1989), 465–503.
- [20] R. Kusner, *Bubbles, conservation laws and balanced diagrams*, *Geometric Analysis and computer graphics* (Berkeley, CA, 1988) pp. 103–108, Mathematical Science Research Institute Publications, **17**, Springer, New York, 1991.
- [21] R. Kusner, R. Mazzeo, and D. Pollack, *The moduli space of complete embedded constant mean curvature surfaces*, Geom. Funct. Anal. **6** (1996), 120–137.
- [22] H.B. Lawson, *Complete minimal surfaces in \mathbb{S}^3* , Ann. Math. **92** (1970), 335–374.
- [23] R. Mazzeo, and F. Pacard, *Constant mean curvature surfaces with Delaunay ends*, Comm. Anal. Geom. **9** (2001), 169–237.
- [24] W.H. Meeks, *The topology and geometry of embedded surfaces of constant mean curvature*, J. Diff. Geom. **27** (1988), 539–552.
- [25] S. Montiel, and A. Ros, *Schrödinger operators associated to a holomorphic map*, *Global differential geometry and global analysis*, Lecture Notes in Mathematics **1481**, Springer, Berlin, 1991, pp. 148–174.
- [26] A.H. Schoen, *Infinite periodic minimal surfaces without selfintersections*, NASA Technical Note TN D-5541 (1970).

TECHNISCHE UNIVERSITÄT DARMSTADT
FB. MATHEMATIK (AG 3)
SCHLOSSGARTENSTR. 7
64289 DARMSTADT
GERMANY
E-mail address: kgb@mathematik.tu-darmstadt.de

MATHEMATICS DEPT.
UNIVERSITY OF MASSACHUSETTS
AMHERST MA 01003
USA
E-mail address: kusner@math.umass.edu

TECHNISCHE UNIVERSITÄT BERLIN
MA 3-2
STR. DES 17. JUNI 136
10623 BERLIN
GERMANY
E-mail address: sullivan@math.tu-berlin.de

RECEIVED NOVEMBER 20, 2006