BOUNDARY PLATEAU LAWS I

CAMILLO DE LELLIS AND FEDERICO GLAUDO

Abstract. We give a complete list of the possible local shapes of soap films where they wet the whole wire. In rigorous terms, given a straight line $\ell \subset \mathbb{R}^3$, we study 2-dimensional cones in $\mathbb{R}^3$ which contain ℓ and are Almgren minimizing with respect to ℓ.

Contents

1. Introduction 1
2. Plan of the paper 8
3. Index of Concepts and Notation 10
4. Stationary nets with two pins 11
5. Derived nets and topological graphs 16
6. Polygonal paths and polygons 23
7. Sieves 26
8. Computer-assisted list of candidate topologies 30
9. Useful inequalities for convex faces 37
10. Relaxed Almgren competitors 39
11. Excluding basic degenerations 46
12. Path-Isometries 56
13. Computer-assisted restriction for metric realizations 60
14. Special cases: extinct nets 66
15. Special cases II: seven difficult families 69
16. From triangulated competitors to Almgren minimality 84
17. Quantitative stability of competitors 95
18. Area-Decreasing Competitors and proof of Theorem 1.3 104
19. Minimality: known cases and open problems 111
Appendix A. Allard's theory and minimizing sets 114
Appendix B. Convex spherical open sets and positively curved arcs 120
Appendix C. Proof of Proposition 6.8 125
Appendix D. Proof of Theorem 6.7 129
Appendix E. Recognizing isomorphic planar maps 130
Appendix F. Computations for the proofs of Propositions 11.2, 11.5, and 11.6 132
Appendix G. Proof of Lemma 15.5 137
References 139

1. Introduction

1.1. **Informal explanation of the main theorem.** Formulated in the nineteenth century by the Belgian physicist Joseph Plateau [23], the classical Plateau laws describe the local

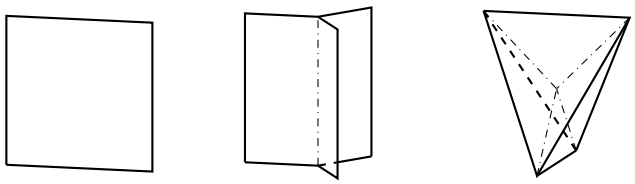


FIGURE 1. The three possible local geometries of a soap film away from the boundary.

geometry of a soap film at points away from the boundary. They state that the film is locally diffeomorphic to one of the following: a plane, three half-planes meeting at 120° angles along a line, or the cone over the 1-skeleton of a regular tetrahedron (see Figure 1). In this context, the "boundary" corresponds to the physical frame – usually a metal or plastic wire – dipped into soapy water. The film attaches to this wire, which dictates the boundary conditions.

The aim of this paper is to classify the possible shapes at a point p located at the boundary Γ, under the assumption that the film wets all of Γ close to p. Assuming that Γ is regular we can move p to the origin of a Cartesian system of coordinates and zoom in. This allows us to treat the boundary as an infinite straight line ℓ going through the origin and the film as an area-minimizing infinite cone containing ℓ (cf. Corollary 1.8 for the rigorous statement). Our list of possible shapes at p comprises six families, depicted in Figures 3, 4, 5 and 2.

List 1.1.

(a) A halfplane with boundary ℓ.
(b) Two halfplanes with boundary ℓ forming a convex angle $\geq \frac{2\pi}{3}$.
(c) Three halfplanes bounded by ℓ forming angles of $\frac{2\pi}{3}$.
(d) The cone over the skeleton of a regular tetrahedron (which contains 3 distinct lines).
(e) The cone over the skeleton Q of a cube.
(f) An element of a family which we call "thin fish", cf. Figure 2 and Definition 1.5.

1.2. **Rigorous statements.** In order to state our results in rigorous terms we first introduce the relevant objects in a more general setting.

Definition 1.2 (Almgren minimizing sets)**.** Let $k < N$ be two given positive integers, $U \subset \mathbb{R}^N$ an open subset, and Γ, E two relatively closed subsets of U. E is a k-dimensional Almgren minimizing set with respect to Γ if E has locally finite $\mathcal{H}^k$ measure in U and the inequality

$$\mathcal{H}^k(E \cap V) \leq \mathcal{H}^k(\Phi(E) \cap V) \tag{1.1}$$

holds for every open set $V \subset\subset U$ and every Lipschitz map Φ with $\overline{W \cup \Phi(W)} \subset\subset V \setminus \Gamma$, where $W = \{x : x \neq \Phi(x)\}$.

With the latter definition in hand, the focus of this paper is the classification of all 2-dimensional cones which are Almgren minimizing with respect to a line ℓ contained in them.

Theorem 1.3. *Assume $\ell \subset \mathbb{R}^3$ is a line and $E \supset \ell$ a closed cone (namely invariant under the homotheties $x \mapsto \lambda x$, $\lambda > 0$) which is Almgren minimizing with respect to ℓ. Then E belongs to List 1.1.*

To complete the statement of Theorem 1.3 we describe in Definition 1.5 the cones in List 1.1(f). Theorem 1.3 has been achieved through a three year project, during which many steps were completed with the assistance of computers in a "traditional sense", namely

through numerical simulations to guide our intuition and through the implementation of suitable computer algorithms to check automatically a large number of cases, after developing theoretical tools which allow a reduction to a finite case analysis. Except for helping with some visualization tools and for a few small fine-tuning of some arguments, we did not use any AI tool to complete this work.

Remark 1.4. Up until a few days before the public release of this manuscript, we were convinced that all the elements of List 1.1 were in fact Almgren minimizing, with the exception of those elements in (e) for which the line ℓ is close to the midpoint of an edge of Q. Our opinion was motivated by both real life and numerical experiments. For instance we dipped wires shaped as the links in List 1.1(f) and in List 1.1(e) in soapy water on several occasions and we invariably saw the films forming cones (with the exceptions mentioned above among the links in List 1.1(e)). The same behavior has been witnessed also by Guy David, who ran independent experiments with different models of the most symmetric element of List 1.1(f), see [8]. Similarly, all the numerical simulations run by us using Brakke's surface evolver confirmed what we saw in the real life experiments.

A few days before releasing this manuscript we doublechecked one of our conjectures, prompting an LLM with the case of the most symmetric element in List 1.1(f). The LLM suggested a relatively simple way to construct a competitor for it and, after understanding the relevant construction and the computations, we then concluded that the same idea could be applied to prove that no elements in List 1.1(e)&(f) are Almgren minimizing. Since we can prove that all the others in List 1.1 are minimizing, this in fact gives a *complete* classification. We also realized that a combination of this new insight and some of the ideas of this paper would give a much shorter proof which is entirely "pen and paper", namely with no need for a computer assisted part. We then asked LLMs to complete a second manuscript with this proof, giving a rather detailed plan of work for the overall strategy and closely steering them during the writing process. In this way, we were able to produce the companion paper [11] in a few hours.

Why then releasing this manuscript? First of all, several parts of it have independent interest and can be used to prove other statements. For instance we go here a long way in the direction of classifying certain type of stationary cone, and in fact some of the tools introduced here are used crucially in [11]. Furthermore, it may have some historical interest as a record of human mathematical work compared to LLM-assisted work during this transitional period. But there is also another, perhaps more important, reason. The stronger classification in paper [11] is at odds with what can be observed experimentally. It does seem that the boundary Plateau laws for soap films need to include shapes as in List 1.1(e)&(f). How do we reconcile this with the main theorem of [11] which excludes their Almgren minimality? We propose a way in Section 1.3. If our suggestions are correct, then the much more complicated tools introduced in this manuscript do provide a better picture of the physical behavior of soap films than the shorter and more elegant arguments given in [11].

Definition 1.5 (The thin fish family)**.** A cone $C \subset \mathbb{R}^3$ is called a *thin fish* if its link $C \cap \partial B_1$ is the union of geodesic arcs γ_i of ∂B_1 arranged in the following way.

- γ_1, γ_2, γ_3, and γ_4 are congruent and form a spherical quadrilateral[1] F with angles $\frac{2\pi}{3}$.
- γ_5 is the longest geodesic arc connecting two opposite vertices of F (say V_1 and V_3).

[1]Note that F is the projection on the sphere ∂B_1 of the face of a cube inscribed in ∂B_1.

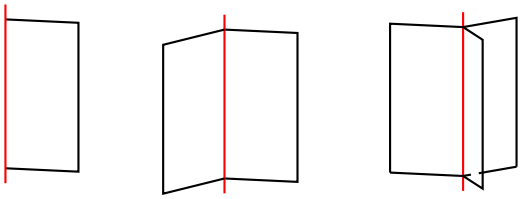


FIGURE 3. The local boundary geometries (a), (b), and (c). The line ℓ is in red.

- γ_6 and γ_7 are two arcs which connect the other vertices, V_2 and V_4, to two antipodal points P_1 and P_2, lying outside F. γ_6 can be trivial (i.e. $P_1 = V_2$ is allowed).
- Any pair of arcs with a common endpoint forms an angle of $\frac{2\pi}{3}$.
- The reference line ℓ is the one passing through P_1 and P_2.

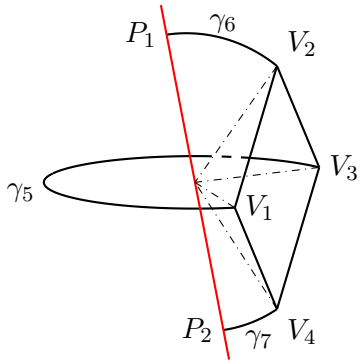


FIGURE 2. A cone C in the thin fish family. In the picture we can see the spherical arcs γ_5, γ_6, and γ_7 of Definition 1.5, while in place of the arcs $\gamma_1, \gamma_2, \gamma_3$, and γ_4 the vertices V_1, V_2, V_3, V_4 are joined by straight edges (which form a square). The length of the arc γ_7 is allowed to vary in a range between 0 (when $V_4 = P_2$) and a maximal length m_0 (when $V_2 = P_1$ and the arc γ_6 vanishes).

Note that the Almgren minimality property is invariant under addition of an $\mathcal{H}^k$-null set; thus, to be precise, the statement of Theorem 1.3 must be appropriately modified. This is done by imposing on E the following "normalization":

$$\mathcal{H}^k(E \cap B_r(x)) > 0 \qquad \forall x \in E \quad \text{and} \quad \forall r > 0\,. \tag{1.2}$$

Closed sets as in Definition 1.2 are actually called $(\mathbf{M}, 0, \infty)$-minimal by Almgren and in much of the subsequent literature (cf. [3] and [25]). In the rest of this work we will simplify our terminology further and just use the term "minimizing". Furthermore it is a theorem by Almgren, cf. [3, Theorem II.3, Conclusion 9], that such sets are in fact necessarily (countably) rectifiable. To streamline our exposition, we henceforth adopt the standing assumption that "minimizing" refers to a rectifiable set as in Definition 1.2 which satisfies the condition (1.2).

Note that the problem of classifying cones which are minimizing with respect to a line ℓ can be stated more generally without assuming $\ell \subset E$. However, since E is a cone, there are only two other possibilities: $\ell \cap E = \{0\}$ and $\ell \cap E$ is a halfline. The first case essentially

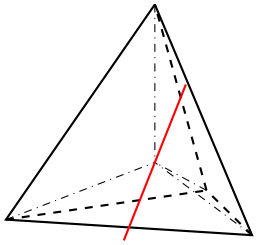


FIGURE 4. The local boundary geometry (d). The line ℓ is going through the midpoints of two opposite sides of the 1-skeleton of the tetrahedron.

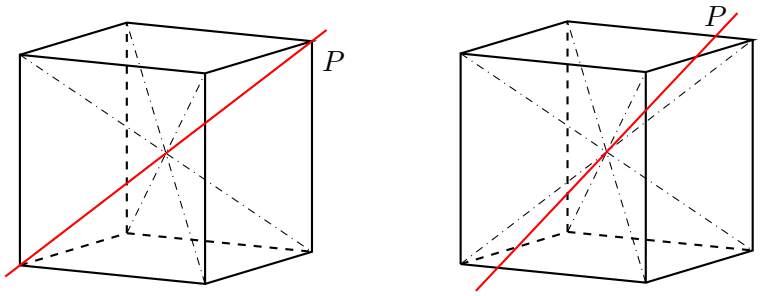


FIGURE 5. The local boundary geometry (e). If P and $-P$ are the intersections of the line ℓ with the skeleton of the cube, then either P is a vertex or it has to be sufficiently close to a vertex.

reduces to the classical theorem of Taylor in [25], cf. the discussion below, while we plan to study the remaining one some time in the future, as the methods introduced in this work can be adapted to address it.

1.3. **Almgren minimality and physical Plateau Laws.** Taylor's theorem [25] immediately implies that the cones in (c) or (d) in List 1.1 are indeed minimizing, whereas the minimality of (a) and (b) is elementary. On the other hand, as noted in Remark 1.4, the cones in (e) and (f) are not minimizing.

As explained in the previous section, our study is motivated by the related problems of understanding the local structure of soap films and area-minimizing sets. We refrain from discussing whether minimizing sets are a good approximation of the shape of a soap film and rather refer the reader to other works, such as [4] and [16]. Concerning the abstract variational problem of minimizing the area subject to appropriate boundary conditions, we just point out that the minimizers of all variational frameworks which we know of (and aim to capture the behavior of soap films) are in fact subsets of the class introduced in Definition 1.2, see for instance [9]. Moreover, real-world experiments of soap films attaching to several printed models of a boundary Γ (thanks to the assistance of the research group of Michael Niedermeier at the University of Passau, cf. Section 1.6) have been of great importance in proving our theorems and have shown a remarkable agreement among experiments, theory, and numerical simulations (performed using Brakke's surface evolver, see [6]). The only exceptions are that, while we can disprove the minimality of cones in (e) and (f) (cf. [11]), we do observe them in real life and numerical experiments. How can we explain such discrepancy?

One obvious possibility is that our experiments are too crude. While our wires were 3d printed in the lab of Michael Niedermeier's research group, these structures are formed by plastic tubes. In fact we used strings of nylon for the line ℓ, upon a suggestion of David that the thickness might have an effect. Ideally our findings should be confirmed by much more accurate experiments.

A second explanation, the one we propose here, is that when the wire has the shape of one of the links in List 1.1(e)&(f), even though the actual minimum is different than the corresponding cone, it is however too close to it to be observed in an experiment (with the exception of some cases in List 1.1(e)).

First of all, it is possible to use the tools of this paper to show the following: for any stationary cone (cf. Definition 4.2) which is not listed in Theorem 1.3 there is a competitor which gains a "substantial amount" in area and coincides with the cone outside the unit ball. Without altering the structure of the manuscript, but scaling appropriately some of the needed numerical computations, this substantial amount can be quantified as being around 10^{-3}. The thickness of a soap film in a classroom experiment ranges between 0.1 and 2 micrometers, namely between 10^{-5} cm and $2 \cdot 10^{-4}$ cm, while the film has a typical size of a few centimeters. This substantial gain in area can also be shown when E is as in List 1.1(e) and the line ℓ contains the midpoint of an edge of Q, cf. Section 19. And in this case we can see experimentally that the actual film will not form a cone, while Brakke's surface evolver does show also a good competitor.

On the other hand, consider the case of the most symmetric element in List 1.1(f). The authors tried themselves and asked an LLM to optimize the ideas in the competitor construction of [11]. The best area gain achieved in the unit ball is approximately $2 \cdot 10^{-15}$. When we consider that the size of a water molecule is roughly $3 \cdot 10^{-8}$ cm, such a gain would be absolutely meaningless even if Almgren minimizers were extremely accurate models of real soap films.

Let E be a cone as in List 1.1(e)&(f) and let $\mathscr{C}(E)$ be the family of closed sets $\Phi(E)$ for any Φ as in Definition 1.2 with $U = B_1$ and $\Gamma = \ell$. Define

$$m(E) := \mathcal{H}^2(E \cap B_1) - \inf\{\mathcal{H}^2(E' \cap B_1) : E' \in \mathscr{C}(E)\}\,. \tag{1.3}$$

We conjecture that $m(E)$ is a very small number, unless E is an element of List 1.1(e) for which ℓ is close to the midpoint of an edge of Q. We therefore propose the following problem

Problem 1.6. Give an upper bound for $m(E)$ for E in List 1.1(e)&(f).

A related question is estimating the distance of the "actual minimizer" with the same boundary data from the cone E. In order to define things properly, consider any minimizing sequence $\{E_k\} \subset \mathscr{C}(E)$, define the measures $\mu_k := \mathcal{H}^2 \llcorner (E_k \cap B_1)$, and let μ be the weak* limit of any subsequence. According to [9], the closed set $E_m := \operatorname{spt}(\mu)$ is an Almgren minimal set. We expect that the Hausdorff distance $d_H(E, E_m)$ between E_m and $E \cap \overline{B}_1$ is very small in the cases under discussions.

Problem 1.7. Give an upper bound for $d_H(E, E_m)$ for E in List 1.1(e)&(f) and any E_m as above.

In particular, let E be the most symmetric element of List 1.1(f). If it were possible to prove that $d_H(E, E_m) \leq 10^{-6}$, then clearly E_m and E would be indistinguishable at the physical scales which are of interest for soap films, given that the ratio between the thickness and the size of the film is in fact larger.

1.4. **Tangent cones.** The relation between Theorem 1.3 and the local structure of minimizing sets is exemplified by the seminal work [25]. In that paper Taylor proved that:

(i) If we replace ℓ with the empty set $\emptyset$ in the statement of Theorem 1.3, then there are precisely three possibilities: E is either a 2-dimensional plane, or the union of three half-planes meeting at 120° angles along a common line, or the cone over the 1-skeleton of a regular tetrahedron;
(ii) if E is minimizing with respect to a boundary Γ and $x \in E \setminus \Gamma$, then there is a neighborhood U of x in which E is diffeomorphic to one of the three cones in (i).

While we do not address the problem of showing further regularity properties (e.g. uniqueness of the tangent cone), our classification theorem, combined with Allard's classical boundary monotonicity formula and the compactness arguments of [1] (or more recent and elementary ones, see e.g. [9]), immediately implies the following corollary (cf. Corollary 4.6 and Corollary 4.5).

Corollary 1.8. *Assume $U \subset \mathbb{R}^3$ is an open set, $\Gamma \subset U$ a smooth arc containing the origin, and $E \subset U$ a minimizing set containing $\Gamma \cap U$. For every sequence $r_k \downarrow 0$ there is a subsequence, not relabeled, such that the rescaled sets*

$$E_{0,r_k} := \{x : r_k x \in E\}$$

converge, locally in the Hausdorff distance, to a cone E which contains the line ℓ tangent to Γ at 0 and is minimizing with respect to it. Therefore E belongs to List 1.1.

1.5. **Previous results.** The problem of finding "boundary" area-minimizing cones is a very natural one and is mentioned in several places in the literature, see e.g. [6, 7, 19, 20]. It is well known that a half-plane with boundary ℓ is in fact a minimizing cone with boundary ℓ. Allard's classical regularity theory in [2] implies then that, if Γ is of class C^2, $x \in \Gamma$, and one tangent cone to E at x is a half-plane (case (a) in List 1.1), then E is diffeomorphic to it in some neighborhood of x. The only other result of this type is due to David in [7]. For a notion of minimizing set which is stronger than the one of Almgren, David gives a complete description of the local behavior of E at x if one tangent cone is as in List 1.1(b) (one cannot always conclude that E is diffeomorphic to the cone in a neighborhood of x since it can be proved that there are some other possible behaviors, as it is explained in detail in [7]). He also proves that if one tangent cone at x is as in List 1.1(c) then the tangent cone is unique (i.e. the cone is the limit of the whole family $r^{-1}(E - x)$ as $r \downarrow 0$); in this case, however, he does not fully characterize the local topology of E. The paper [19] correctly identifies the geometries in (a), (b), (c), and (d) as minimizing, while the case of the cone over the 1-skeleton of the cube with the line ℓ going through two vertices is suggested in [7] as very likely to be minimizing. The remaining cases of (e) are not mentioned anywhere, while the family in (f) appears to be a novel discovery.

1.6. **Acknowledgments.** The first author acknowledges the support of the Simons Foundation through the Simons Initative on the Geometry of Flows (Grant Award ID BD-Targeted-00017375, CDeL). Both authors are very grateful to Michael Niedermeier and Dörte Rüweler, who printed several models of candidate minimizing cones. The experiments involving dipping these models in soapy water have been pivotal in guiding our intuition. The first author is also grateful to Guy David, Jonas Hirsch, and Simone Steinbrüchel for several discussions on the present work.

2. Plan of the paper

The remainder of this paper is devoted to the proof of Theorem 1.3, with the exception of Section 19, where we discuss which cones in List 1.1 are proven to be area-minimizing and motivate our conjectures for the remaining ones. The proof of Theorem 1.3 is split into several steps, many of which rely on computer-assisted verification. In this section, we outline these steps and their logical interdependencies.

Fix a line $\ell \subset \mathbb{R}^3$ and a 2-dimensional cone E containing ℓ, which is minimizing with respect to it. Let $\mathcal{N} = E \cap \partial B_1$ denote the link of the cone. In Section 4 we use the interior regularity theory for Almgren-minimizing sets and a classical blow-up argument to prove the following structural properties of $\mathcal{N}$ (cf. Proposition 4.3):

- $\mathcal{N}$ is the union of finitely many closed geodesic arcs γ_i.
- Distinct edges γ_i can only intersect at their endpoints.
- If p is an endpoint of an arc γ_i, then either $p \in \{P_1, P_2\} = \ell \cap \partial B_1$, or it is a junction where three arcs meet at equal (120°) angles.
- P_1 and P_2 are both endpoints of at least one arc and at most three.
- Any pair of arcs sharing P_1 (or P_2) as an endpoint form an angle in the range $[\frac{2\pi}{3}, \pi]$.

Given their distinct role, we call the points P_1 and P_2 "pins". Sets $\mathcal{N}$ satisfying these conditions are called *stationary nets with two antipodal pins* (cf. Definition 4.1).

Any such $\mathcal{N}$ is naturally associated to an abstract graph embedded in $\mathbb{S}^2$ (cf. Definitions 5.8 and 5.7): the vertices correspond to the endpoints of the geodesic segments, and the edges correspond to the segments themselves. Two nets are said to have the same topology if their underlying graphs are isomorphic and their embeddings in $\mathbb{S}^2$ are homeomorphically equivalent (cf. Definition 5.9).

Roughly speaking, our proof of Theorem 1.3 relies on the following three main tasks:

(I) "Topological classification". Identify meaningful restrictions on the topology of $\mathcal{N}$ and list all topologies satisfying them. These are the *candidate topologies*.

(II) "Geometric realization". For each candidate topology, classify all its possible embeddings as stationary nets with two antipodal pins. These will be called *realizations* of the candidate topologies.

(III) "Area minimization". For each realization $\mathcal{N}$ not included in List 1.1, show that there is a suitable competitor decreasing its area.

2.1. **Task (I).** For any net $\mathcal{N}$, we construct a spherical embedding $\mathcal{N}''$ of a cubic graph (e.g., all vertices have degree three). This purely topological object is called "twice-derived spherical graph", see Definition 5.11 (the attentive reader may observe that in this introduction we are skipping the concept of once-derived net, which is an important intermediate step between the original net and the twice-derived spherical graph, cf. Definition 5.5).

Theoretical considerations in Section 5 imply that the number of valid twice-derived spherical graphs is finite, albeit rather large. In order to classify them, we describe in Section 8 an algorithm which produces a computer-assisted enumeration, cf. Proposition 8.2. Once we have classified the twice-derived spherical graphs, it is a relatively easy job to deduce a finite list of candidate topologies (because, given a twice-derived spherical graph, it is simple to enumerate the possible candidate topologies that could have generated it): this is also explained in Section 8. Some additional geometric "sieves" – simple criteria based on spherical geometry that filter some (geometrically non-realizable) topologies (see Section 7) – are

however used to further restrict the list of candidate topologies. These criteria are based on some important considerations in spherical geometry, which are also used in many arguments in the subsequent parts of the paper, collected in Section 6. The final list of possibilities, to which we arrive at the conclusion of this part of our paper, is given in Proposition 8.5.

2.2. **Task (II).** Finding realizations of a fixed topology can be reduced to a system of algebraic equations where the unknowns are suitable trigonometric functions of the edge lengths. This system is generally overdetermined; thus, most topologies have no realization. However, the system's high degree makes a direct resolution computationally unfeasible and forces us to proceed differently.

The realizations of some exceptional topologies will be studied with pen and paper using considerations in spherical geometries, and some Gröbner basis techniques, cf. Section 14 and Section 15. All the topologies originating from the six geometries in Theorem 1.3 are in fact included in the special cases of Section 14 and Section 15.

In order to study the realizations of the cases not covered in Section 14 and Section 15 but listed in Proposition 8.5 we employ metric, instead of algebraic, arguments. As a first step, thanks to the area-minimizing property of the cone over the net, we show that the edges of the net must all satisfy certain lower bounds on their edge lengths. This simple fact can be proved by compactness, but in our case getting (good) *effective* lower bounds is of fundamental importance, cf. Section 11. The arguments rely on the construction of good competitors. Since, however, the admissibility constraints for competitors in Definition 1.2 are difficult to work with, we will dedicate some efforts to show that they can be relaxed in several ways, cf. Section 10

Thanks to the bounds in Section 11 and to a number of geometric tools (cf. Section 9 and Section 12), we are then able to devise an efficient computer algorithm which, given an assigned precision $\boldsymbol{\delta}$ and a fixed topology:

- Either proves that the topology does not have a geometric realization;
- Or, for each edge of the topology, returns an interval of size at most $\boldsymbol{\delta}$. The algorithm proves that, for any geometric realization of the net, the length of any edge belongs to the corresponding interval.

The details of the algorithm are explained in Section 13 and the theoretical tools employed in it are the lemmata and propositions proved in Section 9 and Section 12. The precision $\boldsymbol{\delta}$ is set (approximately, depending on the net to 10^{-4}. By running our algorithm for a few hours on a standard workstation over the topologies not covered by Section 14 and Section 15, we are then able to show that only a handful of topologies, more precisely 26, are realizable. For these we have a suitable approximation, up to the precision $\boldsymbol{\delta}$, of any possible realization.

2.3. **Task (III).** We are left with 26 geometric realizations, some found in Section 13 and some found in Section 15, which do not belong to List 1.1. These must be excluded on "energetic" grounds, more precisely we shall show that the associated cones, which are in fact stationary, are not minimizing. The strategy is to fix a finite number of "approximations" of such nets $\mathcal{N}$. In the cases found in Section 13 we in fact only know the lengths of the edges up to an error $\boldsymbol{\delta}$. But even for the cases found in Section 15, for which we have an *exact* knowledge, we introduce suitable approximations because they come as infinite (one-parameter and two-parameters) families. We must therefore prove a sufficiently high area saving on the approximation to guarantee that there is a similar competitor for the hypothetical exact net. Section 17 provides the rigorous estimates required to ensure the gain

outweighs the approximation error. Akin to what happens with the lower bound of Section 11, a compactness argument gives a proof of concept of the latter idea, but in reality we need an explicit and sufficiently accurate estimate.

At this stage, we thus have a finite collection $\mathcal{C}$ of cones (which provide a sufficiently good approximation of any cone over any possible existing realization of a candidate topology), and for each element of the collection we shall show the existence of a competitor (in the sense of Definition 1.2) which gains a certain specified amount $\boldsymbol{\eta}$ in area. The competitors are stored in `data/competitors.zip` in [13] (and can be visualized at [12]) and each consists of a collection of triangles. More rigorously, each stored object is a suitable 3-dimensional piecewise linear immersion of a finite 2-dimensional simplicial complex, cf. Definition 16.1 and Definition 16.4. These have been generated with the help of Brakke's surface evolver (with some appropriate modifications). The evolver was intialized with a suitable topological deformation of the corresponding cone $E \in \mathcal{C}$, which has been guessed using real world experiments. We then let it run with some specific choices of parameters (in some cases we needed to try more than one choice) for a suitable amount of time. It is however not necessary to know how Brakke's surface evolver produces the competitor; in particular our proof does not rely on Brakke's code being error-free. Rather, another appropriate code, written by us and described in Section 18, checks the following two relevant conditions for each competitor E' in the database:

(i) E' gains $\boldsymbol{\eta}$ in area over the respective element $E \in \mathcal{C}$.
(ii) E' is valid in the sense of Definition 1.2.

Point (i) has a straightforward implementation. Point (ii) is instead the most difficult aspect of this part of the proof. Recall that according to Definition 1.2, a competitor for E must be of the form $\Phi(E)$ for a Lipschitz map Φ satisfying some specific requirements. In Section 16, we derive a list of combinatorial properties for E' which guarantees the existence of a map Φ as in Definition 1.2. These properties can be described in finite terms, which enables us to check them through a computer code. In fact, strictly speaking, it is not E' which is *directly* a valid competitor for E, but rather a suitable modification of E', cf. Section 16. Our arguments show, however, that this suitable modification gains at least as much area as E'.

3. Index of Concepts and Notation

3.1. **Index of Concepts.**

Almgren minimizing set	Definition 1.2
Approx. metrized marked spherical graph (AMMSG)	Section 13
Boundary distance / Boundary path	Definition 7.1
Children (and Nice children)	Section 8.4
Collapse (of a simplicial complex)	Definition 16.6
Combinatorial map	Definition 8.8
Extinct (net or graph)	Definition 8.1
Fat fish	Definition 15.1
Goggles	Definition 15.2
Graph Embedding	Definition 5.8

Heppes net	Section 5.1
Marked spherical graph	Definition 8.3
Metrized marked spherical graph (MMSG)	Section 13
Once-derived net	Definition 5.5
Path-isometry	Definition 12.3
Pins	Definition 4.2
Planar map	Section 8.3
Polygon / Polygonal path (Spherical)	Definitions 6.1 and 6.2
Positively curved oriented arc	Definition B.2
Simplicial complex / Piecewise linear map	Definitions 16.1 and 16.4
Stationary geodesic net (with pins)	Definition 4.2
Thin fish	Definition 1.5
Twice-derived spherical graph	Definition 5.11
Wedge	Definition 10.4

3.2. **Index of Notation.**

$C_\theta(\ell)$	Section 10
$\mathrm{Darts}(G)$	Definition 8.6
$E(v,t)$	Section 12
$\exp_t(v)$	Section 12
$F_{\mathcal{N}}(d,d')$	Definition 13.5
$F_{(G,\mathrm{len})}$ / $F_{(G,\mathcal{I})}$	Definitions 13.6 and 13.7
$\Gamma(X)$	Section 5.3
$\iota(\sigma,F)$	Section 13
$\mathrm{next}(d)$	Definition 8.6
$R(v,\alpha)$	Section 12
$\mathrm{rev}(d)$	Section 5.3
$\mathrm{rot}_\alpha(v)$	Section 12
$T\mathbb{S}^2$ / $T^1\mathbb{S}^2$	Definition 12.1

4. Stationary nets with two pins

In this first step we use well-known tools in geometric measure theory to infer some basic structural properties for the link of a cone which contains a line ℓ and is minimizing with respect to it. The main statement is Proposition 4.3 below. We start with some relevant definitions.

Definition 4.1. A geodesic net on the 2-dimensional sphere $\mathbb{S}^2 \subset \mathbb{R}^3$ consists of a finite set (possibly empty) of distinct points $V_1, \dots, V_n \in \mathbb{S}^2$ (called *vertices*) and a finite set of distinct *edges* $\sigma_1, \dots, \sigma_m$ such that

(a) Each σ_i is either an equator or a closed geodesic arc whose endpoints are two distinct vertices V_i and V_j;
(b) Distinct edges σ_i can only intersect at their endpoints;
(c) Each V_i is the endpoint of at least one edge.

The support of the net is the subset of $\mathbb{S}^2$ which is the union of the edges. We say that the net is connected if the support is connected.

This definition allows for the trivial empty net and for the net consisting of an empty set of vertices and a single equator as its only edge. Moreover, two edges can form an angle of π at a common endpoint, even if they are the only two edges with that endpoint. This means that the net is not entirely determined by its "support". For instance, if the support of the net consists of a single equator, the net could have, as the example just mentioned, an empty set of vertices and a single edge, but also N vertices and N edges with any $N \geq 1$.

Definition 4.2. A geodesic net is said to have k pins if there is a subset of the vertices (called *pins*) of cardinality k such that:

(a) Any vertex which is not a pin is the endpoint of exactly three edges;
(b) A pin is the endpoint of at most three edges.

In addition we say that the net is stationary if

(c) Any pair of edges with a common endpoint V form an angle which is at least $\frac{2\pi}{3}$.

Note the following consequence of the above definition: at any non-pinned vertex the three edges of a stationary net must form three angles of $\frac{2\pi}{3}$, and of course this occurs also at a pin which is the endpoint of three edges. Moreover, given condition (c) in Definition 4.1, we know that a pin is the endpoint of at least one and at most three edges. Finally, note that since the only vertices which might have degree two are the pins, only at these special points there can be two incoming edges forming a "flat" angle π. The following proposition explains why we consider this possibility: the pins must be "marked" even in the latter case, because they serve a particular purpose.

Proposition 4.3. *Assume $E \subset \mathbb{R}^3$ is a cone containing a line ℓ which is minimizing with respect to it. The link of E, namely $E \cap \partial B_1$, is a stationary net with two antipodal pins, which are the two points forming the set $\ell \cap \partial B_1$.*

Note that we are not claiming that the net $E \cap \partial B_1$ is connected. We will however prove this fact in a later stage. The rest of this section is devoted to the proof of Proposition 4.3 and we will show, in passing, how Corollary 1.8 can be derived from Theorem 1.3. The first important tool will be the following general compactness result.

Proposition 4.4. *Assume $U \subset \mathbb{R}^N$ is a bounded open set and $\{\Gamma_j\}$ a sequence of C^2 $(k-1)$-dimensional submanifolds of U which converges in the C^2 topology to a $(k-1)$-dimensional submanifold Γ. Assume $\{E_j\}$ is a sequence of relatively closed subsets of U such that*

- *E_j is minimizing with respect to Γ_j in U;*
- *$\sup \mathcal{H}^k(E_j) < \infty$;*
- *E_j converges locally in the Hausdorff distance to a closed set E.*

Then,

- *E is minimizing with respect to Γ;*
- *$\lim_j \mathcal{H}^k(V \cap E_j) = \mathcal{H}^k(V \cap E)$ for every open or closed $V \subset U$ with $\mathcal{H}^k(E \cap \partial V) = 0$.*

The proposition was first proved by Almgren in [3] (even though it does not appear there in this precise form), but for the reader's convenience we will give a proof in Section A using Allard's compactness theorem for stationary integral varifolds (other methods can also be used, see e.g. [9], but Allard's theory seems to be the quickest route). In fact, two further ingredients will be inferred from Allard's theory: the monotonicity formula (both at the boundary and in the interior, cf. Corollary 4.5 below) and some density lower bounds (cf. Lemma A.2). These tools can be applied in our case because minimizing sets (in the sense of Almgren) induce integral stationary varifolds, thanks to the simple Lemma A.1.

In what follows we denote by A_Γ the second fundamental form of Γ.

Corollary 4.5. *Let $U \subset \mathbb{R}^N$ be an open set, $\Gamma \subset U$ a C^2 $(k-1)$-dimensional submanifold of U, and E a minimizing set in U with respect to Γ. Fix a point $p \in U$. If $p \notin \Gamma$, then*

$$\frac{d}{dr}\frac{\mathcal{H}^k(E \cap B_r(p))}{r^k} \geq 0 \qquad \text{for } r \in (0, \operatorname{dist}(p, \Gamma \cup \partial U)). \tag{4.1}$$

If $p \in \Gamma$ and $\|A_\Gamma\|_\infty = \sup_\Gamma |A_\Gamma|$, we have

$$\frac{d}{dr}\left(\exp(C\|A_\Gamma\|_\infty r)\frac{\mathcal{H}^k(E \cap B_r(p))}{r^k}\right) \geq 0 \qquad \text{for } r \in (0, \operatorname{dist}(p, \partial U)), \tag{4.2}$$

where C is a dimensional constant.

Moreover, if $\frac{\mathcal{H}^k(E \cap B_{r_0}(p))}{r_0^k} = \frac{\mathcal{H}^k(E \cap B_{r_1}(p))}{r_1^k}$ and

(i) either $p \notin \Gamma$ and $0 < r_0 < r_1 < d(p, \Gamma \cup \partial U)$,
(ii) or Γ is an affine space including p and $0 < r_0 < r_1 < d(p, \partial U)$,

then $E \cap B_{r_1}(p)$ is a cone with vertex p and base $E \cap \partial B_{r_1}(p)$.

The first part of the Corollary can be found in [1] and [2]. For the second part, the conclusion under assumption (i) can be found in [1]; under assumption (ii) it is folklore but not explicitly stated in [2]: in the appendix we show how to derive it quickly from the tools developed in that paper.

We next register a consequence of the results stated thus far, which is just a generalization to all dimensions and codimensions of the first statement of Corollary 1.8.

Corollary 4.6. *Let $U \subset \mathbb{R}^N$ be open and $\Gamma \subset U$ a $(k-1)$-dimensional C^2 submanifold. Assume that E is minimizing with respect to Γ in U and fix $p \in \Gamma$. Then, for every sequence $r_j \downarrow 0$, there is a subsequence, not relabeled, such that the following holds:*

(a) the sets $E_{p,r_j} = r_j^{-1}(E - p)$ converge, locally in the Hausdorff distance, to a closed set E' of $\mathbb{R}^N$ which is a cone with vertex at the origin;
(b) E' is minimizing with respect to $T_p\Gamma$;
(c) If $V \cap \Gamma \subset E$ for some neighborhood V of p, then $T_p\Gamma \subset E'$.

Proof. By compactness (and a diagonal argument), we can extract a subsequence of radii (not relabeled) such that $E_{p,r_j} \to E'$ locally in the Hausdorff distance, where E' is a closed set. We want to apply Proposition 4.4 to show that E' is minimizing with respect to $T_p\Gamma$. To do so, we just need to prove an upper bound on the area of E_{p,r_j}. Observe that, for every fixed

$R > 0$,

$$\mathcal{H}^k(E_{p,r_j} \cap B_R) = R^k \frac{\mathcal{H}^k(E \cap B_{Rr_j}(p))}{(Rr_j)^k}$$

is uniformly bounded by Corollary 4.5.

To prove that E' is a cone, we prove that the mass-ratio in balls centered at the origin is (almost) constant in the radius and then we apply Corollary 4.5. Let

$$\Theta(p) := \lim_{r \downarrow 0} \frac{\mathcal{H}^k(E \cap B_r(p))}{\omega_k r^k},$$

which exists by Corollary 4.5. Since E' is rectifiable, except for at most a countable number of exceptional radii r, we have $\mathcal{H}^k(E' \cap \partial B_r) = 0$. Thus, Proposition 4.4 implies that, for every such r,

$$\mathcal{H}^k(E' \cap B_r) = \omega_k r^k \lim_{j \to \infty} \frac{\mathcal{H}^k(E \cap B_{rr_j}(p))}{\omega_k (rr_j)^k} = \omega_k r^k \Theta(p).$$

Therefore, by the second part of Corollary 4.5, we infer that E' is a cone with vertex at the origin.

Finally, observe that $\Gamma_{p,r_j} := r_j^{-1}(\Gamma - p) \subseteq E_{p,r_j}$. Moreover, Γ_{p,r_j} converge, locally in the Hausdorff distance, to $T_p\Gamma$. Since E_{p,r_j} converge, locally in the Hausdorff distance to E', we conclude that $T_p\Gamma \subseteq E'$. □

Before coming to the proof of Proposition 4.3, we state the following simple lemma.

Lemma 4.7. *Let $F \subset \mathbb{R}^2$ be a one-dimensional rectifiable set with locally finite measure such that $E = F \times \mathbb{R} \subset \mathbb{R}^3$ is minimizing with respect to $\ell = \{(0,0)\} \times \mathbb{R}$. Assume $R > 0$ and $\Psi : \mathbb{R}^2 \to \mathbb{R}^2$ is a Lipschitz map for which there is a closed convex set $K \subset B_R(0) \setminus \{0\}$ contanining $\{x : \Psi(x) \neq x\}$ and its image through Ψ. Then $\mathcal{H}^1(\Psi(F)) \geq \mathcal{H}^1(F)$.*

Proof. We argue by contradiction and assume that $\Psi : \mathbb{R}^2 \to \mathbb{R}^2$ is a Lipschitz map satisfying the properties above and violating the inequality. We next fix a large number L (to be chosen later) and define

$$\Phi(x_1, x_2, x_3) = \begin{cases} (\Psi(x_1, x_2), x_3) & \text{if } |x_3| \leq L \\ ((|x_3| - L)(x_1, x_2) + (L + 1 - |x_3|)\Psi(x_1, x_2), x_3) & \text{if } L \leq |x_3| \leq L + 1 \\ (x_1, x_2, x_3) & \text{otherwise.} \end{cases}$$

Φ is clearly Lipschitz and is the identity outside of the open set $U := B_R(0) \times (-L-1, L+1)$ and in a neighborhood of ℓ. Moreover $W := \overline{\{\Phi(x) : x \neq \Phi(x)\}}$ is by construction contained in $K \times [-L-1, L+1]$ and hence its closure is at a positive distance from ℓ. By minimimality of E, we must have

$$\Delta := \mathcal{H}^2(\Phi(E) \cap U) - \mathcal{H}^2(E \cap U) \geq 0. \tag{4.3}$$

Observe that $|\partial_{x_1}\Phi| \leq 1 + |\partial_{x_1}\Psi|$ and $|\partial_{x_2}\Phi| \leq 1 + |\partial_{x_2}\Psi|$. Moreover, on the open set U we can estimate $|\partial_{x_3}\Psi| \leq 1 + R + |\Psi| \leq 1 + 2R$. In particular the Lipschitz costant of Φ can be bounded by a number M which is independent of L. Therefore we can compute

$$\Delta \leq 2L(\mathcal{H}^1(\Psi(F) \cap B_R(0)) - \mathcal{H}^1(F \cap B_R(0)) + 2(M^2 - 1)\mathcal{H}^1(F \cap B_R(0)). \tag{4.4}$$

Clearly, given that M is fixed while we can choose L very large, we would then conclude $\Delta < 0$, which is at odds with (4.3). □

Proof of Proposition 4.3. First of all, observe that, by Taylor's regularity theorem, at any point $p \in E \setminus \ell$, the set E is locally either a smooth submanifold, or it is diffeomorphic to a triple junction, or it is diffeomorphic to the cone over the skeleton of a tetrahedron. However, because E itself is a cone, we also see immediately that any tangent cone to E at p must be invariant under translation along the direction $v = \frac{p}{|p|}$. Since the tetrahedron is not invariant under any translation, we conclude that E at p is either diffeomorphic to a triple junction or it is a smooth submanifold. If it is a smooth submanifold, then $B_r(p) \cap E \cap \partial B_1$ must be a geodesic arc for any r sufficiently small; if it is a triple junction, then $B_r(p) \cap E \cap \partial B_1$ must be the union of three geodesic arcs meeting at p and forming equal angles there (again for r sufficiently small).

Fix now a system of coordinates such that $\ell \cap \partial B_1 = \{e, -e\}$ for some unit vector $e \in \mathbb{R}^3$. We show that at the point e (and $-e$, but by symmetry we will focus on e only) one of the following three things occur for r sufficiently small:

(a) Either $\partial B_1 \cap E \cap B_r(e)$ is a single geodesic segment connecting the point e to some point on $\partial B_r(e) \cap \partial B_1$;
(b) or $\partial B_1 \cap E \cap B_r(e)$ is the union of two geodesic segments forming at e an angle $\alpha \in [\frac{2\pi}{3}, \pi]$;
(c) or $\partial B_1 \cap E \cap B_r(e)$ is the union of three geodesic segments forming at e three angles of $\frac{2\pi}{3}$.

First of all, we use Corollary 4.6 to infer that, up to subsequences, $E_{e,r_j} := r_j^{-1}(E - e)$ converge locally in the sense of Hausdorff, to some set E' containing ℓ which is minimizing with respect to ℓ and is a cone with vertex at the origin. Note however that, since E itself is a cone, E' must be invariant along the direction e, which is tangent to ℓ. So, $E' = F \times \ell$, where $F \subset \ell^\perp$ is a 1-dimensional stationary cone in the punctured plane $(\ell^\perp \setminus \{0\})$. Note that, since F is a cone, for every $x \in \partial B_1 \cap F$ it contains the entire halfline originating in 0 and passing through x. On the other hand, since F has locally finite Hausdorff 1-dimensional measure there can only be a finite number N of these lines.

We can now apply Lemma 4.7 to see that no pair of distinct half-lines among those forming F can make an angle which is (strictly) smaller than $\frac{2\pi}{3}$, otherwise we can use the competitor illustrated in Figure 6 to contradict its minimality (this is a very well known elementary argument, which we report for the reader's convenience).

The half-lines forming F are, therefore, $0 \leq N \leq 3$. Note furthermore that the number N is determined by the density

$$\Theta(e) = \lim_{r \downarrow 0} \frac{\mathcal{H}^2(E \cap B_r(e))}{\pi r^2},$$

indeed we have $\Theta(e) = \frac{N}{2}$. So, either all tangent cones are single half-lines, or they are all formed by two halflines, or they are all formed by three half-lines. Note that N must be positive because of the condition (3) of Definition 1.2.

We can now apply Allard's interior regularity theorem to infer that, if r is sufficiently small, then

$$E \cap B_r(e) \setminus B_{r/8}(\ell)$$

consists of either one, two or three minimal surfaces, with boundaries contained in $\partial(B_r(e) \setminus B_{r/8}(\ell))$, close to either a single half-plane, or to the union of two, or to the union of three (and we can choose all these half-planes so that their boundaries all coincide with the line ℓ). Since E is a cone, this implies that, for all sufficiently small radii r, $E \cap \partial B_1 \cap (B_r(e) \setminus B_{r/8}(e))$

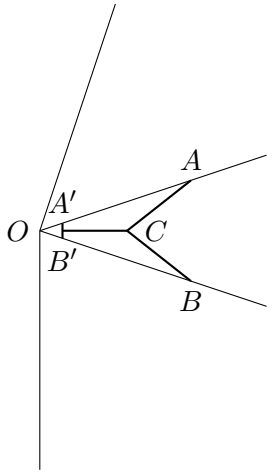


FIGURE 6. Following a well-known argument, we illustrate the competitor disproving the minimality of F if two of its halflines form an angle strictly smaller than $\frac{2\pi}{3}$. W.l.o.g. we can pick the halflines to be a pair of closest ones. We then choose two points A and B on them and connect with the origin O forming a triple junction at some additional point C. Because we cannot change the set in a neighborhood of the origin, we "reconnect" with the segments OA and OB in a neighborhood of O. It is easy to construct a Lipschitz map which deforms the union of the two segments AA' and BB' on the triple junction, while keeping the rest of the set F unchanged. The map can be easily arranged to satisfy the assumptions of Lemma 4.7 and when A' and B' are sufficiently close to the origin we gain in length.

consists of $N = 1, 2$, or 3 geodesic segments, each segment joining one endpoint in $\partial B_r(e)$ with an endpoint in $\partial B_{r/4}(e)$. We can now apply the latter conclusion to r and $\frac{r}{2}$: using that the sets $B_r(e) \setminus B_{r/4}(e)$ and $B_{r/2} \setminus B_{r/8}(e)$ overlap, we infer that the same description holds for the set $E \cap (B_r(e) \setminus B_{r/8}(e)) \cap \partial B_1$. Proceeding inductively, we in fact conclude that the same is true in $(B_r(e) \setminus B_{2^{-N}r}(e)) \cap \partial B_1$. Letting $N \uparrow \infty$, we conclude that $E \cap \partial B_1 \cap (B_r(e) \setminus \{e\})$ consists of either one, or two, or three geodesic segments σ_i, with e being a common endpoint to all of them. On the other hand this also implies that $E \cap B_r(e)$ coincides in fact with $E' \cap B_r(e)$ for its tangent cone E', which is then unique. Therefore any pair of distinct geodesic segments σ_i and σ_j meet at e forming an angle which is at least $\frac{2\pi}{3}$. This completes the proof. □

5. Derived nets and topological graphs

5.1. **Preliminaries on stationary nets with two antipodal pins.** In this section we prove some useful facts on stationary nets with two antipodal pins. First, we recall that stationary nets with 0 pins were classified by Heppes in [14], after a preliminary attempt by Lamarle in [17], who missed one of the nets in the list (and did not provide a proof that the list was complete). We will therefore call them "Heppes nets". Let now $\mathcal{N}$ be the union of the geodesic arcs forming a Heppes net. If this subset of $\mathbb{S}^2$ contains two antipodal points P_1 and P_2, we then obtain easily a stationary net with two antipodal pins. There are in principle three possibilities:

(i) Both P_i's are vertices of the original net; then the net is already also a stationary net with two antipodal pins;
(ii) One point, say P_1, is a vertex, but the other one, i.e. P_2, is internal to some segment σ of the net $\mathcal{N}$. In this case we add P_2 to the collection of vertices and "split" the segment σ into two segments meeting at P_2; naturally, the two new segments "meet" at their common endpoint forming an angle of π;
(iii) Both P_i's are internal points of two segments σ_j and σ_k of $\mathcal{N}$; we then add the pins P_1 and P_2 to the collection of vertices and subdivide each of the segments σ_j and σ_k as described in point (ii) above.

A simple inspection of the Heppes nets shows that (ii) never occurs. On the other hand, both (i) and (iii) occur in a number of cases. The above considerations motivate the following terminology: a net which is obtained from a Heppes net with the above procedure will be called "Heppes with two antipodal pins".

The main goal of this subsection is thus to prove the following proposition. In what follows we will call face of a stationary net $\mathcal{N}$ any connected component of $\mathbb{S}^2 \setminus \mathcal{N}$

Proposition 5.1. *Let $\mathcal{N}$ be a stationary net with two antipodal pins P_1 and P_2. Then:*

(i) At least one of the two pins has degree strictly less than 3 unless $\mathcal{N}$ is Heppes.
(ii) $\mathcal{N}$ is connected; in particular any connected component of $\mathbb{S}^2 \setminus \mathcal{N}$ is an open disk.
(iii) If $\mathcal{N}$ has only two vertices, then it consists of one, two, or three half equators connecting the two pins.
(iv) If $\mathcal{N}$ has at least three vertices, then the two pins cannot belong to the boundary of the same connected component of $\mathbb{S}^2 \setminus \mathcal{N}$ (in particular an edge σ which is incident to P_1 cannot have the other endpoint in common with an edge τ which is incident to P_2).

In the proof we will need the following lemmas and convexity will play a crucial role in both.

Remark 5.2. We warn the reader that a good definition of *general* convex subsets of the sphere presents some challenges, as witnessed by the following example. Let p and $-p$ be two antipodal points and consider two distinct equators passing through them. Their complement consists of four open sectors of the sphere, which are convex according to any possible reasonable definition. It seems natural to call their closures convex. However, if we take the closure of two opposite open sectors, we then find two closed convex sets whose intersection consists of a pair of antipodal points. For this reason we restrict the notion of convexity to open sets of the sphere. In the latter case several equivalent definitions can be given, for instance we can require that open convex subsets of $\mathbb{S}^2$ are intersections of open hemispheres, cf. Section B.

Lemma 5.3. *Let $\mathcal{N}$ be a stationary net with at most two pins. Assume that $\mathcal{N}$ is contained in a (closed) hemisphere. Then, $\mathcal{N}$ consists of two vertices connected by either one or two edges (observe that in the latter case the two vertices must be antipodal).*

Proof. Let Ω be the intersection of all closed hemispheres containing $\mathcal{N}$. Observe that the cone Λ with link Ω and vertex 0 is then the Euclidean convex hull of the cone E over $\mathcal{N}$. Let $p \in \mathbb{S}^2$ be an extremal point of Ω, namely a point of Λ with the property that every line which is going through p but not through the origin intersects Λ only at p.Then $p \in \mathcal{N}$. Since p cannot be in the interior of an edge, it must be a vertex of the net. Furthermore p cannot be a non-pin vertex and if it is a pin, then it can have at most 2 edges coming to

it. Since Ω must have at least two extremal points, we conclude that in fact it has exactly two extremal points, which are the two pins P_1 and P_2 of $\mathcal{N}$. Consequently, Ω is either the geodesic connecting P_1 and P_2 or a sector of $\mathbb{S}^2$, delimited by two half-equators which have P_1 and P_2 as endpoints. If it is the geodesic, we have proven the statement. If it is a sector, then, since no non-pin vertices can belong to the boundary of Ω, the two half-equators are edges of $\mathcal{N}$. To prove that there are no other vertices in $\mathcal{N}$, observe that the other connected components of $\mathcal{N}$ are Heppes nets and cannot belong to a hemisphere. □

Lemma 5.4. *Let $\mathcal{N}$ be a connected stationary net with some pins and let F be one of its faces.*

(1) If the boundary of F does not contain any pins, then F is convex.
(2) If the boundary of F contains one pin, then F is the union of the closures of two open convex sets.

Proof. Let us start with (i). Since $\mathcal{N}$ is connected, F is topologically an open disk with connected boundary. Observe that, a priori, it is not obvious that the boundary of F is topologically a loop. The *combinatorial* boundary of F is a sequence of directed edges of $\mathcal{N}$ connecting a its vertices, with internal convex angles (more precisely, all angles are $2\pi/3$). Hence, we can apply Lemma B.1 and conclude.

We show how a similar argument works also for (ii). Let P be the pin. If P has degree one, we continue, past P, along the equator containing the only edge ending at F, "traveling in the interior" of F until we hit ∂F for the first time in Q. If F has degree 2, we travel along the bisector of the internal angle of P until we hit ∂F for the first time in Q. In both cases, let γ be the line connecting P and Q. Observe that γ is contained in F and it splits F in two connected components which satisfy the assumptions necessary to apply the argument we used in (i). The conclusion follows. □

Proof of Proposition 5.1. (i) is straightforward: if both pins have degree 3, then $\mathcal{N}$ is a Heppes net.

To show (ii) we provide an argument by contradiction. If $\mathcal{N}$ had a connected component without pins, the connected component would be a Heppes net and the other connected components would be contained in the faces of the Heppes net, which are convex. By Lemma 5.3, we conclude that this cannot be true.

Therefore, $\mathcal{N}$ can consist of at most two connected components, each with at least one of the pins. For the sake of contradiction, assume that $\mathcal{N} = \mathcal{N}_1 \sqcup \mathcal{N}_2$, where both $\mathcal{N}_i$ are connected and both contain one pin. Say that $\mathcal{N}_i$ contains the pin P_i. In particular, N_2 belongs to a face of F of N_1. If $P_1 \notin \partial F$, then by Lemma 5.4 we know that F is convex and by Lemma 5.3 we have the desired contradiction. Otherwise, $P_1 \in \partial F$ and $P_2 \in \mathring{F}$, which leads to a contradiction again thanks to Lemma 5.4.

In order to prove (iii) and (iv) we proceed as follows. Assume that the two pins P_1 and P_2 belong to the boundary of the same connected component of $\mathbb{S}^2 \setminus \mathcal{N}$. Thus, there is a path $\gamma \subset \mathcal{N}$ which has nonnegative geodesic curvature and connects P_1 and P_2. Lemma B.3 implies that this path is half of an equator.

Now, rotate this half equator in the counterclockwise direction keeping its two endpoints fixed, until we meet another point of $\mathcal{N}$. Let γ' be the rotated meridian. $\gamma' \cap \mathcal{N}$ cannot contain any vertex except for P_1 and P_2 (because any other vertex will have degree 3). If $\gamma = \gamma'$, then we have conclude that the net consists of just one arc connecting P_1 and P_2.

Otherwise, we repeat the procedure on γ' finding γ''. If $\gamma'' = \gamma$, then the net is made of just two arcs connecting the two pins. Otherwise, we have found three arcs connecting the two pins and thus the net is a triple junction. □

5.2. **Once-derived net.** Given a stationary net with two pins, any pair of edges which meet at a non-pinned vertex form an angle of $\frac{2\pi}{3}$: only those pairs meeting at a pin might form a different angle, in case the degree of that specific pin is different from three. If we remove the pins with degree smaller than three and the edges which contain them as vertices, we get, instead, a geodesic net with the property that all pairs of edges meeting at some vertex form an angle of $\frac{2\pi}{3}$: this "reduced net" is an important stepping stone for a better understanding of the original one. Note however that the resulting net might have 0, 1, or 2 nonconvex faces: in order to have always exactly two nonconvex faces, we "artificially" remove also one edge from each pin of degree 3.

Definition 5.5 (Once-derived net)**.** Take a stationary net with two antipodal pins. We obtain a *once-derived net* by performing the following operations:

- from the set of vertices, we remove the pins with degree 1 or 2;
- from the set of edges we remove the edges with at least one endpoint which is a pin removed in the previous step;
- from the set of edges we remove one edge connecting two pins of degree 3, if such an edge exists, otherwise we remove an edge from every pin of degree 3.

Observe that the once-derived net $\mathcal{N}'$ could very well be the empty set: in the latter case we call it "extinct". We will later include more cases in the term *extinct*, cf. Remark 5.12 and Definition 8.1.

Proposition 5.6 (Properties of once-derived nets)**.** *Consider a non-extinct once-derived net $\mathcal{N}'$ derived from a stationary net $\mathcal{N}$ with two antipodal pins. Then $\mathcal{N}$ is connected and has the following properties:*

(i) All vertices have degree equal to 2 or 3.
(ii) All its faces are polygons homeomorphic to open disks.
(iii) Exactly two faces are nonconvex.
(iv) Every convex angle of any face is $\frac{2\pi}{3}$.
(v) Every nonconvex face has either one or two concave angles equal to $\frac{4\pi}{3}$.
(vi) The area of a face with q concave angles is[2]

$$\frac{\pi}{3}\Big(6 - \#sides + 2q\Big). \tag{5.1}$$

So, if $q = 0$ then $\#sides \le 5$, if $q = 1$ then $\#sides \le 7$, and if $q = 2$ then $\#sides \le 9$.

The convex faces of $\mathcal{N}'$ are one of the following types:

(vii) A face consisting of two half equators connecting two antipodal points;
(viii) An equilateral triangle;
(ix) A parallelogram, determined by the length of one side;
(x) A pentagon, determined by the lengths of two sides.

[2]Note that, when counting the number of sides in a face F, every side which is not in common with another face is in fact counted twice for F.

Proof. (i) follows immediately from Proposition 5.1(iv). Thanks to Proposition 5.1(ii) we know that $\mathcal{N}$ is connected and thanks to Proposition 5.1(iv) we can deduce that also $\mathcal{N}'$ is connected. Then, (ii) follows.

In order to show (iii), we prove that each pin is associated to exactly one nonconvex face in the once-derived net and that any nonconvex face is associated to exactly one pin. Observe that if a face of $\mathcal{N}$ does not have any pin on its boundary, then it is convex and remains the same in $\mathcal{N}'$. Furthermore, thanks to Proposition 5.1(iv) we know that any face of $\mathcal{N}$ can contain at most one pin in its boundary. This allows us to treat the two pins independently. So, fix a pin P. We shall prove that there is exactly one face of $\mathcal{N}'$ which contains P inside or on its boundary and is nonconvex. This is sufficient to conclude. Let us consider the various cases depending on the degree of P in $\mathcal{N}$:

(a) If P has degree 1, then the face of $\mathcal{N}$ which has P on its boundary is nonconvex and becomes a nonconvex face of $\mathcal{N}'$.
(b) If P has degree 2, then the two faces of $\mathcal{N}$ which have P on their boundary are joined to become a nonconvex face of $\mathcal{N}'$.
(c) If P has degree 3, then two of the faces of $\mathcal{N}$ that have P on their boundary are joined to become a nonconvex face of $\mathcal{N}'$.

Note that the above argument implies also (v). The removal of a pin of degree 1 creates a nonconvex face with one nonconvex angle. The removal of a pin of degree 2 eliminates a convex face and creates a nonconvex face with two nonconvex angles of $\frac{4\pi}{3}$ merging a nonconvex face with the removed convex face. The removal of an edge from a degree 3 pin creates a nonconvex face with two concave angles of $\frac{4\pi}{3}$.

(iv) is clear. (vi) follows from Gauss-Bonnet Formula. The points from (vii) to (x) are already discussed in Heppes' paper [14]. □

5.3. **Graphs and embeddings.** In order to handle the topological properties of once-derived net we will borrow basic concepts from the theory of embeddings of graphs, and some classical facts in topological graph theory will be used to devise our algorithm; as a reference we will use [18]. First of all we recall the classical definition of graphs, cf. [18, Definition 1.3.1].

Definition 5.7 (Graph)**.** Let V be a finite set of vertices. An edge is an object that has two (unordered) endpoints in V (which may coincide). Two edges with the same endpoints are not assumed to be equal. We say that an edge is *adjacent* to its two endpoints. A graph is a pair (V, E), where V is a finite set of *vertices* and E is a finite collection of edges.

We are now ready to introduce the concept of embedding of a graph in a topological space. In order to do so, given a topological space X we let $\Gamma(X)$ be the set of continuous embedded arcs. More precisely:

- we first consider the set Y of continuous maps of $[0,1]$ into X which restrict to an embedding of $(0,1)$,
- we introduce the equivalence relation for pairs of elements of Y in which one element can be obtained from the other through composition with an increasing continuous map of $[0,1]$ onto itself,
- and we define $\Gamma(X)$ as the quotient of Y modulo that relation.

Moreover, given an arc in $\Gamma(X)$ and a continuous map $\gamma : [0,1] \to X$ representing its equivalence class, we call "interior of the arc" the set $\gamma((0,1))$. The following definition of graph embedding is equivalent to [18, Definition 1.3.7].

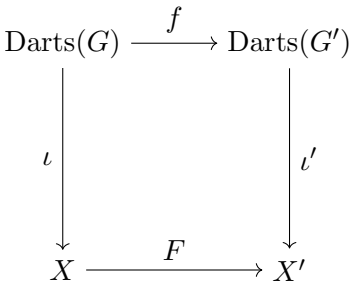


FIGURE 7. The commutative diagram for isomorphic embeddings.

Definition 5.8 (Graph Embedding)**.** For a graph $G = (V, E)$, an embedding of G in a topological space X is a map $\iota : V \to X$ and $\iota : E \to \Gamma(X)$ such that:

- Distinct edges are mapped into arcs with disjoint interiors;
- Distinct vertices are mapped into distinct points;
- If $A, B \in V$ are the two vertices adjacent to the edge e, then either $\iota(A) = \iota(e)(0)$ and $\iota(B) = \iota(e)(1)$, or viceversa;
- If A is any vertex, e any edge, and t any real in $(0, 1)$, then $\iota(e)(t) \neq \iota(A)$.

It is in fact convenient to describe a graph embedding after "directing the edges". More precisely, given a graph $G = (V, E)$ we can associate to it a set $\mathrm{Darts}(G)$ which contains two copies of each edge in E, and admits a canonical involution $\mathrm{rev} : \mathrm{Darts}(G) \to \mathrm{Darts}(G)$ mapping one copy of the edge in the other. We understand the darts as "orientations" for the edge: in particular if the two vertices A and B adjacent to an edge are distinct we have a "dart going from A to B" and a "dart going from B to A". If $A = B$, i.e. the edge is actually a loop, then we simply have two copies of it. The involution reverses the orientation of the dart. For a description we refer to [18, Section 1.3.3]: note that the map rev is denoted by α in there.

Given an embedded arc γ, its "reverse" is given by the class of embeddings obtained from composing γ with a *decreasing* homeomorphism of $[0, 1]$ onto itself. We will moreover abuse our notation and denote it $\mathrm{rev}(\gamma)$. The embedding ι in the above definition can therefore be extended to a map $\iota : \mathrm{Darts}(G) \to \Gamma(X)$, with the property that $\mathrm{rev}(\iota(e)) = \iota(\mathrm{rev}(e))$, and from now on we will refer to the latter.

Finally, the union of the images of the curves $\iota(d)$ with d varying over $\mathrm{Darts}(G)$, which is a closed subset of X, will be called the image of ι and denoted by $\iota(\mathrm{Darts}(G))$. We next define when two embeddings are "equivalent", cf. [18, Definition 1.3.7].

Definition 5.9 (Isomorphism between Graph Embeddings)**.** Let $\iota : G \to X$ and $\iota' : G' \to X'$ be two embeddings from two graphs G, G' onto two oriented surfaces X, X'. An isomorphism between these two embeddings is a pair (f, F), where $f : \mathrm{Darts}(G) \to \mathrm{Darts}(G')$ is a graph isomorphism commuting with rev, and $F : X \to X'$ is an orientation preserving homeomorphism such that $F(\iota(d)) = \iota'(f(d))$ for all $d \in \mathrm{Darts}(G)$, cf. Figure 7.

Clearly, the cases which are relevant to us are when X is either the standard sphere $\mathbb{S}^2$ or the Euclidean plane $\mathbb{R}^2$, which are both orientable. In what follows we will often use the terms "spherical graph" and "planar graph" for, respectively, the embedding of a graph in $\mathbb{S}^2$ and in $\mathbb{R}^2$.

Remark 5.10. We observe that the isomorphisms given in Definition 5.9 do not cover all the ones which are pertinent to our problem: for us, two stationary nets with antipodal pins are equivalent also when one is obtained from the other with a mirror symmetry of the two-dimensional sphere. However the latter map is not an orientation-preserving homeomorphism of $\mathbb{S}^2$. In order to avoid "duplicating" our candidate topologies for once-derived nets, this will be taken into account at a certain stage, cf. Remark 8.4.

5.4. **Twice-derived spherical graphs.** We are ready to introduce the concept of "twice-derived spherical graph". The idea is that it is easier to understand the embedded graph coming from a once-derived net if we get rid of the vertices of degree 2.

Definition 5.11 (Twice-derived spherical graphs)**.** Consider the spherical graph associated to a once-derived net and assume it has at least one vertex of degree 3. We obtain from it a *twice-derived spherical graph* by performing the following operation (as many times as possible):

- Remove a vertex of degree 2 and identify its two adjacent edges in a single edge. The embedding of the resulting edge is the concatenation of the embeddings of the two removed edges.

Remark 5.12. Consider the spherical graph associated to a once-derived net which is not empty and has only vertices of degree 2. If we attempt the same procedure as in Definition 5.11, namely we repeat the operation of removing inductively the vertices of degree 2 and merging the edges, the final result would have no vertices. In fact, since we know that the starting graph is connected, we would end up with a single loop. Strictly speaking this does not respect Definition 5.8: a graph with empty set of vertices is for us necessarily the empty graph. The stipulation, in Definition 5.11, that the spherical graph associated to the once-derived net has at least one vertex of degree 3 is just meant to ensure that the twice-derived spherical graph is well-defined. In order to avoid differentiating too many cases, we stipulate that the empty graph and the single loop with no vertices are also twice-derived spherical graphs and we call them both *extinct*. Likewise, a once-derived net which leads to an extinct twice-dervied spherical graph, will be called extinct.

Since the cases of an empty graph and of a graph with only vertices of degree 2 is excluded by Definition 5.11, they will be treated separately. In fact such cases will be completely classified in Section 14, cf. Theorem 14.1.

Proposition 5.13 (Properties of twice-derived spherical graphs)**.** *A twice-derived spherical graph, if it is not extinct, has the following properties:*

(i) It is connected.
(ii) All vertices have degree 3 *and there is at least one vertex.*
(iii) If v*,* e*, and* f *denote the number of vertices, edges, and faces of the once-derived net, then they satisfy the relations* $3v = 2e$ *and* $v - e + f = 2$ *and*[3] $\sum_{faces}(6 - \#sides) = 12$*.*
(iv) Every face has at most 7 *sides.*
(v) There are at most two faces with more than 5 *sides.*
(vi) It has at most 12 *faces and at most* 30 *edges.*

[3]Note that, when counting the number of sides in a face F, every side which is not in common with another face is in fact counted twice for F

Proof. The connectedness follows from the connectedness of the once-derived net, which is proven in Proposition 5.1(ii), and the fact that the procedure which generates the twice-derived spherical graph from a once-derived keeps it. (ii) is also straightforward because the procedure eliminates all vertices of degree 2. Concerning (iii), observe that $v - e + f = 2$ is Euler's formula for graphs embedded in the spheres. The relation $3v = 2e$ is obtained counting the edges in the following way: at each vertex there are 3 edges, but each vertex connects two edges, so $e = \frac{3v}{2}$. The final identity corresponds to $12 = 6f - 2e$ which can be immediately derived from the other two.

To see (iv) let $\mathcal{N}_1$ be the once-derived spherical graph and $\mathcal{N}_2$ be the twice-derived spherical graph. There is a natural bijection between the faces of the two spherical graphs. We thus fix $F \in \mathcal{F}(\mathcal{N}_2)$ and the corresponding face $G \in \mathcal{F}(\mathcal{N}_1)$. F is obtained from G by removing the vertices of G which have degree 2 and merging the corresponding edges, in particular the number of edges of G coincides with that of F minus its vertices of degree 2. It follows from Proposition 5.6 that the sides of G are at most $5+2n(G)$, where $n(G)$ is the number of degree 2 vertices of G. Hence the sides of F are at most most $5+n(G)$ and hence at most 7 because $n(G) \leq 2$ (cf. Proposition 5.6).

Coming to (v), the considerations above show that there are exactly two faces in $\mathcal{F}(\mathcal{N}_2)$ that originate from a nonconvex face of $\mathcal{F}(\mathcal{N}_1)$. So, except for the ones originating from these, all other faces in $\mathcal{F}(\mathcal{N}_2)$ have no more sides than some convex face in $\mathcal{F}(\mathcal{N}_1)$, namely at most 5 by Proposition 5.6.

Finally, with respect to (vi), note that the number of faces in $\mathcal{F}(\mathcal{N}_2)$ equals the number of faces in $\mathcal{F}(\mathcal{N}_1)$. By Proposition 5.6 the area of a face $F \in \mathcal{F}(\mathcal{N}_1)$ is an integer multiple of $\frac{\pi}{3}$ and hence at least $\frac{\pi}{3}$. Since the sum of the area of these faces is the total area of the sphere, namely 4π, their number does not exceed 12. Using (iii) we get

$$\sum_{F \in \mathcal{F}(\mathcal{N}_2)} \#\text{sides}(F) = 6(\#\text{faces} - 2) \leq 60\,.$$

However, given that

$$\sum_{F \in \mathcal{F}(\mathcal{N}_2)} \#\text{sides}(F) = 2e\,,$$

we conclude the second inequality in (vi). □

6. Polygonal paths and polygons

In this section study polygons in the sphere. We start with some definitions, being careful about whether self-intersections are allowed.

Consider a Lipschitz curve in the sphere which consists of a finite number of geodesic arcs γ_i on the sphere $\mathbb{S}^2$, with $i \in \{1, \dots, N\}$, joining vertices V_{i-1} and V_i. The arcs are supposed to have strictly positive lengths and be shorter than 2π. Given these data we can give to the path a canonical orientation: each geodesic arc is oriented so that V_{i-1} is the starting point and we assume further that γ_i is an arc-length parametrization of the geodesic arc with $\gamma_i(0) = V_{i-1}$ and $\gamma_i(\lambda_i) = V_i$. Given the latter convention we define the angle of the path at the vertex V_i (with $i \leq N-1$) as the counterclockwise angle (at V_i) from the tangent vector $-\dot\gamma_i(\lambda_i)$ to the tangent vector $\dot\gamma_{i+1}(0)$. The path will be called *closed* if $V_N = V_0$. We will identify two such closed paths up to choice of the " starting point" $V_0 = V_N$. Two nonclosed paths are congruent if there is an *orientation preserving* isometry mapping one (ordered) collection of

geodesic segments into the other, while for closed paths the congruence is understood in the same way up to choosing the starting vertex.

Definition 6.1. A *(spherical) polygonal path* is a curve γ as described above. The number of sides is the number of geodesic arcs in the description.

Note that we are not requiring that the path does not self-intersect, nor we are requiring that the lengths of the γ_i's are smaller than π (i.e. that the γ_i are the length-minimizing geodesic connecting V_{i-1} with V_i).

Definition 6.2. A (spherical) polygonal path will be called a *(spherical) polygon* if the path is closed and it is an embedding of $\mathbb{S}^1$ into $\mathbb{S}^2$. In these cases the polygon will be called (sphere) triangle, quadrilateral, pentagon, more in general N-gon depending of the number of its sides. The *interior* of the polygon is the topological disk bounded by the polygon with the property that the exterior normal to it forms a counterclockwise angle of $\frac{\pi}{2}$ with the tangent to the path. The polygon will be called convex if the interior is a convex set of the sphere.

Remark 6.3. It follows from Lemma B.1 that a spherical polygon is convex if and only if all its internal angles are smaller than or equal to π.

6.1. **Triangles.** In this section we focus on spherical triangles. When listing the angles of the triangle we will do so by ordering them in counterclockwise fashion as A, B, and C. The length of the segments will then be denoted by a, b, and c, with the convention that the angle A is opposite to the side a, and so on. This is the standard notation in textbooks on spherical trigonometry, see for instance [27].

Theorem 6.4. *Fix $A, B, C \in (0, \pi)$. There is a convex spherical triangle with angles A, B, C if and only if:*

- $\pi < A + B + C < 3\pi$,
- $\pi + 2\min\{A, B, C\} > A + B + C$.

Under these assumptions, there exists a unique (up to congruence) convex spherical triangle with angles A, B, C. The lengths of its sides a, b, c can be found using the (polar of the spherical) cosine rule:

$$\cos(a) = \frac{\cos(A) + \cos(B)\cos(C)}{\sin(B)\sin(C)}\,.$$

Viceversa the sides determine the angles via the formula:

$$\cos(A) = \frac{\cos(a) - \cos(b)\cos(c)}{\sin(b)\sin(c)}\,.$$

Finally the triangle is also determined by the measure of one side and its two adjacent angles and by the measure of an angle and its adjacent sides.

The theorem above is a classical fact in spherical trigonometry known since the nineteenth century (see [27]), but the reader can also derive it using linear algebra and elementary (planary) trigonometry. It will be useful for us to get a generalization of the theorem above to polygonal paths with three sides. We observe that, according to our definition, the three sides of the polygonal path must then belong to three *distinct* equators. We follow the same convention as above for labeling the measures of the angles of the polygonal path and the lengths of its sides

Remark 6.5. Note that when passing from convex triangles to polygonal paths with three sides, or even just to non-convex triangles, the conclusions of Theorem 6.4 must be handled with some care. For instance, if we know one angle A and the two measures b and c of the adjacent sides, we have determined, up to congruences, the other two vertices of the triangle. In particular one of the two sides must be different than π, otherwise the two vertices would coincide. We have thus determined the equator in which the third side must lie. However there are exactly two possibilities for this arc, and hence the knowledge of the triple of numbers A, b, and c does not determine the polygonal path, but it rather gives exactly two possibilities.

Theorem 6.6. *Fix $A, B, C \in (0, 2\pi) \setminus \{\pi\}$. If there is an odd number of values in (A, B, C) smaller than π, let $(A', B', C') := (A, B, C)$ and otherwise let $(A', B', C') := (2\pi - A, 2\pi - B, 2\pi - C)$. Let $A'' := (A' \mod \pi)$ and cyclic.*

There is a polygonal path with three angles (A, B, C) if and only if there is a convex triangle with angles (A'', B'', C'') (see Theorem 6.4).

Under this assumption, there are exactly two (up to congruence) polygonal paths with (A, B, C) as angles. The lengths of their sides can be obtained applying the following set of rules.

- *If (a, b, c) is a valid triple of side lengths, then the other triple is $(2\pi-a, 2\pi-b, 2\pi-c)$.*
- *The valid triples of side lengths for the angles (A, B, C) coincide with the valid triples of side lengths for the angles $(2\pi - A, 2\pi - B, 2\pi - C)$.*
- *(a, b, c) is a valid triple of side lengths for the angles (A, B, C) (with $B, C < \pi$) if and only if $(2\pi - a, b, c)$ is a valid triple of side lengths for the angles $(A, \pi + B, \pi + C)$.*
- *If $0 < A, B, C < \pi$, then a valid triple of side lengths can be obtained as in Theorem 6.4.*

The above theorem is an exercise which we leave to the reader.

6.2. **Congruence criterion for closed polygonal paths.** The following fact is a corollary of the above considerations for spherical triangles and will be particularly useful to analyze the special nets of Section 14 and Section 15.

Theorem 6.7. *Assume $V_1V_2 \dots V_N$ and $W_1W_2 \dots W_N$ are the vertices of two closed polygonal paths such that:*

(i) The angles $\hat{V}_i$ are all congruent to the angles $\hat{W}_i$, and all ddiferent than π;
(ii) The geodesic arcs V_iV_{i+1} are congruent to the edges W_iW_{i+1} for all $i \in \{1, \dots, N-3\}$;
(iii) The last three edges of both paths are all shorter than π.

Then the two polygonal paths are congruent.

The proof is given in Section D.

6.3. **Convex spherical polygons with angles equal $\frac{2\pi}{3}$.** The explicit formulas in Proposition 6.8 below will be of fundamental importance throughout the rest of the paper, as they give a simple way to relate the lenghts of the sides in the four type of poligons which have angles all equal to $\frac{2\pi}{3}$.

Proposition 6.8. *A convex spherical triangle with angles all equal to $\frac{2\pi}{3}$ is equilateral and its side measures*

$$\lambda_T = \arccos(-\tfrac{1}{3})\,. \tag{6.1}$$

In a convex spherical quadrilateral with angles equal to $\frac{2\pi}{3}$ opposite sides are congruent. Moreover, the measures a and b of two adjacent sides satisfy the relation

$$\tan\tfrac{a}{2}\tan\tfrac{b}{2}=\tfrac{1}{2}\,. \tag{6.2}$$

If two adjacent sides are equal (which means that all sides are congruent) then the diagonal measures λ_T.

A convex spherical pentagon $ABCDE$ with angles equal to $\frac{2\pi}{3}$ is determined by the length of any two of its sides and every side is strictly shorter than λ_T. Moreover, if the vertices are labeled consecutively, then

$$2\cos|DE|-\tfrac{1}{3}=f(|AB|,|BC|)\,. \tag{6.3}$$

where

$$f(x,y):=\cos x+\cos y+\cos x\cos y-2\sin x\sin y\,. \tag{6.4}$$

The sidelength of an equilateral triangle with angles $\frac{2\pi}{3}$ is well known. (6.3) is instead taken from [25] (where, however, it is not proved) and the other statements appear in [14], except for (6.2), which does not seem to appear anywhere. Anyway, we provide proofs for all the statements in Section C, based on the trigonometric identities for triangles given in Theorem 6.4.

7. Sieves

Eventually, in the following section, we plan to describe an algorithm which will produce a finite list of spherical graphs[4] such that any once-derived net, when interpreted as a spherical graph, belongs to the list.

In order to shorten our list, we show here a number of conditions that any once-derived net must satisfy in addition to those already found in Proposition 5.6. We will first state all the conditions: each condition will be stated separately in a lemma and will be assigned a name; the proofs will then follow in separate subsections, each bearing the name of the lemma.

We start by introducing the following concepts. Recalling Proposition 5.6, the area of a face depends only on the number of concave angles and the number of edges of the face.

Definition 7.1 (Boundary distance)**.** Given two vertices A and B in the boundary of a face F we introduce the following.

- A *boundary path* from A to B is an ordered list of $N+1$ vertices $A=V_0,V_1,\dots,V_N=B$ with the property that, for each $i\in\{0,\dots,V_{N-1}\}$, V_i and V_{i+1} are connected by an edge σ_i of F and the union of the edges σ_i form a positively curved oriented arc.
- The *boundary distance* of A and B in the face F is the minimal N for which there is boundary path as in (a) either from A to B or from B to A.
- The boundary distance of A and B *avoiding a specific set* $\mathcal{E}$ of edges is the minimal N for which there is a boundary path with the additional property that none of the edges σ_i belongs to $\mathcal{E}$.

We are now ready to state all the sieves. In the next statements, $\mathcal{N}$ denotes a non-extinct once-derived net.

Lemma 7.2 (No convex triangles next to convex pentagons)**.** *A convex triangle of $\mathcal{N}$ cannot share an edge with a convex pentagon.*

[4]Actually, we will keep track of a bit more information than just a spherical graph, see Definition 8.3.

Lemma 7.3 (No concave biangles)**.** *$\mathcal{N}$ cannot contain a concave biangle.*

Lemma 7.4 (Simple boundary)**.** *The boundary of any face of $\mathcal{N}$ is topologically an $\mathbb{S}^1$.*

The proof of Lemma 7.4 also implies easily the following fact, which is not used as a sieve but will be useful in other contexts.

Corollary 7.5. *Two convex faces of the net $\mathcal{N}$ have at most one edge in common.*

Lemma 7.6 (Biangle)**.** *Let $\mathcal{N}$ be such that one face is a biangle B. Then*

(i) B is the only biangular face.
(ii) The two faces adjacent to B are distinct and nonconvex.
(iii) The convex faces (except for the biangle) are all strictly *contained in a single hemisphere.*
(iv) Any convex triangle in $\mathcal{N}$ cannot share a side with another convex face.

Moreover, the sum of the areas of the convex faces (excluding the biangle), plus the contribution *of the nonconvex faces is at most $\frac{5\pi}{3}$, where the contribution of a nonconvex face is defined as*

(a) 0 if the face has a single concave angle;
(b) $\frac{\pi}{3}\max\{0, 4-d\}$, where d is the boundary distance of the two concave angles avoiding the edges of B, if the face has two concave angles.

Lemma 7.7 (Long edge)**.** *Assume that there is an edge σ of $\mathcal{N}$ which is strictly longer than π or with length π which is not the edge of a biangle. Then all the conclusions of Lemma 7.6 hold, specifically:*

(i) The two faces adjacent to σ are distinct and nonconvex;
(ii) The convex faces are all strictly *contained in a single hemisphere;*
(iii) The sum of the areas of the convex faces plus the contribution from the nonconvex faces is at most $\frac{5\pi}{3}$, where the contribution of a nonconvex face is defined as
  (a) 0 if the face has a single concave angle;
  (b) $\frac{\pi}{3}\max\{0, 4-d\}$, where d is the boundary distance of the two concave angles avoiding σ, if the face has two concave angles.

Lemma 7.8 (Concave angles are far)**.** *The boundary distance between the two concave angles of a face of $\mathcal{N}$ is at least $s-5$, where s is the number of sides of the face.*

Proof of Lemma 7.2. This is proved by Heppes in [14], at any rate it follows immediately from Proposition 6.8. □

Proof of Lemma 7.3. If $\mathcal{N}$ contains a concave biangle, then the two vertices of the biangle have degree 2 and thus the two edges adjacent to them are precisely the edges of the biangle. In particular, the boundary of the biangle is a connected component of $\mathcal{N}$ and, because $\mathcal{N}$ is connected, it is the whole net. But then $\mathcal{N}$ would be an extinct net. □

Proof of Corollary 7.5. The statement follows immediately from Proposition B.4(v) □

Proof of Lemma 7.4. Let F be the face we are considering. If F is a convex face there is nothing to prove.

Consider the case of a nonconvex face with a single concave angle A. Start from A and follow the bisector of A inside the face F until we hit the first point B of the boundary. The resulting geodesic arc η subdivides the face F into parts F_1 and F_2. Using Lemma B.1 we

conclude that both are convex. If we consider the equator γ containing η, we conclude that, since in a neighborhood of A there are points of γ which do not belong do F_1 nor F_2, $F_1 \cap F_2$ contains η, but is not γ. It must therefore be a geodesic arc which contains η. But note that there are points in γ arbitrarily close to B which do not belong to F_1 nor to F_2. So $F_1 \cap F_2 = \eta$. In particular, $\partial F_1 \setminus \eta$ and $\partial F_2 \setminus \eta$ are two disjoint embedded open arcs with endpoints A and B. Since ∂F is the union of the latter arcs and of the points A and B, we conclude that ∂F is topologically $\mathbb{S}^1$.

Consider now the case of a nonconvex face F with two concave angles A and B. In the stationary net that generated $\mathcal{N}$ there is a pin P which is connected to A and B with two edges s and t, contained in the face F. If s and t form an angle π at P then we argue as above decomposing F into two convex faces F_1 and F_2 whose intersection consists only of $s \cup t$. Otherwise we take the bisector of the concave angle formed at P by the edges connecting P at A and B and follow it until it hits the first point C in ∂F. The three geodesic segments PA, PB, and PC decompose F into the closures of three convex open sets F_1, F_2, and F_3 and, again, the intersection of any pair of them is the shortest geodesic connecting P with one of the three points A, B, and C. We thus conclude that ∂F can be decomposed as the union of three embedded open arcs with empty intersection plus their endpoints A, B, and C, forming a topological $\mathbb{S}^1$. □

Proof of Lemma 7.6. First of all denote by V and U the two vertices of the biangle and observe that it must be a convex biangle by Lemma 7.3. If one vertex is a pin the other is a pin too and the biangle would be a connected component of $\mathcal{N}$, hence the whole net, which in turn would be elementary. So both V and U have degree 3. In particular the biangle is also a face of the net $\mathcal{O}$ from which $\mathcal{N}$ is derived. Now rotate the two half meridians forming the biangle keeping V and U fixed, moving them into the complement of the biangle, until they intersect the first points of $\mathcal{O} \setminus \partial B$. If the two meridians meet then the net would be a triple junction, with the two pins lying on the edges, which is impossible. So each rotated meridian must meet a distinct vertex of $\mathcal{O}$ and these two vertices must both be pins. In particular, once they stop, the two rotated meridians form a single equator. The latter subdivides the sphere into two hemispheres. One of them contains the biangle B and the other one contains all the other convex faces of the net $\mathcal{N}$. So the two faces adjacent to B must both be nonconvex.

Next, consider the hemisphere H which contains all the convex faces of $\mathcal{N}$ and fix a pin P, which is contained in the boundary of the hemisphere. If P has degree 2 then both edges adjacent to the pin P either touch ∂H coming from H, or they lie in ∂H. In particular this implies that P is the vertex of a convex face of $\mathcal{O}$ which was removed when deriving the net $\mathcal{N}$. This face was contained in the hemisphere H. We thus conclude that the sum of the areas of the convex faces of $\mathcal{N}$ plus the eventual convex faces of $\mathcal{O}$ which were removed (one for each pin of degree 2) can be at most 2π.

Consider any of the convex faces F removed from $\mathcal{O}$, with P the corresponding pin of degree 2. The internal angle of F at P at is at least $\frac{2\pi}{3}$. If $s(F)$ is the number of sides of F, then by Gauss-Bonnet the area of F is at least $(6 - s(F))\frac{\pi}{3}$. Observe that F is contained in one concave face G of $\mathcal{N}$. The pin P has degree 2 if and only if this face G has two concave angles. In fact if the pin had degree 1, then G would have one concave angle. If the pin had degree 3 then G would need to be the union of two convex polygons. Note however that one would share an edge with a biangle, so would need to be a biangle, implying that the other one would also need to be a biangle. But then the entire net $\mathcal{O}$ would consist of the two pins connected by three meridians and $\mathcal{N}$ would be empty.

Having established that each "removed" face F accounted corresponds to a nonconvex faces G of $\mathcal{N}$ with two concave angles containing it, we now show that $(6-s(F))\frac{\pi}{3}$ coincides with our definition of contribution of the face G. In fact F has $s(F)-2$ sides in common with G. These sides form a boundary path connecting the two concave angles of G and avoiding the side σ of G in common with the biangle: $s(F)-2$ is the boundary distance of the two concave angles avoiding σ. Hence the area of F is $(4-(s(F)-2))\frac{\pi}{3}$, precisely the "contribution of the nonconvex face G". Since the contribution of G has been defined as 0 when there is a concave angle, the argument given so far shows that the sum of the areas of the convex faces of $\mathcal{N}$ plus the contributions of the nonconvex faces is at most the area of the hemisphere H, i.e. 2π. Note however that:

- If one of the two pins has degree 1, then the union of the faces considered has area strictly smaller than H;
- If both pins have degree 2, the only way in which the union of the faces considered in the computation could have the same area as H is that the two angles at the pin are exactly π; but in this case the "contribution" of each nonconvex face G is strictly less than the area of the corresponding convex face F which was removed.

The sum of the areas of the convex faces plus the contribution of the nonconvex ones must therefore be strictly smaller than 2π. Since however this number is an integer multiple of $\frac{\pi}{3}$, it can be at most $\frac{5\pi}{3}$.

Finally, (iv) follows because a convex triangle cannot share an edge with a convex pentagon and if it shares an edge with a convex quadrilateral, then their union is a set of diameter strictly larger than π. This verification is left to the reader, after observing that Proposition 6.8 completely determines the diameter of the union. □

Remark 7.9. We remark here an implicit effect of the proof of Lemma 7.6: if the derived-once net is nonextinct and it contains a biangle, then no pin of the net can have degree 3.

Proof of Lemma 7.7. The proof is achieved using the same idea of the proof of the biangle lemma. If the length of the edge is strictly larger than π we choose two interior antipodal points and rotate the corresponding half meridian in both directions keeping its endpoints fixed. If the length of the edge is exactly π but it is not an edge of a biangle, then we rotate it in both directions. Like in the proof of Lemma 7.6 this operation "sweeps" an hemisphere in the original net $\mathcal{O}$ from which $\mathcal{N}$ was derived. If H is the complementary hemipshere, then ∂H contains the two pins. The considerations in the proof of Lemma 7.6 apply now verbatim to this case as well. □

Proof of Lemma 7.8. If the face F has $s \leq 6$ sides, then the statement is trivial. Thus, we can assume that F is a face with two concave angles and $s \geq 7$ sides. Its area is $(10-s)\frac{\pi}{3}$. Let d be the boundary distance between the concave angles in F. Let P be the pin in the net $\mathcal{O}$ which is contained in the face F.

The pin P together with the $d+1$ vertices on the boundary of the face which form the path in ∂F connecting the two concave angles constitute a simple polygon Q strictly contained in F. The polygon Q has $d+1$ angles equal to $\frac{2\pi}{3}$, and one (the one at the pin) greater or equal than $\frac{2\pi}{3}$. By Gauss-Bonnet, the area of Q is at least $(6-(d+2))\frac{\pi}{3}$.

Since Q is strictly contained in F, the area of Q must be strictly smaller than the area of F. Therefore $6-(d+2) < 10-s$ which is equivalent to $s-5 \leq d$. □

8. Computer-assisted list of candidate topologies

In this section, we describe the algorithms we use to obtain a finite list of candidates for the twice-derived spherical graphs and for the spherical embedded graphs induced by once-derived nets. Let us recall in the following definition the special cases which we will neglect in this section.

Definition 8.1. A twice-derived spherical graph is *extinct* if it is either empty or a loop with no vertices. A once-derived net is extinct if the corresponding twice-derived spherical graph is extinct.

The extinct cases include most of the geometries in Theorem 1.3. We will therefore analyze them directly in Section 14. We can thus focus on the remaining cases, namely the twice-derived spherical graphs in which there is at least one vertex. The first computer algorithm described in this section, combined with the statements of the previous section, proves:

Proposition 8.2. *Any twice-derived spherical graph (as in Definition 5.11) which is not extinct belongs to the list given in* `data/twice_derived_human.txt` *in* [13].

Given a once-derived net $\mathcal{N}$ we regard it as a particular embedding in $\mathbb{S}^2$ of a graph G, carrying, however one additional piece of information: for each vertex of degree 2, we record which angle is concave.

Definition 8.3 (Once-derived marked spherical graph)**.** A marked spherical graph is a graph G embedded in $\mathbb{S}^2$ which has only vertices of degree 3 and 2, where for each vertex of degree 2 we also mark one edge arriving at it. Observe that any once-derived net produces a marked spherical graph when we use the concavity of the angle to mark the edge. More precisely, we adopt the convention that the counterclockwise angle from the marked edge to the next one is concave. The marked spherical graph produced by a once-derived net is a *once-derived marked spherical graph*.

Remark 8.4. Following Definition 5.9 such two marked spherical graphs are isomorphic if they can be mapped one into the other via an orientation-preserving homeomorphism of $\mathbb{S}^2$ which also "respects" the concave angles, namely if the homeomorphism ι maps a degree 2 vertex A of one graph in the vertex A' of the other, then it must map the marked edge arriving at A to the marked edge arriving at A'. Though, for our purposes, we shall also identify mirror symmetric images: in other words we will drop the requirement that ι is orientation preserving. We warn the readers that the list mentioned in Proposition 8.2 will instead not distinguish between a twice-derived spherical graph and its mirror image, cf. Remarks 5.10 and E.3.

The second algorithm described in this section takes as input the list of Proposition 8.2 and produces a corresponding list of possibilities for once-derived marked spherical graphs, avoiding any "double counting" through isomorphic copies (including mirror symmetries, as detailed in the remark above).

Proposition 8.5. *Any once-derived marked spherical graph belongs to the list given in* `once_derived_human.txt` *in* [13] *(see also* [12]*), up to isomorphism as described in Remark 8.4.*

8.1. **High-level illustration of the algorithms.** Here we give a quick illustration of the main ideas, especially suited for readers who already have a working knowledge of concepts in graph theory. The actual algorithms will be more intricate in order to be computationally

feasible (namely ensuring that the algorithms can be run in a reasonable time on common hardware).

Recall that, as shown in Proposition 5.13, any twice-derived spherical graph is a spherical graph where each vertex has degree 3 and the total number of vertices is at most 30. In particular, there are finitely many twice-derived spherical graphs. We want to generate all of them, starting from a cycle and adding one edge at a time.

We generate the list of possible twice-derived spherical graphs by recursively building families of planar graphs, denoted as $\mathcal{C}_k$, where k is the number of edges. While twice-derived spherical graphs are in fact spherically embedded, we can easily regard them as planar (see Section 8.3). The algorithm builds $\mathcal{C}_{k+1}$ from $\mathcal{C}_k$ as follows. The first element of $\mathcal{C}_{k+1}$ is the cyclic graph with $k+1$ edges. All the other elements are created as "nice children" of elements in $\mathcal{C}_k$: the parent graph in $\mathcal{C}_k$ is modified by either attaching a new vertex externally or connecting two external vertices with an edge. Along this process we force the property that graphs in $\mathcal{C}_k$ have at most one leaf. One can show that this algorithm generates all planar graphs without leaves.

To keep the search feasible (i.e. to prevent a fast growth in the size of $\mathcal{C}_k$), the algorithm applies many pruning criteria based on topological constraint derived from the geometric information (cf., Lemma 8.16). Moreover, the algorithm discards isomorphic copies so that C_k does not contain two isomorphic planar graphs. In order to perform this step efficiently some care is required (e.g., we employ a polynomial-time isomorphism check and a sophisticated hashing technique for planar graphs).

Once we have produced a list of spherical graphs that contains all twice-derived spherical graphs, we employ a second (simpler) algorithm to obtain from it a list of marked spherical graphs that contains all once-derived marked spherical graphs. The algorithm iterates through the twice-derived candidates and systematically splits edges to insert between 2 and 4 vertices of degree two. The reconstruction process also keeps track of where the "concave angles" can be placed, cf. Definition 8.3). The resulting objects are checked against the properties of once-derived nets which have been proved in the previous sections (cf. Proposition 5.1 and Section 7). This step also accounts for mirror symmetries.

8.2. **Combinatorial maps.** The material contained in this section is well known. We start with the key observation that equivalence classes of graph embeddings into surfaces can be described through combinatorial data. We recall the notions of darts and darts graph $\mathrm{Darts}(G)$, given in Section 5.3.

Definition 8.6 (Ordering induced by an embedding)**.** Let $G = (V, E)$ be a graph and let $\iota : \mathrm{Darts}(G) \to \Gamma(X)$ be an embedding of G into an *oriented 2-dimensional surface* X. The embedding ι canonically induces a function $\mathrm{next} : \mathrm{Darts}(G) \to \mathrm{Darts}(G)$ as follows. Take a dart e starting from $v \in V$. Look at a neighborhood of $\iota(v)$ in X and assume that, if we start from a point on $\iota(e)$ close to v and move counterclockwise around v (which makes sense because X is an oriented surface), the first edge we encounter is $\iota(e')$. Then, we set $\mathrm{next}(e) = e'$.

The reader can compare with [18, Section 1.3.3], where the map next is denoted by σ. The following lemma summarizes [18, Lemma 1.3.15 & Proposition 1.3.16] (note that the notation φ for the map $\mathrm{next} \circ \mathrm{rev}$ is not completely consistent with the one in [18] as in there the same symbol denotes the inverse of our φ; this change is however immaterial for the conclusion).

Lemma 8.7. *Let $G = (V, E)$ be a graph, and let $\iota : \mathrm{Darts}(G) \to \Gamma(X)$ be an embedding of G into an* oriented surface *X. Let* next $: \mathrm{Darts}(G) \to \mathrm{Darts}(G)$ *be the associated function. We have:*

- *For any $e \in \mathrm{Darts}(G)$, the starting points of e and* next(e) *coincide.*
- *If $e, e' \in \mathrm{Darts}(G)$ have the same starting point, there is a $k \in \mathbb{N}$ such that* next$^k(e) = e'$.
- *Consider the map $\varphi = \mathrm{next} \circ \mathrm{rev} : \mathrm{Darts}(G) \to \mathrm{Darts}(G)$. The image, through ι, of an orbit of φ is (one connected component of) the boundary of a connected component of $X \setminus \iota(\mathrm{Darts}(G))$, traversed in the clockwise direction. The orbits of φ (as well as their images in X through ι) are called* faces. *A vertex belongs to a face if it is a vertex of a dart forming the face.*

Since we will only consider embedded graphs and their embeddings ι into the simply connected spaces $X = \mathbb{R}^2$ and $X = \mathbb{S}^2$, to each "face" as defined above will correspond a unique connected component Ω of the complement in X of the image of the graph, with the property that $\partial\Omega$ is its topological boundary. These will also be called faces. When $X = \mathbb{S}^2$, all of them are topological disks. When $X = \mathbb{R}^2$, all the bounded ones are topological disks, and there is precisely one unbounded one.

We can now introduce the primary combinatorial object which we will use to give our list of possibilities for twice-derived spherical graphs. The reader can compare with [18, Definition 1.3.23 & Remark 1.3.24].

Definition 8.8 (Combinatorial Map)**.** A combinatorial map is a graph G endowed with a function next $: \mathrm{Darts}(G) \to \mathrm{Darts}(G)$ such that

- For any $e \in \mathrm{Darts}(G)$, the starting points of e and next(e) coincide.
- If $e, e' \in \mathrm{Darts}(G)$ have the same starting point, there is $k \in \mathbb{N}$ such that next$^k(e) = e'$.

For our purposes, the most relevant fact of the theory is the following one; for the proof we refer the reader to [18, Construction 1.3.20].

Theorem 8.9. *Any combinatorial map is induced by an embedding of the underlying graph into an oriented surface such that all faces are simply connected.*

In our case the relevant combinatorial maps are the ones induced by embeddings in the sphere. On the other hand the topology of the oriented surface given by the above theorem is completely determined by the Euler characteristic of the graph. Another important point in our classification is that we wish to avoid repetitions: two combinatorial maps inducing isomorphic embeddings of the same graph should not appear in our list. There is in fact a simple way of recognizing isomorphic embeddings through their combinatorial maps (cf. again [18, Definition 1.3.23]).

Definition 8.10 (Isomorphism of Combinatorial Maps)**.** Let (G, next) and (G', next) be two combinatorial maps. An isomorphism between these two combinatorial maps is a bijection $f : \mathrm{Darts}(G) \to \mathrm{Darts}(G')$ that commutes with rev and next.

For the following theorem we refer to [18, Section 1.3.3] (the theorem is not explicitly stated in the reference, but stated informally in [18, Bottom of p. 37]).

Theorem 8.11. *Let $\iota : \mathrm{Darts}(G) \to X$ and $\iota' : \mathrm{Darts}(G') \to X$ be two embeddings of two graphs into two oriented surfaces such that*

- *G and G' are connected, and they both have at least one edge.*
- *X and X' are connected.*
- *Every face of ι and ι' is simply connected.*

Then, the two embeddings ι and ι' are isomorphic if and only if the associated combinatorial maps are isomorphic.

From the theorem we infer immediately the following corollary.

Corollary 8.12. *Let (G, next) be a connected combinatorial map with at least one edge. There is a unique (up to isomorphism) embedding of G into a a 2-dimensional oriented surface, compatible with the combinatorial map, such that its faces are simply connected. The 2-dimensional surface is determined by the Euler characteristic $v - e + f$ of the combinatorial map.*

8.3. **Spherical and planar maps.** Even though we are interested in embeddings of graphs in the sphere, the algorithm that proves that all twice-derived spherical graphs belong to a finite list (cf., Proposition 8.2) works with embeddings of graphs in the plane. Let us spend some time to explain how to pass from a spherical graph to a planar one and vice versa.

We can pass from a spherical embedding with simply connected faces to a planar embedding by removing from the sphere one point which is not in the image of the embedding. Vice versa, to pass from the planar embedding to the spherical embedding it suffices to take the one-point compactification of the plane.

At the level of combinatorial maps this can be described as follows.

Definition 8.13 (Spherical and Planar Maps)**.** A *spherical map* is a combinatorial map for a graph with the Euler characteristic of the sphere. A *planar map* is a spherical map with a special face, namely a triple (G, next, Y), where (G, next) is a combinatorial map, and Y one of its faces, namely an orbit of $\varphi = \mathrm{next} \circ \mathrm{rev}$.

Once we fix an embedding corresponding to the underlying spherical map, we get a corresponding planar map by simply choosing a face in the spherical embedding. In geometric terms, this identifies a connected component in the complement in $\mathbb{S}^2$ of the image of the graph: by removing a point from this connected component, we find a planar embedding. The complement in $\mathbb{R}^2$ of the image of this embedding consists of one unbounded connected component and finitely many bounded connected components which are topological disks. Moreover the unbounded connected component's first homotopy group is $\mathbb{Z}$.

Conversely, given a connected graph and a planar embedding of it, if we compactify $\mathbb{R}^2$ to $\mathbb{S}^2$ by adding a point at infinity we recover a spherical embedding of the graph. The operations just describe give therefore a correspondence between spherical and planar maps (which clearly is not one-to-one). Note that it is also natural to use the term "unbounded face" for the special face of the planar combinatorial map.

8.4. **Generating the list of candidates.** The algorithm for generating the list of candidates twice-derived spherical graphs is recursive over the topological complexity of the graph. In order to describe the recursive procedure we introduce the following terminology. Recall that, in graph theory, a "leaf" of a graph is a vertex of degree 1.

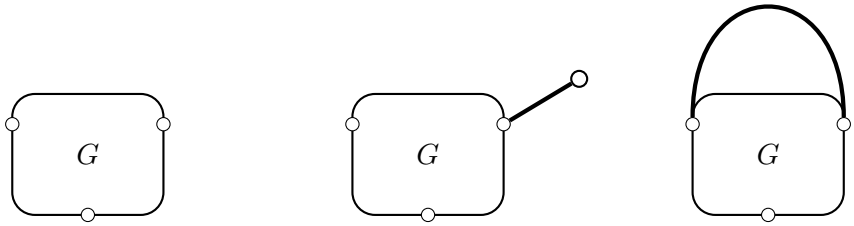
Definition 8.14 (Children of a Planar Map)**.** Let (G, next, Y) be a planar map. Its *children* are the planar maps that can be obtained from (G, next, Y) with one of the following two operations (cf. Figure 8):

- Attach a new vertex "externally" to a vertex in Y (the special face of the child consists of Y and the two additional darts coming from the added edge).
- Connect two vertices (the two vertices can coincide) in Y "externally" with a new edge (the special face of the child consists then of a proper subset of Y and one of the darts coming from the added edge).

The *nice children* are the children such that:

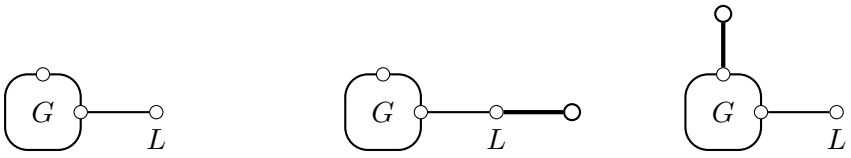
- Any leaf of G is not a leaf of the child (cf. Figure 9).

A *descendant* of a planar map is either the planar map itself or a child of a descendant of the planar map. A *nice descendant* of a planar map is either the planar map itself or a nice child of a nice descendant of the planar map.



(a) Parent. (b) Operation 1 (c) Operation 2

FIGURE 8. The operations defining the children of a planar map. (a) A base planar map (G, next, Y) where the rectangle represents the boundary Y. (b) The first operation: attaching a new vertex (thick circle) to the boundary via a new edge (thick line). (c) The second operation: connecting two existing boundary vertices via a new edge (thick curve) drawn through the outer face.



(a) Map with Leaf L (b) Nice Child (c) Not Nice Child

FIGURE 9. Illustration of "Nice" vs. "Not Nice" children. (a) A map $X = (G, \text{next}, Y)$ containing a leaf vertex L and another boundary vertex on top. (b) A *Nice Child*: The new edge is attached to L. Consequently, L becomes an internal vertex (degree 2) and is no longer a leaf. (c) A *Not Nice Child*: The new edge is attached to the other boundary vertex, so L remains a leaf in the child map.

The following lemma allows us to consider only nice children during our search.

Lemma 8.15. *Any connected planar map with at least one edge that either has no leaves, or has only one leaf contained in the unbounded face, is a nice descendant of a cycle.*

Proof. First of all, we observe that the statement is obvious if the map itself is a cycle. Then we prove the general case by induction on the number of edges. If there is only one edge, then the two endpoints of the edge must coincide, and thus the map is a cycle.

Assume now that the statement holds when the number of edges is $k \geq 1$ and fix a planar map X as in the statement with $k+1$ edges. If there is one leaf, we can remove it to get a parent planar map which satisfies the assumption of the lemma, from which we conclude that X is a nice descendant of a cycle from the inductive hypothesis. If there are no leaves, then all vertices have degree at least 2. Consider the unbounded face. If all the vertices adjacent to it have degree exactly two, then the face is a connected component of the planar map, and thus the entire map is the face itself. In particular it is a cycle. Otherwise there is a vertex V of degree 3. If we now remove one of the edges in the unbounded face which is adjacent to V, the resulting parent planar map X^p has k edges and is still connected. Consider the other vertex V' adjacent to the edge which has been removed. If $V' = V$ then either V had degree > 3 in Y, and hence X^p has no leafs, or it had degree exactly 3, and then X^p contains only one leaf, which is contained in its unbounded face. If $V' \neq V$, then V has degree at least two in X^p, while V', which belongs to the unbounded face of X^p, is the only vertex which could be a leaf in X^p. In all these cases X^p falls under the assumptions of the lemma, and thus by the induction assumption it is a nice descendant of a cycle, in turn implying the same property for X. □

A first algorithm to prove that twice-derived spherical graphs belong to a finite list could go as follows. We define recursively families $\mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, \dots$ of planar maps such that $\mathcal{F}_k$ contains only planar maps with exactly k edges.

- We define $\mathcal{F}_0$ as the empty set.
- $\mathcal{F}_{k+1}$ consists of all the nice children of the planar maps in $\mathcal{F}_k$ and of the cycle of length $k+1$.

Because of Lemma 8.15, a connected planar map without leaves and with k edges belongs to $\mathcal{F}_k$. Since we know that a twice-derived spherical graph has at most 30 edges, we can examine all the generations $\mathcal{F}_k$ with $k \leq 30$ and check which planar maps in these lists satisfy the restrictions of Proposition 5.13. However the cardinality of the families $\mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, \dots$ grows quickly (cf. e.g. [22]), making the above algorithm unfeasible. Hence, the following criteria help us making $\mathcal{F}_k$ sensibly smaller, since they allow us to throw away elements of $\mathcal{F}_k$ for which we are sure that none of their descendants will satisfy the properties of Proposition 5.13.

Lemma 8.16. *Let (G, next, Y) be a planar map. Denote by $\mathcal{V}(G)$, $\mathcal{E}(G)$, $\mathcal{F}(G)$ the vertices, edges, and faces of G. Denote by $\mathcal{BV}(G)$ the vertices on the boundary and by $\mathcal{IF}(G)$ the* bounded *faces of G. Assume that a descendant of (G, next, Y) comes from a twice-derived spherical graph after removing a point from the spherical embedding (cf., Section 8.3). The following properties hold.*

(1) Every $F \in \mathcal{IF}(G)$ has at most 7 sides.
(2) Let $\gamma := \#\{F \in \mathcal{IF}(G) : |F| > 5\}$. Then $0 \leq \gamma \leq 2$.
(3) $\deg_G(w) \leq 3$ for any $w \in \mathcal{V}(G)$ and $\deg_G(w) = 3$ for $w \in \mathcal{V}(G) \setminus \mathcal{BV}(G)$.

(4) *Let* $Q_i(j) := 6q_i(j) - j$, *where* $q_i(j)$ *is the smallest size of a collection of positive integers in* $[1,7]$, *with sum equal to* j, *with at most* i *numbers larger than* 5. *Set* $a := \sum_{v\in\mathcal{BV}(G)}(3-\deg_G(v))$ *and*

$$\tilde a = \begin{cases} a & \text{if } a \text{ is even,}\\ a+3 & \text{if } a \text{ is odd.}\end{cases}$$

We have

$$\min_{\substack{\tilde a\le\alpha\le 60-|\mathrm{Darts}(G)|\\ 6|\alpha-\tilde a}} Q_{2-\gamma}\big(|Y|+\alpha\big) + \sum_{F\in\mathcal{IF}(G)}(6-\mathit{sides}(F)) \le 12\,.$$

Proof. Let (G',next',Y') be the descendant of (G,next,Y) that is a twice-derived spherical graph (up to arbitrarily choosing a face to obtain a planar graph).

Note that a bounded face of G is necessarily also a bounded face of G'. Recalling Proposition 5.13 no face of G' can have more than 7 sides and this proves statement (1). Moreover there can be at most two faces of G' with more than 5 sides, which proves (2). Next, a vertex in $\mathcal{V}(G)\setminus\mathcal{BV}(G)$ has the same degree in G as in G', which proves the second part of (3). In general the degree of a vertex can only increase in G', which shows the first part of (3).

Coming to (4), let $a' = |\mathrm{Darts}(G')\setminus\mathrm{Darts}(G)|$. Since G' is a cubic graph, we have

$$a' = a + 3|\mathcal{V}(G')\setminus\mathcal{V}(G)|.$$

Observe also that a' is even because $a' = 2|\mathcal{E}(G')\setminus\mathcal{E}(G)|$. Thus, we deduce $6 \mid a'-\tilde a$. Moreover, $a' + \mathrm{Darts}(G) = \mathrm{Darts}(G') \le 60$ by Proposition 5.13, and thus $\tilde a \le a' \le 60 - \mathrm{Darts}(G)$.

Again by Proposition 5.13, we have

$$\sum_{F\,\in\,\mathcal{F}(G')\setminus\mathcal{IF}(G)}(6-\mathrm{sides}(F)) + \sum_{F\in\mathcal{IF}(G)}(6-\mathrm{sides}(F)) = \sum_{F\in\mathcal{F}(G')}(6-\mathit{sides}(F)) = 12\,.$$

Our observations reduce the desired statement to proving

$$\sum_{F\,\in\,\mathcal{F}(G')\setminus\mathcal{IF}(G)}(6-\mathrm{sides}(F)) \ge Q_{2-\gamma}\big(|Y|+a'\big)\,.$$

The sum of the sides of F for F running in $\mathcal{F}(G')\setminus\mathcal{IF}(G)$ is exactly the number a' of new darts introduced to get G' and the darts in the unbounded face of G, in other words

$$\sum_{F\,\in\,\mathcal{F}(G')\setminus\mathcal{IF}(G)}\mathrm{sides}(F) = |Y| + a'\,.$$

The sides of each face in $\mathcal{F}(G')\setminus\mathcal{IF}(G)$ are at least 1 and at most 7 and there are at most $2-\gamma$ such faces with more than 5 sides. Hence, we deduce that

$$|\mathcal{F}(G')\setminus\mathcal{IF}(G)| \ge q_{2-\gamma}(|Y|+a'),$$

We finally conclude

$$\begin{aligned}\sum_{F\,\in\,\mathcal{F}(G')\setminus\mathcal{IF}(G)}(6-\mathrm{sides}(F)) &= 6|\mathcal{F}(G')\setminus\mathcal{IF}(G)| - (|Y|+a')\\ &\ge 6q_{2-\gamma}(|Y|+a') - (|Y|+a') = Q_{2-\gamma}(|Y|+a').\end{aligned}$$

□

We can use this Lemma to improve the initial algorithm. We define $\mathcal{C}_0$ as the empty set and then $\mathcal{C}_{k+1}$ is defined recursively as the union of the cycle of length $k+1$ (provided $k+1 \leq 7$) and the set of all nice children of the planar maps in $\mathcal{C}_k$ satisfying the constraints (1), (2), (3), (4). Because of Lemma 8.16, we conclude that a twice-derived net with k edges belongs necessarily to $\mathcal{C}_k$.

However, this new modified algorithm is not yet the one used in our implementation. Without further considerations the families $\mathcal{C}_k$ would contain multiple copies of the same spherical graph, since we are not accounting for isomorphisms of the maps. In theory we could still run the algorithm as described and only deal with the multiple counting when checking which elements of $\mathcal{C}_k$ are twice-derived spherical graphs. However this would be computationally too costly. In our actual algorithm we therefore avoid these multiple counts in every class $\mathcal{C}_k$. How this is done will be explained in Section E.

8.5. **Reconstruction of once-derived candidates.** We describe here how, given the list of Proposition 8.2, we produce the list of Proposition 8.5. The goal is just to give the appropriate rules to "add back" the vertices of degree 2 which were removed, recalling that in a once-derived marked spherical graph there are at most four and at least two degree 2 vertices.

(1) We fix a graph $G = (V, E)$ and its spherical embedding ι, coming from the list of Proposition 8.2.
(2) We fix the number $2 \leq n \leq 4$ of degree two vertices which we wish to add.
(3) We fix an edge $\sigma \in E$ and let A and B be the two vertices connected by it. We then remove σ, add a new vertex C, add two edges (adjacent to A and C and to B and C, respectively), and mark one of them. At the same time we produce an embedding $\iota_{(1)}$ of this new graph $G_{(1)}$ from the embedding ι of G in the obvious way, namely "splitting" the curve connecting $\iota(A)$ and $\iota(B)$ in two pieces.
(4) We repeat step (3) the necessary number of times to get to $G_{(n)}$ and $\iota_{(n)}$.
(5) We check whether $G_{(n)}$ and $\iota_{(n)}$ satisfy the conditions stated in Proposition 5.6 and in Section 7.

The procedure above is repeated over all possible choices of: elements in the list of Proposition 8.2, natural numbers n between 2 and 4, edges selected in step (3), and markings decided in the same step. The algorithm described above can be easily expressed in terms of combinatorial maps and we leave the task to the reader. As is the case for the algorithm of Section 8, we would incur into multiple counting of the same marked spherical graph, if we were to ignore the isomorphic copies which might occur in the list. This point is addressed as described in Section E.

9. Useful inequalities for convex faces

We prove here some elementary inequalities for the sides of convex triangles, quadrilaterals, and, more importantly, pentagons with angles equal to $\frac{2\pi}{3}$, which are derived from Proposition 6.8. These will be the backbone of the algorithm explained in Section 13, which will allow us to effectively search for approximate metrizations as once derived nets of the marked spherical graphs listed in Proposition 8.5. The inequalities can be thought as quantitative versions of the monotonicity considerations of Heppes in [14].

Lemma 9.1. *Let $ABCD$ be a convex spherical quadrilateral with angles of $\frac{2\pi}{3}$. Then*

$$|AB| + |BC| \geq 2 \arctan 2\sqrt{2} \approx 2.46\ldots .$$

Lemma 9.2. *Let $ABCDE$ be a convex spherical pentagon with angles of $\frac{2\pi}{3}$. Then*

$$|AB| + |BC| \geq 1\,.$$

Proposition 9.3. *Let $ABCDE$ be a convex spherical pentagon with angles of $\frac{2\pi}{3}$. Denote by λ_T the length of the side of the equilateral triangle with angles $\frac{2\pi}{3}$ and by f the function*

$$f(x,y) := \cos x + \cos y + \cos x \cos y - 2 \sin x \sin y \tag{9.1}$$

(cf. Proposition 6.8). The following holds for any choice of the numbers x, y, z in the range $(0, \lambda_T)$:

- *If $|AB| \geq x$ and $|BC| \geq y$, then*
$$2\cos(|DE|) - \frac{1}{3} \leq f(x,y)\,.$$
- *If $|AB| \leq x$ and $|BC| \leq y$, then*
$$2\cos(|DE|) - \frac{1}{3} \geq f(x,y)\,.$$
- *If $|AB| \leq x$ and $|DE| \geq z$, then*
$$2\cos z - \frac{1}{3} \geq f(x, |BC|)\,.$$
- *If $|AB| \geq x$ and $|DE| \leq z$, then*
$$2\cos z - \frac{1}{3} \leq f(x, |BC|)\,.$$

9.1. **Proof of Lemma 9.1.** Let $a = |AB|/2$ and $b = |BC|/2$. We have

$$\tan a \cdot \tan b = \frac{1}{2}\,,$$

with $0 < a, b < \frac{\pi}{2}$. Using the tangent addition formula, we get

$$\tan(a+b) = \frac{\tan a + \tan b}{1 - \tan a \cdot \tan b} = 2(\tan a + \tan b) \geq 4\sqrt{\tan a \cdot \tan b} = 2\sqrt{2}\,.$$

Therefore,

$$|AB| + |BC| = 2(a+b) \geq 2 \arctan 2\sqrt{2}\,.$$

9.2. **Proof of Lemma 9.2.** Since $\sqrt{2} \geq 1$, we can assume that $|AB|, |BC| \leq \sqrt{2}$.
For $0 \leq x \leq \sqrt{2}$, we have

$$0 \leq \sin x \leq x \quad \text{ and } \cos x \geq 1 - \frac{x^2}{2} \geq 0\,.$$

Therefore, we have

$$2 \geq 2\cos|DE| - \frac{1}{3} = f(|AB|, |BC|) \geq 3 - |AB|^2 - |BC|^2 - 2|AB| \cdot |BC| = 3 - (|AB| + |BC|)^2\,,$$

which is equivalent to the statement.

9.3. **Proof of Proposition 9.3.** The proposition follows immediately from the formula (6.3) and the following elementary monotonicity property of f.

Lemma 9.4. *The function f of (6.4) is decreasing with respect to both its variables in the square $(0,\lambda_T)^2$.*

Proof. Observe that $\frac{\pi}{2} < \lambda_T < \pi$. Let us compute the derivative of f with respect to x:

$$\partial_x f(x,y) = -\sin x(1+\cos y) - 2\cos x \sin y = -\sin x \sin y\Big(\frac{1+\cos y}{\sin y} + \frac{2\cos x}{\sin x}\Big)\,.$$

Both $\theta \mapsto \frac{1+\cos\theta}{\sin\theta}$ and $\theta \mapsto \frac{2\cos\theta}{\sin\theta}$ are decreasing functions for $\theta \in (0,\pi)$. Thus, for $x, y \in (0,\lambda_T)$, we deduce

$$\partial_x f(x,y) \le -\sin x \sin y\Big(\frac{1+\cos\lambda_T}{\sin\lambda_T} + \frac{2\cos\lambda_T}{\sin\lambda_T}\Big) = -\sin x \sin y \frac{1+3\cos\lambda_T}{\sin\lambda_T} = 0\,.$$

Hence, f is decreasing with respect to x and, by symmetry, also with respect to y. □

10. Relaxed Almgren competitors

Several arguments in the rest of the paper will take advantage of the minimizing assumption of E in Theorem 1.3 by exhibiting suitable competitors. However the simpler "ideal comparison surfaces" will typically not comply with the constraints required in Definition 1.2 and we will need to argue that they can be suitably approximated with competitors which instead comply with them. This section proves three very useful general propositions which will allow us to relax considerably the requirements of Definition 1.2. The first one shows that it suffices to define the map Φ in Definition 1.2 on the cone E.

Proposition 10.1. *Assume that $\ell \subset \mathbb{R}^3$ is a line through the origin and $\ell \subset E \subset \mathbb{R}^3$ is the cone over a link $\mathcal{N} \subset \mathbb{S}^2$ which is a connected geodesic net (cf. Definition 4.1). Let $H : E \to \mathbb{R}^3$ be a Lipschitz map such that:*

(i) $H(x) = x$ outside some ball B_{r_0};
(ii) If $W := \{x \in E : H(x) \neq x\}$, then $\overline{W \cup H(W)} \subset \mathbb{R}^3 \setminus \ell$.

Then there is an extension G of H to the whole Euclidean space $\mathbb{R}^3$ which complies with the requirements of Definition 1.2.

The next propositions will give substantial more leeway, because we will allow the set $\overline{W \cup \Phi(W)}$ in Definition 1.2 to "touch" ℓ in a controlled way. For the first proposition we introduce the following terminology.

Definition 10.2. Given $\theta \in (0, \frac{\pi}{2})$ we denote by $C_\theta(\ell)$ the closed "ice-cream" cone

$$C_\theta(\ell) := \{x \in \mathbb{R}^3 : \operatorname{dist}(x,\ell) \le |x| \sin\theta\}\,.$$

We will also use the notation $B_\delta(A)$ for the open set $\{x : \operatorname{dist}(x,A) < \delta\}$, namely the open δ-neighborhood of a set A.

Proposition 10.3 (Ice-cream relaxed condition)**.** *Assume $\ell \subset \mathbb{R}^3$ is a line through the origin and $\ell \subset E \subset \mathbb{R}^3$ the cone over a connected geodesic net $\mathcal{N} \subset \mathbb{S}^2$. Let $F : E \to \mathbb{R}^3$ be a Lipschitz map such that:*

(i) $F(x) = x$ outside some ball B_{r_0};
(ii) $F(x) = x$ for $x \in C_\theta(\ell)$ and $F^{-1}(\ell) \subset \ell$.

Then, for every $\varepsilon > 0$, there is a Lipschitz map $G : \mathbb{R}^3 \to \mathbb{R}^3$ such that:

(a) $G(x) = x$ *outside* B_{r_0}*;*
(b) upon setting $W := \{x : G(x) \neq x\}$*, we have* $\overline{W \cup G(W)} \subset \mathbb{R}^3 \setminus \ell$*;*
(c) $\mathcal{H}^2(G(E \cap \overline{B}_{r_0})) \leq \mathcal{H}^2(F(E \cap \overline{B}_{r_0})) + \varepsilon$*.*

The next one will require slightly stronger conditions in a structurally different cone.

Definition 10.4. We will call *wedge* any open convex cone of $\mathbb{R}^3$ which is invariant under translations along the direction of ℓ.

Note that such a wedge V can be written as $V' + \ell$, where V' is an open convex cone in the plane orthogonal to ℓ. Furthermore observe that V' is either the whole plane or an open convex subset of the plane bounded by two half-lines emanating from the origin.

Proposition 10.5 (Wedge relaxed condition)**.** *Assume that* $\ell \subset \mathbb{R}^3$ *is a line through the origin,* $\ell \subset E \subset \mathbb{R}^3$ *the cone over a connected geodesic net* $\mathcal{N} \subset \mathbb{S}^2$*. Let* $F : E \to \mathbb{R}^3$ *be a Lipschitz map such that:*

(i) $F(x) = x$ *outside some ball* B_{r_0}*;*
(ii) upon setting $U := \{x : F(x) \neq x\}$*, we have* $\overline{U \cup F(U)} \subset \mathbb{R}^3 \setminus (V \cap B_\delta(\ell))$ *for some wedge* V *and some* $\delta > 0$*;*

Then, for every $\varepsilon > 0$*, there is a Lipschitz map* $G : \mathbb{R}^3 \to \mathbb{R}^3$ *such that:*

(a) $G(x) = x$ *outside* B_{r_0}*;*
(b) upon setting $W := \{x : G(x) \neq x\}$*, we have* $\overline{W \cup G(W)} \subset \mathbb{R}^3 \setminus \ell$*;*
(c) $\mathcal{H}^2(G(E \cap \overline{B}_{r_0})) \leq \mathcal{H}^2(F(E \cap \overline{B}_{r_0})) + \varepsilon$*.*

10.1. **Proof of Proposition 10.1.** We will need the following technical lemma, whose proof can be achieved with a regularization procedure.

Lemma 10.6. *Let* $\Omega \subset \mathbb{R}^k$ *be a bounded open set and let* $H : \overline{\Omega} \to \mathbb{R}^N$ *be a continuous map whose restriction to* $\partial\Omega$ *is Lipschitz. Then, for every* $\eta > 0$*, there is a Lipschitz map* $G : \overline{\Omega} \to \mathbb{R}^N$ *such that* $G = H$ *on* $\partial\Omega$ *and* $\|G - H\|_{C^0} \leq \eta$*.*

Proof. Let $\tilde{H} : \overline{\Omega} \to \mathbb{R}^N$ be a Lipschitz function that coincides with H on $\partial\Omega$. Let $A : \mathbb{R}^k \to \mathbb{R}^N$ be defined as $A = H - \tilde{H}$ in Ω and $A = 0$ in $\mathbb{R}^k \setminus \Omega$. The function A is continuous and compactly supported. Fix a standard smooth convolution kernel φ, and, for any $\varepsilon > 0$, let $A_\varepsilon := A * \varphi_\varepsilon$. Observe that $A_\varepsilon \to A$ uniformly as $\varepsilon \to 0$ and furthermore A_ε is Lipschitz. Define $G_\varepsilon : \overline{\Omega} \to \mathbb{R}^N$ as

$$G_\varepsilon(x) = \tilde{H}(x) + \min\left\{1, \tfrac{\operatorname{dist}(x,\partial\Omega)}{\varepsilon}\right\} A_\varepsilon(x)\,.$$

For any $\varepsilon > 0$, the function G_ε is Lipschitz and $G_\varepsilon(x) = \tilde{H}(x) = H(x)$ for $x \in \partial\Omega$. It remains only to bound uniformly $|G_\varepsilon(x) - H(x)|$ for $x \in \overline{\Omega}$. If $\operatorname{dist}(x, \partial\Omega) \geq \varepsilon$, then

$$|H(x) - G_\varepsilon(x)| = |A(x) - A_\varepsilon(x)| \leq \|A - A_\varepsilon\|_{C^0}\,,$$

while if $\operatorname{dist}(x, \partial\Omega) < \varepsilon$, then

$$|H(x) - G_\varepsilon(x)| \leq |A(x)| + |A_\varepsilon(x)| \leq \|A - A_\varepsilon\|_{C^0} + 2 \sup_{x:\, \operatorname{dist}(x,\partial\Omega) < \varepsilon} |A(x)|\,.$$

Since $A|_{\partial\Omega} \equiv 0$ and A is uniformly continuous, the last two inequalities imply that, if we choose ε small enough compared to η, the map G_ε satisfies the requirements of the statement. □

Proof of Proposition 10.1. First of all, without loss of generality we will assume that $r_0 = 1$. Next extend H on the complement of B_1 simply by setting it equal to the identity. The main task becomes therefore to further extend to $\overline{B}_1 \setminus E$. Let

$$W' := \{x \in E : x \neq H(x)\}$$

and choose δ so that $2\delta < \theta$ and

$$2\delta < \min\{\text{dist}\,(y, \ell) : y \in W' \cup H(W')\}\,.$$

Impose further that 2δ is strictly smaller than the minimum distance between any pair of vertices of the net $\mathcal{N}$.

We then extend H as the identity on the closed set

$$K := (\mathbb{R}^3 \setminus B_{1-2\delta}) \cup \overline{B}_{2\delta}(0) \cup C_{2\delta}(\ell)\,.$$

This extension is clearly Lipschitz on $K \cup E$. Now, let U be the open set

$$U := B_{1-2\delta} \setminus (K \cup E)\,.$$

Denote by P and $-P$ the two pins of the net $\mathcal{N}$. Each connected component Ω of U corresponds to a connected component Λ of $V := \mathbb{S}^2 \setminus (\mathcal{N} \cap \overline{B}_{2\delta}(P) \cup \overline{B}_{2\delta}(-P))$ through the relation

$$\Omega = \{tx : x \in \Lambda, 2\delta < t < 1 - 2\delta\}\,.$$

In particular, since each Λ is homeomorphic to the open disk of $\mathbb{R}^2$, each Ω is homeomorphic to the open 3-ball. We can then extend the map H *continuously* in each connected component Ω so that the values of the latter extension over each Ω stay at distance 2δ from ℓ: this follows from the fact that the second homotopy group of $\mathbb{R}^3 \setminus \overline{B}_{2\delta}(\ell)$ is trivial. We denote by H' such extension.

We can now use Lemma 10.6 to perturb H' in Ω to a new Lipschitz map H'' close to H' in the C^0 norm and coinciding with H on $\partial\Omega$: the map G is defined to be H on $\mathbb{R}^3 \setminus \Omega$ and H'' on $\overline{\Omega}$. G is then Lipschitz and the closure of the set $W := \{x : G(x) \neq x\}$ is contained in $B_1 \setminus \ell$. If we impose the inequality $\|H - H''\|_{C^0} \leq \delta$ then we also ensure $G(W \cap \Omega) \subset \mathbb{R}^3 \setminus B_\delta(\ell)$. Since however we know that $G(W \setminus \Omega) \subset\subset \mathbb{R}^3 \setminus \ell$, we conclude that $G(W) \subset\subset \mathbb{R}^3 \setminus \ell$. Therefore G satisfies all the requirements of the Proposition. □

10.2. **Proof of Proposition 10.3.** Since we can invoke Proposition 10.1, we really need to define G on E. Moreover, without loss of generality we will assume $r_0 = 1$.

Step 1. Definition of G on E. First of all, since E is a cone and F is the identity on $\partial B_1 \cap E = \mathcal{N}$, we know that F on $E \setminus \overline{B}_1$ as a 1-homogeneous map, i.e. we have

$$F(x) = |x| F\left(\tfrac{x}{|x|}\right) \qquad \forall x \in E \setminus \overline{B}_1\,.$$

We then fix $\eta > 0$, whose choice will be specified later, and define the following Lipschitz function λ on the positive real line:

$$\lambda(t) = \begin{cases} \frac{1}{\eta} & \text{if } 0 < t \leq \eta \\ \exp(-\eta^{-1}(2\eta - t)\log\eta) & \text{if } \eta < t < 2\eta \\ 1 & \text{if } t \geq 2\eta \end{cases}\,.$$

Note that $1 \leq \lambda \leq \frac{1}{\eta}$.

At this point it is notationally convenient to use $d(x)$ for the dist (x, ℓ). We then introduce

$$\Phi(x) := \lambda(d(x))x$$

and finally define explicitly the map G for $x \in E$, by setting

$$G(x) = \begin{cases} x & \text{if } d(x) < \eta \\ \\ \frac{|x|}{|\Phi(x)|} F(\Phi(x)) & \text{otherwise} \end{cases} .$$

We show below that this map satisfies all the requirements of the proposition.

Step 2. Topological conditions. In this step we check that, upon choosing η small enough, G complies with the requirements of Definition 1.2. The fact that (when $d(x) \neq 0$) $x = \frac{|x|}{|\Phi(x)|}\Phi(x)$ will be used repeatedly. By construction, the map G is the identity both when $d(x) < \eta$ and when $|x| \geq 1$. Moreover, it is continuous and locally Lipschitz. Hence, it is globally Lipschitz. We also have that the set $W := \{x : x \neq G(x)\}$ has positive distance from ℓ and is contained in B_1.

Observe moreover that $d(\Phi(x)) = \lambda(d(x))d(x) \geq d(x)$ for every x. In particular $\Phi(W)$ has positive distance from ℓ. Given that $F^{-1}(\ell) \subset \ell$, we immediately conclude that $\overline{F(\Phi(W))} \subset F(\Phi(\overline{W}))$ is a compact subset of $\mathbb{R}^3 \setminus \ell$. Observe next that, since $\lambda \leq \frac{1}{\eta}$,

$$\inf\{d(G(x)) : x \in W\} \geq \eta \inf\{d(F(\Phi(x))) : x \in W\} > 0 .$$

Next, since W is bounded and G Lipschitz, $G(W)$ is bounded. Summarizing, we have proved that G is Lipschitz, that $\overline{W} \subset \overline{B}_1 \setminus \ell$, and that $G(W) \subset\subset \mathbb{R}^3 \setminus \ell$.

Step 3. Area estimate. In order to complete the proof we need to prove estimate (C) for $\mathcal{H}^2(G(E \cap \overline{B}_1))$. First of all we know that $G(x) = F(x)$ over $E \cap \overline{B}_1 \setminus B_{2\eta}(\ell)$. In particular $G(E \cap \overline{B}_1 \setminus B_{2\eta}(\ell))$ is contained in the set $F(E \cap \overline{B}_1)$.Thus we can write

$$\mathcal{H}^2(G(E \cap \overline{B}_1)) \leq \mathcal{H}^2(G(E \cap \overline{B}_1 \cap B_{2\eta}(\ell))) + \mathcal{H}^2(F(E \cap \overline{B}_1)) .$$

Observe that $\Phi(C_\theta(\ell)) \subset C_\theta(\ell)$. In particular G is the identity on $C_\theta(\ell) \cap E$ and thus

$$\mathcal{H}^2(G(C_\theta(\ell) \cap E \cap \overline{B}_1 \cap B_{2\eta}(\ell))) = \mathcal{H}^2(C_\theta(\ell) \cap E \cap \overline{B}_1 \cap B_{2\eta}(\ell)) \leq C_0\eta$$

for some constant C_0 which depends on E but not on η. Hence we can write

$$\mathcal{H}^2(G(E \cap \overline{B}_1)) \leq \mathcal{H}^2(F(E \cap \overline{B}_1)) + C_0\eta + \mathcal{H}^2(G((E \setminus C_\theta(\ell)) \cap \overline{B}_1 \cap B_{2\eta}(\ell))) . \tag{10.1}$$

We now observe further that

$$(E \setminus C_\theta(\ell)) \cap \overline{B}_1 \cap B_{2\eta}(\ell) \subset E \cap B_{3\theta^{-1}\eta}$$

and, in order to estimate $\mathcal{H}^2(G(E \cap B_{3\theta^{-1}\eta}))$, we introduce the notation $J_\sigma(A)$ for the Jacobian of the restriction of a linear map $A : \mathbb{R}^3 \to \mathbb{R}^3$ to a 2-dimensional plane σ and define

$$J = \sup\{J_{T_x E}(dG|_x) : x \in E \cap \overline{B}_{3\theta^{-1}\eta}\} .$$

In particular, using the area formula, we can estimate

$$\mathcal{H}^2(G((E \setminus C_\theta(\ell)) \cap \overline{B}_1 \cap B_{2\eta}(\ell))) \leq \mathcal{H}^2(G(E \cap \overline{B}_{3\theta^{-1}\eta})) \leq C_1 J \eta^2 , \tag{10.2}$$

where C_1 is a constant which depends on the cone E and on θ.

We recall that, given any orthonormal basis $\{e_1, e_2\}$ of the 2-dimensional plane σ, we have

$$J_\sigma(A) = |A[e_1] \times A[e_2]| \leq |A[e_1]||A[e_2]| . \tag{10.3}$$

Evidently in the region where G is the identity, for every 2-plane σ we have that $J_\sigma(dG|_x) = 1$. We are thus interested in the region where $G(x)$ equals

$$\frac{|x|}{|\Phi(x)|} F(\Phi(x)) = (\lambda(d(x)))^{-1} F(\lambda(d(x))x)\,.$$

There we can compute

$$dG_x = dF_{\Phi(x)} + \frac{\lambda'(d(x))}{\lambda(d(x))}\left(-\frac{F(\Phi(x))}{\lambda(d(x))} + dF_{\Phi(x)}[x]\right) \otimes \nabla d(x)\,.$$

We recall that $d(x)$ is a 1-Lipschitz function and it is differentiable on the complement of ℓ, with $|\nabla d(x)| = 1$. Given an arbitrary $x \in B_{3\theta^{-1}\eta} \setminus \ell \subset \overline{B}_1 \setminus \ell$ (so we assume that $\eta \leq \frac{\theta}{3}$) and an arbitrary 2-dimensional plane, we can choose an orthonormal basis e_1, e_2 with the property that $e_1 \perp \nabla d(x)$. In particular (10.3) gives the estimate

$$J_\sigma(dG_x) \leq |dF_{\Phi(x)}[e_1]|\left(|dF_{\Phi(x)}[e_2]| + \frac{|\lambda'(d(x))|}{\lambda(d(x))}(|G(x)| + |dF_{\Phi(x)}[x]|)\right)$$

for the specific 2-dimensional plane σ. Recall that F is Lipschitz, that $|G(x)| \leq \|F\|_{C^0(B_1)}$ for $x \in \overline{B}_1$, and that

$$\left|\frac{\lambda'}{\lambda}\right| = |(\log(\lambda))'| \leq \frac{|\log \eta|}{\eta}\,.$$

We thus achieve the estimate $J(dG_x) \leq C_2 \eta^{-1} |\log \eta|$, where the constant C_2 depends on $\mathrm{Lip}\,(F)$, $\|F\|_{C^0(B_1)}$, and θ but not on the parameter η. We can now insert the latter bound in (10.2) and (10.1) to conclude

$$\mathcal{H}^2(G(E \cap \overline{B}_1)) \leq \mathcal{H}^2(F(E \cap \overline{B}_1)) + C_0 \eta (1 + |\log \eta|)\,.$$

While the constant C_0 depends upon $\mathrm{Lip}\,(F)$, E, and θ, it does not depend on η. Clearly if we choose η small enough so that $C_0\eta(1 + |\log \eta|) < \varepsilon$ we achieve the estimate (c) in the statement of the proposition, thus completing the proof.

10.3. **Proof of Proposition 10.5.** Without loss of generality we can assume that $r_0 = 1$ and (since otherwise F itself already satisfies the assumption of Proposition 10.1) that V is not the entire space $\mathbb{R}^3$.

We next fix coordinates so that $\ell = \{(0, 0, t) : t \in \mathbb{R}\}$. In particular $V = V' \times \mathbb{R}$, where V' is an open convex cone of $\mathbb{R}^2$ bounded by two half-lines emanating from the origin. We then fix a Lipschitz deformation retraction $r : [0, 1] \times \mathbb{R}^2 \setminus (B_1 \cap V') \to \mathbb{R}^2 \setminus B_1$ with the additional roperty that $r(t, x) \neq 0$ for every $t > 0$ and define a corresponding Lipschitz retraction $\rho : [0, 1] \times \mathbb{R}^3 \setminus (B_1(\ell) \cap V) \to \mathbb{R}^3 \setminus B_1(\ell)$ as

$$\rho(t, (x', x_3)) = (r(t, x'), x_3)\,.$$

Finally, for every positive η we define further the retraction $\rho_\eta : [0, 1] \times \mathbb{R}^3 \setminus (B_\eta(\ell) \cap V) \to \mathbb{R}^3 \setminus B_\eta(\ell)$ as

$$\rho_\eta(t, x) = \eta\rho\left(t, \tfrac{x}{\eta}\right)\,.$$

Next, for every $x, y \in \mathbb{R}^3 \setminus V$ we let $\gamma_{x,y} \subset \mathbb{R}^3 \setminus V$ be the shortest path connecting x and y in $\mathbb{R}^3 \setminus V$. The path exists because $\mathbb{R}^3 \setminus V$ is closed and arc-connected, but in fact we will show at the end that it is unique and has a simple description. This allows us to define the map $[0, 1] \ni t \mapsto \Gamma(t, x, y)$ as the constant speed parametrization of $\gamma_{x,y}$ and in the last step of the proof we will also show that Γ is locally Lipschitz in all variables.

Assuming the latter properties of Γ, we fix a positive number $\eta < \frac{\delta}{2}$ (whose choice will be specified later) and denote by $d(x)$ the distance of x to ℓ. We first define the map over E and it is convenient to introduce the following partition of $\mathbb{R}^3$ into three sets which are clearly invariant under translations along ℓ (cf. Figure 10):

$$\begin{aligned}\Omega_1 &:= \overline{B}_\eta(\ell) \cup (\overline{B}_{2\eta}(\ell) \cap V)\\ \Omega_2 &:= \overline{B}_{2\eta(\ell)} \setminus (\overline{B}_\eta(\ell) \cup V)\\ \Omega_3 &:= \mathbb{R}^3 \setminus \overline{B}_{2\eta}(\ell)\,.\end{aligned}$$

We then set

$$G(x) := \begin{cases} x & \text{if } x \in E \cap \Omega_1\\ \rho_\eta\left(\frac{d(x)-\eta}{\eta}, \Gamma\left(\frac{d(x)-\eta}{\eta}, x, F(x)\right)\right) & \text{if } x \in E \cap \Omega_2\\ \rho_\eta(1, F(x)) & \text{if } x \in E \cap \Omega_3\end{cases}$$

The role of the function $x \mapsto \frac{d(x)-\eta}{\eta}$ is to interpolate between the value 0 on $\partial B_\eta(\ell)$ and 1 on $\partial B_{2\eta}(\ell)$: any other interpolation with Lipschitz constant comparable to η^{-1} would be suitable.

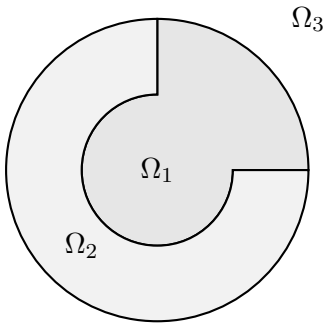


FIGURE 10. A cross section, orthogonal to ℓ, of the partition of $\mathbb{R}^3$ into the regions Ω_i. In this picture the wedge V is given by $\{(x_1, x_2, x_3) : x_1, x_2 \geq 0\}$.

Observe next that, since $F(x) = x$ on $E \setminus B_1$ and thus $G(x) = \rho_\eta(1, x)$ on $E \cap \Omega_3 \setminus B_1$ and $G(x) = \rho_\eta\left(\frac{d(x)-\eta}{\eta}, x\right)$ on $E \cap \Omega_2 \setminus B_1$. Since however $\rho_\eta(t, x) = x$ for ever $x \in \mathbb{R}^3 \setminus B_\eta(\ell)$ and $\Omega_2 \cup \Omega_3 \subset \mathbb{R}^3 \setminus B_\eta(\ell)$, we immediately conclude that $G(x) = x$ on $E \setminus B_1$.

Note next that by construction $W := \{x : G(x) \neq x\}$ is contained in the compact set $K := \overline{B}_1 \setminus B_\eta(\ell)$. Next, in order to check that $G(W) \subset\subset \mathbb{R}^3 \setminus \ell$ it suffices to show that $G(K) \cap \ell = \emptyset$. By definition of ρ, $\rho_\eta(t, x)$ is an element of ℓ if and only if $t = 0$ and $x \in \ell$. It is therefore obvious that $G(K) \cap \ell = \emptyset$. Thus we can apply Proposition 10.1 and extend G to a map on $\mathbb{R}^3$ which complies with the requirements (a) and (b). The rest of the proof is devoted to prove that, upon choosing η sufficiently small, the map satisfies (c) as well.

Step 1. *G is Lipschitz.* We first check the Lipschitz continuity of G. Clearly G is Lipschitz in each of the three sets Ω_i and in fact $G|_{\Omega_i}$ extends to a unique Lipschitz map G_i on $\overline{\Omega}_i$ simply using each of the three explicit formulae. In order to check that it is globally Lipschitz it suffices to show that the map is continuous on $\overline{\Omega}_i \cap \overline{\Omega}_j$ for every $i \neq j$, which is equivalent to showing that $G_i = G_j$ on the latter set. This is easily checked as follows:

- $\overline{\Omega}_1 \cup \overline{\Omega}_2$ consists of the two pieces $\partial B_\eta(\ell) \setminus V$ and $\partial V \cap \overline{B}_{2\eta}(\ell) \setminus B_\eta(\ell)$. On the first we have
$$G_2(x) = \rho_\eta(0, \Gamma(0, x, F(x))) = \rho_\eta(0, x) = x = G_1(x)\,,$$
while on the second we have $F(x) = x$ and $\rho_\eta(t, x) = x$, and therefore
$$G_2(x) = \rho_\eta(\tfrac{d(x)-\eta}{\eta}, \Gamma(\tfrac{d(x)-\eta}{\eta}, x, x)) = \rho_\eta(\tfrac{d(x)-\eta}{\eta}, x) = x = G_1(x)\,;$$
- $\overline{\Omega}_1 \cap \overline{\Omega}_3 = \partial B_{2\eta}(\ell) \cap \overline{V}$. On this set we have $F(x) = x$ and $\rho_\eta(1, x) = x$ and therefore $G_3(x) = x = G_1(x)$.
- $\overline{\Omega}_2 \cap \overline{\Omega}_3 = \partial B_{2\eta}(\ell) \setminus V$. On this set we have
$$G_2(x) = \rho_\eta(1, \Gamma(1, x, F(x)) = \rho_\eta(1, F(x)) = G_3(x)\,.$$

Step 2. Lipschitz constant. We next claim that the Lipschitz constant of G is bounded by a constant L independent of η. In order to see this it suffices to show that $|DG| \leq C$ at all points of differentiability in $B_1 \cap E$. This is simply achieved using the chain rule and observing that

- $|D_x\Gamma(t, x, y)| + |D_y\Gamma(t, x, y)| \leq C_0(|x| + |y|)$;
- $|\partial_t\Gamma(x, y)| \leq C_0|x - y|$;
- $|D_x\rho(t, x)| \leq C_0$;
- $|\partial_t\rho(t, x)| \leq C_0\eta$.

The last two are in fact obvious conclusions, while the first two estimates will be proved in the last step. The bound on the derivative becomes then obvious on the regions Ω_1 and Ω_3.

The bound on Ω_2 is more subtle. First define $\Lambda(x) = \Gamma\left(\frac{d(x)-\eta}{\eta}, x, F(x)\right)$ and observe that
$$\begin{aligned} D_x\Lambda(x) =& \partial_t\Gamma\left(\tfrac{d(x)-\eta}{\eta}, x, F(x)\right) \otimes \tfrac{\nabla d(x)}{\eta} + D_x\Gamma\left(\tfrac{d(x)-\eta}{\eta}, x, F(x)\right) \\ &+ D_y\Gamma\left(\tfrac{d(x)-\eta}{\eta}, x, F(x)\right) D_xF(x)\,. \end{aligned}$$
We can tame the only "dangerous" term, namely the first term, using
$$\left|\partial_t\Gamma\left(\tfrac{d(x)-\eta}{\eta}, x, F(x)\right) \otimes \tfrac{\nabla d(x)}{\eta}\right| \leq \tfrac{C_0}{\eta}|x - F(x)|$$
and observing that, since the above chain rule is applied when $d(x) \leq 2\eta$, the assumption that $F(x) = x$ on ℓ implies $|x - F(x)| \leq d(x)(1 + \operatorname{Lip}(F)) \leq 2\eta(1 + \operatorname{Lip}(F))$. This shows then $|D_x\Lambda| \leq C$. Next we compute
$$D_xG(x) = \partial_t\rho_\eta\left(\tfrac{d(x)-\eta}{\eta}, \Lambda(x)\right) \otimes \tfrac{\nabla d(x)}{\eta} + D_x\rho_\eta\left(\tfrac{d(x)-\eta}{\eta}, \Lambda(x)\right) D_x\Lambda\,.$$
This time the dangerous term is tamed using $|\partial_t\rho| \leq C_0\eta$.

Step 3. Area estimate. Note that the maps G and F coincide on $\overline{B}_1 \setminus (B_{2\eta}(\ell) \cap F^{-1}(B_\eta(\ell))$ and that $G(x) = \rho_\eta(1, F(x))$ for $x \in \overline{B}_1 \setminus B_{2\eta}(\ell)$ and in particular $G = F$ on the set $\overline{B}_1 \setminus (B_{2\eta}(\ell) \cup F^{-1}(B_\eta(\ell)))$. Thus we can estimate
$$\begin{aligned} \mathcal{H}^2(G(E \cap \overline{B}_1 \setminus B_{2\eta}(\ell))) &\leq \mathcal{H}^2(E \cap \overline{B}_1) + \mathcal{H}^2(\rho_\eta(1, B_\eta(\ell) \cap F(E \cap \overline{B}_1))) \\ &\leq \mathcal{H}^2(E \cap \overline{B}_1) + C_0^2\mathcal{H}^2(F(E \cap \overline{B}_1) \cap B_\eta(\ell))\,. \end{aligned}$$
Now, $F(E \cap \overline{B}_1)$ is a compact set with finite Hausdorff measure and therefore
$$\lim_{\eta \downarrow 0} \mathcal{H}^2(F(E \cap \overline{B}_1) \cap B_\eta(\ell)) = 0\,.$$

Given the above estimates, it is obvious that

$$\mathcal{H}^2(G(E\cap \overline{B}_1)) \leq \mathcal{H}^2(F(E\cap \overline{B}_1)) + \varepsilon$$

as soon as η is sufficiently small.

Step 4. Lipschitz continuity of Γ. Consider $x = (x', t)$ and $y = (y', s)$ in $\mathbb{R}^3 \setminus V$. If the segment joining x and y is contained in $\mathbb{R}^3$, then this is the (unique) path $\gamma_{x,y}$ of least length. This happens if and only if one of the following happens:

- x' is the origin;
- $x' \neq 0$ and if we let $\alpha^+ \subset \mathbb{R}^2$ be the open half-space containing V' and bounded by the line $\{\lambda x' : \lambda \in \mathbb{R}\}$, then $y' \notin \alpha^+$.

When $x' \neq 0$ and $y' \in \alpha^+$, upon setting $L := |x'| + |y'|$ and

$$z := \left(0, \tfrac{|x'|}{L} t + \tfrac{|y'|}{L} s\right),$$

the unique shortest path $\gamma_{x,y}$ is the union of the two segments $[x, z]$ and $[z, y]$.

Given this it is easy to check that Γ is locally Lipschitz and that $|D_x\Gamma(t, x, y)| \leq |x|$. It is also easy to check that the length of $\gamma_{x,y}$ is bounded by $C_0|x - y|$, where C_0 depends only upon the opening angle of the cone V', thus giving $|\partial_t\Gamma(t, x, y)| \leq C_0|x - y|$.

11. Excluding basic degenerations

The primary goal of this section is to give effective lower bounds on the length of the edges of a once-derived net: while so far we have mostly leveraged the stationarity of the net, the assumption that it comes from the link of an area-minimizing cone as in Theorem 1.3 will play here a crucial role. These bounds will be of great help in the computer-assisted search for approximate metrizations of the candidate once-derived topologies (cf. Section 13). We will split the lower bound into three cases, summarized in the following theorem. In order to simplify our statements we will use the notation $\mathrm{len}(\sigma)$ for the length of a geodesic arc σ in $\mathbb{S}^2$.

Proposition 11.1. *Consider the link $\tilde{\mathcal{N}} = E \cap \partial B_1$ of an area-minimizing cone as in Theorem 1.3, let $\mathcal{N}$ be the corresponding once-derived net and σ be an edge of $\mathcal{N}$.*

(i) If σ is the side of a convex face of $\mathcal{N}$ or a convex face of $\tilde{\mathcal{N}}$, then $\mathit{len}(\sigma) \geq \frac{1}{100}$.

(ii) If $\mathit{len}(\sigma) < \frac{1}{100}$ then there is an edge τ of $\mathcal{N}$ adjacent to σ such that

$$\mathit{len}(\sigma) \geq \tfrac{1}{10}\mathit{len}(\tau)\,.$$

(iii) Let σ and τ be the two edges adjacent to a vertex of degree 2. If both edges are shorter than $\frac{1}{100}$ then there is a third distinct edge ς, adjacent to either σ or τ, such that

$$\mathit{len}(\sigma) + \mathit{len}(\tau) \geq \tfrac{1}{10}\mathit{len}(\varsigma)\,. \tag{11.1}$$

We include in this section also another type of lower bound, which will require a more sophisticated analysis to produce a rather tight bound, but which nonetheless is achieved with similar methods. Assume $\tilde{\mathcal{N}}$ is the link of a minimizing cone and P one of the pins. Assume the degree of the pin is strictly less than 3. Let F be the nonconvex face of $\mathbb{S}^2 \setminus \tilde{\mathcal{N}}$ whose boundary contains P. A compactness argument shows that P gets arbitrarily close to an edge of ∂F which is *not* adjacent to it, then it in fact gets arbitrarily close to some other vertex. A more precise and effective version of the latter statement is given below in

Proposition 11.2 for pins of degree 1. This proposition will be used in Section 18, when we describe how to discard many stationary links on energetic grounds.

Proposition 11.2. *Let E be as in Theorem 1.3, let $\tilde{\mathcal{N}} = E \cap \partial B_1$ be the link of E and $P \in \ell \cap \partial B_1$ a pin of degree 1 contained in the edge σ. Let $\mathcal{N}$ be the corresponding once-derived net and F the face of $\mathbb{S}^2 \setminus \mathcal{N}$ which contains the pin. Assume that the distance of P to the closest vertex of $\tilde{\mathcal{N}}$ is at least 0.22 and, moving from P on the equator $\gamma \supset \sigma$ away from the edge σ, let Q be the first intersection with ∂F. Assume that the smallest angle formed by γ with any edge containing Q is at least $\frac{\pi}{4}$. Then the distance from P to Q is at least 0.01.*

11.1. **Minimizing criterion.** The crucial ingredient to prove the effective bounds in Proposition 11.1 and Proposition 11.2 is the minimizing criterion of Corollary 11.4, which is shown exhibiting appropriate competitors. We recall that $\#A$ denotes the cardinality of the set A. The basic construction is detailed separately in the following lemma because it will be rather useful at a later stage.

Lemma 11.3. *Consider a cone E as in Theorem 1.3 and let $\tilde{\mathcal{N}} = E \cap \mathbb{S}^2$ be its link. Let $\rho : [0,1] \times \tilde{\mathcal{N}} \to \mathbb{S}^2$ be a Lipschitz homotopy and let Ω be the smallest closed set such that*

$$\{x : \rho(t,x) \neq x\} \cup \{\rho(t,x) : \rho(t,x) \neq x\} \subset \Omega \quad \text{for every } t.$$

Assume that

(i) $\rho(0,x) = x$ for all $x \in \tilde{\mathcal{N}}$;

(ii) Either there is a wedge V as in Definition 10.4 such that $\Omega \subset \mathbb{S}^2 \setminus V$ or $\Omega \subset \mathbb{S}^2 \setminus \ell$.

Then, for every $\delta > 0$ there is $\varepsilon < \delta$ and a Lipschitz map $\Phi : E \cap \overline{B}_1 \to \overline{B}_1$ with the following properties:

(a) $\Phi(x) = x$ for $x \in \tilde{\mathcal{N}}$;

(b) $\Phi(E \cap \overline{B}_{1-\varepsilon}) = E' \cap \overline{B}_{1-\varepsilon}$, where E' is the cone with base $\rho(1, \tilde{\mathcal{N}})$ and vertex 0;

(c) The following inequality holds

$$\mathcal{H}^2(\Phi(E \cap \overline{B}_1 \setminus B_{1-\varepsilon})) \leq \int_{\mathbb{S}^2} \#\rho^{-1}(\{x\})\, d\mathcal{H}^2(x) + \delta\,; \tag{11.2}$$

(d) Φ satisfies the assumptions of Proposition 10.3 in case (i) and the assumptions of Proposition 10.5 in case (ii).

Proof. We let $\varepsilon > 0$ be any positive real number and define Φ_ε on $E \cap \overline{B}_1$ as follows:

- $\Phi_\varepsilon(0) = 0$;
- $\Phi_\varepsilon(x) = |x| \rho(1, \frac{x}{|x|})$ for every $x \in E$ with $0 < |x| < 1 - \varepsilon$;
- $\Phi_\varepsilon(x) = |x| \rho(\frac{1-|x|}{\varepsilon}, \frac{x}{|x|})$ for every $x \in E$ with $1 - \varepsilon \leq |x| \leq 1$.

This map clearly is Lipschitz and satisfies (a), (b) and (d). We therefore only need to show that

$$\limsup_{\varepsilon \to 0} \mathcal{H}^2(\Phi_\varepsilon(E) \cap (\overline{B}_1 \setminus B_{1-\varepsilon})) \leq \int_{\mathbb{S}^2} \#\rho^{-1}(\{x\})\, d\mathcal{H}^2(x)\,. \tag{11.3}$$

Let $A_\varepsilon = E \cap (\overline{B}_1 \setminus B_{1-\varepsilon})$. Parameterize $x \in A_\varepsilon$ by $(r, \omega) \in [1-\varepsilon, 1] \times \tilde{\mathcal{N}}$ with $x = r\omega$ and $d\mathcal{H}^2 = r\, dr\, d\mathcal{H}^1(\omega)$. On A_ε, the map is $\Phi_\varepsilon(r, \omega) = r\rho(\frac{1-r}{\varepsilon}, \omega)$.

By the area inequality for Lipschitz maps (and the change of variable $t = \frac{1-r}{\varepsilon}$), we bound the area by integrating the Jacobian $J\Phi_\varepsilon$:

$$\mathcal{H}^2(\Phi_\varepsilon(A_\varepsilon)) \leq \int_{\tilde{\mathcal{N}}} \int_{1-\varepsilon}^1 J\Phi_\varepsilon(r,\omega)\, r\, dr\, d\mathcal{H}^1(\omega) = o(1) + \int_{\tilde{\mathcal{N}}} \int_0^1 J\rho(t,\omega)\, dt\, dt\mathcal{H}^1(\omega)\,.$$

Passing to the limit and applying the area formula to ρ yields Equation (11.3). □

Corollary 11.4. *Consider a cone E as in Theorem 1.3 and let $\tilde{\mathcal{N}} = E \cap \mathbb{S}^2$. Let ρ be as in Lemma 11.3. Then,*

$$\mathcal{H}^1(\tilde{\mathcal{N}}) - \mathcal{H}^1(\rho(1,\tilde{\mathcal{N}})) \leq 2 \int_{\mathbb{S}^2} \#\rho^{-1}(\{x\})\, d\mathcal{H}^2(x)\,. \tag{11.4}$$

Proof. Fix any δ and consider the map Φ_ε and the number $\varepsilon = \varepsilon(\delta) < \delta$ given by Lemma 11.3. Since we can apply either Proposition 10.3 or Proposition 10.5, we conclude that

$$\tfrac{1}{2}\mathcal{H}^1(\tilde{\mathcal{N}}) = \mathcal{H}^2(E \cap \overline{B}_1) \leq \mathcal{H}^2(\Phi(E \cap \overline{B}_1)\,. \tag{11.5}$$

On the other hand, by (11.2) we get

$$\begin{aligned}\mathcal{H}^2(\Phi(E \cap \overline{B}_1)) &= \mathcal{H}^2(\Phi(E) \cap B_{1-\varepsilon}) + \mathcal{H}^2(\Phi(E) \cap \overline{B}_1 \setminus \overline{B}_{1-\varepsilon}) \\ &\leq \tfrac{(1-\varepsilon)^2}{2}\mathcal{H}^1(\rho(1,\tilde{\mathcal{N}})) + \int_{\mathbb{S}^2} \#\rho^{-1}(\{x\})\, d\mathcal{H}^2(x) + \delta\,.\end{aligned}$$

Combining the latter inequality with (11.5) we then get

$$\mathcal{H}^1(\tilde{\mathcal{N}}) - (1-\varepsilon)^2\mathcal{H}^1(\rho(1,\tilde{\mathcal{N}})) \leq 2 \int_{\mathbb{S}^2} \#\rho^{-1}(\{x\})\, d\mathcal{H}^2(x) + 2\delta\,.$$

Recalling that $\varepsilon < \delta$ and letting $\delta \downarrow 0$ we achieve (11.4) □

11.2. **Suitable homotopies.** From the minimizing criterion in Corollary 11.4 and through a suitable choice of appropriate homotopies we will gain the following two "lower level estimates", from which in turn we will get Proposition 11.1.

Proposition 11.5. *Consider the link $\tilde{\mathcal{N}} = E \cap \partial B_1$ of an area-minimizing cone E as in Theorem 1.3, let α, β, γ be three edges of $\tilde{\mathcal{N}}$ which form a path in this order, and denote by A and B the vertices of β in common, respectively, with α and γ. Assume that:*

- *(i) α, β, and γ are in the boundary of a face F of $\mathbb{S}^2 \setminus \tilde{\mathcal{N}}$;*
- *(ii) the internal angles of F at A and B measure $\frac{2\pi}{3}$;*
- *(iii) let H be the open hemisphere with $\beta \subset \partial H$ containing a neighborhood of β in F; either F is convex or the eventual pin[5] contained in H is at distance at least $\frac{\pi}{4}$ from A.*

Then

$$\mathit{len}(\beta) \geq \tfrac{1}{10}\min\{\tfrac{1}{10}, \mathit{len}(\alpha), \mathit{len}(\gamma)\}.$$

Proposition 11.6. *Consider the link $\tilde{\mathcal{N}} = E \cap \partial B_1$ of an area-minimizing cone E as in Theorem 1.3, let $\alpha, \beta, \gamma, \delta$ be four edges of $\tilde{\mathcal{N}}$ which form a path in this order, and denote by A, B, and C the vertices in common between, respectively, α and β, β and γ, and γ and δ. Assume that:*

- *(i) α, β, γ, and δ are in the boundary of a face F of $\mathbb{S}^2 \setminus \tilde{\mathcal{N}}$;*
- *(ii) the internal angles of F at A, B, and C measure $\frac{2\pi}{3}$.*
- *(iii) denote by H_1 and H_2 the open hemispheres with $\beta \subset \partial H_1$, $\gamma \subset \partial H_2$ and containing, respectively a neighborhood of β and a neighborhood of γ in F; either F is convex or the eventual pin contained in $H_1 \cap H_2$ is at distance at least $\frac{\pi}{4}$ from B.*

[5] H might contain no pin because both pins might lie in its boundary: in this case the assumption is considered as verified.

Then

$$len(\beta) + len(\gamma) \geq \tfrac{1}{10} \min\{\tfrac{1}{5}, len(\alpha), len(\delta)\}\,.$$

Proof of Proposition 11.5. Assume by contradiction that the edge β is strictly shorter than $\frac{1}{100}$, $\frac{\mathrm{len}(\alpha)}{10}$, and $\frac{\mathrm{len}(\gamma)}{10}$. We then mark on α and γ the two points A_2 and B_2 with the property that the lengths of the two segments A_2A and B_2B are $10\mathrm{len}(\beta)$. Furthermore we consider the isosceles triangles AA_2D and BB_2C with angles $\frac{\pi}{6}$ at A, A_2, B and B_2 with the property that the edges AD and BC are orthogonal to AB, cf. Figure 11. Finally we connect C and D with a special arc η which is not geodesic but which rather consists of the points which are equidistant to the edge AB.

We now describe an homotopy ρ to which we wish to apply Corollary 11.4. The homotopy keeps all points of the net $\tilde{\mathcal{N}}$ fixed except for those in the open arcs AA_2 and BB_2. The segment AA_2 is deformed continuously onto the union of AD and DA_2; while BB_2 is deformed onto the union of BA, AD, η, CB_2. The homotopy is defined in such a way that all the points in the interior of the triangles AA_2C, BB_2C, and the region Q bounded by $DA \cup AB \cup BC \cup \eta$ have exactly one preimage, while the points which do not belong to their closures or to $\tilde{\mathcal{N}}$ have no preimages. We also draw the shortest geodesic segment joining A_2 and B_2 and introduce the quadrilateral Ω, which contains $AA_2D \cup BB_2C \cup Q$. We claim that

(a) the following inequality holds

$$|AA_2|+|BB_2| > |A_2D|+|DA|+\mathrm{len}(\eta)+|B_2C|+2(\mathcal{H}^2(AA_2D)+\mathcal{H}^2(BB_2C)+\mathcal{H}^2(Q))\,. \quad (11.6)$$

The claims (a) will be proved in Section F.1.

We need however to check that the homotopy satisfies condition (ii) in Lemma 11.3 in order to apply the Corollary 11.4. If Ω does not contain any pin, then clearly condition (ii) in Lemma 11.3 is satisfied and the proof follows after observing that

$$\mathcal{H}^1(\tilde{\mathcal{N}}) - \mathcal{H}^1(\rho(1,\tilde{\mathcal{N}})) = |AA_2| + |BB_2| - (|A_2D| + |DA| + \mathrm{len}(\eta) + |B_2C|)$$

while

$$\int_{\mathbb{S}^2} \#(\rho^{-1}(x))\, d\mathcal{H}^2(x) = (\mathcal{H}^2(AA_2D) + \mathcal{H}^2(BB_2C) + \mathcal{H}^2(Q))2\mathcal{H}^2(AA_2QM)\,,$$

in particular (11.6) would be at odds with (11.4).

We next show that if Ω contains a pin P, then it is either A or B and thus it follows immediately that condition (ii) in Lemma 11.3 holds. In order to prove that $P \in \{A, B\}$ note first that, let F be the face in the statement of the proposition. If F is convex, the interior ω of the segment A_2B_2 (which is shorter than π and hence length-minimizing) is contained in the interior of F. Therefore also the interior of Ω is contained in the interior of F. Since the interior of F does not contain any pin, then $P \in \partial\Omega \setminus \omega \subset \tilde{\mathcal{N}}$. But P is a vertex of $\tilde{\mathcal{N}}$ and the only two vertices of the net contained in $\partial\Omega \setminus \omega$ are A and B. If F is not convex, consider the open hemisphere H of condition (iii) in the statement. Then $\Omega \setminus \beta \subset H$ and, since it can be easily checked that Ω has diameter strictly smaller than $\frac{\pi}{4}$, we immediately infer $P \in \beta$. As above, since β is an edge of $\tilde{\mathcal{N}}$, P must be either A or B. □

Proof of Proposition 11.6. We argue by contradiction and assume that

$$\mathrm{len}(\beta) + \mathrm{len}(\gamma) < \min\{\tfrac{1}{50}, \tfrac{\mathrm{len}(\alpha)}{10}, \tfrac{\mathrm{len}(\delta)}{10}\}\,.$$

Without loss of generality, we also assume $\mathrm{len}(\beta) \leq \mathrm{len}(\gamma)$. We then draw the equator λ which bisects the internal angle of F at the vertex B. Let M be the point on λ, in the direction of

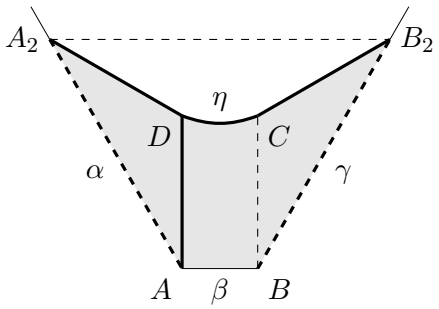


FIGURE 11. The homotopy used to prove Proposition 11.5. All the straight segments drawn in the pictures are length-minimizing geodesics, while the arc η is made of points equidistant from the edge β.

the face F, at distance $5(\mathrm{len}(\beta)+\mathrm{len}(\gamma))$ from B. Then, let $\bar\lambda$ be the equator perpendicular to λ at M. In Section F.1, we will use elementary geometric arguments to show that:

(a) $\bar\lambda$ intersects the edge α at an interior point E and the edge δ at an interior point D;
(b) The pentagon $ABCDE$ is convex;
(c) $|CD| > |DE| + 2\mathcal{H}^2(ABCDE)$;

cf, Figure 12.

Now, we describe the homotopy ρ to be used in Corollary 11.4. First of all ρ moves the points in the the interior of the arc CD, while the rest of the net $\tilde{\mathcal{N}}$ is kept fixed. The arc CD is mapped by $\rho(1,\cdot)$ onto the union of the arcs γ, β, AE, and DE. Moreover the homotopy is defined in such a way that each point in the open pentagon $ABCDE$ has exactly one preimage and the points outside the union of its closure and $\tilde{\mathcal{N}}$ have no preimages. Arguing as in the proof of Proposition 11.5 we identify the closed pentagon $ABCDE$ as the region Ω which satisfies condition (ii) of Lemma 11.3. In particular we claim that, if a pin P belongs to Ω, then $P \in \{A, B, C\}$. In order to show this, we consider first the case in which the face F is convex. Then, since Ω is convex, its interior and the interior ω of the segment DE are both contained in the interior of F. So $P \in \partial\Omega \setminus \omega \subset \tilde{\mathcal{N}}$ and, because P is a vertex of $\tilde{\mathcal{N}}$ and the only vertices of $\tilde{\mathcal{N}}$ contained in $\partial\Omega \setminus \omega$ are A, B, and C, we must necessarily have $P \in \{A, B, C\}$. If F is not convex, consider however that, again by convexity of the interior of Ω, $\Omega \setminus (\beta \cup \gamma) \subset H_1 \cap H_2$, where H_1 and H_2 are the hemispheres of assumption (iii) in the proposition. But because we easily see that the diameter of Ω is strictly smaller than $\frac{\pi}{4}$, this implies that $P \in \beta \cup \gamma$. Again, since P is a vertex of the net $\tilde{\mathcal{N}}$ and the only vertices contained in $\beta \cup \gamma$ are A, B, and C, we infer that P is one of them.

Clearly

$$\mathcal{H}^1(\tilde{\mathcal{N}}) - \mathcal{H}^1(\rho(1,\tilde{\mathcal{N}})) = |CD| - |ED|$$

while

$$\int_{\mathbb{S}^2} \#(\rho^{-1}(x))\, d\mathcal{H}^2(x) = \mathcal{H}^2(ABCDE)\,.$$

But then (c) implies that the inequality (11.4) is violated by this explicit ρ. □

11.3. **Proof of Proposition 11.1. Proof of (i).** Fix the convex face F of the assumption. If the face is a biangle or a triangle, the lengths are rigid and the statement holds by Proposition 6.8. Now, assume the face F is either a convex quadrilateral or a convex pentagon.

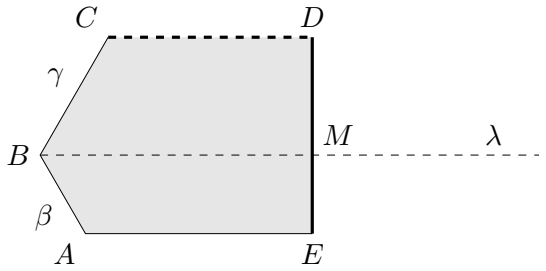


FIGURE 12. The competitor used to prove Proposition 11.6. The thick dashed segment is removed from the net and replaced with the thick segment. The gray area is the region spanned by the homotopy from the original net to the competitor.

First of all observe that, if F is a face of $\mathcal{N}$, then (since it is convex), it is also a face of $\tilde{\mathcal{N}}$. Let a be the length of σ and b, c be the lengths of the adjacent sides of F. Thanks to Lemmas 9.1 and 9.2, we have $a+b, a+c \geq 1$. So, either $a \geq 0.01$ or $b, c \geq 0.99$. In the first case, we have proven the statement. Otherwise, we can apply Proposition 11.5 and deduce $a \geq 0.1 \min\{0.1, b, c\} \geq 0.01$.

Proof of (ii). We can assume that

$$\operatorname{len}(\sigma) < 0.01\,, \tag{11.7}$$

Let A, B be the two vertices of σ. In $\tilde{\mathcal{N}}$, both A and B have degree 3. None of them can be a pin, because otherwise it would be a pin of degree 3 and hence σ would be the edge of a convex face in $\mathbb{S}^2 \setminus \tilde{\mathcal{N}}$: in this case we could apply statement (i) of the Proposition, which is at odds with (11.7). In particular:

(p) if an edge τ of $\tilde{\mathcal{N}}$ adjacent to A or B does not belong to $\mathcal{N}$, then the other endpoint of τ is a pin.

Let H be an open hemisphere bounded the equator containing σ which either does not contain a pin or such that the only pin P in H is at distance at least $\frac{\pi}{4}$ from A. Let α be the edge adjacent to A and γ the edge adjacent to B which enter into H. Observe that if α, resp. γ, does not belong to $\mathcal{N}$, then, by (p), its length is at least $\frac{\pi}{4} - 0.01$. Applying Proposition 11.5 we obtain

$$\operatorname{len}(\sigma) \geq \frac{1}{10} \min\{\tfrac{1}{10}, \operatorname{len}(\alpha), \operatorname{len}(\gamma)\}\,.$$

Since $\operatorname{len}(\sigma) < 0.01$, it must hold $\min\{\operatorname{len}(\alpha), \operatorname{len}(\gamma)\} < \frac{1}{10}$. Without loss of generality, we assume that α is the shortest between α and β. So $\operatorname{len}(\sigma) \geq \frac{1}{10}\operatorname{len}(\alpha)$ and $\operatorname{len}(\alpha) < \frac{1}{10}$. Thus, as observed above, the edge α belongs to $\mathcal{N}$ and this concludes the proof.

Proof of (iii). We can, without loss of generality, assume

$$\max\{\operatorname{len}(\sigma), \operatorname{len}(\tau)\} < 0.01\,. \tag{11.8}$$

Let us label as B the only vertex in common between σ and τ. Observe that:

- the other endpoint A of σ must be distinct from the other endpoint C of τ, otherwise both σ and τ would be long π;

- Neither A, nor B, nor C can be a pin, because it would be a pin of degree 3 and thus either σ or τ would belong to a convex face of $\mathbb{S}^2 \setminus \tilde{\mathcal{N}}$, which, by point (i) of the Proposition, is at odds with (11.8).

Recall next that B has degree 2 in $\mathcal{N}$. Thus there is a third edge σ_B adjacent to B in $\tilde{\mathcal{N}}$, which has been removed because it connects B to a pin P. Clearly $\mathrm{len}(\sigma_B) < \pi$, otherwise σ_B would contain P and its antipodal P', which we know is the other pin of the net.

We next label with σ_A and σ'_A the other two edges of $\tilde{\mathcal{N}}$ adjacent to A and with σ_C and σ'_C the other two edges adjacent to C, cf. Figure 13. The labels are placed so that there are three faces G_A, G_C and G, of $\mathbb{S}^2 \setminus \tilde{\mathcal{N}}$ such that:

$$\sigma_A \cup \sigma \cup \sigma_B \subset \partial G_A$$
$$\sigma_B \cup \tau \cap \sigma_C \subset \partial G_C$$
$$\sigma'_A \cup \sigma \cap \tau \cup \sigma'_C \subset \partial G\,,$$

cf. again Figure 13. Note finally that:

(p') If σ'_A, resp. σ'_C, does not belong to $\mathcal{N}$, then it connects A, resp. C, to the pin P'.

To complete the proof we analyze two possibilities.

Alternative 1: $\mathbf{len}(\sigma_B) \geq \frac{\pi}{4}$. At least one between σ_A and σ_C must belong to $\mathcal{N}$, otherwise $\mathcal{N}$ would have a face with three concave angles. Without loss of generality, assume σ_A is an edge of $\mathcal{N}$. We can apply Proposition 11.5 to the three edges σ_B, σ, σ_A and infer

$$0.01 > \mathrm{len}(\sigma) \geq \tfrac{1}{10} \min\{\tfrac{1}{10}, \mathrm{len}(\sigma_A), \mathrm{len}(\sigma_B)\}\,.$$

In particular, $\min\{\frac{1}{10}, \mathrm{len}(\sigma_A), \mathrm{len}(\sigma_B)\} < \frac{1}{10}$, and, since $\mathrm{len}(\sigma_B) \geq \frac{\pi}{4}$, we conclude

$$\min\{\tfrac{1}{10}, \mathrm{len}(\sigma_A), \mathrm{len}(\sigma_B)\} = \mathrm{len}(\sigma_A)\,.$$

Hence (11.1) holds if we set $\varsigma = \sigma_A$.

Alternative 2. $\mathbf{len}(\sigma_B) < \frac{\pi}{4}$. In this case we have $\mathrm{dist}\,(B, P') > \frac{3\pi}{4}$. Let now γ_A be the equator containing σ and H_1 the hemisphere bounded by γ_A in which the segment σ'_A enters. Clearly P belongs to the other hemisphere bounded by γ_A, in particular its antipodal P' is in H_1. If we denote by γ_C the equator containing τ and by H_2 the hemisphere bounded by γ_C in which σ'_C enters, we conclude likewise that $P' \in H_2$. We then see that we are in the position to apply Proposition 11.6, where the hemispheres of condition (iii) therein are precisely the ones just identified. We then conclude

$$0.02 > \mathrm{len}(\sigma) + \mathrm{len}(\tau) \geq \tfrac{1}{10}\{\tfrac{1}{5}, \mathrm{len}(\sigma'_A), \mathrm{len}(\sigma'_C)\}\,.$$

If we let ς be the shortest segment between σ'_A and σ'_C, we then see that it is necessarily shorter than $\frac{1}{5}$, which means that it must belong to $\mathcal{N}$. In fact (p') implies that, if it were removed, then it would connect either A or C with the pin P', and hence it would be longer than $\frac{3\pi}{4} - 0.01$. Hence we conclude the validity of (11.1) in this case as well.

11.4. **Proof of Proposition 11.2.** As in the cases analyzed in Proposition 11.1, we will argue by contradiction showing the existence of an admissible competitor with lower area. In order to do so we start by fixing a system of coordinates (labeled (x, y, z) as usual) in which the point Q is the "north pole" $Q = (0, 0, 1)$ and the equator containing the edge σ is the plane $\{y = 0\}$. We then intersect E with the plane $\wp = \{z = 1\}$: this intersection is a planar net, i.e. a closed set consisting, locally, of a finite number of straight segments connecting a finite number of vertices (the set is not compact, but this plays no role in our argument,

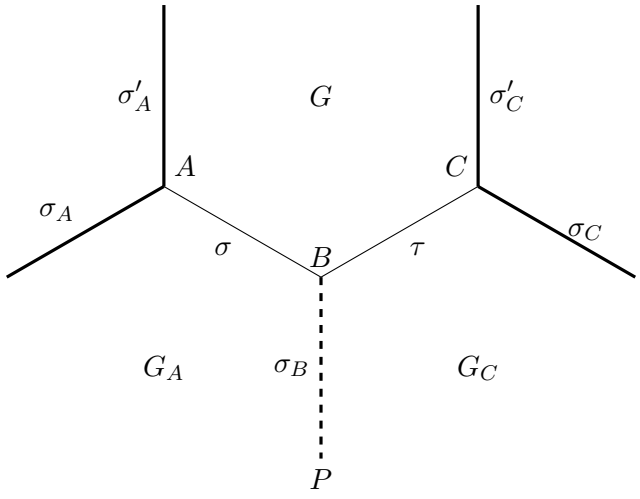


FIGURE 13. The edges and faces used in the proof of Proposition 11.1(iii). The edges σ and τ belong to the once-derived net $\mathcal{N}$. The dashed edge σ_B belongs to $\tilde{\mathcal{N}}$ but certainly does not belong to $\mathcal{N}$. All the other (thick but not dashed) edges belong to $\tilde{\mathcal{N}}$ and might or might not be removed by the procedure to derive $\mathcal{N}$.

which is local around the point Q). In fact we can establish a natural map from $E \cap \wp$ to $\tilde{\mathcal{N}}$ by mapping x into $\frac{x}{|x|}$: this maps vertices of $E \cap \wp$ into vertices of $\tilde{\mathcal{N}}$ and, locally, edges into edges. The map has Lipschitz constant 1 if we endow $\wp$ with the Euclidean distance and $\mathbb{S}^2$ with the geodesic distance. The intersection of ℓ with $\wp$ consists then of the one point, which we label by M, such that $P = \frac{M}{|M|}$. If we let $\bar{\delta}$ be the distance (on the sphere) between the pin P and Q and we assume it is shorter than 0.01, then the length δ of the segment MQ is $\tan \bar{\delta}$ at we therefore conclude that it is smaller than 0.011.

We next consider the plane containing the equator γ and let $\mathfrak{s}$ be its intersection with $\wp$ and denote by $\hbar$ the intersection of the plane $\{y = 0\}$ with $\wp$. Of course $\mathfrak{s} \cap \hbar = \{Q\}$ and we let α be the (smallest) angle formed by $\mathfrak{s}$ and $\hbar$. Note that this is the angle formed by the equator γ and the edge σ at Q, therefore $\frac{\pi}{4} \leq \alpha \leq \frac{\pi}{2}$. In particular, after possibly applying a reflection along the plane $\{x = 0\}$, we conclude that the point M has coordinates

$$M = (\delta \cos \alpha, \delta \sin \alpha, 1)\,.$$

Consider now any vertex V in $E \cap \wp$: having established that the map $V \mapsto \frac{V}{|V|}$ is 1-Lipschitz and that it maps vertices into vertices, we conclude that the segment AM is at least 0.22, in particular the segment AQ is at least $0.22 - 0.011 = 0.209 > 0.2 =: L$. We then mark as A, B, and C the points

$$A = (L \cos \alpha, L \sin \alpha, 1) \tag{11.9}$$

$$B = (L, 0, 1) \tag{11.10}$$

$$C = (-L, 0, 1)\,. \tag{11.11}$$

We remark that:

- The segments AM and BC are contained in $E \cap \wp$;

- Indeed $\overline{B}_{0.2}(M) \cap E \cap \wp = AM \cup BC$, cf. Figure 14.

Consider now the cone

$$\mathcal{C} := \{rx : x \in \overline{B}_{0.2}(M) \cap \wp \quad \text{and} \quad 0 \leq r \leq 1\}\,.$$

The set $\Sigma_0 := E \cap \mathcal{C}$ consists of the two triangles BOC and MOA, where O denotes the origin: this is the cone with vertex 0 over the set $\overline{B}_{0.2}(M) \cap E \cap \wp = AM \cup BC$. We construct a competitor Σ which consists of 10 suitable triangles contained in $\mathcal{C}$. We will show that:

(a) $\mathcal{H}^2(\Sigma_0) - \mathcal{H}^2(\Sigma) > 0$;
(b) while clearly $\mathcal{E} := (E \setminus \mathcal{C}) \cup \Sigma$ will not be an admissible competitor for E in the sense of Definition 1.2, there is a Lipschitz map $F : E \to \mathbb{R}^3$ which maps E into $\mathcal{E}$ and satisfies the assumptions of Proposition 10.5.

Therefore (a) and (b) contradict the minimizing property of E and hence will complete the proof.

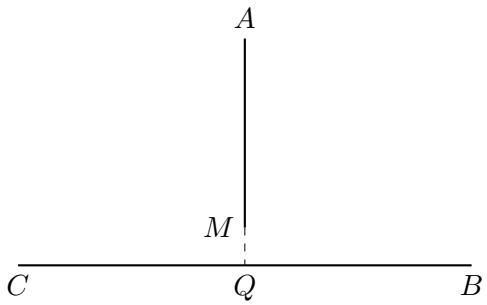


FIGURE 14. The set $E \cap \wp$ in a neighborhood of $M = \ell \cap \wp$. In this case $\alpha = \frac{\pi}{2}$.

Description of Σ. We start by describing the triangles forming Σ. In order to do so we need to introduce:

- The Torricelli point D of A, B, and Q. This is the unique point D of the plane $\wp$ with the property that AD, BD, and QD meet at D forming equal angles;
- The points $A' = rA$, $B' = rB$, $C' = rC$, and $D' = rD$, where $r = 0.8$.

The surface Σ consists then of:

- The triangles BCM and CMO;
- The cone with basis the net $A'D' \cup B'D' \cup M'D'$, which is formed by the three triangles $A'OD'$, $B'OD'$, and $M'OD'$;
- The "transition" surface formed by the five triangles $AA'D'$, AMD', $MM'D'$, BMD', and $BB'D'$.

The geometric idea behind the competitor is that we would like to substitute Σ_0 in $\mathcal{C}$ with a new cone Σ' with vertex 0 and with base given in the Figure 15. But transitioning between the two cones directly on the plane $\wp$ would be energetically costly. We therefore insert rather a scaled copy of the cone $r\Sigma'$ ($r = 0.8$ is the scaling factor) and interpolate at the intermediate scales. We can afford one horizontal interpolation in the triangle BCM because it is very small, the remaining interpolating surface takes advantage also of the direction orthogonal to $\wp$.

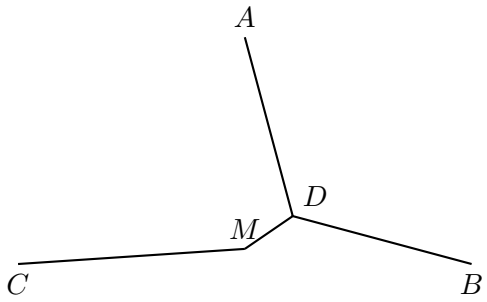


FIGURE 15. This picture shows the basis of the cone Σ' which would be an ideal competitor for Σ_0. Σ contains a scaled copy $r\Sigma'$ of the latter, while the remaining part of Σ interpolates between $\Sigma_0 \cap \wp$ and the basis of $r\Sigma'$.

Proof of (a). We note that the quantity $\mathcal{H}^2(\Sigma_0) - \mathcal{H}^2(\Sigma)$ depends in fact on x and α (indeed it can even be computed through an explicit formula). This gives a function $(\delta, \alpha) \mapsto f(\delta, \alpha)$, defined on the domain $(0, 0.011] \cap [\frac{\pi}{4}, \frac{\pi}{2}]$. We will show in Section F.1 that f is Lipschitz, but in fact it can be immediately seen, either geometrically or from the explicit formula, that it can be extended to the domain $\mathcal{D} = [0, 0.011] \times [\frac{\pi}{4}, \frac{\pi}{2}]$. (a) amounts to showing that the function f is positive on the domain. The Lipschitz constant of f will be estimated explicitly in section Section F.1. We will then partition $\mathcal{D}$ in small boxes using a grid fine enough so to guarantee that, if the value of f is above an appropriate threshold η at the center of a box of the partition, then the function is in fact positive on the whole $\mathcal{D}$. A computer program will then evaluate the function at all the centers and validate the bound.

Proof of (b). It remains to show the existence of the map F. First of all F is defined to be the identity on $E \setminus \Sigma_0$. The objective is thus to define $F : \Sigma_0 \to \Sigma$ so that F is the identity on the boundaries of the two triangles COB and AOM forming Σ_0. F is constructed as the composition of two maps $F = F_2 \circ F_1$ and from the construction it will be clear that it satisfies the assumptions of Proposition 10.5. Both maps are in fact explicit maps from a simplicial complex to another simplicial complex.

F_1 is the identity on AOM and maps the triangle COB onto the union of the triangles $CMO \cup BMO \cup CMB$. This is easily done as follows: we pick a point $\bar{M}$ interior to the triangle COB and then prescribe $F_1(C) = C$, $F_1(B) = B$, $F_1(O) = O$, and $F_1(\bar{M}) = M$. F_1 is then defined to be the unique affine extension in each of the three triangles $C\bar{M}O$, $B\bar{M}O$, and $CB\bar{M}$, cf. Figure 16.

$\Sigma_1 = F_1(\Sigma_0)$ is therefore the union of four triangles, namely the "horizontal" triangle BCM and the three triangles CMO, BMO and AMO. We next define F_2 on Σ_1. F_2 is the identity on the horizontal triangle and on the triangle CMO. It thus has to map the two triangles BMO and AMO on the rest of Σ, which we denote by Σ_2. Recalling that D is the Torricelli point of A, B, and M and that $A' = rA$, $B' = rB$, $M' = rM$, and $D' = rD$, Σ_2 is the union of

- The cone with basis the net $A'D' \cup B'D' \cup M'D'$, which is formed by the three triangles $A'OD'$, $B'OD'$, and $M'OD'$;
- The "transition" surface formed by the five triangles $AA'D'$, AMD', $MM'D'$, BMD', and $BB'D'$.

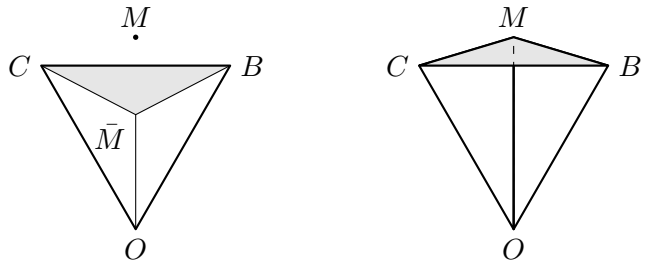


FIGURE 16. The map F_1 is piecewise affine and maps the triangle CBO in the union of the three triangles CMB, CMO, and MBO.

Note that F_2 is required to be the identity on the boundaries of both triangles BOM and AOM, namely on the four segments MO, AO, BO, BM, and AM.

F_2 is in fact defined in very similar fashion to F_1. On the triangle BMO we mark the points B' and M' and the midpoint D_B of the segment $B'M'$. We then prescribe that F_2 is the identity on all the five points O, M, B, M', and B', while $F_2(D_B) = D'$. Joining the five fixed points with D_B we subdivide BMO into five triangles: F_2 is then he unique affine extension on each of these. The procedure on AMO mirrors the one for BMO. More precisely we mark the point A' and let D_A be the midpoint of the segment $A'M'$. We then prescribe $F_2(A) = A$, $F_2(A') = A'$, and $F_2(D_B) = D'$, while of course we already have fixed $F_2(M') = M'$, $F_2(M) = M$, and $F_2(O) = O$. Even in this case we subdivide AMO into the corresponding five triangles and define F_2 to be the unique affine extension in each of them.

12. PATH-ISOMETRIES

In this section, we describe an algorithm that assigns an orientation-preserving isometry of the sphere to each polygonal path. We then provide a neat formula for determining this isometry and we infer an estimate on the distance between the isometries assigned to two different paths, assuming they share the same internal angles but differ in edge lengths. These tools are central in the computer-assisted search for approximate metrizations of candidate once-derived marked spherical graphs, as detailed in Section 13.

12.1. **Path-isometry Formulas.** As usual, orientation-preserving isometries of the sphere will be called *special isometries* and will be denoted (as a group) by $\mathrm{SO}(3)$. $T\mathbb{S}^2$ will denote the tangent bundle of the sphere $\mathbb{S}^2$. It will be useful for us to identify $T\mathbb{S}^2$ with the appropriate subset of pairs (e, w) of vectors in $\mathbb{R}^3 \times \mathbb{R}^3$, as specified in the following definition.

Definition 12.1 (Unit tangent bundle)**.** The tangent bundle $T\mathbb{S}^2$ is the set of pairs $(e, w) \in \mathbb{R}^3 \times \mathbb{R}^3$ such that

$$|e| = 1 \tag{12.1}$$

$$w \cdot e = 0\,. \tag{12.2}$$

Finally, $T^1\mathbb{S}^2$ is the subset of $T\mathbb{S}^2$ consisting of tangent vectors of length 1, namely the pairs (e_1, e_2) which satisfy (12.1), (12.2), and $|e_2| = 1$.

We further introduce two families of elements of $\mathrm{SO}(3)$ related to $T^1\mathbb{S}^2$

Definition 12.2 (Exponentiation and rotation)**.** If $v = (e_1, e_2) \in T^1\mathbb{S}^2$, we let $e_3 = e_1 \times e_2$. We then denote

- by $E(v,t)$ the counterclockwise rotation of angle t around the axis directed by e_3;
- by $R(v,\alpha)$ the counterclockwise rotation of angle α around the axis directed by e_1.

Note that $R(v,\alpha)$ depends only upon the base point e_1 of the element v. We are now ready to define what we will often call *path-isometry*.

Definition 12.3 (Path-isometry)**.** Consider a polygonal path consisting of vertices $V_0,\dots,V_N$ connected by geodesic arcs γ_i (starting at V_{i-1} and ending at V_i), parametrized over $[0,\lambda_i]$ by arclength (note that the arcs are not necessarily length-minimizing). The associated *path-isometry* is the unique special isometry that maps the initial point V_0 into the last point V_N and the initial unit tangent $\dot\gamma_1(0)$ into the last unit tangent $\dot\gamma_N(\lambda_N)$.

Because of the definition of $E(v,t)$ and $R(v,\alpha)$, the following formula for the path-isometry is almost tautological.

Lemma 12.4. *Following the same notation for a path as in Definition 12.3, assume that the counterclockwise angle at V_i from $-\dot\gamma_i(\lambda_i)$ to $\dot\gamma_{i+1}(0)$ is α_i. The* path-isometry *associated to this path is*

$$E(W_{N-1},\lambda_N)\circ R(W_{N-1},\alpha_{N-1}-\pi)\circ E(W_{N-2},\lambda_{N-1})\circ\ldots\circ R(W_1,\alpha_1-\pi)\circ E(W_0,\lambda_1)\,, \quad (12.3)$$

where $W_i := (V_i,\dot\gamma_{i+1}(0)) \in T^1\mathbb{S}^2$.

Our goal is to obtain a formula for path-isometries that is more stable than Equation (12.3) with respect to perturbation in the lengths λ_i. We state the formula immediately, but defer the proof until the necessary technical tools have been introduced.

Proposition 12.5. *Following the same notation for a path as in Lemma 12.4, let $w_0 = (V_0,\dot\gamma_1(0))$ and define $w_i := (V_0, R(w_0,\alpha_1+\cdots+\alpha_i-i\pi)\dot\gamma_1(0))$. Then, the path-isometry associated to this path is*

$$E(w_0,\lambda_1)\circ E(w_1,\lambda_2)\circ\cdots\circ E(w_{N-1},\lambda_N)\circ R\big(w_0,\alpha_1+\cdots+\alpha_{N-1}-(N-1)\pi\big)\,. \quad (12.4)$$

First of all, we note that the group $\mathrm{SO}(3)$ acts naturally on $T^1\mathbb{S}^2$.

Definition 12.6 (Action of $\mathrm{SO}(3)$ on $T^1\mathbb{S}^2$)**.** The differential of $T\in\mathrm{SO}(3)$ gives a map $T_*: T^1\mathbb{S}^2\to T^1\mathbb{S}^2$. Using the identification of Definition 12.1, its action on $(e_1,e_2)\in T^1\mathbb{S}^2\subset\mathbb{R}^3\times\mathbb{R}^3$ is $(e_1,e_2)=v\mapsto T_*(v)=(T(e_1),T(e_2))$.

Remark 12.7 (Transitivity)**.** The above action of $\mathrm{SO}(3)$ on $T^1\mathbb{S}^2$ is transitive: given any two pairs (e_1,e_2), (f_1,f_2) of orthogonal unit vectors of $\mathbb{R}^3$ there is an element T of $\mathrm{SO}(3)$ which maps e_1 into f_1 and e_2 into f_2. In fact T is unique: if we let $e_3=e_1\times e_2$ and $f_3=f_1\times f_2$ it follows necessarily that $f_3=T(e_3)$, thus determining the value of T on a basis of $\mathbb{R}^3$.

We next recast the exponential and rotation introduced previously as differentials.

Definition 12.8 (Exponential and Rotation)**.** Let us define the exponential map

$$\exp:\mathbb{S}^1\times T^1\mathbb{S}^2\to T^1\mathbb{S}^2$$

as $\exp_t((e_1,e_2))=(\gamma(t),\gamma'(t))$, where $\gamma:\mathbb{R}\to\mathbb{S}^2$ is the unit-speed parametrization of the geodesic of $\mathbb{S}^2$ starting at $\gamma(0)=e_1$ with $\gamma'(0)=e_2$.

Let us define the rotation map $\mathrm{rot}:\mathbb{S}^1\times T^1\mathbb{S}^2\to T^1\mathbb{S}^2$ as $\mathrm{rot}_\alpha(v)=R_*(v)$, where $R\in\mathrm{SO}(3)$ is the (counterclockwise) rotation of $\mathbb{R}^3$ around the (oriented) axis directed along the basepoint e_1 (of the element $v=(e_1,e_2)$) with angle α.

The following lemma is a simple exercise left to the reader.

Lemma 12.9 (Exponential and Rotation as Linear Maps)**.** *For any $v \in T^1\mathbb{S}^2$ and any $t, \alpha \in \mathbb{S}^1$, we have $\exp_t(v) = E(v,t)_* v$ and $\mathrm{rot}_\alpha(v) = R(v,\alpha)_* v$.*

Remark 12.10. Observe that both $\exp_t : T^1\mathbb{S}^2 \to T^1\mathbb{S}^2$ and $\mathrm{rot}_\alpha : T^1\mathbb{S}^2 \to T^1\mathbb{S}^2$ commute with the action of $\mathrm{SO}(3)$ on $T^1\mathbb{S}^2$.

We are ready to show Proposition 12.5.

Proof of Proposition 12.5. We prove by induction on $1 \leq k \leq N$, the identity

$$\big(E(w_0,\lambda_1) \circ \cdots \circ E(w_{k-1},\lambda_k)\big)_*(w_{k-1}) = (V_k, \dot\gamma_k(\lambda_k))\,. \tag{12.5}$$

For $k = 1$, the identity boils down to $E(w_0,\lambda_1)_* w_0 = (V_1, \dot\gamma_1(\lambda_1))$, which follows from Lemma 12.9. Now, assuming that Equation (12.5) holds for a given k, we show it for $k+1$. Using Lemma 12.9 and the fact that exp and rot commute with the action of $SO(3)$, we obtain

$$E(w_k,\lambda_{k+1})_*(w_k) = \exp_{\lambda_{k+1}}(w_k) = \exp_{\lambda_{k+1}}(\mathrm{rot}_{\alpha_k-\pi}(w_{k-1})) = \big(\exp_{\lambda_{k+1}} \circ \mathrm{rot}_{\alpha_k-\pi}\big)(w_{k-1})\,.$$

Using the latter identity, we deduce

$$\begin{aligned}
&\big(E(w_0,\lambda_1) \circ \cdots \circ E(w_k,\lambda_{k+1})\big)_*(w_k)\\
&\qquad = \big(E(w_0,\lambda_1) \circ \cdots \circ E(w_{k-1},\lambda_k)\big)_*\big((\exp_{\lambda_{k+1}} \circ \mathrm{rot}_{\alpha_k-\pi})(w_{k-1})\big)\\
&\qquad = (\exp_{\lambda_{k+1}} \circ \mathrm{rot}_{\alpha_k-\pi})\Big(\big(E(w_0,\lambda_1) \circ \cdots \circ E(w_{k-1},\lambda_k)\big)_*(w_{k-1})\Big)\\
&\qquad = (\exp_{\lambda_{k+1}} \circ \mathrm{rot}_{\alpha_k-\pi})(V_k, \dot\gamma_k(\lambda_k)) = \exp_{\lambda_{k+1}}(V_k, \dot\gamma_{k+1}(0)) = (V_{k+1}, \dot\gamma_{k+1}(\lambda_k))\,,
\end{aligned}$$

which is Equation (12.5) when k is replaced by $k+1$.

Applying Equation (12.5) with $k = N$, we obtain

$$\begin{aligned}
&\Big(E(w_0,\lambda_1) \circ \cdots \circ E(w_{N-1},\lambda_N) \circ R\big(w_0, \alpha_1 + \cdots + \alpha_{N-1} - (N-1)\pi\big)\Big)_*(V_0, \dot\gamma_1(0))\\
&\quad = \big(E(w_0,\lambda_1) \circ \cdots \circ E(w_{N-1},\lambda_N)\big)_*(w_{N-1}) = (V_N, \dot\gamma_N(\lambda_N))
\end{aligned}$$

which concludes the proof. □

12.2. **Path-isometry stability.** Let us state the stability estimate for path-isometries. As usual the operator norm $\|S\|_{op}$ of an element $S \in \mathrm{SO}(3)$ is defined as

$$\|S\|_{op} := \{\max |S(e)| : e \in \mathbb{S}^2\}\,.$$

Proposition 12.11 (Path-isometry Stability)**.** *Consider two polygonal paths with the following properties:*

- *The starting vertex and the unit tangent vector of the first edge at it are the same;*
- *The number N of edges is the same;*
- *The internal angles at the i-th vertex are the same for every $i \in \{1,\dots,N-1\}$.*

Let S and $\bar S$ be the path-isometries assigned to the respective paths as in Definition 12.3. If we denote by λ_i and $\bar\lambda_i$ the lengths of the respective edges, then

$$\|S - \bar S\|_{op} \leq \sum_{i=1}^N |\lambda_i - \bar\lambda_i|\,. \tag{12.6}$$

The proof of the estimate in Proposition 12.11 is a direct consequence of the formula in Proposition 12.5 and of the following two simple observations.

Lemma 12.12. *Let $T_1, T_2, \dots, T_k \in \mathrm{SO}(3)$ and $\bar{T}_1, \bar{T}_2, \dots, \bar{T}_k \in \mathrm{SO}(3)$ be two k-uples of (special) isometries. We have*

$$\|T_1 \circ \cdots \circ T_k - \bar{T}_1 \circ \cdots \circ \bar{T}_k\|_{op} \leq \sum_{i=1}^{k} \|T_i - \bar{T}_i\|_{op} .$$

Lemma 12.13. *For any $v \in T^1\mathbb{S}^2$ and any $s, t \in \mathbb{S}^1$, we have*

$$\|E(v,s) - E(v,t)\|_{op} \leq |s - t| .$$

Before coming to their proofs, we show how Proposition 12.11 follows from them.

Proof of Proposition 12.11. Let $V_0, \dots, V_N$ and $\bar{V}_0, \dots, \bar{V}_N$ be the vertices of the two paths and let λ_i and $\bar{\lambda}_i$ be the length of the geodesic segments connecting V_{i-1} with V_i and $\bar{V}_{i-1}$ with $\bar{V}_i$. Observe that the elements $w_i \in T^1\mathbb{S}^2$ appearing in Equation (12.4) coincide for the two paths as the angles coincide. In particular, we have

$$\|S - \bar{S}\|_{op} = \|E(w_0, \lambda_1) \circ \cdots \circ E(w_{N-1}, \lambda_N) - E(w_0, \bar{\lambda}_1) \circ \cdots \circ E(w_{N-1}, \bar{\lambda}_N)\|_{op} .$$

The conclusion follows applying first Lemma 12.12 and then Lemma 12.13. □

Proof of Lemma 12.12. We prove it by induction on k. For $k = 1$, the statement is trivial.

For $k \geq 2$, we have

$$\begin{aligned} &\|T_1 \circ \cdots \circ T_k - \bar{T}_1 \circ \cdots \circ \bar{T}_k\|_{op} \\ &\leq \|T_1 \circ \cdots \circ T_k - \bar{T}_1 \circ \cdots \circ \bar{T}_{k-1} \circ T_k\|_{op} + \|\bar{T}_1 \circ \cdots \circ \bar{T}_{k-1} \circ T_k - \bar{T}_1 \circ \cdots \circ \bar{T}_k\|_{op} \\ &= \|(T_1 \circ \cdots \circ T_{k-1} - \bar{T}_1 \circ \cdots \circ \bar{T}_{k-1}) \circ T_k\|_{op} + \|\bar{T}_1 \circ \cdots \circ \bar{T}_{k-1} \circ (T_k - \bar{T}_k)\|_{op} \\ &= \|T_1 \circ \cdots \circ T_{k-1} - \bar{T}_1 \circ \cdots \circ \bar{T}_{k-1}\|_{op} + \|T_k - \bar{T}_k\|_{op} . \end{aligned} \tag{12.7}$$

In the above computations we have used that the operator norm is a norm, and hence satisfies the triangle inequality, and that it is invariant with respect to the (both left and right) composition with elements of SO(3). Joining Equation (12.7) with the inductive hypothesis, the conclusion follows. □

Proof of Lemma 12.13. Without loss of generality, $v = (e_1, e_2) \in T^1\mathbb{S}^2$ where $(e_1, e_2, e_3 := e_1 \times e_2)$ will be used as the basis of $\mathbb{R}^3$. For any $\lambda \in \mathbb{S}^1$, we have

$$\|E(v,s) - E(v,t)\|_{op} = \leq \int_s^t \left\| \frac{d}{d\lambda} E((e_1, e_2), \lambda) \right\|_{op} d\lambda . \tag{12.8}$$

Observe that $E((e_1, e_2), \theta)$ is represented by the matrix

$$\begin{pmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

and the derivative with respect to θ is

$$\begin{pmatrix} -\sin\theta & \cos\theta & 0 \\ -\cos\theta & -\sin\theta & 0 \\ 0 & 0 & 0 \end{pmatrix} .$$

The latter is the composition of a projection with an isometry, thus its norm is at most 1 (and it is, in fact, equal to 1). This observation, together with Equation (12.8), concludes the proof. □

13. Computer-assisted restriction for metric realizations

In this section we will describe an algorithm which accomplishes the following task. The algorithm takes as input a marked spherical graph G in the list of possibilities obtained in Proposition 8.5 for a once-derived net. As an output

(i) either it proves that there is no cone E as in Theorem 1.3 with a link inducing a once-derived net $\mathcal{N}$ with marked spherical graph G;
(ii) or it assigns a length to each edge of the graph so that the resulting net $\mathcal{N}$ embedded in the sphere satisfies *approximately* the constraints we proved to hold for such an $\mathcal{N}$.

Furthermore, for all marked combinatorial maps listed in Proposition 8.5 but two, the algorithm shows that there is at most one output as in (ii), namely that (up to isometries and up to small perturbations) there is at most one $\mathcal{N}$ "realizing" the given combinatorial map.

In order to make the above statement more precise we introduce the following terminology.

Definition 13.1 (Metrized Marked Spherical Graph)**.** A *metrized* marked spherical graph (MMSG) is a marked spherical graph G (see Definition 8.3) endowed with an assignment $\mathcal{L} : \mathcal{E}(G) \to (0, 2\pi)$ of a positive length to each of its edges.

Given two metrized marked spherical graphs $(G, \mathcal{L}_1)$ and $(G, \mathcal{L}_2)$ on the same underlying marked spherical graph G, their distance is the L^∞ distance between $\mathcal{L}_1$ and $\mathcal{L}_2$.

Remark 13.2. Observe that a once-derived net $\mathcal{N}$ induces a length assignment $\mathcal{L}_\mathcal{N}$ on the associated marked spherical graph G. However, a metrized marked spherical graph is not necessarily induced by a once-derived net: we do not even require that it is induced by a geodesic net in the sense of Definition 4.1.

Let us state the Theorem proved by the algorithm.

Theorem 13.3. *Consider a set E as in Theorem 1.3 and let $\mathcal{N}$ be the net once-derived from its link. If the once-derived marked spherical graph G associated to $\mathcal{N}$ is non-extinct, then it belongs to the following subset of the list give in Proposition 8.5:*

$$\mathscr{L} := \{28, 99, 158, 182, 250, 319, 325, 346, 361, 421, 468, 484, 520, 614,$$
$$681, 712, 736, 779, 827, 909, 920, 938, 1029, 1096, 1100, 1122\}$$

Moreover, if $G \in \mathscr{L} \setminus \{827, 1029\}$, then there is a metrized marked spherical graph $(G, \mathcal{L})$ (see `data/metrized_nets/` *in* [13] *or* [12]*) with the following property:*

(U) let E be as in Theorem 1.3 and $\mathcal{N}$ be the net once-derived from its link; if marked spherical graph associated to $\mathcal{N}$ is G and $\mathcal{L}_\mathcal{N}$ is the corresponding length assignment for its edges, then the distance from $(G, \mathcal{L})$ to $(G, \mathcal{L}_\mathcal{N})$ is at most 10^{-4}.

For none of the cases mentioned in Theorem 13.3 we actually prove either existence or uniqueness of a once-derived net $\mathcal{N}$ realizing G. Nonetheless, we strongly believe that in all of them a once-derived net actually exists. Algebraic methods might be sufficient to upgrade the statement to an existence and uniqueness one, but we do not pursue this goal since it is not necessary for our scope (and might require a lot of additional work).

Before coming to the description of the algorithm we point out the heuristic reason behind the fact that for most topologies the metrization is absent and, when it is present, in all but

two cases (goggles and fat fish families, cf. Section 15.1) the once-derived net is in fact rigid. Indeed, one can check that the lengths $\sigma_1, \dots, \sigma_N$ must satisfy $N+1$ equations: this gives an overdetermined system of equations which hence, in most cases, does not have any solution.

13.1. **High-level presentation of the algorithm.** Let us relax the definition of metrized marked spherical graph by allowing the length to be intervals instead of singletons.

Definition 13.4 (Approximately Metrized Marked Spherical Graph)**.** An *approximately* metrized marked spherical graph (AMMSG) is a marked spherical graph G together with an assignment $\mathcal{I} : \mathcal{E}(G) \to \{I : I \subseteq [0, 2\pi] \text{ is an interval}\}$ of an interval to each of its edges.

We say that a metrized marked spherical graph $(G, \mathcal{L})$ is a *realization* of an approximately metrized marked spherical graph $(G, \mathcal{I})$ if, for each edge $e \in \mathcal{E}(G)$, we have $\mathcal{L}(e) \in \mathcal{I}(e)$.

We say that an AMMSG $(G, \mathcal{I})$ is a refinement of another $(G, \mathcal{I}')$ if $\mathcal{I}(e) \subseteq \mathcal{I}'(e)$ for each edge $e \in \mathcal{E}(G)$.

The core routine `search-valid-realization` of our algorithm takes as input an approximately metrized marked spherical graph $(G, \mathcal{I})$ and either

- Proves that there is no realization of $(G, \mathcal{I})$ which is a net $\mathcal{N}$ once-derived from the link of a minimizing cone E as in Theorem 1.3,
- Or produces a realization of $(G, \mathcal{I})$ which *approximately* satisfies most of the constraints of such an $\mathcal{N}$. Notice that in this case we do not guarantee anything on the metrized marked spherical graph (only that it is a realization of $(G, \mathcal{I})$).

Furthermore, one can specify some additional constraints (e.g., two edges must have the same length) and restrict the search to metrized marked spherical graphs satisfying these additional constraints.

Let us explain how `search-valid-metrization` can produce a proof of Theorem 13.3. For each marked combinatorial graph G in the list of Proposition 8.5 which does not belong to the list $\mathscr{L}$ of Theorem 13.3, we run `search-valid-metrization` starting from the approximately metrized marked spherical graph $(G, \mathcal{I})$ where $\mathcal{I}(e)$ is the interval I defined as follows.

- If e is the edge of a face with 2 sides, then $I = \{\pi\}$.
- If e is the edge of a convex face with 3 sides, then $I = \{\lambda_T\}$ (recall Proposition 6.8).
- If e is the edge of a convex face with 4 sides, then $I = [0.01, \pi - 0.001]$ (recall Proposition 11.1).
- If e is the edge of a convex face with 5 sides, then $I = [0.01, \lambda_T - 0.001]$ (recall Propositions 6.8 and 11.1).
- If e is not the edge of a convex face but we can apply Lemma 7.7 to conclude that its length is at most π, then $I = [0, \pi]$;
- If none of the above applies, then $I = [0, 2\pi]$.

Then, `search-valid-metrization` proves that any metrized marked spherical graph on G fails to satisfy at least some of the necessary conditions found so far to be the once-derived net associated to the link of a minimizing cone E as in Theorem 1.3.

It remains to prove rigidity, namely Theorem 13.3(U), when G belongs to $\mathscr{L} \setminus \{827, 1029\}$. Let G be one such element, and perform the following sequence of operations.

(1) Run `search-valid-realization` on $(G, \mathcal{I})$ where $\mathcal{I}$ is the same described above. This produces a metrized marked spherical graph $(G, \mathcal{L})$.
(2) Rerun `search-valid-realization`, starting again from $(G, \mathcal{I})$, but with the additional constraint that we search only metrized marked spherical graphs with distance

greater than 10^{-4} from $(G,\mathcal{L})$. This shows that no such metrized marked spherical graph is once-derived from the link of a cone E as in Theorem 1.3.

This procedure works in all cases apart from net 736, which requires additional considerations. Indeed, if we run this procedure for net 736, the routine is unable to show the rigidity because there is a 1-parameter family of "valid" metrizations. These are valid only approximately (i.e., they don't come from true once-derived nets) and we are able to show rigidity in the following way. After the first run of the strategy we can restrict the values of some specific edges to small intervals: in other words, we find rigidity for a specific subnet. We then use this information together with synthetic and algebraic arguments to prove that two specific edges must be congruent, cf. Proposition 15.9. Finally we run again `search-valid-realization` subject to this constraint and show rigidity.

The remaining part of this section is devoted to explaining the internal working of the routine `search-valid-realization`.

13.2. **The core routine:** `search-valid-realization`**.** Recall that the input of the routine is an approximately metrized marked spherical graph $(G,\mathcal{I})$. The routine consists then of the following three steps.

1. We run a sequence of checks on $(G,\mathcal{I})$. If $(G,\mathcal{I})$ admits a realization that is a net once-derived from a link of a minimizing set E as in Theorem 1.3, then all the checks would pass. Therefore, if any of the checks fails, the routine has proven that there is no valid realization. These checks can modify $\mathcal{I}$ (some intervals may be reduced). If this happens, it is guaranteed that any valid realization before the change is a realization also after.
2. If $\sum_{e\in\mathcal{E}(G)}|\mathcal{I}(e)|<10^{-5}$, the routine returns the realization of $(G,\mathcal{I})$ that is obtained by assigning to e the midpoint of $\mathcal{I}(e)$ for all $e\in\mathcal{E}(G)$.
3. If we reached this step, it means that we are not able to rule out the existence of a valid realization of $(G,\mathcal{I})$ and $(G,\mathcal{I})$ is not "precise" enough for us to be content. Therefore, we create a collection $\{X_1,\dots,X_k\}$ of refinements of $(G,\mathcal{I})$ such that any realization of $(G,\mathcal{I})$ is a realization of X_i for at least one $1\le i\le k$. Then, we run the routine on each of X_i independently. If any of these subroutines produces a metrized marked spherical graph, we return it. Otherwise, we have proven that no valid realization of X exists.

In this sketch, two crucial details are missing: what are the checks we run during Step 1 and how we produce the refinements of Step 3. These decisions are, arguably, the core of the algorithm and slight variations (in the idea or in the implementation) can easily lead to much longer execution times. In fact slight variations can also lead to the checks being less effective on deciding that some realizations are valid. This clearly would not produce a "wrong" answer, but just an answer which is less efficient in the road towards a proof of Theorem 1.3.

13.3. **The refinements of Step 3.** The refinements $X_1,\dots,X_k$ of $(G,\mathcal{I})$ mentioned in Step 3 are produced as follows. A subset $\mathcal{A}$ of the edges of G is selected; the selection process depends only on G and not on $\mathcal{I}$. We obtain $k=2^{|\mathcal{A}|}$ refinements by bisecting each of the intervals assigned to the edges in $\mathcal{A}$.

Since it is not crucial for the correctness of the algorithm, we give only a brief idea of how $\mathcal{A}$ is chosen (we refer the interested reader to `embed_or_exclude.cpp` in [13]). The set $\mathcal{A}$ is a small subset of edges such that if we knew the length of all the edges in $\mathcal{A}$ then

we could deduce the lengths of all the remaining edges with a simple sequence of inferences (e.g., thanks to Proposition 6.8 if we know two sides of a pentagonal convex face, then we immediately know all the sides). Moreover, we select $\mathcal{A}$ so that the deduction of the lengths of the remaining edges given the lengths of $\mathcal{A}$ is as numerically stable as possible.

13.4. **The checks of Step 1.** We consider a metrized marked spherical graph "valid" if it is induced by a once-derived net (see Definition 5.5) which also satisfies the lower bounds on the edge lengths stated in Proposition 11.1.

A check is a process that takes in an AMMSG and tries to prove that there is no valid realization of the given AMMSG. If it fails, it may still refine the given AMMSG by showing that if a valid realization exists then it is a valid realization also of the refinement it produces.

In order to describe the checks, we shall first explain how to obtain a quantitative approximation of the relative position of two vertices (or darts) in an approximately metrized marked spherical graph.

13.4.1. *Approximate Relative Position for AMMSGs.* Let us begin by introducing the concept of relative position for once-derived nets.

Definition 13.5. Consider a once-derived net $\mathcal{N}$. Identify a dart $d \in \mathrm{Darts}(\mathcal{N})$ with the element of $T^1\mathbb{S}^2$ given by initial vertex and the direction of the dart. Given two darts d and d', let $F_{\mathcal{N}}(d, d')$ be the (special) isometry that maps d onto d'.

Observe that $F_{\mathcal{N}}(d, d')$ encodes the relative position of two darts; in particular it is invariant under rotations of the whole net.

We will now describe how to obtain an analogous object for MMSGs and AMMSGs.

Definition 13.6. Consider a metrized marked spherical graph $(G, \mathcal{L})$. Let $d_1, d_2, \dots, d_k$ be a path of darts in G (i.e., the second vertex of d_i is the first vertex of d_{i+1}).

We know the assigned lengths $\mathcal{L}(d_1), \dots, \mathcal{L}(d_k)$, and we have the information that the clockwise angle from d to $\mathrm{next}(d)$ is $\frac{2\pi}{3}$ for all $d \in \mathrm{Darts}(G)$, with the exceptions of the marked darts, for which the angle is $\frac{4\pi}{3}$. Given these pieces of information, after fixing a point $V_0 \in \mathbb{S}^2$ for the starting vertex of d_1 and a direction for it, we can build a corresponding polygonal path which following the darts $d_1, d_2, \dots, d_k$, cf. Definition 6.1. Then, as already done in Definition 13.5, we can identify each d_i with an element of $T^1\mathbb{S}^2$, and we let $F_{(G,\mathcal{L})}(d_1, [d_2, \dots, d_{k-1}], d_k)$ be the (special) isometry that maps d_1 onto d_k. Observe that this isometry does not depend on the position chosen initially for d_1.

Observe that, differently from the case of once-derived nets, this time we have to choose a path from d_1 to d_k to define their relative position in an MMSG. Moreover, notice that if $(G, \mathcal{L})$ is the MMSG associated to a once-derived net $\mathcal{N}$, then

$$F_{\mathcal{N}}(d_1, d_k) = F_{(G,\mathcal{L})}(d_1, [d_2, \dots, d_{k-1}], d_k)\,,$$

in other words it is really in this case that the relative position is *independent* of the chosen path.

We can now move to the case of AMMSGs.

Definition 13.7. Consider an approximately metrized marked spherical graph $(G, \mathcal{I})$. Let $(G, \mathcal{L})$ be the metrized marked spherical graph obtained by setting $\mathcal{L}(e)$ to be the midpoint of $\mathcal{I}(e)$ for all edges $e \in \mathcal{E}(G)$. Let $d_1, d_2, \dots, d_k$ be a path of darts in G (i.e., the second vertex of d_i is the first vertex of d_{i+1}).

Let $F_{(G,\mathcal{I})}(d_1, [d_2, \dots, d_{k-1}], d_k)$ be the subset of $\mathrm{SO}(3)$ given by

$$\left\{ T \in \mathrm{SO}(3) : \left\| T - F_{(G,\mathcal{L})}(d_1, [d_2, \dots, d_{k-1}], d_k) \right\|_{op} \leq \frac{|\mathcal{I}(d_1)| + \cdots + |\mathcal{I}(d_{k-1})|}{2} \right\}.$$

The crucial fact is the following one.

Lemma 13.8. *Let $(G, \mathcal{I})$ be an approximately metrized marked spherical graph and consider a path $d_1, d_2, \dots, d_k$ of darts in G. If $\mathcal{N}$ is a once-derived net which is a realization of $(G, \mathcal{I})$ (when interpreted as a MMSG), then*

$$F_{\mathcal{N}}(d_1, d_k) \in F_{(G,\mathcal{I})}(d_1, [d_2, \dots, d_{k-1}], d_k).$$

Proof. After unrolling the definition, this is a restatement of Proposition 12.11. □

13.4.2. *List of checks.* We are ready to present the checks. However, we do not provide here a detailed explanation of their implementation. Instead, for each of them, we describe what is the property of the net that they verify and which of the results we proved they are enforcing or applying. We also refer the reader to our `C++` implementation `metrization_constraints.cpp` in [13]. As soon as one of the checks fails, the algorithm terminates as it has proven that no valid realization exists.

Our list of checks may seem arbitrary. We want to emphasize that if any of them is removed, then the list of *plausible* nets produced by the algorithm is larger than the one of Theorem 13.3. The reader should interpret the list of checks as a *meta-proof*: we are describing the main ingredients of the minimal proof we could find. Note also that, even if we do not prove it, we believe that all metrized marked spherical graphs found in Theorem 13.3 correspond to a true stationary geodesic with two antipodal pins. This implies that our list of checks is exhaustive.

Inequalities on convex faces: We apply, on convex faces with 4 or 5 sides, the inequalities in Lemma 9.1, Lemma 9.2, and Proposition 9.3 to restrict the intervals assigned to the edges. The algorithm keeps cycling the inequalities through all the convex quadrilaterals and pentagons until it cannot produce any further restriction (or one of the intervals assigned becomes empty, proving that no valid realization exists).

Closure of nonconvex faces: We check whether there is at least one nonconvex face which is "obstructed from closing". Let us show how one can apply Lemma 13.8 to obtain an obstruction. Let $(G, \mathcal{I})$ be the AMMSG we are working with and let $d_1, d_2, \dots, d_k$ be the darts forming the boundary of a nonconvex face. Then, if $(G, \mathcal{I})$ admits a valid realization $\mathcal{N}$, then

$$\mathbb{1} = F_{\mathcal{N}}(d_1, d_1) \in F_{(G,\mathcal{I})}(d_1, [d_2, \dots, d_k], d_1).$$

In particular, if the identity does not belong to the family of isometries on the right-hand side (which is computable) then there is no valid realization.

Our algorithm implements a more delicate check which is more numerically stable and allows us to also restrict some of the intervals when we are not immediately able to exclude the existence of a valid realization. We refrain from describing it here and refer the reader to the code.

Pins' antipodality I: In a stationary net with two antipodal pins, the two pins must be antipodal. In particular, consider a stationary net with two pins $\mathcal{N}$ whose once-derived version is a realization of the AMMSG we are working on. Consider three darts of $\mathcal{N}$ whose second

vertex is a pin (not necessarily the same pin). Then, the three equators which contain the three darts must all go through the pins. Applying Lemma 13.8 similarly to how we used it in the previous check, we are able to exploit this observation to obtain an obstruction for the AMMSG.

Angle at pin: In a stationary net with two antipodal pins, if one pin has degree at least two then its two edges form an angle which is at least $\frac{2\pi}{3}$. Once again, this condition can be transformed in a possible obstruction for the AMMSG using Lemma 13.8.

Pins are outside convex faces: Consider a once-derived net $\mathcal{N}$. Even though we do not know the exact positions of the pins, we do know the *beginning* of the darts that were removed to obtain the once-derived net and thus, as we observed in *Pins' antipodality I*, we can identify a few equators which contain the pins. Assume that, for $\mathcal{N}$, at least two such equators do not coincide. Thus, the two pins are uniquely identified as the intersection of those two equators. In particular the two intersection points cannot be strictly inside one of the convex faces of $\mathcal{N}$. This statement about once-derived nets, thanks to Lemma 13.8, can be distilled into a possible obstruction for AMMSGs.

Intersection of distinct edges: Any two edges of a once-derived net which are not adjacent to the same vertex are disjoint. In particular, this must hold for two edges belonging to the boundary of a nonconvex face. This statement about once-derived nets, thanks to Lemma 13.8, can be distilled into a possible obstruction for AMMSGs.

Pins' antipodality II: Consider a once-derived net $\mathcal{N}$ and let Q_1 and Q_2 be two degree 2 vertices belonging to two distinct nonconvex faces F_1 and F_2.

Starting from Q_1 and entering in the face F_1 along the bisector of its internal angle at Q_1 formed by the two edges at Q_1 we must first encounter the pin P_1 belonging to F_1 (in the stationary net with two antipodal pins that induced $\mathcal{N}$) and then the boundary of F_1. Let d_1 be the amount of distance we have to go along the bisector to meet the boundary of F_1. We have $d_1 \geq \mathrm{dist}_{\mathbb{S}^2}(P_1, Q_1)$.

Define analogously P_2 and d_2. Then, since the two pins P_1 and P_2 are antipodal, we have

$$\pi = \mathrm{dist}_{\mathbb{S}^2}(P_1, P_2) \leq \mathrm{dist}_{\mathbb{S}^2}(P_1, Q_1) + \mathrm{dist}_{\mathbb{S}^2}(Q_1, Q_2) + \mathrm{dist}_{\mathbb{S}^2}(P_2, Q_2) \leq d_1 + d_2 + \mathrm{dist}_{\mathbb{S}^2}(Q_1, Q_2)\,.$$

Observe that this inequality involves only quantities that are known if we know $\mathcal{N}$ (we do not need to know the position of the pins).

This statement about once-derived nets, thanks to Lemma 13.8, can be distilled into a possible obstruction for AMMSGs.

Boundary test: This last check is probabilistic and is used to rule out some very pathological cases. Because of its probabilistic nature, it is run multiple times to increase the probability of finding an obstruction.

Consider a once-derived net $\mathcal{N}$. Take a random oriented geodesic segment σ on the sphere whose two endpoints do not belong to $\mathcal{N}$. For any dart d of $\mathcal{N}$ (which is an oriented geodesic edge), let $\iota(\sigma, d)$ be the intersection number of σ and d, that is,

- 0 if the two geodesic segment do not intersect;
- 1 if they intersect positively, namely the two tangents $\dot{\sigma}$ and $\dot{d}$ form a positively oriented pair at the intersection point;
- -1 if they intersect negatively.

For a face F of $\mathcal{N}$, we then define

$$\iota(\sigma, F) = \sum_{d \in \partial F} \iota(\sigma, d)\,.$$

We leave the proof of the following lemma to the reader.

Lemma 13.9. *The following holds:*
(i) $\iota(\sigma, F) \in \{-1, 0, 1\}$ *for every face* F *of* $\mathcal{N}$*;*
(ii) $\iota(\sigma, F) \neq 0$ *for at most two faces* F *of* $\mathcal{N}$*;*
(iii) $\sum_F \iota(\sigma, F) = 0$ *(where* F *runs over all faces of* $\mathcal{N}$*).*

This statement about a once-derived net (and a geodesic segment σ) can be distilled into a possible obstruction for AMMSGs thanks to Lemma 13.8.

14. Special cases: extinct nets

In this section we analyze the first batches of special nets. These are the nets which lead to an empty twice-derived spherical graph. Their once-derived nets contain therefore only vertices of degree 2, cf. Definition 5.5 and Definition 8.1. These are the ones which we called "extinct" and they are classified here in Theorem 14.1. Any of the nets found in the latter classification is a link of some element of List 1.1.

Theorem 14.1. *Let* $\mathcal{N}$ *be the link of an area-minimizing cone* E *as in Theorem 1.3, let* $\mathcal{N}'$ *be the corresponding once-derived net, and consider the marked spherical graph* G *induced by the latter. Assume that* $\mathcal{N}'$ *is extinct, namely that the vertices of* G *have degree* 2*. Then* G *is either empty, or it has* 2*,* 3*, or* 4 *vertices. Moreover:*

(i) If G *is the empty graph, then* E *belongs to List 1.1(a)&(b);*
(ii) If G *has two vertices, then* E *belongs to List 1.1(c);*
(iii) If G *has three vertices, then* E *belongs to List 1.1(f);*
(iv) If G *has four vertices, then* E *belongs to List 1.1(d).*

Recalling Proposition 5.6, in any once-derived net there are at most 4 vertices of degree 2, while all the others are of degree 3. Since, under the assumption of Theorem 14.1, $\mathcal{N}'$ has only vertices of degree 2, these vertices are at most 4. We will analyze the five possible cases (which include the case of the empty net) separately.

14.1. G **is empty.** In this case the only vertices of $\mathcal{N}$ must be the two pins, because the other vertices are never removed. But then each edge of $\mathcal{N}$ is an half-equator joining the two pins. In particular the two pins have the same degree $d \in \{1, 2, 3\}$. The degree cannot be 3, because in this case the pins would not have been removed to derive $\mathcal{N}'$. If $d = 1$, then $\mathcal{N}$ consists of a single half-equator joining the two pins, and thus E is as in (a) of List 1.1. If $d = 2$, then $\mathcal{N}$ consists of two half equators joining the two pins and they must form a convex angle $\geq \frac{2\pi}{3}$. This is case (b) of List 1.1.

14.2. G **has one vertex.** This is not possible, because then the net would consist of a loop starting from this vertex and getting back to it: the initial net does not have loops and the edges of the once-derived net are edges of the initial one.

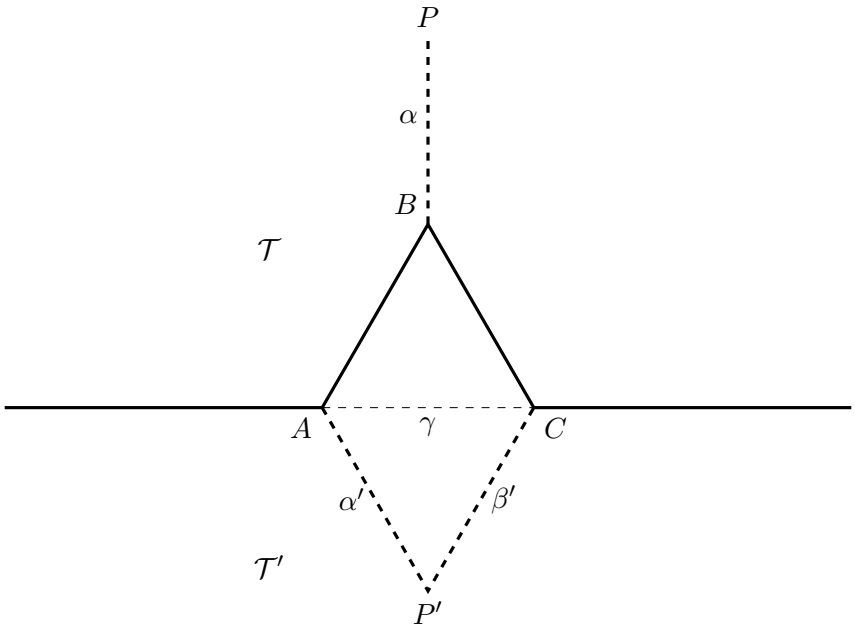


FIGURE 17. The net $\mathcal{N}$ and $\mathcal{N}'$ when $\mathcal{N}'$ is a cycle with three vertices. The thick segments depict the three edges of $\mathcal{N}'$ while the thick dashed segments represent the edges of $\mathcal{N}$ removed in the derivation of $\mathcal{N}'$.

14.3. G **has two vertices.** Let us denote them by A and B. There must then be two edges of $\mathcal{N}'$ connecting them, because they have degree 2 and we have already argued that G contains no loops. In particular A and B would need to be antipodal and the two edges would need to be two half equators connecting them and meeting at $\frac{2\pi}{3}$ angles. The complement of $\mathcal{N}'$ consists then of a convex biangle and a concave biangle. The two pins of the original net $\mathcal{N}$ would need to be in the nonconvex face (or on its boundary if they are A or B) and would also need to belong to the bisector of one of the two angles. These bisectors are actually the same equator and its intersection with the nonconvex face has length exactly π. Since the pins are antipodal, we conclude that A and B are the two pins. Since they have not been removed, they both have degree 3 in $\mathcal{N}$. But then $\mathcal{N}$ consists of the two pins joined by three half-equators meeting at equal angles, namely E is as in (c) of List 1.1.

14.4. G **has three vertices.** The net $\mathcal{N}'$ forms a cycle which splits the sphere into two triangles, which we denote by $\mathcal{T}$ and $\mathcal{T}'$. We assume, without loss of generality, that the internal angles of $\mathcal{T}$ measure $\frac{2\pi}{3}$, $\frac{2\pi}{3}$, and $\frac{4\pi}{3}$ at three vertices which we denote, respectively, by A, C, and B. $\mathcal{T}$ is isosceles and using Theorem 6.6 and Theorem 6.4 we conclude that the lengths of its edges are determined. Let λ_T and λ_Q be, respectively, the side lengths of the equilateral triangle and equilateral quadrilateral with angles equal to $\frac{2\pi}{3}$. We leave to the reader to check that

$$d_{\mathbb{S}^2}(A,B) = d_{\mathbb{S}^2}(B,C) = \lambda_Q \text{ and } d_{\mathbb{S}^2}(A,C) = 2\pi - \lambda_T\,.$$

Moreover, if γ is the equator containing the edge connecting A and C, then $\mathcal{T}$ is contained in one of the two hemispheres bounded by γ, cf. Figure 17.

Observe that there must be a pin P' in (or on the boundary of) $\mathcal{T}'$ connected to the two vertices A and C by two edges α' and β' which have been removed from the net $\mathcal{N}$: P', A, and C form then a convex triangle congruent to the convex triangle formed by A, B, and C. In particular P' is symmetric to B with respect to γ. The other pin P is the antipodal of P' and $\mathcal{N}$ contains an additional edge α joining it to C. We have completely identified the net up to rigid motions and, by Definition 1.5, is is the (link of) a thin fish with six edges.

14.5. G **has four vertices.** In this final case the goal is to show that $\mathcal{N}$ is in fact the support of the link of List 1.1(d), which we call *tetrahedral net*. First of all, the derived-once net $\mathcal{N}'$ splits the sphere into two quadrilaterals, which we denote by Q and Q'. Each quadrilateral has two (internal) angles measuring $\frac{2\pi}{3}$ and two (internal) angles measuring $\frac{4\pi}{3}$. In particular we have two possibilities:

(a) The vertices at which the internal angles of Q measure $\frac{4\pi}{3}$ are consecutive;
(b) The internal angles of Q at any pair of two consecutive vertices are different.

The first case can be excluded by the following lemma.

Lemma 14.2. *If $\mathcal{N}$ is as in Theorem 14.1(iv), then alternative (a) is ruled out.*

It is possible to give a direct argument with pen and paper, but we have not been able to find one which is shorter than a couple of pages. On the other hand we can quickly prove it running the code described in Section 13, cf. `processImpossibleQuadrilateralExtinct` in `embed_or_exclude.cpp` in [13].

From now on, we treat case (b). Therefore, in both quadrilaterals Q and Q' the internal angles alternate in concavity.

Step 1. Let us denote the vertices of the face Q as A, B, C, D; with A and C having internal angles equal to $\frac{2\pi}{3}$, while B and D having internal angles equal to $\frac{4\pi}{3}$. Both B and D are adjacent to an edge of $\mathcal{N}$ which was removed when deriving $\mathcal{N}'$. If this is the same edge (which would happen when both B and D have degree 3 in $\mathcal{N}$), then it subdivides Q into two equiangular triangles with angle $\frac{2\pi}{3}$, thereby determining that all of the edges of $\mathcal{N}'$ are long λ_T. From this we immediately see that $\mathcal{N}$ is the tetrahedral net. Moreover the same consideration can be done on the quadrilateral Q'. We can therefore, without loss of generality, assume that:

(c) B and D are connected to a pin P internal to Q by two edges β and δ of $\mathcal{N}$;
(d) A and C are connected to a pin P' internal to Q' by two edges α and γ of $\mathcal{N}$.

Note in addition that, if any of these pairs of edges form a flat angle at one of the two pins, then we again conclude that $\mathcal{N}$ is the tetrahedral net. We can therefore also assume

(e) the smallest angle formed by β and δ at P is strictly smaller than π.

From now on we want to show that (c), (d), and (e) lead to a contradiction.

Step 2. We next claim that all the edges of $\mathcal{N}'$ are shorter than π. For the sake of contradiction, assume that, e.g., the edge AB is π or longer. Then the quadrilateral $ABDP$ is not convex, cf. Figure 18, and its area is, by Gauss-Bonnet, at most $\frac{4\pi}{3}$ (recall that the internal angle at P is at most $\frac{4\pi}{3}$). Now, on the edge AB let χ be the half-equator which has vertices B and B'. We can rotate χ inside $ABPD$ until we hit a point of the boundary which is distinct from B and B'. We must be able to rotate until we hit the edge AP. Otherwise we would hit some other point on the boundary, but this is not possible because P is the only concave angle of the quadrilateral. So we must have rotated χ by $\frac{2\pi}{3}$, the angle at B, forming

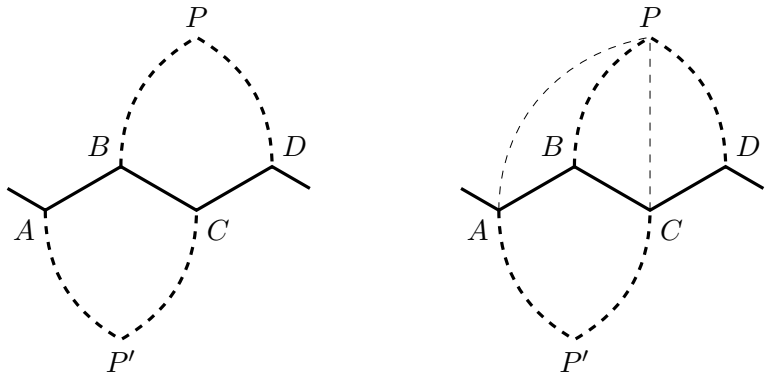


FIGURE 18. Final arguments for the rigidity of the tetrahedral case. On the left the quadrilaterals $ABCP'$ and $BPDC$ are used in Step 2 to prove that all edges of $\mathcal{N}'$ are shorter than π. Thus Theorem 6.7 shows that AB is congruent to CD and BC is congruent to AD. On the right, in Step 3 the triangles CPD and APB are shown to be congruent, and likewise the triangles BPC and ADP. As a result we conclude that the angle BPD is flat.

a biangle of area $\frac{4\pi}{3}$ contained strictly inside $ABPD$. Since the latter has area at most $\frac{4\pi}{3}$ we have reached a contradiction.

We now orient the closed loop formed by $\mathcal{N}'$ in two ways: one going from A to B to C to D, the other going from B to A to D to C. In doing so we discover two polygonal paths which have a common edge and equal internal angles: since all the edges are shorter than π, we can apply Theorem 6.7 and conclude that BC is congruent to DA. We then repeat apply the argument with the two loops B to C to D to A and C to B to A to D, concluding that AB and CD are also congruent.

Step 3. In this (final) step we will use that the pins P and P' are antipodal. Consider in particular the equator contanining $P'A$: this equator necessarily contains P given that it is antipodal to P'. Hence its intersection with Q contains a geodesic segment PA which bisects the internal angle of Q at A. Likewise for C, cf. Figure 18. So, the triangles ABP and CDP have the two edges AB and CD congruent and the two angles adjacent to the respective edges congruent (measuring $\frac{\pi}{3}$ and $\frac{2\pi}{3}$). By Theorem 6.6, the two triangles are congruent. So the angles $A\hat{P}B$ and $C\hat{P}D$ are congruent. An analogous argument implies that $B\hat{P}C$ and $D\hat{P}A$ are also congruent. Looking at the pin P we realize however that $2A\hat{P}B + 2B\hat{P}C = 2\pi$, in particular the internal angle at P of the quadrilateral $BADP$ is π, contradicting (e).

15. Special cases II: seven difficult families

Here we will analyze seven marked spherical graphs among the ones listed in Proposition 8.5. The overall conclusions are given in Theorem 15.3 and Proposition 15.9. The link of any element in List 1.1 generates a marked spherical graph which is either analyzed in Theorem 14.1 or in Theorem 15.3. However, Theorem 15.3 involves also two further families of links which are described in detail in Section 15.1 and depend upon one and two parameters. These are not in List 1.1: for these Theorem 15.3 must be considered just a step before they will be finally excluded (on energetic grounds) in the next sections. Proposition 15.9 concerns the marked spherical graph numbered as 736 in the list of Proposition 8.5. As it

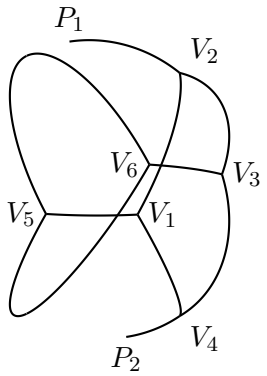


FIGURE 19. An element in the fat fish family. The picture should be compared to Figure 2: a portion of γ_5 in Figure 2 has been removed and replaced by the half-equators which in the picture above join V_5 and V_6 (another difference is that in Figure 2 the quadrilateral $V_1V_2V_3V_4$ is made of straight segments, while in the above picture we have drawn the corresponding projections on the sphere).

was already mentioned in Section 13, a more subtle interaction of theoretical pen-and-paper considerations with the code described in Section 13 allows ultimately to complete the proof of Theorem 13.3 for the links of minimizing cones E which induce the marked spherical graph 736: this is explained in the proof of Proposition 15.9.

15.1. **Fat fish and goggles.** Our next goal is to analyze metric realizations of the marked spherical graphs numbered as 158, 681, 827, 920, 1029, and 1100, (see Figures 22, 23, 27, 26, 21 and 24 for diagrams depicting them) in the list of Proposition 8.5. In order to proceed with our analysis we need to introduce two families of stationary geodesic nets with two pins. At a later stage we will show that none of these nets is the link of a minimizing cone, but in this section we will prove that these are the only possible realizations of the marked spherical graphs 827 and 1029.

Definition 15.1 (Fat fish)**.** Consider any geodesic net $\mathcal{N}'$ which is an element of the thin fish family, cf. Definition 1.5. Let γ_5 be the (only) geodesic arc in $\mathcal{N}'$ which is longer than π and let V_5 and V_6 be two arbitrary antipodal points in (the interior of) γ_5. Eliminate from the net the half equator connecting V_5 and V_6 which is a subset of γ_5 and add the two half-equators connecting the same vertices which meet the remaining portions of γ_5 forming angles of $\frac{2\pi}{3}$ at V_5 and V_6, cf. Figure 19. If one of the two half-equators joining V_5 and V_6 intersects one of the two edges γ_6 and γ_7, then we discard the net, otherwise it is a *fat fish*.

Once we mod out the symmetries of the sphere, the fat fish family is a 2-parameter family: the lengths of the arcs V_5V_1 and V_2P_1 both vary in an interval. All the other edges have instead a fixed length.

We next define a second family, which we call "goggles". We recall that λ_T denotes the length of the sides of an equiangular convex spherical triangle with internal angles equal $\frac{2\pi}{3}$. Note that, given any pair of points A and C on the sphere at distance λ_T, there are exactly two points B and D with the property that $ABCD$ is a convex equilateral equiangular quadrilateral with angles $\frac{2\pi}{3}$: this can be proved observing that the two diagonals of the

quadrilateral AC and BD must measure λ_T, must be perpendicular, and must intersect at the respective midpoints.

Definition 15.2 (Goggles)**.** Let γ be any equator on $\mathbb{S}^2$ and mark on it four consecutive points V_1, V_2, V_3, V_4 with the property that the geodesic distances between V_1 and V_2 and V_3 and V_4 are both λ_T. Remove the respective arcs from γ, so that what remains consists of two length-minimizing geodesic arcs joining V_2 with V_3 and V_4 with V_1. Next add the four vertices V_5, V_6, V_7, V_8 such that $V_1V_2V_5V_6$, resp. $V_3, V_4V_7V_8$, are the four vertices of an equilateral quadrilateral with angles $\frac{2\pi}{3}$. The vertices are labeled in such a way that V_5 and V_7 lie on the same connected component of $\mathbb{S}^2 \setminus \gamma$, cf. Figure 20.

Add to the net the four sides of each quadrilateral. Furthermore, let P_1 and P_2 be the two centers of the two hemispheres forming $\mathbb{S}^2 \setminus \gamma$, labeled so that P_1 is in the same hemisphere as V_5 and V_7, and add the length-minimizing geodesic arcs joining P_1 with V_5 and V_7 and P_2 with V_6 and V_8. We say that the net obtained in this fashion is a *goggle* if and only if the convex angle formed by V_5P_1 and V_7P_1 at P_1 is at least $\frac{2\pi}{3}$.

Note that the nets in the goggles family consist of 10 vertices and twelve edges. Each of these nets is symmetric with respect to γ. There are (up to isometries) two extremal elements in the family. One is given when the convex angle at the pins is exactly $\frac{2\pi}{3}$. If this is the angle internal to the pentagon $V_2V_3V_7P_1V_5$, then the latter and the symmetric pentagon $V_2V_3V_8P_2V_6$ are equiangular with angles equal $\frac{2\pi}{3}$. The other extremal case is the one in which the angle at the pin is π, namely when P_1 is the midpoint of the length-minimizing geodesic arc joining V_5 and V_7 (which we will denote as V_5V_7) and P_2 is the midpoint of the corresponding geodesic arc V_6V_8. In this case the 8 vertices $V_1V_2V_3V_4V_5V_6V_7V_8$ are in fact the vertices of a cube inscribed in the unit sphere and the line ℓ joining P_1 and $P_2 = -P_1$ goes through the midpoints of two opposite edges of the cube. This extremal net is thus in common with the link of a cone in the cubical family described in Theorem 1.3(e). Note however that this is one case for which we will prove that the corresponding cone is not Almgren minimizing.

15.2. **Classification of six special nets.** We are now ready to state the main theorem classifying the following six special nets.

Theorem 15.3. *Let $\mathcal{N}$ be the link of a cone E as in Theorem 1.3 and denote by $\mathcal{N}'$ the corresponding once-derived net.*

Let id be the identifier of the marked spherical graph associated to $\mathcal{N}'$ in the list of Proposition 8.5. Then, for some special values of id we have the following statements:

id = 158*: This is impossible.*
id = 681*: This is impossible.*
id = 827*: $\mathcal{N}$ is an element of the fat fish family.*
id = 920*: $\mathcal{N}$ is an element of the thin fish family (see List 1.1(f)).*
id = 1029*: $\mathcal{N}$ is either an element of the cubical family (see List 1.1(e)) or the goggles family.*
id = 1100*: $\mathcal{N}$ is an element of the cubical family (see List 1.1(e)).*

The next subsections will analyze the various cases of Theorem 15.3.

15.3. **1029 is a cube or a goggle.** In this step we will prove the following fact. Assume $\mathcal{N}$ and $\mathcal{N}'$ are as in Theorem 15.3, and that the marked spherical graph associated to $\mathcal{N}'$ is numbered 1029. Then $\mathcal{N}$ is necessarily an element of the goggle family or of the cube family.

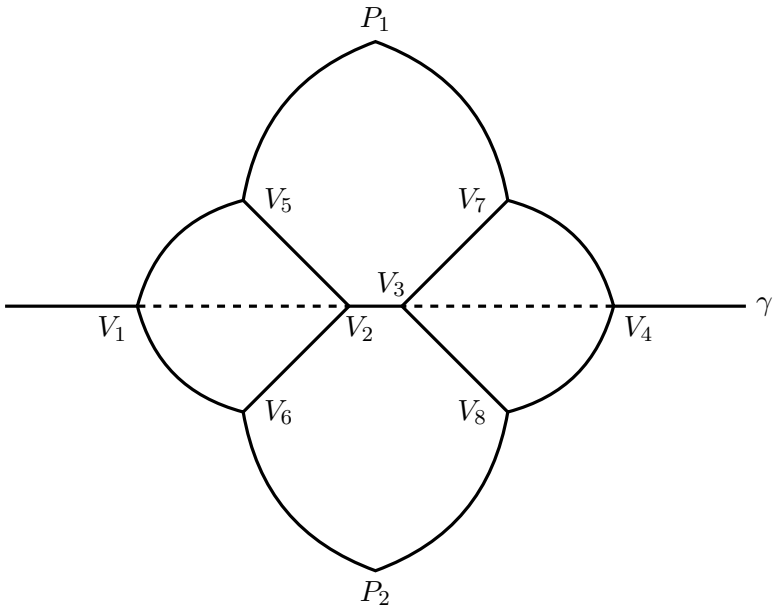


FIGURE 20. A schematic picture of the goggle family. The long straight line represent the equator γ: the dashed lines are taken away from it, so that the remaining portion consists of two segments connecting V_2 and V_3 and V_4 and V_1. The two quadrilaterals $V_1V_5V_2V_7$ and $V_3V_6V_4V_8$ are equilateral and have internal angles $\frac{2\pi}{3}$. The pins P_1 and P_2 are the two "poles", namely the centers of the hemispheres forming the complement of γ.

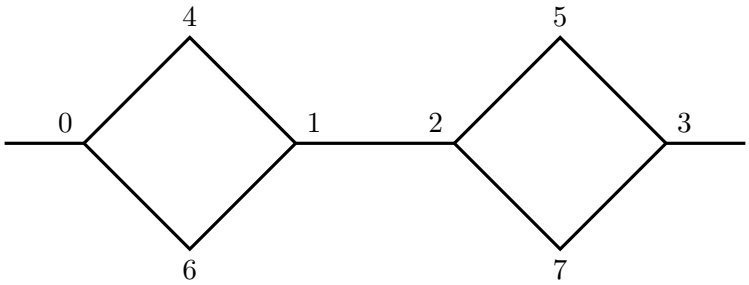


FIGURE 21. A picture of the marked spherical graph 1029 for a once-derived net $\mathcal{N}'$. The picture is on the plane, the graph becomes spherically embedded by adding a point at infinity so that the vertices 0 and 3 are connected by an edge. The two quadrilateral 0416 and 2537 are convex.

We start by drawing a picture schematizing the properties of the once-derived net $\mathcal{N}'$, cf. Figure 21.

Since the most important part of the proof will also be used to deal with other cases, we isolate its consequences in the following lemma.

Lemma 15.4. *Consider a geodesic net such that:*

(i) Its spherical graph is depicted in Figure 21 (from now on, we identify vertices of the graph with vertices of the net);
(ii) Convex angles in Figure 21 correspond to angles of $\frac{2\pi}{3}$ in the net;
(iii) The edge 12 is shorter than π.

Then, one of the following two statements hold:

(a) The two quadrilaterals 0416 and 2537 are equilateral and the four vertices 0, 1, 2, and 3 belong to the same equator.
(b) The geodesic segments of length π starting from 4 and 5 and bisecting the respective concave angles do not intersect.

In order to apply the lemma to the case at hand we need to check condition (iii): this is, however, a consequence of Lemma 7.7. Moreover, because of the existence of the pins, we can exclude case (b) of the conclusion. Thus, (a) holds and we deduce that $\mathcal{N}$ is either an element of the goggle family or an element of the cube family. The rest of this section is devoted to a proof of Lemma 15.4.

Proof of Lemma 15.4. Let Q and Q' be the equiangular quadrilaterals 2537 and 0416. Observe that, if any of the two is equilateral, then (a) holds. Indeed, without loss of generality, assume that Q is equilateral. Then, the diagonal 23 of Q is a bisector of the angles in 2 and 3 and thus 0, 1, 2, and 3 all lie on the same equator γ. Take now the diagonal 10 of Q'; it must also lie on γ and thus the latter diagonal bisects the two angles at 1 and 0. So, the convex triangles 106 and 104 are both equiangular and equilateral. We conclude that Q' is also equilateral.

From now on, we assume that neither Q nor Q' is equilateral. We draw the two diagonals of the quadrilaterals Q and Q', which we denote by 01 and 23 and consider:

- the closed polygonal path $\mathcal{P}$ which start at 0 and proceeds to 1, 2, 3;
- the closed polygonal path $\mathcal{P}'$ from 1 to 0 to 3 to 2.

Since Q and Q' are not equilateral, none of the angles of these loops are flat. Furthermore, all edges of the loops, apart from 03, are shorter than π thanks to assumption (iii). Note moreover that the four angles of the loops are congruent (here we are using that opposite sides in an equiangular convex quadrilateral are congruent). We can thus apply Theorem 6.7 and conclude that the two loops are congruent. This implies that the edges 03 and 21 are congruent.

Next, we observe that the two equators containing the two edges 03 and 21 are at the same time bisectors of opposite angles of both quadrilaterals. We will then use the following lemma, which can be checked via a direct computation in spherical trigonometry: the explicit computations are given in Section G.

Lemma 15.5. *Consider two equiangular spherical quadrilaterals Q_1 and Q_2 with angles of $\frac{2\pi}{3}$. Assume that the two equators bisecting two opposite angles of Q_1 form the same angle as the two equators bisecting two opposite angles of Q_2. Then Q_1 and Q_2 are congruent.*

The lemma implies that the quadrilaterals Q and Q' are congruent. So either 41 is congruent to 52 or 41 is congruent to 27. If we now look again at the polygonal path $\mathcal{P}$ (where each pair of edges which do not have a vertex in common consists of congruent edges) the two alternatives amount to the angles at 1 and 2 in this path being equal or summing to 2π. If the angles were equal, then the diagonals of the two quadrilaterals Q and Q' would lie in the same closed hemisphere H bounded by the equator passing through 1 and 2. In particular 3 and 0 would belong to the interior of H and so would segment edge 30, since it is shorter than

π. But then one of the two quadrilaterals Ω bounded by the polygonal path $\mathcal{P}$ would have area strictly smaller than 2π. But Gauss-Bonnet applied to Ω shows that its area is exactly 2π, leading to a contradiction.

Summarizing, we have shown that the two convex quadrilaterals are congruent and, moreover:

- The edges 37, 16, 40, and 52 are all congruent;
- The edges 72, 60, 41, and 53 are all congruent.

Let $u_4, u_5 \in T^1\mathbb{S}^2$ be the unit tangent vectors directed along the bisectors of the concave angles in 4 and 5.

Let O be the unique orientation-preserving isometry of the sphere which maps the point 1 in the point 3 and the edge 16 to the the edge 37 (cf. Figure 21). The map O maps the geodesic net into itself, and in particular it maps u_4 to u_5 and viceversa (recall the action of $SO(3)$ on $T^1\mathbb{S}^2$ introduced in Definition 12.6).

Now, let us consider three cases depending on the relative position of u_4 and u_5.

(1) Assume that u_4 and u_5 identify the same equator (via the exponential map) with the same orientation. Then, given that $O_*(u_4) = u_5$ and $O_*(u_5) = u_4$, the points 4 and 5 must be antipodal, which is a contradiction because 4 and 5 lay on the same side of the equator identified by 12.

(2) Assume that u_4 and u_5 identify the same equator (via the exponential map) with opposite orientations. Then, consider the quadrilateral 1254 where the edge 54 is the shortest geodesic connecting 5 to 4. Observe that the edges of this quadrilateral are non-intersecting because 4 and 5 lay on the same side of the equator identified by 12. The edge 54 is either the bisector of the concave angles in 5 and 4 or it is the bisector of the convex angles in 5 and 4.
   - If it is the bisector of the concave angles, then the quadrilateral 1254 is a convex equiangular quadrilateral with angles of $\frac{2\pi}{3}$ and thus 14 is congruent to 25; which is a contradiction because it implies that Q and Q' are equilateral.
   - If it is the bisector of the convex angles, then Gauss-Bonnet implies that its area is $\frac{8\pi}{3}$, which is impossible because the quadrilateral is contained in a hemisphere bounded by the equator identified by 12.

(3) The equators identified (via the exponential map) by u_4 and u_5 are distinct. If the conclusion (b) of the statement holds, we do not have to prove anything. Otherwise, there is a point V such that $\exp_s(u_4) = V = \exp_t(u_5)$ with $0 < s, t < \pi$. Since O swaps u_4 and u_5, we have $O(V) = V$ and in particular $s = t$, which is equivalent to the congruence of $4V$ and $5V$. Applying Theorem 6.7 to the pentagon $4125V$ we deduce that 41 is congruent to 25 which is a contradiction because it implies that Q and Q' are equilateral.

In particular $O(V) = V$, which in turn implies the congruence of $6V$ and $7V$. Thanks to Theorem 6.7, we deduce that the pentagon $6127V$ is isosceles, from which in turn we infer that 16 and 27 are also congruent. □

15.4. **158 and 681 are not possible.** The arguments are very similar and we will therefore carry them in parallel. First of all we draw schematic pictures of both topologies, cf. Figure 22 and Figure 23.

Observe that, by Lemma 7.7 any once-derived net $\mathcal{N}'$ with either marked spherical graph has all edges which are strictly shorter than π. We now see that in both cases there is a

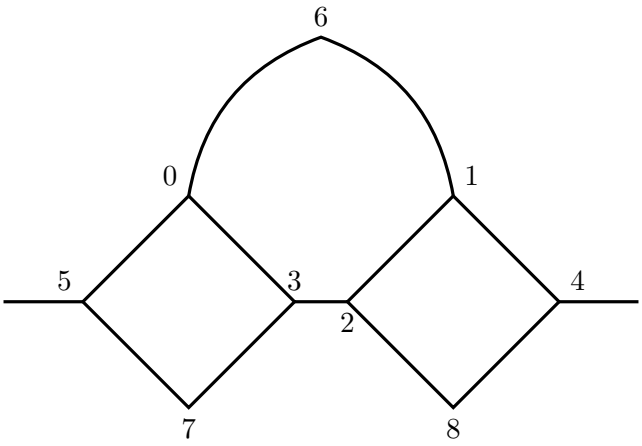


FIGURE 22. A picture of the marked spherical graph 158 for a once-derived net $\mathcal{N}'$. The quadrilaterals 5037 and 2148 and the pentagon 30612 are convex.

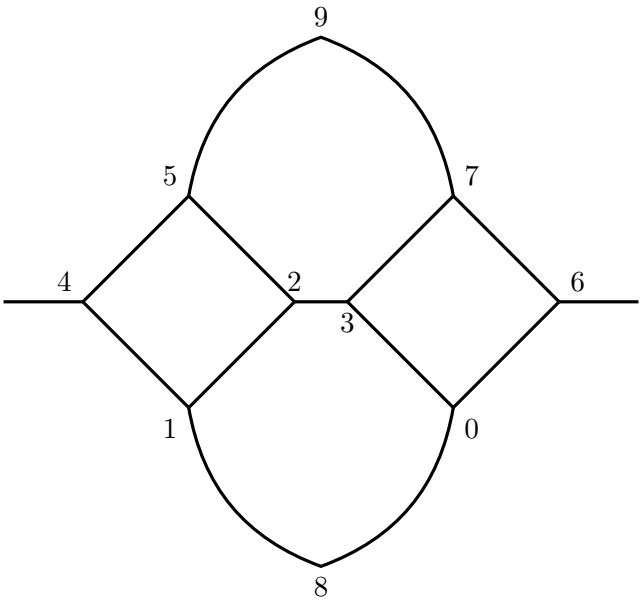


FIGURE 23. A picture of the marked spherical graph 681 for a once-derived net $\mathcal{N}'$. The quadrilaterals 4521 and 3760 and the pentagons 25973 and 30812 are convex.

subnet $\mathcal{N}''$ to which we can apply Lemma 15.4. In the case of 158 the subnet is obtained by removing the vertex 6 and the edges 60 and 61, while in the case of 681 we obtain it by removing the vertices 9 and 8 and the edges 95, 97, 80 and 81. Since conclusion (b) is false for both nets, Lemma 15.4 guarantees that the two convex quadrilaterals in the respective nets are both equilateral. Note that we have uniquely determined the convex pentagons of both nets, since they have to be isosceles and they have two (nonconsecutive) congruent sides with length which equals the sidelength of the equailateral quadrilaterals. In particular such

pentagon is the convex pentagon which appears as two faces of the one element of the goggles family for which the angle at the pins is $\frac{2\pi}{3}$ (this is also the same pentagon of the pentagonal faces of one of Heppes' nets in [14]). We therefore conclude that:

- the marked spherical graph 681 has a unique metrization which must be precisely this element of the goggles family; on the other hand this would imply that the vertices 8 and 9 are the pins of the net, but the rules for the once-derivation would have removed both, and thus we would have been left with another marked spherical graph (in fact the marked spherical graph 1029);
- 158 has a unique metrization which coincides with the latter element of the goggles family with only one pin (and the corresponding two edges) removed; but again this is in contradiction with the derivation rules, which require *both* pins to be removed.

15.5. **1100 is a cube.** Here we analyze the marked spherical graph numbered 1100 and prove the following proposition.

Proposition 15.6. *Assume $\mathcal{N}$ is the link of a cone E as in Theorem 1.3 and $\mathcal{N}'$ a net once-derived from $\mathcal{N}$ whose underlying marked spherical graph is numbered 1100 in the list of Proposition 8.5. Then:*

(a) E is an element in List 1.1, namely the cone over the 1-skeleton of a cube;
(b) ℓ goes through two vertices of the cube, namely the two pins have degree 3;
(c) $\mathcal{N}'$ is obtained by removing from $\mathcal{N}$ two edges σ and σ'. The two edges σ and σ' are adjacent to the pins P and P', and to two other vertices which are not antipodal.

We start by drawing the marked spherical graph of the net $\mathcal{N}'$ in Figure 24 below. Let us define:

- the point F' to be the midpoint of the edge 20;
- the equator γ to be the one passing through F' and orthogonal to the edge 20;
- the point F to be the antipodal point to F' on the equator γ.

Note that in the picture the point F is also the midpoint of the edge 57 and that the whole net is central-symmetric with respect to F. These two facts are proven in the following lemma. Before stating it, we recall our convention of denoting by λ_Q the length of the edges of an equilateral equiangular spherical quadrilateral with angles $\frac{2\pi}{3}$. A posteriori we must conclude that the length of all the edges of the nets $\mathcal{N}'$ and $\mathcal{N}$ is, in fact, precisely λ_Q.

Lemma 15.7. *If $\mathcal{N}'$ is as in Proposition 15.6 then the following holds:*

(i) if λ is the length of an edge of $\mathcal{N}'$, then $|\lambda - \lambda_Q| < 2 \cdot 10^{-4}$;
(ii) the edges 20, 63, *and* 14 *are congruent;*
(iii) the edges 12, 40, 03 *and* 26 *are congruent;*
(iv) γ intersects 14 *and* 63 *orthogonally in the respective midpoints;*
(v) The central symmetry s with respect to the line through F and F' maps 6 *into* 4, 1 *into* 3, 2 *into* 0, 7 *into* 5*;*
(vi) F is the midpoint of the edge 57*;*
(vii) If 26 *and* 63 *are congruent, then the conclusion of Proposition 15.6 follows.*
(viii) If Proposition 15.6 is false, then the pins of $\mathcal{N}$ are symmetric with respect to the line through F and F'. In particular their distance to F is $\frac{\pi}{2}$.

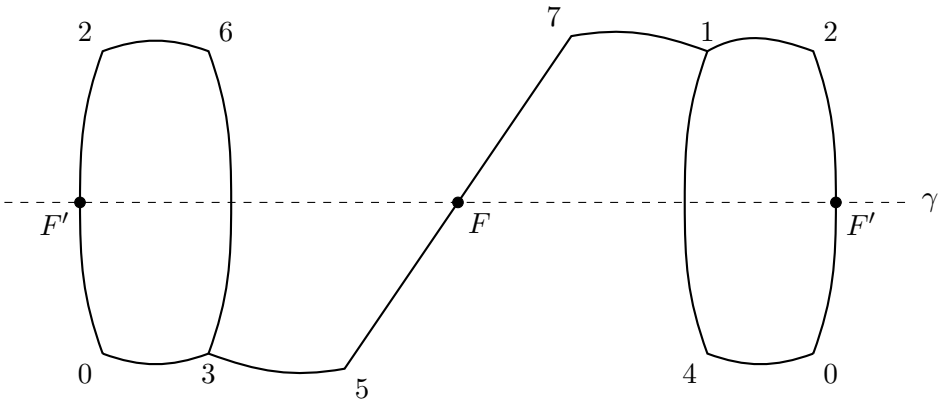


FIGURE 24. A diagram describing the marked spherical graph 1100. In the metric realization the angles at 6 and 7 facing upwards measure $\frac{4\pi}{3}$, as the angles at 5 and 1 facing downwards.

Proof. The statement (i) of the lemma is a consequence of Theorem 13.3. The statements (ii), (iii), and (iv) are immediate from the fact that the two quadrilaterals 1204 and 2603 are equiangular with angles $\frac{2\pi}{3}$ and share the edge 20.

Concerning (v), the fact that $s(3) = 1$, $s(6) = 4$, and $s(2) = 0$ follows immediately from (ii) and (iii). Note that in order to show that $s(5) = 7$ we just need to show that the edges 35 and 71 are congruent. In order to prove the latter, consider the hexagons with vertices 126357 and 304175. Note that the angles at the vertices (in the given order) form the same sequence $\frac{2\pi}{3}$, $\frac{2\pi}{3}$, $\frac{4\pi}{3}$, $\frac{2\pi}{3}$, $\frac{2\pi}{3}$, $\frac{4\pi}{3}$. Moreover 30 is congruent to 12, 04 to 26, and 63 to 41, thanks to (ii) and (iii). In particular, by Theorem 6.7 (which is applicable because by (i) all the segments in the polygonal paths are shorter than π) 35 is congruent to 17. We have therefore shown (v).

As for (vi), since $s(5) = 7$ and $s(7) = 5$, and the segment 57 is shorter than π, s must map it into itself. In particular the midpoint of 57 is a fixed point of the symmetry s. The latter has only two fixed points: F and F'. But F' does not belong to 57.

Under the assumption of (vii), namely that 26 and 63 are congruent, we immediately see that 26, 63, and 16 are all congruent and their length is λ_Q.

Take then the spherical link of the cone over a cube and remove two edges σ and σ' from this link to obtain a spherical net $\tilde{\mathcal{N}}$. We impose, however, if P and A are the vertices connected by σ and P' and A' the ones connected by σ', then P' is antipodal to P, but A' is not antipodal to A. We see immediately that there is a face F of $\mathbb{S}^2 \setminus \tilde{\mathcal{N}}$ which is an hexagon with the same internal angles as 126357. Having proven that 36, 62, and 21 have all length λ_Q allows us to apply Theorem 6.7 to infer that in fact 126357 is congruent to F. This shows that the rigid motion which maps 126357 onto F maps in fact the net $\mathcal{N}'$ onto $\tilde{\mathcal{N}}$, thus showing Proposition 15.6.

Finally, coming to (viii), note first that the pins belong to the four equators bisecting the angles in 4, 5, 6, and 7. At least the ones bisecting 5 and 6 must be distinct because of (i). Therefore the pins P_1 and P_2 are the only pair of points common to all four bisectors mentioned above. But the symmetry s maps the bisector of 6 in that of 4 and the bisector of 7 in that of 5. So $s(\{P_1, P_2\}) = \{P_1, P_2\}$. But since s is an involution, either $s(P_1) = P_2$, or

$s(P_1) = P_1$. The latter option is however not possible, because the only fixed points of s are F and F', which are contained in edges of $\mathcal{N}'$, while the pins cannot. □

We are now ready to close the argument of the proof by contradiction in Proposition 15.6. We derive six polynomial equations in the sines and cosines of 6 unknown geometric quantities, which are lengths of edges or measurements of angles. The list of the geometric quantities is given below in (15.1)–(15.6), while the six equations are given by (15.7), (15.8), (15.9), (15.10), (15.11), and (15.12). Using that $\sin^2 + \cos^2 = 1$, we obtain a system of 12 polynomial equations in 12 unknowns. Next note that Lemma 15.7(i) specifies a certain range for the 12 unknowns. Moreover, if we assume by contradiction that Proposition 15.6 does not hold, this range is further restricted to exclude the cases in which the segments 26 and 63 are congruent. Using a computer algebra software, we will see that there is no solutions of the system in the range.

15.5.1. *Trigonometric identities.* In order to describe the six geometric quantities and the triangles used to relate them we refer to Figure 25 below. The solid segments are all part of the net $\mathcal{N}'$, cf. Figure 24. The dashed segments are all geodesics.

While we keep the notation F for the same point in both figures, in the sequel we use B, C and X for the three vertices labeled 0, 3 and 5 in Figure 24. We extend the edges 20 and 63 along the respective equators on the sides of 0 and 3 until they meet in the point A.

Note that, since the edges 20 and 63 intersect the equator γ orthogonally, the point A is the pole of the lower hemisphere with boundary γ, more precisely it is the point equidistant from all points of γ which lies in the same connected component of $\mathbb{S}^2 \setminus \gamma$ of the vertices 0 and 3. The dashed segments with extrema 5 and 6 are on the bisectors of the corresponding of the angles at the corresponding vertices of $\mathcal{N}'$. Their second common point, labeled E, is thus one of the pin, in fact the pin lying in the northern hemisphere (this is justified recalling Lemma 15.7: our net $\mathcal{N}$ is close to the cubical net and so the pin is close to the edge joining 6 and 7, which lies on the northern hemisphere). The segment AE intersects γ orthogonally at some point O.

In order to justify the accuracy of Figure 25 we need to ensure that the segment AE intersects the segment CX in a point D which is interior to both. Recall that we are working under the assumption that the segments 26 and 63 have nearly the same length. When they are congruent the point E coincides with 6, the point D coincides with C and XF is orthogonal to γ. In particular it follows that in our situation E is close to 6 and D is close to C. Moreover, observe that E cannot be inside the quadrilateral (because it is a pin) and cannot be on its boundary (because then it would coincide with 6, which would be a pin, and the net $\mathcal{N}$ would be the cubical net). Hence D, which is the symmetric of E with respect to the plane containing γ, is strictly inside CX

The six geometric quantities involved in the system are:

$$\text{(the size of) the angle } B\hat{A}C; \tag{15.1}$$

$$\text{(the length of) the segment } AC; \tag{15.2}$$

$$\text{the angle } A\hat{D}C; \tag{15.3}$$

$$\text{the segment } AD; \tag{15.4}$$

$$\text{the segment } EX; \tag{15.5}$$

$$\text{the angle } D\hat{E}X. \tag{15.6}$$

By the cosine rule on the triangle $0AC$, we have

$$\cos\frac{\pi}{3} = -\cos\frac{\pi}{3}\cos B\hat{A}C + \sin\frac{\pi}{3}\sin B\hat{A}C\cos AC\,. \tag{15.7}$$

By the sine rule on the triangle ACD, we have

$$\frac{\sin A\hat{D}C}{\sin AC} = \frac{\sin\frac{\pi}{3}}{\sin AD}\,. \tag{15.8}$$

By Lemma 15.7(viii), $EF = \frac{\pi}{2}$ and also $FA = \frac{\pi}{2}$. In particular, the geodesic connecting A and E is a subset of the equator at distance $\frac{\pi}{2}$ from F and thus $OF = OF' = \frac{\pi}{2}$. Moreover, $0\hat{A}D = OF'$, because γ consists is the equator at distance $\frac{\pi}{2}$ from A. Hence, $C\hat{A}D = \frac{\pi}{2} - B\hat{A}C$, in particular $\cos C\hat{A}D = \sin B\hat{A}C$ and $\cos C\hat{A}D = \sin B\hat{A}C$. By the cosine rule on the triangle ACD, we have

$$\cos\frac{\pi}{3} = -\cos B\hat{A}C\sin A\hat{D}C\cos AD + \sin B\hat{A}C\cos A\hat{D}C\,. \tag{15.9}$$

Observe that $E\hat{D}X = A\hat{D}C$ and $DE = \pi - 2\,AD$ so $\sin DE = 2\sin AD\cos AD$. By the sine rule on the triangle DEX, we have

$$\frac{\sin A\hat{D}C}{\sin EX} = \frac{\sin\frac{\pi}{3}}{2\sin AD\cos AD}\,. \tag{15.10}$$

By the cotangent rule on the triangle DEX, we have

$$(\sin^2 AD - \cos^2 AD)\cos D\hat{E}X = 2\cot XE\sin AD\cos AD - \cot A\hat{D}C\sin D\hat{E}X\,. \tag{15.11}$$

Since $X\hat{E}F = \frac{\pi}{2} - D\hat{E}X$, by the cotangent rule on the triangle EFX, we have

$$\cos EX\sin D\hat{E}X = -\cot\frac{\pi}{3}\cos D\hat{E}X\,. \tag{15.12}$$

15.5.2. *Conclusion via Gröbner basis.* The statements that follow are supported by the python script `scripts/almost_cube_computation.py` in [13] (the output of the script can be found in `scripts/almost_cube_computation_output.txt`), whose execution takes a few seconds.

Observe that the three equations Equation (15.10), Equation (15.11), and Equation (15.12) can be reduced to a polynomial system in the sines and cosines of $A\hat{D}C$, EX, AD, $D\hat{E}X$. Through the computation of a Gröbner basis, this system is shown to imply

$$xy(x-1)(x+1)(y-1)(y+1)(4x^2y^2 - 4x^2 + 1)(4x^2y^2 - 4x^2 - 4y^2 + 3) = 0, \tag{15.13}$$

where $x = \cos AD$ and $y = \cos A\hat{D}C$. Thanks to Lemma 15.7(i), we know that $AD \approx \frac{\lambda_T}{2}$ and $A\hat{D}C \approx \frac{2\pi}{3}$. Therefore, the only factor of Equation (15.13) that can be zero is

$$4x^2y^2 - 4x^2 + 1 = -4\cos^2 AD\sin^2 A\hat{D}C + 1 = -(2\cos AD\sin A\hat{D}C - 1)(2\cos AD\sin A\hat{D}C + 1)$$

Using again Lemma 15.7(i), we deduce that

$$2\cos AD\sin A\hat{D}C = 1\,. \tag{15.14}$$

Now, observe that the four equations Equation (15.7), Equation (15.8), Equation (15.9), Equation (15.14) can be reduced to a polynomial system in the sines and cosines of AD, $A\hat{D}C$, $B\hat{A}C$, AC. Through the computation of a Gröbner basis, this system is shown to imply $\cos B\hat{A}C = 0$. This implies that the quadrilateral 2630 is equilateral, which is sufficient to conclude thanks to Lemma 15.7(vii).

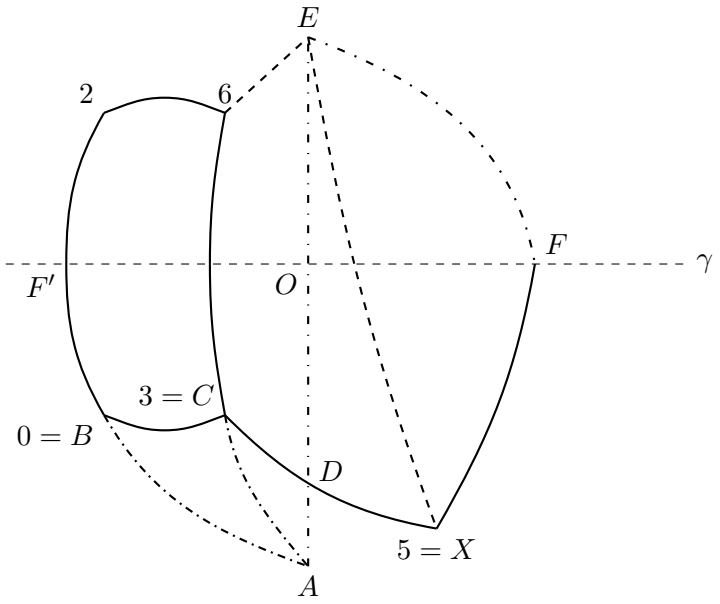


FIGURE 25. The picture referenced in the concluding argument for Proposition 15.6.

15.6. **920 is a thin fish and 827 is a fat fish.** In this section we conclude the proof of Theorem 15.3 analyzing the two remaining marked spherical graphs, those numbered 920 and 827 in the list of Proposition 8.5. The statement is summarized in the following proposition

Proposition 15.8 (Classification of thin and fat fish)**.** *The marked spherical graph 920 is only compatible with once-derived nets of the thin fish family, while the marked spherical graph 827 is only compatible with once-derived nets of the fat fish family.*

Proof. We examine first the marked spherical graph 920 and we refer to the schematic picture in Figure 26. Observe that the the triangle 031 has angles $\frac{2\pi}{3}$, $\frac{2\pi}{3}$ and $\frac{4\pi}{3}$. Note moreover that the two edges 31 and 30 must be shorter than π because they are in common with a convex quadrilateral. From Theorem 6.6 the lengths of the edges of the triangle are uniquely determined. Note moreover that the edges 31 and 30 are congruent. This in particular implies that the convex quadrilateral Q with vertices 3210 is equilateral. This determines completely the metric of the net and it can be easily checked that it is the metric of the once-derived net which one obtains from any member of the thin fish family for which both pins have degree 1.

We analyze next the marked spherical graph 827 and refer to the schematic picture in Figure 27. Observe that the vertices 2 and 3 are connected by two edges. These two edges must therefore be long exactly π, while the vertices 2 and 3 are antipodal. The edges connecting 2 and 3 form therefore a convex biangle $\mathcal{B}$ of opening $\frac{2\pi}{3}$. Observe in particular that the edges 12 and 30 must lie on the bisector γ of the angles of $\mathcal{B}$. If we remove the two edges of the biangle and add the half equator $\gamma \cap \mathcal{B}$, the resulting net induces the marked spherical graph 920 and therefore we can follow the argument above to prove that the only nets for which the once-derivation procedure would produce the net 827 are the members of the fat fish family for which both pins have degree 2. □

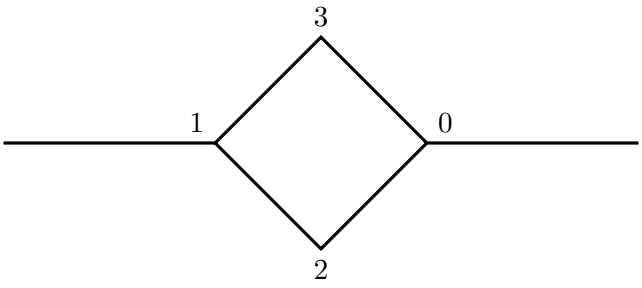


FIGURE 26. A picture of the marked spherical graph 920 for a once-derived net $\mathcal{N}'$. The quadrilateral 1302 is convex.

.

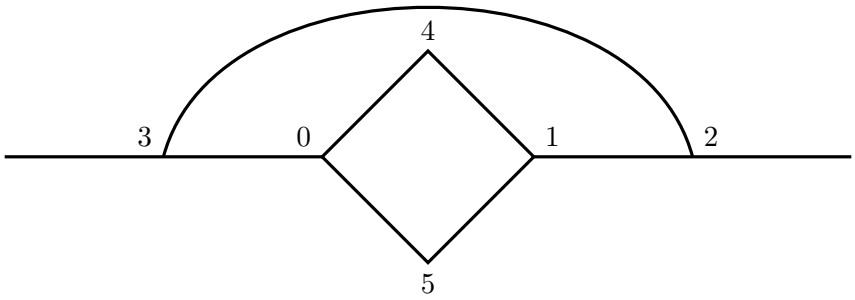


FIGURE 27. A picture of the marked spherical graph 827 for a once-derived net $\mathcal{N}'$. The quadrilateral 0415 is convex and the two vertices 2 and 3 are joined by two edges and must, therefore, be antipodal points.

.

15.7. **The marked spherical graph 736.** In this section we prove Proposition 15.9. The proof is computer assisted: first we use the computer program explained in Section 13 to show that the lengths of *some* specific edges are close to some values; then synthetic arguments, helped by some algebraic computations performed with a computer algebra software, will finish the proof. We recall also that, as it is explained in Section 13, with the additional information gained from Proposition 15.9 an additional run of the code described Section 13 will then determine the lengths of *all* the edges of any valid metric realization, up to a small error.

Proposition 15.9. *Let $\tilde{\mathcal{N}}$ be the link of a minimizing cone as in Theorem 1.3 and assume that the once-derived marked spherical graph associated to it is numbered 736, as represented in Figure 28. Then the edges* 56 *and* 72 *are congruent.*

Proof. We denote by $\mathcal{N}'$ the once-derived net obtained from $\tilde{\mathcal{N}}$.

Step 1. We denote by λ_Q the length of the edges of a spherical equilateral quadrilateral with angles $\frac{2\pi}{3}$. Consider a pentagon with angles $\frac{2\pi}{3}$ and vertices labeled clockwise as A, B, C, D, and E. If the length of DE and BC is λ_Q, Proposition 6.8 shows that the lengths of the remaining three edges are determined and that the sides AE and AB are congruent. We

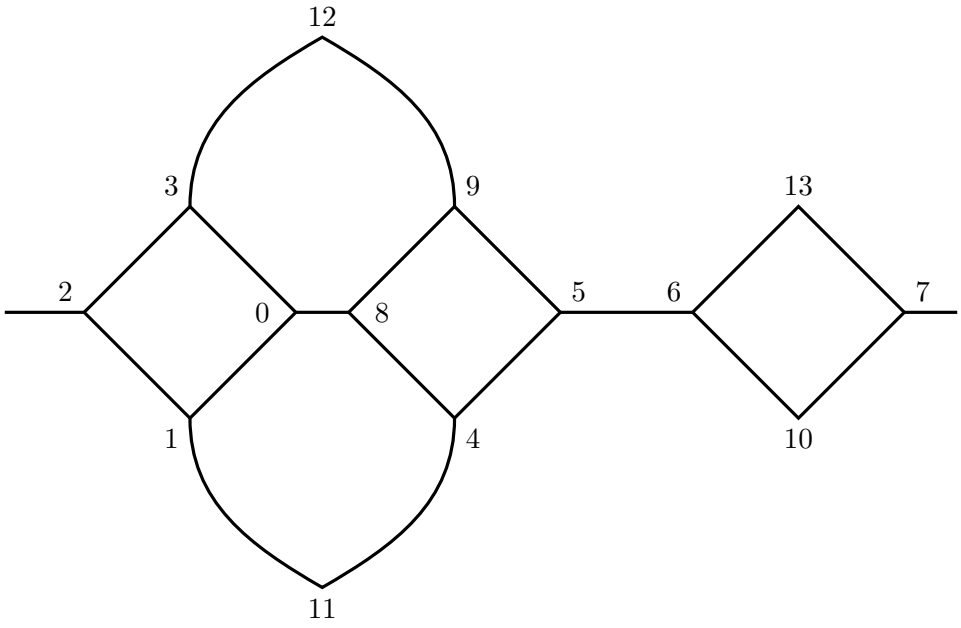


FIGURE 28. A picture of the marked spherical graph 736 for a derived-once net $\mathcal{N}$. Thanks to the algorithm described in Section 13.1 we have an *almost* rigidity for all edges apart from 56 and 72.

denote by λ_B the length of CD and by λ_U the length of AE and AB. Upon setting $\boldsymbol{\delta} = 2{\cdot}10^{-4}$, we use the algorithm explained in Section 13.1 (cf. `embed_or_exclude.cpp` in [13]) to prove:

(a) The lengths of the twelve edges 32, 30, 01, 12, 89, 95, 54, 48, 6(13), (13)7, 7(10), and (10)6 belong to the interval $[\lambda_Q - \boldsymbol{\delta}, \lambda_Q + \boldsymbol{\delta}]$.
(b) The lengths of the four edges 3(12), (12)9, 4(11), and (11)1 belong to the interval $[\lambda_H - \boldsymbol{\delta}, \lambda_H + \boldsymbol{\delta}]$;
(c) The length of the edge 08 belongs to $[\lambda_B - \boldsymbol{\delta}, \lambda_B + \boldsymbol{\delta}]$.

We also record a simple consequence of the above estimates and the smallness of $\boldsymbol{\delta}$:

(d) The distance between the vertices 5 and 2 is strictly less than π and the unique shortest geodesic segment joining them does not intersect the 4 convex faces 1032, 8459, 089(12)3, 084(11)1.

This happens because the diagonal 20 and 85 would be aligned with the edge 08 if $\boldsymbol{\delta}$ would be zero, making the union of these three segments a single geodesic arc of length $\lambda_B + 2\lambda_T > \pi$ (here λ_T denotes the length of the diagonal of the equilateral equiangular quadrilateral with angles $\frac{2\pi}{3}$, which equals the length of the sides of the equiangular triangle with angles $\frac{2\pi}{3}$).

Step 2. We will use geometric considerations and algebraic computations to show that

(e) 89, 45, 01, and 23 are congruent;
(f) 95, 84, 30, and 21 are congruent;
(g) 3(12) and 4(11) are congruent;
(h) 9(12) and 1(11) are congruent.

In fact, it suffices to show that 30 is congruent to 84, as the remaining congruences are then due to the following facts, recorded in Proposition 6.8: opposite edges in an equiangular

quadrilateral are congruent; a quadrilateral with internal angles all equal to $\frac{2\pi}{3}$ is determined by the length of one edge; a pentagon with internal angles all equal to $\frac{2\pi}{3}$ is determined by the lengths of two consecutive edges.

In order to show our claim, we denote by x, y, and z the values of $\tan\frac{\lambda}{2}$ when λ is the length of, respectively, 30, 08, and 95. Recalling Proposition 6.8, $\tan\frac{\lambda}{2}$ equals then $\frac{1}{2x}$ and $\frac{1}{2z}$ if λ is the length of, respectively, 01, and 84. The congruence of 30 and 84 is thus equivalent to

$$2xz = 1\,. \tag{15.15}$$

Using the formulae in Proposition 6.8 for the edges of the pentagon, some classical trigonometric identities, and computing the Gröbner basis of an algebraic system we deduce the following lemma (see `736_computation.py` and `736.sing` in [13] for the details).

Lemma 15.10. *Let a, b, c be the length of three consecutive (in this order) sides of a convex spherical pentagon. Then $Q(\tan\frac{a}{2}, \tan\frac{b}{2}, \tan\frac{c}{2}) = 0$, where Q is the polynomial*

$$Q(x,y,z) = x^2y^2z^2 - 2y^2(x^2+z^2) + x^2z^2 - 2(x^2+z^2) - 6xy^2z + 6y(x+z) + 6xz + y^2 - 2\,.$$

Applying this lemma to the two pentagons 089(12)3 and 084(11)1 we infer that $Q(x,y,z) = 0$ and $Q(\frac{1}{2x}, y, \frac{1}{2z}) = 0$. The second identity can also be recast as a polynomial equation and this enables us to eliminate the variable y and compute the resultant polynomial equation $R(x,z) = 0$ in x and z. The polynomial R can be factored as

$$R(x,z) = 9(2xz-1)^2 S(x,z)\,,$$

where $S(x,z)$ is a polynomial of degree 12 with integer coefficients (see again `736_computation.py` in [13] for the details). Either $2xz - 1$, which is what we want, or the factor $S(x,z)$ vanishes. Thanks to (a), we know that x, z are (quantitatively) close to $\tan\frac{\lambda_Q}{2} = \frac{1}{\sqrt{2}}$. We have $S(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}) = -1944$ and furthermore $S(x', z') \neq 0$ whenever $|x' - \frac{1}{\sqrt{2}}|, |z' - \frac{1}{\sqrt{2}}| < 0.05$. Thus, thanks to (a), we deduce that $S(x,z) \neq 0$ and therefore $2xz = 1$ as desired.

Step 3. We prove that either the conclusion of the Proposition holds, or

(i) The points 5, 6, 7, and 2 are all on the same equator, and the quadrilateral (13)7(10)6 is equilateral.

Consider the diagonal of the quadrilateral joining the vertices 6 and 7, which we denote by 67. (i) is equivalent to the edge 67 forming a flat angle with 56. Notice that the angle formed by the polygonal path $5 \to 6 \to 7$ at 6 is congruent to that formed by $2 \to 7 \to 6$ at 7. So either they are both π, and (i) holds, or they are both distinct from π. In the second case, consider the closed polygonal paths $2 \to 3 \to 0 \to 8 \to 9 \to 5 \to 6 \to 7 \to 2$ and $5 \to 4 \to 8 \to 0 \to 1 \to 2 \to 7 \to 6$. Note that they have the same internal angles, all distinct from π, and the first five edges are congruent. If we know that all the edges are strictly shorter than π, we then conclude from Theorem 6.7 that the last three are also congruent, which however implies the desired conclusion that 56 and 72 are congruent.

Note that from (a), (b) and (c) we know that the quadrilateral 6(13)7(10) is close to be equilateral, so the polygonal path $5 \to 6 \to 7 \to 2$ is close to a geodesic arc connecting 5 and 2, which by (d) must be shorter than π. So the fact that the lengths of 56 and 72 are shorter than π is guaranteed.

It remains to show the congruence of the segments 56 and 72 under assumption (i). We begin by proving that

(j) the quadrilaterals 3021 and 8951 are also equilateral.

Consider a new net in which the vertices $6, 10, 13$, and 7 are eliminated and all the edges 56, 6(13), (13)7, 7(10), (10)6, and 72 are replaced by the geodesic segment joining 5 and 2 (and containing 6 and 7). Since the edges 3(12) and 9(12) are both strictly shorter than π, we can apply Lemma 15.4 and conclude that (j) holds.

Given all the information (a) – (j) we are ready to conclude. The three convex quadrilaterals in the net $\mathcal{N}'$ have been proved to be equilateral. In particular, their diagonals 31, 94, and (10)(13) identify three concurrent equators all passing through 12 and 11. We infer that 12 and 11 are the two pins (with degree 3) and thus the three equators must form angles of $\frac{2\pi}{3}$ with each other. We have completely identified the net (which coincides with the Heppes' net $N_6^{(3)}$ in [14]) and the congruence between 56 and 72 follows. □

16. From triangulated competitors to Almgren minimality

In Section 18 we will complete the proof of Theorem 1.3 excluding all the possible nets allowed by Theorem 13.3 and Theorem 15.3 which are not listed in List 1.1. In order to achieve this goal we will produce a database of suitable "triangulated 2-dimensional sets" which save area compared to the respective cones. Since the set of possibilities to exclude is infinite, we need a suitable criterion to reduce the problem to the existence of a finite database: this will be accomplished in Section 17, where we show that the existence of a competitor, which saves a certain amount of area, for some specific cone E guarantees the existence of a competitor saving area for every sufficiently close cone E'. The material contained in this section is instead meant to free us from the task of producing a map F satisfying the conditions of Definition 1.2 from the cone to the competitor. In particular, Theorem 16.7 gives us a set of sufficient conditions, to be checked on the (combinatorial data) of the triangulated sets, which guarantees the existence of a map as in Definition 1.2. The reader who is familiar with the language of simplicial complexes can skip directly to the statement of Theorem 16.7.

16.1. **Preliminaries on simplicial complexes.** We start here by detailing our description of the "triangulated sets", which are simply images, under piecewise linear maps, of abstract finite 2-dimensional simplicial complexes.

Definition 16.1 (Simplicial complex)**.** An abstract finite simplicial complex Σ is a collection of subsets of a finite set $\mathcal{V}(\Sigma) = \{V_1, \dots, V_N\}$ with the properties that:

- All singletons are in Σ;
- if $S \in \Sigma$, then all the subsets of S belong to Σ.

The singletons will be called vertices, while any element $S \in \Sigma$ of cardinality $k+1$ will be called a k-dimensional face of Σ, and any simplex which is a subset of S will be called a face of S. A face of Σ is called maximal if it is not a face of a face of Σ of higher dimension. Σ is said to be k-dimensional if it contains at least a k-dimensional face and no $(k+1)$-dimensional faces.

For one-dimensional and two-dimensional faces it is common to use the terms *edges* and *triangles*. In particular, since we will only deal with one-dimensional and two-dimensional complexes, we will only encounter vertices, edges, and triangles.

Definition 16.2 (Connected simplicial complex)**.** The complex is connected if for each pair of vertices A and B there is a collection of vertices $A = V_0, V_1, \dots, V_L = B$ such that $[V_i, V_{i+1}]$ is an edge of the complex for every $0 \le i \le L-1$.

If a face S of Σ consists of the vertices $A_0, \dots, A_k$, we will denote it by $[A_0, A_1, \dots, A_k]$.

Definition 16.3 (Geometric realization)**.** For a simplicial complex Σ, we define its geometric realization $|\Sigma|$ as the subset of $\mathbb{R}^{\mathcal{V}(\Sigma)}$ defined as

$$|\Sigma| := \bigcup_{S\in\Sigma} \operatorname{conv}(\{e_V : V \in S\}) .$$

We equip $|\Sigma|$ and all its subcomplexes with the restriction of the Euclidean metric of $\mathbb{R}^{\mathcal{V}(\Sigma)}$.

For any face $S \in \Sigma$, we denote by $|S|$ the subspace of $|\Sigma|$ given by $\operatorname{conv}(\{e_V : V \in S\})$.

We remark that the choice of the metric on $|\Sigma|$ is not particularly important: we will only use maps which are Lipschitz and thus any other biLipschitz equivalent metric (e.g. the path metric as subset of $\mathbb{R}^{\mathcal{V}(\Sigma)}$) would serve our purposes.

Definition 16.4 (Piecewise linear map)**.** A piecewise linear map Φ from a simplicial complex Σ is a map $\Phi : \mathcal{V}(\Sigma) \to \mathbb{R}^N$ for some N. Given a face $S = [V_0, \dots, V_k]$ we define its image $\Phi(S)$ as the subset of $\mathbb{R}^N$ consisting of the convex combination of the points $\Phi(V_i)$. $\Phi(\Sigma)$ will denote the union of all $\Phi(S)$ for $S \in \Sigma$.

We can also interpret Φ as a map from $|\Sigma|$ into $\mathbb{R}^N$ by extending linearly the map $\Phi(e_V) := \Phi(V)$. With a slight abuse of notation, we will use the same symbol Φ for the three maps on the domains $\mathcal{V}(\Sigma)$, Σ, $|\Sigma|$.

For any face S, we denote by $\Phi^\circ(S)$ the image of the *interior* of S. The latter consists therefore of the *strict* convex combination of the points $\Phi(V_i)$, i.e.,

$$\Phi^\circ(S) := \{\lambda_0\Phi(V_0) + \dots + \lambda_k\Phi(V_k) : \lambda_0 + \dots + \lambda_k = 1 \text{ and } \lambda_0, \lambda_1, \dots, \lambda_k > 0\} \subseteq \mathbb{R}^N .$$

Definition 16.5 (Piecewise linear embedding)**.** A piecewise linear map Φ from a simplicial complex Σ is called an *embedding* if any of the following equivalent conditions holds:

- for every pair of faces $S, S' \in \Sigma$,
$$\Phi(S) \cap \Phi(S') = \Phi(S \cap S'),$$
and, for every k-dimensional face $S \in \Sigma$, the set $\Phi(S)$ is a k-dimensional simplex;
- $$\Phi^\circ(S) \cap \Phi^\circ(S') = \emptyset$$
whenever $S, S' \in \Sigma$ are distinct faces;
- $\Phi : |\Sigma| \to \mathbb{R}^N$ is injective.

If $K \subset \mathbb{R}^N$ and $\Phi(\Sigma) \subset K$, we will use the shorthand notation $\Phi : \Sigma \to K$. Note that, when K is convex, we only need to check that every vertex of Σ is mapped into K.

The database of competitors will consist of (computer-generated) images of 2-dimensional simplicial complexes Σ through piecewise linear maps Φ into $\mathbb{R}^3$. Finally, we will need the concept of "collapse" of an abstract simplicial complex Σ onto a subcomplex $\tilde{\Sigma}$.

Definition 16.6 (Collapse)**.** Consider a simplicial complex Σ. A subcomplex $\tilde{\Sigma}$ of Σ is obtained from the elementary collapse of a cell if there is a k-dimensional simplex $S \in \Sigma$ with the following properties:

- S is maximal;
- there is a $(k-1)$-dimensional face F of S which is free, namely it is not a face of any other simplex of Σ;
- $\tilde{\Sigma}$ is obtained from Σ after removing S and F from Σ.

We say that Σ can be collapsed on a subcomplex $\tilde\Sigma$ if there is a sequence of subcomplexes $\tilde\Sigma = \Sigma_N \subset \Sigma_{N-1} \subset \ldots \subset \Sigma_0 = \Sigma$ with the property that, for every $i \geq 1$, Σ_i is obtained from the elementary collapse of a cell of Σ_{i-1}.

16.2. **Existence of Almgren maps.** We are finally ready to state the main theorem of this section, which will enable us to produce a "competitor map" in the sense of Almgren's minimality (cf. Definition 1.2) from images of 2-dimensional simplicial complexes through piecewise linear maps.

Theorem 16.7. *Let $\mathcal{N} \subset \partial B_1$ be a geodesic net (not necessarily stationary) with two antipodal pins P and $-P$ on the line ℓ, E be the cone over $\mathcal{N}$, $K \subset B_1$ be a closed convex polyhedron containing the origin in its interior, Σ be a connected abstract 2-dimensional simplicial complex and $\Phi : \Sigma \to \mathbb{R}^3$ a piecewise linear map. Assume in addition that*

(i) $\Phi(\Sigma) \cap \ell = K \cap \ell$ and there is a one-dimensional subcomplex $\Gamma \subset \Sigma$ such that Φ is an embedding on Γ and $\Phi(\Gamma) = (\ell \cap K) \cup (E \cap \partial K)$;
(ii) There is a subcomplex L of Γ that consists of all faces F of Σ such that[6] $\Phi^\circ(F) \cap \ell \neq \emptyset$;
(iii) Σ can be collapsed on the subcomplex L.

Then for every $\varepsilon > 0$ there is a Lipschitz map $F : \mathbb{R}^3 \to \mathbb{R}^3$ such that

(a) If $W := \{x : F(x) \neq x\}$ then $\overline{W \cup F(W)} \subset B_1 \setminus \ell$;
(b) the following inequality holds:

$$\mathcal{H}^2(F(E \cap B_1)) \leq \mathcal{H}^2((B_1 \setminus K) \cap E) + \mathcal{H}^2(\Phi(\Sigma)) + \varepsilon\,. \tag{16.1}$$

The rest of this section is devoted to proving Theorem 16.7. Hence from now on we fix Σ, Φ, and ε as in the statement and our goal is to produce the map F. The proof consists of three parts. First, we modify the pair (Σ, Φ) near the axial subcomplex L. The resulting complexes retain the collapse property, their images gain only a small amount of area, and provide more control on the combinatorial/geometric structure in a neighborhood of L. This is accomplished in Section 16.3. Second, in Section 16.4 we show how collapsibility implies the existence of a controlled homotopy from a complex onto a subcomplex it can be collapsed into. We construct this homotopy for the simplicial complex built in Section 16.3 and use it, after some nontrivial modification, to define the map F required in the statement of Theorem 16.7.

16.3. **New simplicial complexes and new piecewise linear maps.** In this subsection we modify the simplicial complex Σ and the map Φ in two steps. Let us denote by Γ_{bd} the subcomplex of those vertices and edges of Γ which are mapped by Φ into $E \cap \partial K$. Let Q_+, Q_- be the two "pins", i.e., the vertices belonging to both L and Γ_{bd}.

We produce two new simplicial complexes (Σ^1, Φ^1) and (Σ^2, Φ^2).

(1) (Σ^1, Φ^1) is obtained from (Σ, Φ) as a subdivision that guarantees that all edges with one vertex on the line are short.
(2) (Σ^2, Φ^2) is obtained from (Σ^1, Φ^1) adding a thin collar along the boundary Γ^1_{bd} to have a better control of the combinatorial structure around the two pins.

First perturbation: refinement to Σ^1 and Φ^1. The construction depends on a small parameter $\eta > 0$ which will be chosen at the end. Consider the cylinder $B_\eta(\ell) = \{p : \mathrm{dist}(p, \ell) < \eta\}$, namely the tubular neighborhood of ℓ of radius η. Notice that $\partial B_\eta(\ell)$ consists of the points at distance exactly η from ℓ.

[6] Equivalently $|L| = \Phi^{-1}(\ell)$, considering Φ as a map on $|\Sigma|$.

Consider the set $\mathcal{E}_\eta$ which contains the edges σ of Σ such that $\Phi(\sigma)$ intersects $\partial B_\eta(\ell)$. The set $\mathcal{E}_\eta$ is the same for all sufficiently small values of η and contains only edges with exactly one vertex in L. In fact consider the set $\mathcal{E}_L$ of edges with exactly one vertex in L and one vertex not in L. Any other edge is either an edge of L, and hence is mapped in $\ell \subset B_\eta(\ell)$ or an edge σ with $\Phi(\sigma) \cap \ell = \emptyset$. Since there are finitely many edges, it suffices to choose η sufficiently small to ensure that $\mathcal{E}_\eta = \mathcal{E}_L$. Furthermore, under the latter assumption the intersection $\Phi(\sigma) \cap \partial B_\eta(\ell)$ consists of a single point which belongs to $\Phi^\circ(\sigma)$.

To proceed, let us introduce a standard subdivision for simplicial complexes.

Definition 16.8 (Stellar Subdivision)**.** Consider a simplicial complex Σ and a face F of Σ. The stellar subdivision of Σ at F is the new simplicial complex Σ' built as follows:

- The vertices of Σ' are the vertices of Σ and an additional new vertex v_F (which is, intuitively, *inside* F).
- The faces of Σ' are the faces of Σ which are not a superset of F and the following additional family of faces (all containing the new vertex v_F):

$$\left\{ \{v_F\} \sqcup B \setminus A : \emptyset \neq A \subseteq F \subseteq B \in \Sigma \right\}.$$

Let us build Σ^1 by performing successive stellar subdivisions on the edges σ of Σ such that $\Phi(\sigma)$ intersects $\partial B_\eta(\ell)$. The order does not matter for our purposes (even though the actual resulting Σ^1 depends on the order). Observe that, once the order is specified independently of η, for every sufficiently η the abstract complex Σ^1 is independent of η.

Let us build the map Φ^1 so that it coincides with Φ on the shared vertices, while for each new vertex v_σ, where $\sigma \in \mathcal{E}_\eta$, we set $\Phi^1(v_\sigma)$ to be the only element of the singleton $\Phi(\sigma) \cap \partial B_\eta(\ell)$.

Observe that a stellar subdivision of a simplicial complex acts on any subcomplex as a stellar subdivision (or as an empty operation). Therefore, we can define canonically $L^1, \Gamma^1, \Gamma^1_{bd}$ as the subcomplexes of Σ^1 corresponding to L, Γ, Γ_{bd} respectively.

Then the following lemma holds.

Lemma 16.9. *The pair (Σ^1, Φ^1) satisfies all the properties satisfied by (Σ, Φ) in the assumptions of Theorem 16.7, the image $\Phi^1(\Sigma^1)$ coincides with the image $\Phi(\Sigma)$, and:*

- *The number of edges of Σ^1 adjacent to a vertex in L^1 does not depend on η (if η is sufficiently small).*
- *Every vertex of Σ^1 that does not belong to L^1 but is adjacent to a vertex of L^1 is mapped by Φ^1 into $\partial B_\eta(\ell)$.*

Proof. We use the classical fact that simplicial collapses are preserved under stellar subdivision: if $K \searrow K_0$ and σK is obtained from K by a finite sequence of stellar subdivisions, then

$$\sigma K \searrow \sigma K_0;$$

see [28, §4, Theorem 4, p. 252]. In the present case all the subdivided edges belong to $\mathcal{E}_L = \mathcal{E}_\eta$ and hence are not faces of L. Therefore the induced subdivision of L is trivial, and we conclude that

$$\Sigma^1 \searrow L^1 = L.$$

The remaining claims are obvious given the above discussion. Indeed, since each new vertex is placed on the image of the edge being subdivided and Φ^1 is affine on the resulting simplices,

we have

$$\Phi^1(\Sigma^1) = \Phi(\Sigma).$$

The induced subdivision of Γ preserves the embedding property, and the natural identification of $|\Sigma^1|$ with $|\Sigma|$ intertwines Φ^1 and Φ. Hence the characterization of L^1 as the inverse image of ℓ is also preserved. □

Second perturbation: Σ^2 and Φ^2. Σ^2 is constructed as follows. First of all, it contains all the vertices, segments, and triangles of Σ^1.

Then, for every $V \in \Gamma^1_{bd}$ adjacent (in Γ^1_{bd}) to Q_+ or Q_-, we add a vertex $A(V)$ and the following faces (these additions are depicted in Figure 29):

- The segment $[V, A(V)]$;
- For each V' so that $[V, V']$ is an edge in Γ^1_{bd}, we add the segment $[A(V), V']$ and the triangle $[A(V), V, V']$ (in particular we add the triangle $[A(V), V, Q]$ and the segment $[A(V), Q]$).

We let:

- $L^2 = L^1$;
- Γ^2_{bd} be the 1-dimensional simplicial complex of Σ^2 which coincides with Γ^1_{bd} apart from the fact that we replace V with $A(V)$ for all V adjacent (in Γ^1_{bd}) to Q_+ or Q_-;
- $\Gamma^2 = \Gamma^2_{bd} \cup L^2$.

For $Q \in \{Q_+, Q_-\}$, let $N(Q)$ be the set of vertices $U \in \Gamma^2_{bd}$ adjacent (in Γ^2_{bd}) to Q. Observe that $N(Q)$ coincides with the set of vertices $A(V)$ for all V adjacent (in Γ^1_{bd}) to Q.

The construction of Σ^2 guarantees the validity of the following collapsibility result, which is the whole purpose of introducing Σ^2.

Lemma 16.10. *Σ^2 can be collapsed onto L^2 with a sequence of collapses with the following property. For any $U \in N(Q)$ with $Q \in \{Q_+, Q_-\}$, the edge $[Q, U]$ is collapsed onto Q.*

Proof. We describe the sequence of collapses explicitly.

For all $A(V) \in N(Q)$, with $Q \in \{Q_+, Q_-\}$, we perform the following collapses (in this order):

(1) For all edges $[V, V']$ in Γ^1_{bd} with $V' \notin \{Q_+, Q_-\}$, collapse the triangle $[A(V), V, V']$ onto $[A(V), V]$ and $[V, V']$.
(2) Collapse the triangle $[A(V), Q, V]$ on $[A(V), Q]$ and $[Q, V]$;
(3) Collapse the segment $[A(V), Q]$ onto Q (as required in the statement).

Using the definition of collapse pairs, we thus perform the following collapses:

$$\big([A(V), V, V'], [A(V), V']\big), \qquad \big([A(V), Q, V], [A(V), V]\big), \qquad \big([A(V), Q], A(V)\big).$$

At this point, we have collapsed Σ^2 on Σ^1, which we already know by Lemma 16.9 is collapsible on $L^1 = L^2$. □

The map $\Phi^2 : \Sigma^2 \to \mathbb{R}^3$ is defined as:

- $\Phi^2(V) = 0$ for any $V \in L^2$,
- $\Phi^2(V) = \Phi^1(V)$ for any $V \in \Sigma^1 \setminus L^2$,
- $\Phi^2(A(V)) = \Phi^1(V)$ for all $A(V) \in N(Q)$ with $Q \in \{Q_+, Q_-\}$,

Let us record in the following statement some properties of the image of Σ^2 via Φ^2.

Lemma 16.11. *We have:*

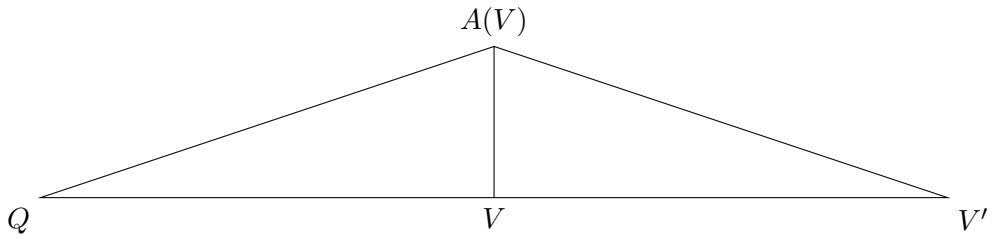


FIGURE 29. The vertices, edges, and triangles added to Σ^1 to form the simplicial complex Σ^2. Here Q represents either Q_+ or Q_- and belongs to both L^2 and Γ^2_{bd}. The edge $[Q,V]$ belongs to Γ^1_{bd}, but V does not belong to Γ^2_{bd} as it is replaced by $A(V)$. The edge $[V,V']$ belongs to Γ^1_{bd} and V' belongs to both Γ^1_{bd} and Γ^2_{bd}.

(1) *For all $U \in N(Q)$, where $Q \in \{Q_+, Q_-\}$, the point $\Phi^2(U)$ is at distance exactly η from ℓ.*
(2) *Let $\tilde{\Gamma}^2_{bd}$ be the subcomplex obtained by removing Q_+ and Q_- (and all the edges adjacent to them) from Γ^2_{bd}. Φ^2 is an embedding on $\tilde{\Gamma}^2_{bd}$ with image contained in $E \cap \partial K$.*
(3) *We have the decomposition*

$$E \cap \partial K = \Phi^2(\tilde{\Gamma}^2_{bd}) \cup \bigcup_{Q \in \{Q_-, Q_+\},\, V \in N(Q)} [\Phi^1(Q), \Phi^2(V)]\,.$$

(4) $\Phi^2(L^2) = \{0\}$.
(5) $(\Phi^2)^{-1}(\ell) = |L^2|$; *equivalently L^2 consists of all faces F of Σ^2 such that $(\Phi^2)^\circ(F) \cap \ell \neq \emptyset$.*
(6) *There is a constant $C > 0$ (independent from η) such that*

$$\mathcal{H}^2(\Phi^2(\Sigma^2)) \leq \mathcal{H}^2(\Phi(\Sigma)) + C\eta.$$

Proof. Parts (1), (2), (3), and (4) are immediate consequences of the construction of Φ^2.

Now, let us show (5). Consider a face F such that $(\Phi^2)^\circ(F) \cap \ell \neq \emptyset$ and assume, for the sake of contradiction, that F does not belong to L^2.

If $F \subseteq \mathcal{V}(L^2)$, we know by assumption (ii) of Theorem 16.7 that $F \in L^2$. Hence, we can assume that at least one vertex of F does not belong to $\mathcal{V}(L^2)$.

Let $\tilde{F} := F \setminus \mathcal{V}(L^2)$. Since every vertex of L^2 is mapped to the origin and ℓ is a linear subspace, we have $(\Phi^2)^\circ(\tilde{F}) \cap \ell \neq \emptyset$. If $\tilde{F}$ does not contain any vertex in $N(Q)$ (with $Q \in \{Q_+, Q_-\}$), then $\tilde{F}$ is a face of $\Sigma^1 \setminus L^1$, and thanks to Lemma 16.9 we have a contradiction. Otherwise, let $U \in N(Q)$ be a vertex of $\tilde{F}$, and consider $V \in \Sigma^1$ such that $U = A(V)$. Let $\hat{F} := \tilde{F} \cup \{V\} \setminus \{A(V)\}$. Since $\Phi^2(U) = \Phi^2(V)$, $(\Phi^2)^\circ(\tilde{F}) = (\Phi^2)^\circ(\hat{F})$. Moreover, by construction of Σ^2, $\hat{F}$ is a face of Σ^1 and thus we have already proven that $(\Phi^2)^\circ(\hat{F})$ is disjoint from ℓ.

We prove (6). The faces of Σ^2 that were not faces of Σ^1 are mapped into sets with zero area by Φ^2 (because $\Phi^2(A(V)) = \Phi^2(V)$). Moreover, for any face F with no vertex from $\mathcal{V}(L^2)$, we have $\Phi^2(F) = \Phi^1(F)$. Therefore, we have

$$\mathcal{H}^2(\Phi^2(\Sigma^2)) \leq \mathcal{H}^2(\Phi^1(\Sigma^1)) + \sum_{\{F \in \Sigma^1 :\, F \cap \mathcal{V}(L^2) \neq \emptyset\}} \mathcal{H}^2(\Phi^2(F))$$

Thanks to Lemma 16.9, $\mathcal{H}^2(\Phi^1(\Sigma^1)) = \mathcal{H}^2(\Phi(\Sigma))$. Moreover, again thanks to Lemma 16.9, the summation is over a number of faces which is independent of η and, for each face F considered in the summation, $\Phi^2(F) \subseteq \overline{B}_\eta(\ell) \cap K$; in particular $\mathcal{H}^2(\Phi^2(F)) \leq C_0\eta$, where C_0 depends only on K. Thus, the right-hand side can be bounded by $\mathcal{H}^2(\Phi(\Sigma)) + C\eta$, for some constant C independent of η. □

16.4. **Controlled homotopies.** In this section we wish to produce a controlled homotopy which encodes the collapse of a complex onto a subcomplex. We will then apply this trick on the complex Σ^2 (which collapses on L^2) constructed in the previous subsection.

Let us remark that property (iii) of the following statement will never be fully used (it would have sufficed to know the weaker "$H(t,x) = x$ for t sufficiently large") but it is required in a technical step of the inductive proof of Lemma 16.12.

Lemma 16.12. *Assume Σ is a finite connected simplicial complex which can be collapsed onto a subcomplex Δ. Then there is a Lipschitz map $H : [0,\infty) \times |\Sigma| \to |\Sigma|$ such that*

(i) $H(0, |\Sigma|) \subset |\Delta|$;
(ii) $H(t,\cdot)^{-1}(|\Delta|) = |\Delta|$ for all $t > 0$;
(iii) $H(t,x) = x$ for every $t \geq \operatorname{dist}(x, |\Delta|)$.

Furthermore, if $[V, V'] \in \Sigma$ is an edge with $V \in \Delta$ that, during the collapsing sequence, is collapsed onto V, then we may require that $H(t, |[V,V']|) \subseteq |[V,V']|$ for all $t \geq 0$.

Proof. First, we construct the homotopy H explicitly when Δ is obtained from the elementary collapse of a single cell, then we show how to concatenate the construction to prove the full statement.

Controlled homotopy for the elementary collapse. Assume that Δ is obtained from Σ by a single elementary collapse of a maximal simplex $S = [v_0, \dots, v_k]$ with free face $F = [v_1, \dots, v_k]$. So, $|\Sigma| = |\Delta| \cup |S|$.

Let $\lambda_0, \dots, \lambda_k$ denote the barycentric coordinates on S, i.e., for any point $x \in |S|$ we let $\lambda_i(x)$ be the only real numbers such that $\sum_i \lambda_i(x) = 1$ and

$$x = \lambda_0(x) e_{v_0} + \cdots + \lambda_k(x) e_{v_k}\,.$$

Define $\rho : |S| \to [0,1]$ as

$$\rho(x) := \min_{1 \leq i \leq k} \lambda_i(x)$$

and extend it outside to be identically 0. Observe that $|S| \cap |\Delta| = \{x \in |S| : \rho(x) = 0\}$. Set $c := \max\{1, \operatorname{Lip}(\rho)\}$.

We define $H(t,x) = x$ for $x \in |\Delta|$, while for $x \in |S|$ we define $H(t,x)$ as

$$H(t,x) := x + \max\{0, \rho(x) - ct\}\Big[k e_{v_0} - (e_{v_1} + \cdots + e_{v_k})\Big].$$

Since $\rho(x) \leq \lambda_i(x)$ for all $1 \leq i \leq k$, for $t \geq 0$ the expression for $H(t,x)$ yields an element of $|S|$ (i.e., a convex combination of $e_{v_0}, \dots, e_{v_k}$). Moreover, since $\rho(x) = 0$ on $|S| \cap |\Delta|$, the two definitions agree on $|S| \cap |\Delta|$.

For any $x \in |S|$, we have $\rho(x) = \lambda_i(x)$ for some $1 \leq i \leq k$ and thus $\lambda_i(H(0,x)) = 0$ which implies $\rho(H(0,x)) = 0$. It follows that

$$H(0, |\Sigma|) \subset |\Delta|.$$

On the other hand, if $x \notin |\Delta|$ and $t > 0$, then $\rho(x) > 0$ and so

$$\lambda_i\big(H(t,x)\big) = \lambda_i(x) - \max\{0, \rho(x) - ct\} > \lambda_i(x) - \rho(x) \geq 0 \qquad \text{for every } i = 1, \dots, k.$$

Consequently $H(t,x) \notin |\Delta|$, and hence

$$H(t,\cdot)^{-1}(|\Delta|) = |\Delta| \qquad \text{for every } t > 0.$$

Since ρ vanishes on $|S| \cap |\Delta|$, for any $x \in |S|$, we have

$$\rho(x) \leq \mathrm{Lip}\,(\rho)\mathrm{dist}(x,|\Delta|) \leq c\,\mathrm{dist}(x,|\Delta|).$$

Thus, whenever $t \geq \mathrm{dist}(x,\Delta)$, the defining formula gives $H(t,x) = x$.

The Lipschitz continuity of H follows immediately from its definition. This proves the elementary-collapse case.

General case by induction. The general case follows then by induction once we can show that, if there is such a map H_1 for two complexes $\Delta' \subset \Sigma$ and H_2 for $\Delta \subset \Delta'$, then there is a map H for $\Delta \subset \Sigma$. In fact we define $t_0(x) := \mathrm{dist}\,(H_1(0,x),\Delta)$ and define first

$$H'(t,x) := \begin{cases} H_2(t, H_1(0,x)) & \text{if } t < t_0(x) \\ H_1(t - t_0(x), x) & \text{otherwise} \end{cases}$$

and then $H(t,x) = H'(Ct,x)$. We now check that H satisfies all the conditions up to a choice of a suitably large scaling constant C.

Lipschitz continuity in t. We need to estimate $|H'(t,x) - H'(s,x)|$ for $t > s$. Clearly we have

$$|H'(t,x) - H'(s,x)| \leq \mathrm{Lip}\,(H_1)|t-s| \quad \text{if } t > s \geq t_0(x).$$

On the other hand, since

$$H_2(t_0(x), H_1(0,x)) = H_1(0,x) = H'(t_0(x),x)\,, \tag{16.2}$$

we also have

$$|H'(t,x) - H'(s,x)| \leq \mathrm{Lip}\,(H_2)|t-s| \quad \text{for } s < t \leq t_0(x)\,.$$

In turn, the two inequalities above give immediately

$$|H'(t,x) - H'(s,x)| \leq \max\{\mathrm{Lip}\,(H_1), \mathrm{Lip}\,(H_2)\}|t-s| \quad \text{when } s < t_0(x) < t.$$

Lipschitz continuity in x. We need to estimate $|H'(t,x) - H'(t,y)|$ and we can assume, without loss of generality, $t_0(x) \geq t_0(y)$. First of all

$$|H'(t,x) - H'(t,y)| \leq \mathrm{Lip}\,(H_2)\mathrm{Lip}\,(H_1(0,\cdot))|x-y| \qquad \text{if } t \leq t_0(y) \leq t_0(x). \tag{16.3}$$

Since t_0 is a Lipschitz function itself, we also get

$$|H'(t,x) - H'(t,y)| \leq \mathrm{Lip}\,(H_1)(1 + \mathrm{Lip}\,(t_0))|x-y| \qquad \text{if } t \geq t_0(x) \geq t_0(y) \tag{16.4}$$

It remains to consider the situation $t_0(y) < t < t_0(x)$. In this case we estimate

$$\begin{aligned} &|H'(t,x) - H'(t,y)| \\ \leq &|H'(t,x) - H'(t_0(x),x)| + |H'(t_0(x),x) - H'(t_0(x),y)| + |H'(t_0(x),y) - H'(t,y)|\,. \end{aligned}$$

The first and third terms are estimated with $C_0|t - t_0(x)| \leq C_0|t_0(y) - t_0(x)| \leq C_0\mathrm{Lip}\,(t_0)|x-y|$, using the Lipschitz continuity in t of H' from the previous step, and the Lipschitz continuity of t_0. To estimate the second summand we can use (16.4).

We now check the claims (i), (ii), and (iii) in the Lemma.

(i). If $0 < t_0(x)$, then clearly $H(0,x) = H'(0,x) = H_2(0, H_1(0,x)) \in |\Delta|$. If $t_0(x) = 0$, then by definition of t_0 we have $H_1(0,x) \in |\Delta|$ and in particular (16.2) implies $H(0,x) = H'(0,x) = H_1(0,x) \in |\Delta|$.

(ii). We first show $H'(t,\cdot)^{-1}(|\Delta|) \subset |\Delta|$ for positive t. Fix $t > 0$ and $x \in |\Sigma| \setminus |\Delta|$. We need to show that $H'(t,x) \notin |\Delta|$. If $0 < t < t_0(x)$, then $H_1(0,x) \notin |\Delta|$, because $t_0(x) > 0$. Property (ii) for H_2 gives

$$H'(t,x) = H_2(t, H_1(0,x)) \notin |\Delta|.$$

If $t > t_0(x)$, set $s := t - t_0(x) > 0$. Then

$$H'(t,x) = H_1(s,x).$$

If $x \notin |\Delta'|$, property (ii) for H_1 gives

$$H_1(s,x) \notin |\Delta'|,$$

and hence $H'(t,x) \notin |\Delta|$. If instead $x \in |\Delta'| \setminus |\Delta|$, then $\operatorname{dist}(x, |\Delta'|) = 0$, so property (iii) for H_1 implies

$$H_1(s,x) = x \notin |\Delta|.$$

Thus $H'(t,x) \notin |\Delta|$ in either case.

It remains to analyze the case $t = t_0(x) > 0$. Under the latter assumption $H'(t,x) = H_1(0,x)$ and, since $t_0(x) = \operatorname{dist}(H_1(0,x), |\Delta|) > 0$, we have $H_1(0,x) \notin |\Delta|$.

Finally, to show $|\Delta| \subset H'(t,\cdot)^{-1}(|\Delta|)$, note that for $x \in |\Delta|$ and any $t \geq 0$ we have $H_1(t,x) = x$, and hence $H'(t,x) = x$.

(iii). Note that, since $H_1(0,x) = x$ for every $x \in |\Delta'|$, we have

$$t_0(x) \leq \operatorname{Lip}(H_1) \operatorname{dist}(x, |\Delta|). \tag{16.5}$$

We now choose $C \geq \operatorname{Lip}(H_1) + 1$ and show that the claim holds with the latter choice.

Indeed, let $t \geq \operatorname{dist}(x, \Delta)$. In particular (16.5) implies

$$\begin{aligned} Ct - t_0(x) &\geq C\operatorname{dist}(x, |\Delta|) - \operatorname{Lip}(H_1)\operatorname{dist}(x, |\Delta|) \\ &\geq \operatorname{dist}(x, |\Delta|) \geq \operatorname{dist}(x, |\Delta'|), \end{aligned}$$

where the last inequality follows from $\Delta \subset \Delta'$. Note that the last chain of inequalities also implies $Ct \geq t_0(x)$, so that the second branch in the definition of H' applies, namely $H(t,x) = H'(Ct,x) = H_1(Ct - t_0(x), x)$. Property (iii) for H_1 therefore gives

$$H(t,x) = H_1(Ct - t_0(x), x) = x\,.$$

This proves property (iii).

The additional property on the image of edges contained in the statement is a direct consequence of the construction of the homotopy. Indeed

- before the edge is collapsed, it lies in every remaining subcomplex, so all preceding elementary homotopies fix it;
- at the edge collapse, the elementary formula moves it inside itself;
- every subsequent homotopy fixes its surviving endpoint V.

This obviously yields $H(t, |[V,V']|) \subset |[V,V']|$. . □

16.5. **Construction of the map and proof of Theorem 16.7.** We can now construct a map F satisfying the requirements of Theorem 16.7. We will do this in two steps:

- Using Lemma 16.12 on Σ^2, we will build a function $G : E \cap K \to \mathbb{R}^3$ whose image is *almost* $\Phi^2(\Sigma^2)$ and which satisfies all the assumptions necessary to apply Proposition 10.3 apart from the Lipschitz continuity in the origin.

- We will then perturb G into $\tilde{G}$ to obtain the Lipschitz continuity in the origin (while keeping the image $\tilde{G}(E \cap K)$ under control). This allows us to apply Proposition 10.3 and conclude.

Let us recall that $\tilde{\Gamma}^2_{bd}$ is the 1-dimensional simplicial complex obtained by removing from Γ^2_{bd} the two vertices Q_+ and Q_- as well as all the edges adjacent to them. By Lemma 16.11, we know that Φ^2 is an embedding from $\tilde{\Gamma}^2_{bd}$ into a subset Z of $E \cap \partial K$. Let Z' be the union of segments that appear in Lemma 16.11. Let us recall the decomposition

$$E \cap \partial K = Z \cup Z' .$$

For notational simplicity, let $\{P_+, P_-\} := \ell \cap \partial K$ (where P_+ lies on the same ray of ℓ as the pin P) and let $M_\pm := \{\Phi^2(V) : V \in N(Q_\pm)\}$. Observe that

$$Z' = \left(\bigcup_{U \in M_+} [P_+, U] \right) \sqcup \left(\bigcup_{U \in M_-} [P_-, U] \right) .$$

We will define $G : E \cap K \to \mathbb{R}^3$ differently in the cone generated by Z and in the cone generated by Z'.

Construction of G on the cone over Z. Thanks to Lemma 16.10, we can apply Lemma 16.12 on Σ^2 and obtain a Lipschitz function $H : [0,1] \times |\Sigma^2| \to |\Sigma^2|$ such that:

- $H(0, |\Sigma^2|) \subseteq |L^2|$.
- $H(t, \cdot)^{-1}(|L^2|) = |L^2|$ for all $t > 0$.
- $H(1, \cdot)$ is the identity.
- For every $V \in N(Q)$, where $Q \in \{Q_+, Q_-\}$, $H(t, V) \in |[V, Q]|$ for all $t \in [0,1]$.

In fact the Lemma 16.12 yields an homotopy $\tilde{H}$ which satisfies the first, second, and fourth requirement, while in place of the third we know that there is a $T > 0$ such that $\tilde{H}(t, x) = x$ for all $x \in |\Sigma^2|$ and all $t \geq T$. Thus, in order to achieve the third we simply set $H(t, x) := \tilde{H}(Tt, x)$.

Observe that an arbitrary nonzero point in the cone over Z can be expressed (uniquely) as $t\Phi^2(p)$ with $0 < t \leq 1$ and $p \in |\tilde{\Gamma}^2_{bd}|$. We define

$$G(t\Phi^2(p)) := \Phi^2(H(t, p)) .$$

Since $\Phi^2(H(0, p)) = 0$, extending it to $G(0) = 0$ preserves the above identity no matter what p in $|\tilde{\Gamma}^2_{bd}|$ we choose. Since $H(0, |\Sigma^2|) \subseteq |L^2|$ and $\Phi^2(|L^2|) = \{0\}$, the map G is continuous in the origin and maps the origin into the origin. Furthermore, G is locally Lipschitz away from the origin. By construction, we have that the image of G is contained in $\Phi^2(\Sigma^2)$ and thus, thanks to Lemma 16.11-(6),

$$\mathcal{H}^2(G(\{tz : 0 \leq t \leq 1,\, z \in Z\})) \leq \mathcal{H}^2(\Phi(\Sigma)) + C\eta$$

for some constant $C > 0$ independent from η. Moreover, because of Lemma 16.11-(5),

$$G^{-1}(\ell) = G^{-1}(\{0\}) = \{t\Phi^2(p) : H(t, p) \in |L^2|\} = \{t\Phi^2(p) :\ t = 0 \text{ or } p \in |L^2|\} = \{0\} .$$

Finally, for any $U \in M_\pm$ (with $U = \Phi^2(V)$ and $V \in N(Q)$) and for all $0 \leq t \leq 1$, we have

$$G(tU) = G(t\Phi^2(V)) = \Phi^2(H(t, |V|)) \in \Phi^2([V, Q]) = [\Phi^2(V), \Phi^2(Q)] = [U, 0],$$

which implies $G([0, U]) \subseteq [0, U]$.

Construction of G on the cone over Z'. Now, we extend G on the cone over Z' as follows. We define it independently on the cone over $[P_+, U]$ for each $U \in M_+$ (and analogously for P_-). Let $W \in \mathbb{R}^3$ be the midpoint of the segment $[P_+, U]$.

- On the cone over $[P_+, W]$, G is the identity.
- An arbitrary point in the cone over $[W, U]$ can be expressed as $t(\lambda W + \mu U)$ with $0 \leq t \leq 1$ and $\lambda, \mu \geq 0$ so that $\lambda + \mu = 1$. We define

$$G\big(t(\lambda W + \mu U)\big) := \lambda t W + \mu G(tU)\,.$$

The finitely many definitions agree on their common edges and hence give a locally Lipschitz map away from 0.

The fact that $G([0, U]) \subseteq [0, U]$ guarantees that this extension has two crucial properties:

- $G(\{tz' : 0 \leq t \leq 1,\ z' \in [P_+, U]\}) \subseteq \{tz' : 0 \leq t \leq 1,\ z' \in [P_+, U]\}$;
- $G^{-1}(\ell) \cap \{tz' : 0 \leq t \leq 1,\ z' \in [P_+, U]\} \subseteq \ell$.

Thanks to Lemma 16.11-(1), we deduce that

$$G(\{tz' :\ 0 \leq t \leq 1,\ z' \in Z'\}) \subseteq E \cap B_\eta(\ell),$$

The latter set consists of the intersection with the tube $B_\eta(\ell)$ of finitely many planar conical sectors adjacent to the two pin rays and a small region around the origin of radius η. The intersection of the sectors with the tube contribute $O(\eta)$ to the area, while the small region around the origin contributes only $O(\eta^2)$. So the total area is bounded bounded by $C\eta$ for some constant $C > 0$ (independent from η).

Properties of G and construction of $\tilde{G}$. By construction, and thanks to the observations above, we know that the map $G : E \cap K \to \mathbb{R}^3$ is continuous and it is locally Lipschitz away from the origin. Moreover, it satisfies

(i) $G(x) = x$ for $x \in E \cap \partial K$.
(ii) $G(x) = x$ for $x \in C_\theta(\ell) \cap K$ (with θ depending on $\eta > 0$) and $G^{-1}(\ell) \subset \ell$.
(iii) $\mathcal{H}^2(G(E \cap K)) \leq \mathcal{H}^2(\Phi(\Sigma)) + C\eta$ for some constant $C > 0$ independent from η.
(iv) The image of G is a subset of $E \cup \Phi^2(\Sigma^2)$.

These assumptions are *almost* sufficient to directly apply Proposition 10.3. The missing bit is the Lipschitz continuity of G at the origin. We will now use (iv) to perturb G into $\tilde{G}$ in a way that guarantees the desired Lipschitz continuity.

First of all, we note that there is a cone $\mathcal{T} \subseteq \mathbb{R}^3$ with locally finite area such that, for any small enough $r > 0$, $G(E \cap B_r(0)) \subset \mathcal{T}$. Indeed let $\mathscr{S}_0 := \{S \in \Sigma^2 : 0 \in \Phi^2(S)\}$. By Lemma 16.11-(5), every simplex in $\mathscr{S}_0$ has a face contained in L^2, and this face is mapped to 0. Hence, if we set,

$$\mathcal{T} := E \cup \bigcup_{S \in \mathscr{S}_0} \{ty : t \geq 0,\ y \in \Phi^2(S)\}$$

then $\mathcal{T}$ is a finite union of cones of dimension at most two and therefore is itself a cone with locally finite $\mathcal{H}^2$-measure. Since there are only finitely many simplices not belonging to $\mathscr{S}_0$, their images have positive distance from the origin. Hence $\Phi^2(\Sigma^2) \cap B_r \subset \mathcal{T}$ for every sufficiently small $r > 0$.

Choose a small radius $r > 0$ such that $G(E \cap B_r(0)) \subset \mathcal{T} \cap B_\eta(0)$ and so that $r < \operatorname{dist}(0, \partial K)$. We define $\tilde{G} : E \cap K \to \mathbb{R}^3$ as

$$\tilde{G}(x) = \begin{cases} G(x) & \text{if } |x| \geq r, \\ \frac{|x|}{r} G(r \frac{x}{|x|}) & \text{if } 0 < |x| < r \\ 0 & \text{if } x = 0. \end{cases}$$

The map $\tilde{G}$ is Lipschitz. Indeed the restriction of G to $E \cap \partial B_r$ is Lipschitz, since G is locally Lipschitz away from the origin, and since E is a finite union of planar cones, the positively homogeneous extension of this boundary map is Lipschitz on $E \cap B_r$. Clearly $\tilde{G}$ retains property (i). To see that property (ii) is retained note that:

- in the directions where G is the identity, $\tilde{G}$ is also the identity;
- if $\tilde{G}(x) \in \ell$ and $x \neq 0$, then $G(rx/|x|) \in \ell$ and therefore $x \in \ell$.

Thanks to property (iii) and (iv), we have

$$\mathcal{H}^2(\tilde{G}(E \cap K)) \leq \mathcal{H}^2(G(E \cap K)) + \mathcal{H}^2(\mathcal{T} \cap B_\eta(0)) \leq \mathcal{H}^2(\Phi(\Sigma)) + C\eta + C'\eta^2,$$

for some constant $C' > 0$ independent of η.

We are now almost ready to apply Proposition 10.3. First of all we extend it to a map $\widehat{G}$ on E by setting it equal to the identity outside of K. Since $\tilde{G}$ is the identity on $E \cap \partial K$, such extension is Lipschitz. Moreover it is the identity on $C_\theta(\ell)$ and $\widehat{G}^{-1}(\ell) \subset \ell$. Furthermore,

$$\begin{aligned} \mathcal{H}^2\big(\widehat{G}(E \cap B_1)\big) &\leq \mathcal{H}^2\big((B_1 \setminus K) \cap E\big) + \mathcal{H}^2\big(\widetilde{G}(E \cap K)\big) \\ &\leq \mathcal{H}^2\big((B_1 \setminus K) \cap E\big) + \mathcal{H}^2\big(\Phi(\Sigma)\big) + C\eta + C'\eta^2. \end{aligned}$$

Choose $\eta > 0$ sufficiently small that $C\eta + C'\eta^2 < \frac{\varepsilon}{2}$. Applying Proposition 10.3 to $\tilde{G}$ with error $\varepsilon/2$ gives a Lipschitz map $F : \mathbb{R}^3 \to \mathbb{R}^3$ satisfying (a) and (b) in the statement of the theorem.

17. Quantitative stability of competitors

In this section we prove two propositions which, in combination with the theory developed in Section 16, will allow us to use the code described in Section 18 to produce effective competitors for the remaining topologies. In Section 16 we addressed the fact that the code produces competitors which are triangulated surfaces, while Definition 1.2 requires additionally the existence of suitable Lipschitz maps. In this section we address a different issue, namely, that the competitors produced by the code do not exactly match the desired boundary conditions, in other words they "span" a net which is only an approximation of the one which we want to exclude. The first proposition will deal with the fat fish family, while the second one will be used for the other nets.

17.1. **Fat fish.** First, recall Definition 15.1 and Figure 19. In order to parametrize the fat fish family, we denote by ξ the length of the edge P_1V_2 and by ζ the length of the edge V_5V_1 (cf. Definition 15.1). Moreover, without loss of generality, we will assume that P_1V_2 is shorter than (or congruent to) P_2V_4 and that V_5V_1 is shorter than (or congruent to) V_3V_6. Recall (cf. Definition 15.1) that an element of the fat fish family is required to be an embedding of the corresponding spherical graph. In particular $\zeta > 0$ and the edges P_1V_2 and P_2V_4 cannot intersect the two half-equators connecting V_5 and V_6. A direct computation shows that this implies $\xi > 0$. These and other useful pieces of information are summarized in the following lemma.

Lemma 17.1. *The pins of any element $\mathcal{N}$ of the fat fish family have degree one. Moreover there is a rigid motion R of the sphere such that:*

- *the pins of $R(\mathcal{N})$ are the north and south poles $(0,0,1)$ and $(0,0,-1)$;*
- *the edges adjacent to the pins lie in the halfplane $\{x_2 = 0, x_1 \geq 0\}$;*
- *the edge adjacent to $(0,0,1)$ is shorter than (or congruent to) the one adjacent to $(0,0,-1)$ (and its length is ξ);*
- *between the two edges adjacent to the two edges of length π, the one in the half-space $\{x_2 < 0\}$ is no longer than the one lying in $\{x_2 > 0\}$ (and its length is ζ).*

We call this the canonical position *of $\mathcal{N}$. Furthermore:*

(a) Two elements in canonical position always coincide in a neighborhood of the pins;
(b) Given two elements in canonical position, the four vertices forming the convex quadrilateral face coincide if and only if their ξ-coordinates coincide;
(c) Two elements in canonical position are the same net if and only if they have the same coordinates (ξ, ζ).

We are now ready to state the stability proposition.

Proposition 17.2. *Consider two elements $\mathcal{N}$ and $\mathcal{N}'$ of the fat fish family with corresponding cones E and E', and coordinates (ξ, ζ) and (ξ', ζ'). Assume that $|\xi - \xi'| + |\zeta - \zeta'|$ is strictly smaller than the distance of any pin of $\mathcal{N}$ to any point of $\mathcal{N}$ not belonging to an edge adjacent to the pin in question. Let $F : \mathbb{R}^3 \to \mathbb{R}^3$ be a Lipschitz map such that*

(i) F is the identity outside B_1 and on a neighborhood of ℓ;
(ii) $\overline{\{F(x) : x \neq F(x)\}} \cap \ell = \emptyset$.

Then, for every positive ε, there is a Lipschitz map G which satisfies (i) and (ii) and such that

$$\begin{aligned}\mathcal{H}^2(G(E \cap \overline{B}_1)) &\leq \mathcal{H}^2(F(E' \cap \overline{B}_1)) + 4(|\xi - \xi'| + |\zeta - \zeta'|) + 4|\xi - \xi'|\left(\tfrac{2\sqrt{3}}{3} + \sqrt{\tfrac{3-\sqrt{6}}{6}}\right) + \varepsilon \\ &\leq \mathcal{H}^2(F(E' \cap \overline{B}_1)) + 10|\xi - \xi'| + 4|\zeta - \zeta'| + \varepsilon\,.\end{aligned}$$

17.2. **Adjusting triangulated competitors.** In this section we deal with the other cases. However, rather than assuming that we have a "mapped competitor" for the approximating net, we assume that we have a triangulated competitor.

We will need a suitable "controlled change of topology", since the approximating net might have a slightly different topology than the real one.

Definition 17.3. Let $\mathcal{N}$ and $\tilde{\mathcal{N}}$ be two geodesic nets with the same antipodal pins, not necessarily stationary, and let G and $\tilde{G}$ be the corresponding embedded spherical graphs. We say that there is a controlled change of topology from $\tilde{\mathcal{N}}$ to $\mathcal{N}$ if the graph G can be obtained from $\tilde{G}$ by contracting at most two edges, each incident to a different pin.

Note that the contraction eliminates at most one vertex adjacent to each pin of $\tilde{G}$. Such a vertex will be called the "special vertex associated with the pin". We introduce the following maps $\tau : \mathcal{V}(\tilde{G}) \to \mathcal{V}(G)$ and $\iota : \mathcal{V}(G) \to \mathcal{V}(\tilde{G})$:

- τ maps each pin $\tilde{P}$ of $\tilde{G}$ into the corresponding pin P of G and the special vertex associated with $\tilde{P}$ (if it exists) in P;
- After contracting the edges (and identifying the corresponding special vertices with the corresponding pins), τ is an isomorphism of embedded graphs.

- ι maps any vertex V of G which is not a pin into its unique preimage $\tau^{-1}(V)$. For a pin P of G, if $\tau^{-1}(P)$ is a singleton $\{\tilde{P}\}$, then $\iota(P) = \tilde{P}$; otherwise we let $\iota(P) = \tilde{S}$, the special vertex associated with $\tilde{P}$.

Therefore τ is a surjective map, ι an injective map, and $\tau \circ \iota$ is the identity.

We are now ready to state the second proposition.

Proposition 17.4. *Let $\mathcal{N}'$ be a stationary geodesic net with two antipodal pins $\pm P$, let ℓ be the line through the pins, and let E' be its cone. Let $\mathcal{N}$ be a geodesic net (not necessarily stationary) with the same pins such that*

(a) There is a controlled topological change from $\mathcal{N}'$ to $\mathcal{N}$ and we let τ and ι be the corresponding maps in Definition 17.3.
(b) All the edges of $\mathcal{N}$ and all the edges of $\mathcal{N}'$ are shorter than π;
(c) For every pair (U, V) of adjacent vertices of $\mathcal{N}$ which are not pins, there is an open half-space[7] H_+ with $\partial H_+ \supset \ell$ and $U, V, \iota(U), \iota(V) \in H_+$;
(d) No vertex V of $\mathcal{N}$ is adjacent to both pins.

Define, for every edge $\sigma = [U, V]$ of $\mathcal{N}$, the function

$$D(\sigma) := \max\left\{\mathcal{H}^2([U, V, \iota(V)]) + \mathcal{H}^2([U, \iota(U), \iota(V)]), \mathcal{H}^2([V, U, \iota(U)]) + \mathcal{H}^2([V, \iota(V), \iota(U)])\right\}, \tag{17.1}$$

where $[A, B, C]$ denotes the triangle of $\mathbb{R}^3$ with vertices A, B, and C. Let K and K' be the convex hulls of the vertices of $\mathcal{N}$ and $\mathcal{N}'$ respectively, while K_c is the convex hull of the union of their vertices.

We have the following two conclusions:

- *Assume that Σ is a 2-dimensional simplicial complex and $\Phi : \Sigma \to \mathbb{R}^3$ is a piecewise linear map satisfying assumptions (i), (ii), (iii) of Theorem 16.7 with respect to $(\mathcal{N}, E, K)$ (i.e., a valid competitor for E). Then, there is a simplicial complex Σ' and a piecewise linear map $\Phi' : \Sigma' \to \mathbb{R}^3$ satisfying the assumptions (i), (ii), (iii) of Theorem 16.7 with respect to $(\mathcal{N}', E', K_c)$ (i.e., a valid competitor for E') such that*

$$\mathcal{H}^2(\Phi'(\Sigma')) \leq \mathcal{H}^2(\Phi(\Sigma)) + \sum_{\sigma \in \mathcal{E}(\mathcal{N})} D(\sigma)\,. \tag{17.2}$$

- *Assume that Σ' is a 2-dimensional simplicial complex and $\Phi' : \Sigma' \to \mathbb{R}^3$ is a piecewise linear map satisfying assumptions (i), (ii), (iii) of Theorem 16.7 with respect to $(\mathcal{N}', E', K')$ (i.e., a valid competitor for E'). Then, there is a simplicial complex Σ and a piecewise linear map $\Phi : \Sigma \to \mathbb{R}^3$ satisfying the assumptions (i), (ii), (iii) of Theorem 16.7 with respect to $(\mathcal{N}, E, K_c)$ (i.e., a valid competitor for E) such that*

$$\mathcal{H}^2(\Phi(\Sigma)) \leq \mathcal{H}^2(\Phi'(\Sigma')) + \sum_{\sigma \in \mathcal{E}(\mathcal{N})} D(\sigma)\,. \tag{17.3}$$

17.3. **Proof of Proposition 17.2.** Lemma 17.1 has a straightforward proof, which we leave to the reader.

We next come to the proof of Proposition 17.2. Fix two elements of the fat fish family $\mathcal{N}$ and $\mathcal{N}'$. Without loss of generality we can assume that $\mathcal{N}$ and $\mathcal{N}'$ are both in canonical position (this only requires conjugating the map F with isometries, which does not change

[7] An open half-space of $\mathbb{R}^3$ is any set H_+ of the form $H_+ := \{x \in \mathbb{R}^3 : \langle x, v\rangle > 0\}$ for some nonzero vector $v \in \mathbb{R}^3$.

any of the assumptions on F). We will produce a Lipschitz homotopy $\rho : [0,1] \times \mathcal{N} \to \mathbb{S}^2$ with $\rho(1,\mathcal{N}) = \mathcal{N}'$ which satisfies requirement (ii) in Lemma 11.3 and the inequality

$$\int_{\mathbb{S}^2} \#(\rho)^{-1}(\{x\})\, d\mathcal{H}^2(x) \leq 4(|\xi - \xi'| + |\zeta - \zeta'|) + 4|\xi - \xi'|\left(\tfrac{2\sqrt{3}}{3} + \sqrt{\tfrac{3-\sqrt{6}}{6}}\right). \tag{17.4}$$

In order to define the homotopy ρ it is convenient to use the labels V_i for the non-pin vertices of $\mathcal{N}$ and corresponding labels V_i' for the non-pin vertices of $\mathcal{N}'$, where we refer to Figure 19.

The homotopy ρ is obtained by concatenating two homotopies. The first one, ρ^1, adjusts the coordinate ξ to ξ' and it is defined as follows:

- Denote by R_θ the rotation of angle θ around the axis x_2. If x belongs to any edge of $\mathcal{N}$ which is not adjacent to the pins, then $\rho^1(t,x) = R_{t(\xi'-\xi)}(x)$. Note that, by Lemma 17.1, $\rho^1(1,V_i) = V_i'$ for $i \in \{1,2,3,4\}$.
- Let W_2 and W_4 be the two points adjacent to the pins at distance $a = \frac{1}{2}\min\{\xi,\xi'\}$. ρ^1 is then the identity on P_1W_2 and P_2W_4. On W_2V_2 we extend ρ^1 so as to take values on the equator $\mathbb{S}^2 \cap \{x_2 = 0\}$. Likewise on the other edge adjacent to the pin P_2. Moreover we extend it so that it is arclength-affine along the relevant equatorial arc at every t.

Note that $\mathcal{N}'' := \rho^1(1,\mathcal{N})$ is also an element of the fat fish family, with coordinates (ξ',ζ). We now wish to find a second homotopy ρ^2 which maps $\mathcal{N}''$ onto $\mathcal{N}'$. Denote now by V_i'' the vertices $\rho^1(1,V_i)$ and note that $V_i'' = V_i'$ for $i \in \{1,\dots,4\}$. So we define ρ^2 to be the identity on the edges P_1V_2'', P_2V_4'', $V_1''V_2''$, $V_2''V_3''$, $V_3''V_4''$, and $V_4''V_1''$. For points in the two half-equators connecting V_5'' and V_6'' we let S_θ be the rotation of angle θ around the axis x_3 and define $\rho^2(t,x) = S_{t(\zeta'-\zeta)}(x)$. On the edges $V_1''V_5''$ and $V_3''V_6''$ we then extend the map so that it takes values on the equator $\mathbb{S}^2 \cap \{x_3 = 0\}$, for instance we can take it arclength-affine along the relevant equatorial arc at every time t.

Note that the concatenation ρ of the homotopies ρ^1 and ρ^2 satisfies (ii) in Lemma 11.3:

- by construction it is the identity in a neighborhood of the pins;
- by construction the points on the two edges adjacent to the pins which are displaced by the homotopy are anyway all mapped to a positive distance from the pins;
- all the other points, belonging to the other edges, are moved at most at a distance $|\xi - \xi'| + |\zeta - \zeta'|$ from themselves, and hence they all stay at a positive distance from the pins.

Therefore it remains to show (17.4), for which we will invoke the following lemma.

Lemma 17.5. *Denote by O_θ the rotation around an axis r of $\mathbb{R}^3$ of angle θ and let σ be a geodesic segment on the sphere $\mathbb{S}^2$ of length at most π. Consider the map $\tau : [0,1] \times \sigma \to \mathbb{S}^2$ given by $\tau(t,x) = O_{t\theta}(x)$. Then*

$$\int_{\mathbb{S}^2} \#\tau^{-1}(\{x\})\, d\mathcal{H}^2(x) \leq 2|\theta| \sin\left(\tfrac{\mathit{len}(\sigma)}{2}\right). \tag{17.5}$$

We will prove Lemma 17.5 below, while for the moment we assume its validity. Note first that the interpolated portions of the pin-adjacent edges remain in $\mathbb{S}^2 \cap \{x_2 = 0\}$, while the interpolations on $V_1''V_5''$ and $V_3''V_6''$ remain in $\mathbb{S}^2 \cap \{x_3 = 0\}$. Their images are therefore one-dimensional, so their two-dimensional Jacobians vanish almost everywhere. This observation puts us in the position of estimating the left-hand side of (17.4) for all the other edges.

Note that the two edges connecting V_5 and V_6 have length π and thus contribute each $2(|\xi - \xi'| + |\zeta - \zeta'|)$, while the four edges V_1V_2, V_2V_3, V_3V_4, and V_4V_1 have all length λ_Q, which

is the sidelength of a spherical equilateral quadrilateral with angles $\frac{2\pi}{3}$. Since $\sin\frac{\lambda_Q}{2} = \frac{1}{\sqrt{3}}$ we can then estimate

$$\int_{\mathbb{S}^2} \#\rho^{-1}(\{x\})\, d\mathcal{H}^2(x) \le 4(|\xi-\xi'|+|\zeta-\zeta'|)+\tfrac{8}{\sqrt{3}}|\xi-\xi'|+2|\xi-\xi'|\left(\sin\tfrac{\mathrm{len}(V_1V_5)}{2}+\sin\tfrac{\mathrm{len}(V_3V_6)}{2}\right).$$

Now, observe that $\mathrm{len}(V_1V_5)+\mathrm{len}(V_3V_6) = \pi-\lambda_T$, where $\lambda_T = \arccos(-\frac{1}{3})$ is the sidelength of a spherical equilateral triangle with angles $\frac{2\pi}{3}$. We can thus estimate

$$\sin\tfrac{\mathrm{len}(V_1V_5)}{2}+\sin\tfrac{\mathrm{len}(V_3V_6)}{2} \le \max_{0\le x\le \pi-\lambda_T}\left(\sin\tfrac{x}{2}+\sin\tfrac{\pi-\lambda_T-x}{2}\right).$$

The latter maximum can be easily computed to be at $x = \frac{\pi-\lambda_T}{2}$, thus giving the bound

$$\sin\tfrac{\mathrm{len}(V_1V_5)}{2}+\sin\tfrac{\mathrm{len}(V_3V_6)}{2} \le 2\sin\tfrac{\pi-\lambda_T}{4} = 2\sqrt{\tfrac{3-\sqrt{6}}{6}}.$$

Putting all these inequalities together gives us the estimate (17.4).

We now apply Lemma 11.3 to F and the homotopy ρ, and using (17.4), for every $\varepsilon>0$ we obtain a Lipschitz map G satisfying conditions (i) and (ii) of Proposition 17.2 such that

$$\begin{aligned}\mathcal{H}^2\big(G(E\cap\overline{B}_1)\big) &\le \mathcal{H}^2\big(F(E'\cap\overline{B}_1)\big)+\int_{\mathbb{S}^2}\#\rho^{-1}(\{x\})\,d\mathcal{H}^2(x)+\varepsilon\\ &\le \mathcal{H}^2\big(F(E'\cap\overline{B}_1)\big)+4\big(|\xi-\xi'|+|\zeta-\zeta'|\big)\\ &\quad+4|\xi-\xi'|\left(\tfrac{2\sqrt{3}}{3}+\sqrt{\tfrac{3-\sqrt{6}}{6}}\right)+\varepsilon.\end{aligned}$$

This proves the first claimed inequality. Moreover,

$$4+4\left(\tfrac{2\sqrt{3}}{3}+\sqrt{\tfrac{3-\sqrt{6}}{6}}\right)<10.$$

Consequently,

$$\mathcal{H}^2\big(G(E\cap\overline{B}_1)\big) \le \mathcal{H}^2\big(F(E'\cap\overline{B}_1)\big)+10|\xi-\xi'|+4|\zeta-\zeta'|+\varepsilon.$$

This completes the proof of Proposition 17.2 and we are thus left with the proof of Lemma 17.5.

Proof of Lemma 17.5. Let us parametrize σ over $[0,\alpha]$ at unit speed, where clearly α is the length of σ. We can then write $\tau(t,s) = O_{t\theta}(\sigma(s))$: if we consider it as a map from the planar domain $[0,1]\times[0,\alpha]$ into $\mathbb{R}^3$ we can use the area formula to compute

$$\int_{\mathbb{S}^2}\#\tau^{-1}(\{x\})\,d\mathcal{H}^2(x) = \int_0^1\int_0^\alpha \left|\tfrac{\partial\tau}{\partial t}\times\tfrac{\partial\tau}{\partial s}\right| ds\,dt.$$

Consider the unit vector v orienting the axis of the rotation O so that O_θ rotates counterclockwise around v (i.e. $\frac{\partial}{\partial t}(O_{t\theta}(x)) = \theta v\times O_{t\theta}(x)$). We then have

$$\tfrac{\partial\tau}{\partial t}(t,s) = \theta v\times\tau(t,s) = \theta v\times O_{t\theta}(\sigma(s)).$$

On the other hand

$$\tfrac{\partial\tau}{\partial s}(t,s) = O_{t\theta}(\dot\sigma(s)).$$

Since $O_{t\theta}(v) = v$ and rotations of $\mathbb{R}^3$ commute with the cross product, we can write

$$\tfrac{\partial\tau}{\partial t}\times\tfrac{\partial\tau}{\partial s} = \theta O_{t\theta}((v\times\sigma(s))\times\dot\sigma(s))$$

and thus

$$\int_0^1 \int_0^\alpha \left|\tfrac{\partial \tau}{\partial t} \times \tfrac{\partial \tau}{\partial s}\right| ds\, dt = |\theta| \int_0^\alpha |(v \times \sigma(s)) \times \dot\sigma(s)|\, ds\,.$$

Observe next that

$$(v \times \sigma) \times \dot\sigma = \sigma(v \cdot \dot\sigma) - v(\sigma \cdot \dot\sigma) = \sigma(v \cdot \dot\sigma)\,,$$

given that $\dot\sigma$ is orthogonal to σ. Therefore

$$\int_{\mathbb{S}^2} \#\tau^{-1}(\{x\})\, d\mathcal{H}^2(x) = |\theta| \int_0^\alpha |v \cdot \dot\sigma(s)|\, ds\,.$$

Consider now the plane containing the geodesic σ and observe that we can choose an orthonormal basis $\{w_1, w_2\}$ for it with the property that w_1 is orthogonal to v. $\sigma(s)$ is now equal to $w_1 \cos(s-\beta) + w_2 \sin(s-\beta)$ for some angle β. Then $\dot\sigma(s) = -w_1 \sin(s-\beta) + w_2 \cos(s-\beta)$, and hence $|v \cdot \dot\sigma(s)| = |v \cdot w_2|\, |\cos(s-\beta)| \leq |\cos(s-\beta)|$. Changing variables in the integral we can therefore estimate

$$\int_{\mathbb{S}^2} \#\tau^{-1}(\{x\})\, d\mathcal{H}^2(x) \leq |\theta| \int_{-\beta}^{\alpha-\beta} |\cos(s')|\, ds'\,.$$

Maximizing the right-hand side in β (and using $0 \leq \alpha \leq \pi$) we easily see

$$\int_{\mathbb{S}^2} \#\tau^{-1}(\{x\})\, d\mathcal{H}^2(x) \leq |\theta| \int_{-\frac{\alpha}{2}}^{\frac{\alpha}{2}} |\cos(s')|\, ds' = 2|\theta| \sin \tfrac{\alpha}{2}\,.$$

□

17.4. **Proof of Proposition 17.4.** We show only the first conclusion; the proof of the second is entirely analogous.

We argue that there are exactly three possible topological behaviors for $\mathcal{N}'$ at a pin P compared to $\mathcal{N}$. Since the nets have the same pins, we will use the notation $\deg_{\mathcal{N}'}(P)$ and $\deg_{\mathcal{N}}(P)$ for the degrees in the corresponding nets. The first possibility is that

(a0) No edge incident to P is contracted. In this case $\deg_{\mathcal{N}'}(P) = \deg_{\mathcal{N}}(\tau(P))$ and τ maps the neighbors of P bijectively onto the neighbors of $P = \tau(P)$.

Otherwise there is an edge $\sigma = [P, S]$ which is contracted. Since S is a non-pin vertex of the net $\mathcal{N}'$, it has degree three. Moreover, every pin of $\mathcal{N}$ has degree at most three. Hence

$$\deg_{\mathcal{N}}(P) = \deg_{\mathcal{N}'}(P) + 1 \leq 3,$$

and therefore $\deg(P) \in \{1, 2\}$. We thus have exactly two possibilities:

(a1) $\deg_{\mathcal{N}}(P) = 1$ and $\deg_{\mathcal{N}}(P') = 2$. Then S is the only vertex adjacent to P, $\tau(S) = P$ and the two neighbors of S distinct from P are mapped by τ bijectively onto the two neighbors of P, cf. Figure 30.

(a2) $\deg_{\mathcal{N}'}(P) = 2$ and $\deg_{\mathcal{N}}(P) = 3$. Then S is one of the two vertices adjacent to P. The other two neighbors of S are mapped by τ bijectively onto two of the three neighbors of P. The other neighbor of P is mapped to the remaining neighbor of P, cf. Figure 31.

Definition of Σ' and Φ'. Recall the subcomplexes Γ and L of Σ as in Theorem 16.7, and denote by Γ_{bd} the subcomplex of those vertices and edges of Γ which are mapped by Φ into $E \cap \partial K$. Observe that L and Γ_{bd} have no edges in common and the vertices they have in common are mapped onto the two pins, which we denote by Q_+ and Q_- (with the convention that $\Phi(Q_\pm) = \pm P$). First of all, we impose that Σ' contains Σ as a subcomplex. Next, we add

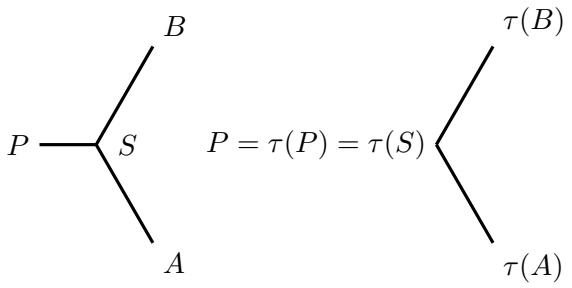


FIGURE 30. The change in topology at the pin of case (a1).

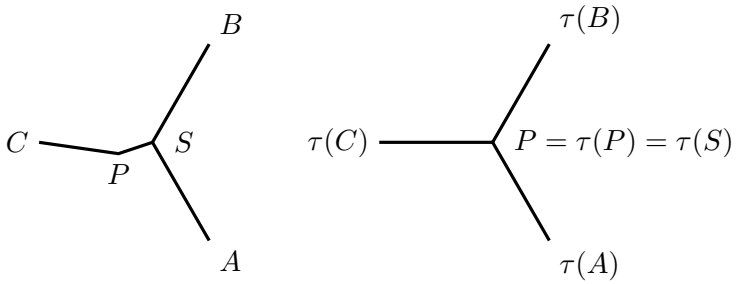


FIGURE 31. The change in topology at the pin of case (a2).

the following simplices to Σ'. As it will become clear, it is actually convenient to introduce immediately also the map Φ'. Being piecewise linear, it will be sufficient to define it on all the vertices. On the ones common to Σ it coincides with Φ.

New vertices. For every vertex $V \in \Gamma_{bd}$ which is not a pin we add V' and we prescribe that $\Phi'(V') = \iota(\Phi(V))$. If at the pin $P = \Phi(Q_+) \in \mathcal{N}'$ there is a change of topology, then $\mathcal{N}'$ has an additional special vertex S_+ in $\mathcal{N}'$ which is mapped by τ to $P = \Phi(Q_+)$. In this case we add one vertex to Σ' as well, we denote it by Z_+ and we set $\Phi'(Z_+) = S_+$. The same is applied at the pin $-P = \Phi(Q_-)$ and, if there is a change of topology, the corresponding vertices are denoted by S_- and Z_-.

New edges. For every $V \in \Gamma_{bd}$ which is not a pin we add the edge $[V, V']$.

Next, for every edge $[V_1, V_2]$ in Γ_{bd} such that neither V_1 nor V_2 is a pin, we add the edge $[V_1', V_2']$ and one of the two diagonals $[V_1, V_2']$ and $[V_1', V_2]$: which diagonal we choose is irrelevant.

If at the pin $P = \Phi(Q_+)$ there is no change in topology, then, for every $V \in \Gamma_{bd}$ adjacent to it we add the edge $[V', Q_+]$. If there is a change in topology we add $[Z_+, Q_+]$. Furthermore we include the following additional edges, depending on the two model alternatives (a1) and (a2).

- In alternative (a1), we let A and B be the two vertices of Γ_{bd} adjacent to Q_+ and we add the edges $[Z_+, A']$, $[Q_+, A']$, $[Z_+, B']$, and $[Q_+, B']$.
- In the alternative (a2), we let A, B, and C be the three vertices of Γ_{bd} adjacent to Q_+. Then we denote by A', B', and C' the corresponding vertices added to form Σ', and we ensure that the labels are assigned so that $\Phi'(A')$ and $\Phi'(B')$ are adjacent

to S_+ in $\mathcal{N}'$, while $\Phi'(C')$ is adjacent to $P = \Phi'(Q_+)$. Then we add the edges $[A', Z_+], [B', Z_+], [A', Q_+]$, $[B', Q_+]$, and $[C', Q_+]$.

We proceed in the same way at the pin $-P = \Phi(Q_-)$, with Q_- and Z_- replacing Q_+ and Z_+ in the above constructions.

New triangles. Consider any pair $A, B \in \Gamma_{bd}$ of adjacent vertices which are not pins and not special vertices. We label them so that $[A, B']$ is the "diagonal edge" added to Σ' in the previous step. Then we add the two triangles $[A, A', B']$ and $[A, B, B']$.

At the pin Q_+ for every vertex $V \in \Gamma_{bd}$ adjacent to Q_+ we add the triangle $[Q_+, V, V']$. If there is a change of topology we add the triangles $[Q_+, Z_+, A']$ and $[Q_+, Z_+, B']$. We proceed in the same way at the pin Q_+, with Q_- and Z_- replacing Q_+ and Z_+ in the above constructions.

Definition of Γ'_{bd} and Γ'. Γ'_{bd} is a one-dimensional simplicial complex. Its vertices are Q_+, Q_- and all the new vertices added to Σ'. Moreover, given an edge $[A, B]$ of Σ' connecting two vertices of Γ'_{bd} we include it in the complex Γ'_{bd} if the points $\Phi'(A)$ and $\Phi'(B)$ are adjacent in the net $\mathcal{N}'$.

The subcomplex Γ' is finally defined as the union of the subcomplexes L and Γ'_{bd}.

Verification that Theorem 16.7 is applicable. First of all, Φ' maps the vertices of Σ in the convex hull K of the vertices of $\mathcal{N}$, and the new vertices in the vertices of $\mathcal{N}'$. Hence, by piecewise linearity, Φ' maps Σ' into K_c, which is the convex hull of $\mathcal{V}(\mathcal{N}) \cup \mathcal{V}(\mathcal{N}')$. The fact that Φ' is an embedding of Γ'_{bd} onto $\partial K' \cap E'$ is obvious from the construction. The fact that it is an embedding of L onto $\ell \cap K_c$ is also obvious, since it coincides with Φ and $\ell \cap K_c = \ell \cap K$. Finally, the only faces in common between Γ'_{bd} and L are the points Q_+ and Q_- and these are mapped onto P and $-P$: since these two points form the intersection of $\partial K' \cap E'$ with $\ell \cap K_c$, we conclude immediately that Φ' is an embedding of Γ' onto $(\partial K' \cap E') \cup (\ell \cap K')$ and that $\Phi'(\Sigma') \cap \ell = K' \cap \ell$.

Having verified Theorem 16.7(i) we proceed to verify Theorem 16.7(ii). Consider a face F with the property that $(\Phi')^\circ(F) \cap \ell \neq \emptyset$. If F is a face of Σ, then F is by assumption a face of L. If F is not a face of Σ and it is a vertex, then we have already proved above that it cannot be mapped into ℓ. If F is a triangle T we then distinguish two cases.

First of all consider a T which does not have one of the pins $Q_\pm$ among its vertices. Then this triangle is of type $[A, B, B']$ or $[A, A', B']$, where both A and B are neither a pin nor a special vertex of $\mathcal{N}$ and A', B' are the corresponding vertices added to Σ' following the description given above. In particular we can apply condition (c) in Proposition 17.4 to the four points $U = \Phi'(A) = \Phi(A)$, $V = \Phi'(B) = \Phi(B)$, $U' = \Phi'(A')$, and $V' = \Phi'(B')$; by definition $U' = \iota(U)$ and $V' = \iota(V)$. Condition (c) in Proposition 17.4 provides an open half-space H_+ with $\ell \subset \partial H_+$ and $U, V, U', V' \in H_+$. Since the three vertices of $\Phi'(T)$ are a subset of $\{U, V, U', V'\}$ we can use therefore convexity to establish that the closed convex set $\Phi'(T)$ cannot intersect ℓ.

Consider next a triangle T which has a pin among its vertices, and assume without loss of generality that it is Q_+. Then T can be of two types:

- $T = [Q_+, A, A']$, where A is a vertex adjacent to Q_+ in Γ_{bd} which is not a special vertex. Then the segment $[\Phi'(A), \Phi'(A')]$ does not intersect ℓ, since it belongs to a triangle which has already been examined: thanks to (d) A must be adjacent to a vertex which is neither a pin nor a special vertex. In particular the triangle $\Phi'(T)$ intersects ℓ only at the point $\Phi'(Q_+) = P$.

- $T = [Q_+, Z_+, A']$. But then $S_+ := \Phi'(Z_+)$ and $W := \Phi'(A')$ are two adjacent vertices of the net $\mathcal{N}'$ which are not pins. In particular the edge connecting them in the net is a length-minimizing geodesic of the sphere strictly shorter than π, which does not contain any pin. We conclude that the corresponding segment $[S_+, W]$ does not intersect ℓ. In particular, for the triangle $\Phi'(T)$, which is the convex hull of $[S_+, W]$ and P, we conclude that its only intersection with ℓ is the point $\Phi'(Q_+) = P$.

We have thus verified that, for every new triangle T, either Φ' maps it into $\mathbb{R}^3 \setminus \ell$, or the only point of T which is mapped in ℓ is a vertex. In particular, since every edge added to form Σ' is the side of some new triangle T, this verifies condition (ii) in Theorem 16.7 also for the new edges.

We finally check condition (iii) which is about the collapsibility of Σ' onto L. We start noticing that $([A, A', B'], [A', B'])$ is a collapse pair for any edge $\sigma = [A, B]$ of Γ_{bd} such that A and B are not pins. We thus collapse all of these triangles and arrive at a second simplicial complex Σ''. Furthermore, for every such edge, $([A, B, B'], [A, B'])$ is now a collapse pair and, after collapsing all of these triangles we arrive at a third complex Σ'''. At this point, we can collapse the pair $([A, A'], A')$ for every A which is not adjacent to a pin in Γ_{bd}. After performing all these collapses we arrive at a fourth complex $\Sigma^{(4)}$. All the triangles of $\Sigma^{(4)}$ which are not triangles of Σ have now a pin among their vertices, while every edge and every vertex of $\Sigma^{(4)}$ which are not in Σ are contained in one such triangle. Consider first the triangles which have Q_+ as a vertex. Among them, for those of the form $[Q_+, A, A']$, the pair $([Q_+, A, A'], [A, A'])$ is a collapse pair, because, thanks to condition (d) in Proposition 17.4, the vertex A cannot be adjacent to Q_- and in particular any other triangle of Σ' which had $[A, A']$ as a side has already been collapsed. We therefore collapse all such triangles and all analogous triangles at Q_-. We are now left with the triangles of the form $[Q_+, Z_+, A']$ (which exist only if there was a change of topology at $\Phi(Q_+) = P$). Since $([Q_+, Z_+, A'], [Z_+, A'])$ is a collapse pair, we can collapse all these triangles and likewise all the analogous triangles with one vertex in Q_-. Compared to Σ, the resulting new complex $\Sigma^{(5)}$ has no additional triangles. If there are additional edges of the form $[Q_+, V]$ and $[Q_-, V]$ where V is a vertex of Γ'_{bd}, the pairs $([Q_\pm, V], V)$ are all collapse pairs, and after collapsing them we arrive at the complex Σ.

Area estimate. We now need to verify Equation (17.2). We count the triangles $\Phi'(T)$ added to $\Phi(\Sigma)$. Each triangle $\Phi'(T)$ is associated to an edge $\sigma \in \mathcal{E}(\mathcal{N})$ in the following way:

- If $T = [A, B, B']$ or $[A, A', B']$ and neither A nor B is a pin in Γ_{bd}, then $\sigma = [\Phi(A), \Phi(B)]$;
- otherwise $T = [Q_\pm, A, A']$ or $[Q_\pm, Z_\pm, A']$ for some $A \in \Gamma_{bd}$ adjacent to the pin $Q_\pm$; in this case we set $\sigma = [\Phi(Q_\pm), \Phi(A)]$.

Under this count, for any $\sigma \in \mathcal{E}(\mathcal{N})$ which is not incident to a pin, there are exactly two triangles, while for an edge incident to a pin, there may be two triangles or just one.

If $\sigma \in \mathcal{E}(\mathcal{N})$ is not adjacent to a pin then the sum of the areas of the two added triangles is obviously bounded by $D(\sigma)$. Assume next that $\sigma = [U, P]$ and that there is no change of topology at P. Then we have $\iota(P) = P$ and it follows immediately that

$$D(\sigma) = \mathcal{H}^2([U, \iota(U), P])$$

which is the area of the only triangle added. In the other case, $\iota(P)$ is a special vertex S and we have added at most the two triangles $[P, U, \iota(U)]$ and $[P, S, \iota(U)] = [P, \iota(P), \iota(U)]$: the sum of their areas is still bounded by $D(\sigma)$.

18. Area-Decreasing Competitors and proof of Theorem 1.3

In this section we employ all the results proven up to now to conclude the proof of Theorem 1.3. Also this final part is computer-assisted. Let us begin by summarizing what needs to be proven and what we already know.

Consider an Almgren minimizing cone E with respect to a line ℓ as in Theorem 1.3, let $\mathcal{N}$ be its link and denote by $\mathcal{N}'$ the corresponding *once-derived* net. We want to show that E belongs to the List 1.1.

If $\mathcal{N}'$ is extinct, we have shown in Theorem 14.1 that the cone E belongs to the List 1.1. So, if $\mathcal{N}'$ is extinct there is nothing left to prove.

On the other hand, if $\mathcal{N}'$ is not extinct, we have proven in Theorem 13.3 that its topology is encoded by one of the marked spherical graphs of Proposition 8.5 with identifiers

$$\begin{aligned}\mathscr{L} := \{&28, 99, 158, 182, 250, 319, 325, 346, 361, 421, 468, 484, 520, 614,\\ &681, 712, 736, 779, 827, 909, 920, 938, 1029, 1096, 1100, 1122\}\,.\end{aligned}$$

Moreover, we have also proven *approximate metric rigidity* for all these topologies apart from 827 and 1029. Thanks to Theorem 15.3, if the once-derived topology is 827 then $\mathcal{N}$ is an element of the fat fish family (see Definition 15.1), while if the once-derived topology is 1029 then $\mathcal{N}$ is an element of either the cubical family or the goggles family. If $\mathcal{N}$ is an element of the cubical family, which belongs to the List 1.1, we do not have anything to prove.

In Theorem 15.3, we have also proven that the once-derived topology cannot be 158 or 681 and that if the once-derived topology is one of 920 or 1100 then E belongs to the List 1.1.

Summing up what we said, we are left to show the following proposition.

Proposition 18.1. *None of the cones whose links belong to the following list of stationary nets with two antipodal pins is Almgren minimizing with respect to the line ℓ passing through the pins:*

- *The nets whose once-derived topology belongs to*

$$\mathscr{L}_{rigid} := \{182, 250, 319, 325, 361, 421, 468, 484, 520, 614, 712, 736, 909, 938, 1122\}\,.$$

- *The nets whose once-derived topology belongs to*

$$\mathscr{L}_{flexible} := \{28, 99, 346, 779, 1096\}\,.$$

- *Any element of the goggles family (corresponding to the once-derived topology* 1029*).*
- *Any element of the fat fish family (corresponding to the once-derived topology* 827*).*

In the above schematic plan, we have split the once-derived topologies to rule out into two seemingly arbitrary lists $\mathscr{L}_{rigid}$ and $\mathscr{L}_{flexible}$. Let us explain why we came up with this partition. Any of the *approximately rigid* once-derived topologies we are considering is associated to an Approximately Metrized Marked Spherical Graph (cf. Definition 13.4). Our presentation will be informal and will assume that the reader understands the content of Section 13.4. For the details about the omitted arguments, we refer the reader directly to `find_competitor.cpp` in [13].

Observe that if $(G, \mathcal{I})$ is an AMMSG containing the once-derived net $\mathcal{N}'$, one could approximately compute (up to isometries of the whole net) the positions of the two pins that were removed from $\mathcal{N}$ to obtain $\mathcal{N}'$: it is sufficient to intersect all the equators containing edges adjacent to a pin (recall Lemma 13.8). There are two cases:

- If the equators do not all (approximately) coincide, then we can intersect them and (approximately) identify the position of the pins and so (approximately) determine the location of all the vertices of $\mathcal{N}$. This is the case for the topologies in $\mathscr{L}_{rigid}$. One exception is the topology 736, whose approximate metrization allows us to (approximately) identify the two pins even if all the edges adjacent to a pin belong to the same equator.
- If the equators are all (approximately) the same, then the position of the pins is not (approximately) identified. This is the case for the topologies in $\mathscr{L}_{flexible}$. In this case, to (approximately) determine the location of the pins of $\mathcal{N}$ we must introduce an additional parameter which determines the length of one (arbitrarily chosen) edge adjacent to a pin.

Finally, let us remark that we employ Proposition 11.2 to conclude that, if the respective cone is area-minimizing, the parameters that define a net in 1096 and in the fat fish family must be (quantitatively) away from the extreme points of their domains (this is the only place where we use Proposition 11.2). Being able to state a-priori that every pin P cannot be too close to an edge that is not adjacent to an edge adjacent to P allows us to avoid some pathological cases that would make the rest of our strategy fail.

18.1. **Ambiguity in the topology of $\mathcal{N}$.** Let us understand whether the topology of $\mathcal{N}'$ determines the topology of $\mathcal{N}$. Let us assume that we are given an AMMSG $(G, \mathcal{I})$ and, if necessary, the additional parameter that is necessary to determine the (approximate) location of the pins. We refer the reader to Definition 5.5 to refresh the exact method to obtain $\mathcal{N}'$ from $\mathcal{N}$.

If the position of a pin is distinct from the position of any vertex of $\mathcal{N}'$ (observe that we know the positions only approximately, so we can only determine whether they are surely distinct, not whether they are surely the same), then there is no ambiguity and the topology of $\mathcal{N}$ is uniquely determined from the topology of $\mathcal{N}'$. If we cannot rule out the possibility that a pin of $\mathcal{N}$ coincides with a vertex of $\mathcal{N}'$ using the data of $(G, \mathcal{I})$, then the topology of $\mathcal{N}$ is not uniquely determined. The two possible ambiguous cases are represented in Figures 30 and 31.

18.2. **Reducing all cases to the same subproblem.** It remains to prove a list of statements all of the form "no net in this family generates an Almgren-minimizing cone". We will now explain how to reduce this list to a (much longer) list of more structured subproblems.

We have observed that the nets whose once-derived topology belongs to $\mathcal{L}_{flexible}$ are (approximately) parametrized by a single real parameter (and their topology is defined up to the ambiguities described in the previous section). Similarly, the family of goggles is parametrized by a single parameter. The fat fish family instead is parametrized by two real parameters.

Observe that a geodesic net with antipodal pins can be encoded by a triple $(G, \{p_1, p_2\}, e)$, where G is a spherical graph, $\{p_1, p_2\}$ are two vertices of G (the pins), and $e : \mathcal{V}(G) \to \mathbb{S}^2$ is an embedding of the vertices into the unit sphere such that any two adjacent vertices are not antipodal[8] and $e(p_1) = -e(p_2)$.

[8]If two vertices are antipodal, then the triple $(G, \{p_1, p_2\}, e)$ does not contain enough data to reconstruct the edge connecting the two vertices. In this case, one can add an artificial vertex in the middle of the edge and the resulting geodesic net, which coincides with the original one as a set, can be fully encoded by this framework.

Definition 18.2 (Approximate Geodesic Net with two pins)**.** An approximate geodesic net with two pins is a triple $(G, \{p_1, p_2\}, \tilde{e})$, where G is a spherical graph, $\{p_1, p_2\}$ are two vertices of G (the pins), and $\tilde{e} : \mathcal{V}(G) \to \mathcal{P}(\mathbb{S}^2)$ is a map from the vertices of G into the subsets of the unit sphere.

A realization $(G, \{p_1, p_2\}, e)$ of an approximate geodesic net $(G, \{p_1, p_2\}, \tilde{e})$ is a geodesic net with antipodal pins on the same graph (with the same pins) such that $e(v) \in \tilde{e}(v)$ for all $v \in \mathcal{V}(G)$.

Now, assume that the once-derived net associated to $\mathcal{N}$ has a topology in $\mathcal{L}_{rigid}$. Then, applying Lemma 13.8, we can construct an Approximate Geodesic Net $(G, \{p_1, p_2\}, \tilde{e})$ with two pins such that, up to isometry of $\mathbb{S}^2$, the net $\mathcal{N}$ is a realization of $(G, \{p_1, p_2\}, \tilde{e})$ and $\tilde{e}(v)$ is a *small* set for all $v \in \mathcal{V}(G)$. What we have just said is not correct, because the topology of $\mathcal{N}$ cannot always be inferred from the topology of its once-derived net $\mathcal{N}'$. We solve this annoying issue as follows. We choose G to be the more *complicated* topology that is consistent with the once-derived topology. Then, $\mathcal{N}$ is either a realization of $(G, \{p_1, p_2\}, \tilde{e})$ or it is approximated by a realization of $(G, \{p_1, p_2\}, \tilde{e})$. More precisely, in the latter case, for any $\varepsilon > 0$ there is a realization $\tilde{\mathcal{N}}$ of $(G, \{p_1, p_2\}, \tilde{e})$ such that we can apply Proposition 17.4 to $\mathcal{N}$ and $\tilde{\mathcal{N}}$ with the sum $\sum_{\sigma \in \mathcal{E}(\tilde{\mathcal{N}})} D(\sigma)$ appearing in Equation (17.2) being smaller than ε. In particular, if we can find a competitor (in the sense of Theorem 16.7) for all realizations of $(G, \{p_1, p_2\}, \tilde{e})$ that decreases the area by a fixed amount, then Proposition 17.4 and Theorem 16.7 guarantee that $\mathcal{N}$ is not Almgren minimizing. Given that we will show the existence of such a competitor, we are allowed to assume that G represents the *most complicated* topology, and from now on we will not discuss this matter anymore.

Analogously, let us instead assume that the once-derived net associated to $\mathcal{N}$ has a topology in $\mathcal{L}_{flexible}$ and let I be a *small* interval such that the length of an (arbitrarily fixed) edge adjacent to a pin belongs to I. Then, using only the information on the topology of $\mathcal{N}'$ together with the interval I, we can construct an Approximate Geodesic Net with two pins satisfying the same constraints of the previous paragraph.

The same can be done also if $\mathcal{N}$ is a goggle and we know that the parameter identifying it belongs to a small interval I. If it is a fat fish, and we know that the two parameters defining it belong to two small intervals I, J, the argument is equivalent but we apply Proposition 17.2 instead of Proposition 17.4.

In all cases, if we allow ourselves to discretize the zero, one, or two-dimensional space of parameters associated to a topology in sufficiently small regions[9], we can reduce the proof of Theorem 1.3 to many instances of the following problem:

Given an Approximate Geodesic Net with two pins $(G, \{p_1, p_2\}, \tilde{e})$, prove that the cone over any of its realizations is not Almgren minimizing.

18.3. **Solving the subproblem.** We describe how to treat one instance of the above-mentioned problem. Consider a geodesic net with antipodal pins $\mathcal{N} := (G, \{p_1, p_2\}, e)$ chosen so that, for all $v \in \mathcal{V}(G)$, the quantity

$$\sup_{p \in \tilde{e}(v)} |p - e(v)|$$

is *small*. In our implementation, $\tilde{e}(v)$ is always a parallelepiped (intersected with $\mathbb{S}^2$) and we choose $e(v)$ to be its center (with an exception at the pins, to guarantee antipodality).

[9]We refer the reader to the code for details on how we do this.

Here again, there is an exception related to the fact that the topology is ambiguous when the location of the pin falls very close to the location of a vertex of the once-derived net. Let us focus on the pin p_1, the analogous argument shall be repeated verbatim for the other pin p_2. If there is a vertex $v \notin \{p_1, p_2\}$ such that the minimum distance between $\tilde{e}(p_1)$ and $\tilde{e}(v)$ is below a given threshold, then $\mathcal{N}$ will be induced by $(\tilde{G}, \{p_1, p_2\}, e)$ where the vertices p_1 and v are collapsed onto a single vertex in $\tilde{G}$. So, in this case, the topology of $\mathcal{N}$ and the topology of the realizations of $(G, \{p_1, p_2\}, \tilde{e})$ do not match but there is a controlled topological change (cf. Definition 17.3) from any $\mathcal{N}$ to any realization. Treating explicitly this case would lead to some very cumbersome notation without yielding any further understanding, so we will assume from now on that $\mathcal{N}$ has the topology of G and we will just relieve the worried reader by making it clear that the only place where this would actually play a role is in the application of Proposition 17.4 which conveniently supports nets which are obtained one from the other via a controlled topological change.[10]

Take any realization $(G, \{p_1, p_2\}, e')$ of $(G, \{p_1, p_2\}, \tilde{e})$. Observe that there exists an isometry T (depending on the realization) such that

- $\|T - \mathbb{1}\|_{op} \leq \sup_{p \in \tilde{e}(p_1)} |p - e(p_1)| =: \lambda$,
- $T(e'(p_1)) = e(p_1)$ and $T(e'(p_2)) = e(p_2)$.

In particular, the geodesic net with antipodal pins encoded by $(G, \{p_1, p_2\}, T \circ e')$ is a realization of the approximate geodesic net $(G, \{p_1, p_2\}, f)$ where f is defined as

- $f(p_1) = \{e(p_1)\}$ and $f(p_2) = \{e(p_2)\}$,
- $f(v) = B(\tilde{e}(v), \lambda)$ for all $v \in \mathcal{V}(G) \setminus \{p_1, p_2\}$.

Since being Almgren minimizing is invariant under isometries, if we are able to prove that none of the cones over a realization of $(G, \{p_1, p_2\}, f)$ is Almgren minimizing, then we have solved the problem.

What we have achieved in the last paragraph is to restrict our study on a family of nets with the two pins fixed in the same position as the pins of the special net $(G, \{p_1, p_2\}, e)$ that we had previously chosen. The price that we had to pay was the enlargement of $\tilde{e}(v)$ to $f(v) = B(\tilde{e}(v), \lambda)$ for a positive $\lambda > 0$ (which is as small as the precision with which we initially knew the position of the pins). This trick is necessary because later we will want to apply Proposition 17.4.

Let us assume that for any pair (U, V) of adjacent vertices of G that are not pins, there is a half-space H_+ with $e(p_1), e(p_2) \in \partial H_+$ and $e(U), e(V) \in H_+$ and $f(U), f(V) \subset H_+$. This assumption is mild given that all geodesic edges have length strictly less than π and we assume to have an instance such that $f(V)$ is a rather *small* set containing $e(V)$ for any $V \in \mathcal{V}(G)$.

Let K be the convex hull of $\{e(v) : v \in \mathcal{V}(G)\}$. Let E be the cone over the geodesic net $(G, \{p_1, p_2\}, e)$.

Now, imagine that we are able to create an abstract 2-dimensional simplicial complex Σ and a piecewise linear map $\Phi : \Sigma \to \mathbb{R}^3$ satisfying the assumptions of Theorem 16.7. Thanks to Theorem 16.7, if

$$\mathcal{H}^2(\Phi(\Sigma)) < \mathcal{H}^2(E \cap K), \tag{18.1}$$

[10]More delicate, and more important, is the matter of *why* we decide to opt for this complicated choice. The main reason is that, eventually, we will want to construct a competitor (in the sense of Theorem 16.7) for $\mathcal{N}$ and to do so we will employ Brakke's surface evolver, which tends to behave erratically when one vertex is very close to a pin. Furthermore, if we did not employ this strategy, then proving Proposition 17.4(c) would become much more cumbersome, or impossible, when a vertex falls very close to a pin.

then E is not Almgren minimizing. We will soon notice that we need a slightly stronger inequality than Equation (18.1) for our purposes.

Fix a realization $(G, \{p_1, p_2\}, e')$ of $(G, \{p_1, p_2\}, f)$ and let K' be the convex hull of $\{e'(v) : v \in \mathcal{V}(G)\}$. Let E' be the cone over the geodesic net $(G, \{p_1, p_2\}, e')$.

By applying[11] Proposition 17.4 and Theorem 16.7, for any $\varepsilon > 0$, there is a map $F : \mathbb{R}^3 \to \mathbb{R}^3$ which is the identity outside of B_1, satisfies the requirements of Definition 1.2, and

$$\mathcal{H}^2(F(E' \cap B_1)) \leq \mathcal{H}^2(E' \cap B_1 \setminus K') + \mathcal{H}^2(\Phi(\Sigma)) + \varepsilon + \sum_{\sigma \in \mathcal{E}(G)} D_{e,e'}(\sigma),$$

where $D_{e,e'}((U,V))$ is defined as

$$\begin{aligned} D_{e,e'}((U,V)) := \max\{&\mathcal{H}^2([e(U), e(V), e'(V)]) + \mathcal{H}^2([e(U), e'(U), e'(V)]), \\ &\mathcal{H}^2([e(V), e(U), e'(U)]) + \mathcal{H}^2([e(V), e'(V), e'(U)])\}. \end{aligned}$$

Therefore, if we could show that

$$\mathcal{H}^2(\Phi(\Sigma)) + \sum_{\sigma \in \mathcal{E}(G)} D_{e,e'}(\sigma) < \mathcal{H}^2(E' \cap K') \tag{18.2}$$

we would deduce immediately that E' is not Almgren minimizing as desired. First of all, observe that this condition is *almost* equivalent to Equation (18.1). More importantly, this condition can be transformed in a statement which does not involve E' or K' at all arguing as follows.

Define

$$D_e(\sigma) := \sup_{e' \colon e'(v) \in f(v) \forall v \in \mathcal{V}(G)} D_{e,e'}(\sigma),$$

and observe that

$$\mathcal{H}^2(E' \cap K') = \sum_{(U,V) \in \mathcal{E}(G)} \mathcal{H}^2([0_{\mathbb{R}^3}, e'(U), e'(V)]) \geq \sum_{(U,V) \in \mathcal{E}(G)} \inf_{p \in f(U), q \in f(V)} \mathcal{H}^2([0_{\mathbb{R}^3}, p, q]).$$

So, Equation (18.2) is implied by

$$\mathcal{H}^2(\Phi(\Sigma)) < \sum_{(U,V) \in \mathcal{E}(G)} \inf_{p \in f(U), q \in f(V)} \mathcal{H}^2([0_{\mathbb{R}^3}, p, q]) - \sum_{\sigma \in \mathcal{E}(G)} D_e(\sigma).$$

In this latter inequality, all terms are computable (or at least approximately computable) and thus we have reduced the subproblem to finding a good enough competitor (Σ, Φ) for one particular antipodal net $(G, \{p_1, p_2\}, e)$.

This is, in fact, the core of the problem. As we will describe next, we are able to find a *good enough* competitor in all the necessary cases and this concludes the proof. The competitors for all the instances of the subproblem can be found in `data/competitors.zip` of [13] or in the website [12].

[11] For simplicity, we describe only the details of the strategy for nets that are not in the fat fish family. For nets in the fat fish family, the argument is in fact simpler because one can directly apply Proposition 17.2 and skip some of the computations described.

18.4. **Constructing a Competitor.** Even though the proof itself does not depend on *how* we have constructed these competitors, as far as we can verify that they respect all the required constraints, we think that it is important to document how they were constructed.

We are given a geodesic net with antipodal pins $\mathcal{N}$ and we want to construct a competitor, in the sense of Theorem 16.7, for the cone over $\mathcal{N}$. We may assume that the geodesic net $\mathcal{N}$ is not pathological in any way, in particular we may assume that no edge of the net is very short (we have spent quite some effort in the previous subsections to obtain these assumptions).

The construction proceeds in two steps: first we build a competitor that has a larger area but has a topology different from the topology of the cone, then we run Brakke's (constrained) mean curvature flow. At the end, we are just left to verify that the result is a valid competitor with small enough area, which are straight-forward checks described in Theorem 16.7.

Cone-popping. To alter the topology of the cone, we perform an operation that we call *popping the cone* (this terminology is taken from Brakke, who introduced a pop command in his code; however our cone-popping differs from his). The standard cone over a net $\mathcal{N}$ consists of triangles connecting every edge of $\mathcal{N}$ to the origin. Popping the cone replaces the portion of the surface inside a small ball $B(0, r)$ centered at the origin with a different simplicial structure.

Definition 18.3 (Popped cone)**.** Let $\mathcal{N}$ be a geodesic net on ∂B_1 with two antipodal pins P and $-P$.

Assume that we are provided with a pair of combinatorial data $(\mathcal{Z}, \gamma)$:

- $\mathcal{Z}$ is a set of faces of $\mathcal{N}$, which we refer to as the *excluded faces*.
- γ is a 1-dimensional subgraph of $\mathcal{N}$ that contains both pins.
- The 2-dimensional simplicial complex associated to $\mathbb{S}^2 \setminus \mathcal{Z}$ can be collapsed on its 1-dimensional subcomplex γ.

The *popped cone* associated to $(\mathcal{Z}, \gamma)$ is an abstract 2-dimensional simplicial complex Σ mapped into $\mathbb{R}^3$ via a piecewise linear map $\Phi : \Sigma \to \mathbb{R}^3$. The geometric realization $\Phi(\Sigma)$ is an immersed simplicial complex constructed as the union of three components, depending on a parameter $0 < r < 1$:

- **The outer cone.** For every edge $[u, v]$ of $\mathcal{N}$, $\Phi(\Sigma)$ contains a quadrilateral (triangulated into two faces) that connects the boundary vertices $\Phi(u)$ and $\Phi(v)$ to their scaled counterparts $r\Phi(u)$ and $r\Phi(v)$.
- **The capped faces.** For every face F of $\mathcal{N}$ that is *not* in the excluded set $\mathcal{Z}$, we add a (not necessarily flat) 2-dimensional (triangulated) polygonal face with vertices $\{r\Phi(v) : v \in F\}$. If the face contains a pin, the triangulation is pivoted at that pin (otherwise, any triangulation generates a valid popped cone).
- **The inner cone.** For every edge $[u, v]$ in γ, we add the triangle $[0_{\mathbb{R}^3}, r\Phi(u), r\Phi(v)]$ (note that this is a small triangle if we choose r small).

The following lemma motivates our choice, even though it ultimately does not play a role in the actual argument for Proposition 18.1.

Lemma 18.4 (Popped cones are valid competitors)**.** *Let $\mathcal{N}$ be a geodesic net with two antipodal pins $P, -P$ on the line ℓ. Let K be the convex hull of the vertices of $\mathcal{N}$.*

Given $(\mathcal{Z}, \gamma)$, as in Definition 18.3, the associated popped cone (Σ, Φ) is a valid competitor for the cone over $\mathcal{N}$ in K in the sense of Theorem 16.7.

Proof. Let $\iota : \mathcal{V}(\mathcal{N}) \to \mathcal{V}(\Sigma)$ be the map such that $\Phi(\iota(V)) = V$. Moreover, for any vertex $V \in \mathcal{V}(\mathcal{N})$, let (with a slight abuse of notation) $v = \iota(V)$ and let $v' \in \mathcal{V}(\Sigma)$ be the vertex such that $\Phi(v') = rV$. Let $o \in \mathcal{V}(\Sigma)$ be the vertex mapped into the origin by Φ.

We must verify the conditions (i), (ii), (iii) of Theorem 16.7.

Let us identify L as the subcomplex connecting, in this order, the vertices $p, p', o, -p', -p$ (with a slight abuse of notation we denote by p and $-p$ the vertices of Σ mapped onto P and $-P$).

Given the definition of L and the fact that (Σ, Φ) and $\mathrm{cone}(\mathcal{N}) \cap K$ coincide in $K \setminus rK$, it is a tedious but straightforward task to verify that (Σ, Φ) satisfies Theorem 16.7(i). We skip this check.

To check the validity of (ii) we shall verify that, for any face F of Σ not in L, $\Phi^\circ(F)$ is disjoint from the line ℓ.

Let us consider a few cases:

- If F is part of the **outer cone**, then it is tedious but straightforward to check that $\Phi^\circ(F)$ is disjoint from ℓ.
- Similarly, if F is part of the **inner cone**, then it is tedious but straightforward to check that $\Phi^\circ(F)$ is disjoint from ℓ.
- The hard case is when F is part of the **capped faces**. Here we must distinguish between the case in which F is part of a face that contained the pin or not. If the original face of $\mathcal{N}$ that is being capped did not contain the pin, then it was convex. Therefore $\Phi(F)$ must belong (independently of the triangulation that was chosen) to the convex cone generated by the face that is being capped and since this face does did not contain either P or $-P$, its cone is disjoint (up to the origin) from ℓ. On the other hand, if the original face of $\mathcal{N}$ that is being capped contained a pin then we know, thanks to the special triangulation we have chosen, that either F contains p' or $-p'$ or it is an edge connecting $[u', v']$ where $[U, V]$ is an edge of $\mathcal{N}$. In both cases, it is easy to see that $\Phi^\circ(F)$ is disjoint from ℓ.

Now, we conclude by showing a collapsing sequence that proves (iii).

- Consider a face $[u, v, v', u']$ (which is made of two triangles) of Σ, where $[U, V]$ is an edge of $\mathcal{N}$. Then, collapse it (in two steps) onto $[u, u'], [u', v'], [v, v']$. Do this for all edges $[U, V]$ of $\mathcal{N}$.
- For each vertex $V \in \mathcal{N}$ different from P and $-P$, collapse $[v, v']$ onto v'.
- Let $\mathcal{A}$ be the subcomplex of Σ whose vertices are $\{v' : V \in \mathcal{N}\}$ and its faces are the **capped faces**. Let $\mathcal{B}$ be the 1-dimensional subcomplex of $\mathcal{A}$ whose vertices are $\{v' : V \in \mathcal{V}(\gamma)\}$ and whose edges are the ones induced naturally by γ. We have that Φ is an embedding restricted on $\mathcal{A}$ and maps it onto a subset of $\mathbb{R}^3$ equivalent (in the right sense) to $\mathbb{S}^2 \setminus \mathcal{Z}$, where $\mathcal{Z}$ is the set of excluded faces. Therefore, by assumption, we deduce that we can collapse $\mathcal{A}$ onto $\mathcal{B}$ and that is exactly what we decide to do.
- For each edge $[U, V]$ in γ, we collapse $[o, u', v']$ onto $[o, u']$ and $[o, v']$.
- Finally, for each vertex V in γ that is not P or $-P$, we collapse $[v', o]$ onto o.

□

Brakke's Surface Evolver. Once we have obtained a popped cone, we evolve it using Brakke's surface evolver [6]. More precisely, we set the line ℓ as a parametrized boundary, and we *fix* the vertices and the edges on $\partial K \cap \Sigma$ (where Σ is the cone and K is the convex hull of the vertices of the net).

Brakke's surface evolver runs a discrete mean curvature flow which is designed to let the surface converge towards a constrained minimum of the area. Let us remark that, during the evolution, Brakke's surface evolver is allowed to change the topological structure of the surface but should, in theory, guarantee that the resulting surface is a valid competitor for the initial surface, given that we started from a valid one by Lemma 18.4. We do not rely on the correctness of Brakke's surface evolver as we do simply check that the final surface satisfies the assumptions of Theorem 16.7 (and so, we do not rely on Lemma 18.4 either).

A large number of ad-hoc tricks are necessary to make the algorithm converge for all our surfaces. Let us mention just a few difficulties that shall be taken care of.

- Brakke's surface evolver has some bugs that lead to bad behaviors in otherwise non-pathological situations. We fixed some of these bugs; our version of Brakke's surface evolver can be found here `https://gist.github.com/dario2994/25b6b268ffd76eeef8e8d36904a12c09`
- If initially one of the edges of the net $\mathcal{N}$ is very short, the surface evolver has the (reasonable) tendency to collapse the two vertices of the boundary into one. This operation is not valid as it breaks the property that the final surface is a valid competitor for the first one. We avoid this issue by proving a-priori that no edges are short in Proposition 11.1.
- The fat fish family presents two idiosyncratic difficulties. It is a 2-parameter family and to obtain a proof that none of its nets generates a minimizing cone we need to discretize the family into approximately 50 000 nets that shall be treated independently. So, we need to run a very fast flow otherwise the task becomes computationally unfeasible. Unluckily, some of these nets are very close to being minimizers (the competitor has area that is just 0.1% better than the initial cone; which is also why we need to discretize the family so much to obtain a proof) and to find it we need a refined flow. Therefore, for each of these nets, we run first a very rough routine which is fast and, if it fails to find a competitor, we then run a much more refined routine (in the sense that the time-evolution steps are smaller and the grid of triangles given to the computer to describe the competitor is finer). Moreover, sometimes the optimal competitor produced by the surface evolver exhibits some self-intersections on the line ℓ (we do not fully understand why this issue does not appear for the other nets we treat). Any such self-intersection makes the resulting surface an invalid competitor for the initial cone (cf. Theorem 16.7-(ii)). To tackle this problem, we consider three different initial *popped cones* and we verify that, for each of the $\approx 50\,000$ nets, at least one of them leads to a valid competitor.

19. Minimality: known cases and open problems

In this section we will discuss the minimizing properties of the cones in List 1.1.

19.1. **The four easy cases.** The cones in (a), (b), (c), and (d) are all minimizing.

Proposition 19.1. *All the elements listed in the cases (a), (b), (c), and (d) are Almgren minimizing with respect to ℓ.*

Proof. First of all, for (c) and (d) the statement follows immediately from [25]. Indeed, consider a cone E as in (c) or (d). According to [25] the inequality (1.1) is satisfied for every bounded open set V and every Lipschitz map $\Phi : \mathbb{R}^3 \to \mathbb{R}^3$ such that $\{x : \Phi(x) \neq x\}$ is compact. On the other hand the Almgren-minimizing property with respect to ℓ requires that the same inequality holds for a *strict* subset of such maps Φ.

Consider next a cone E as in case (a) and (b). In case (a) E consists of a half-plane τ with boundary ℓ, while in case (b) it consists of two half-planes τ and τ' with boundary ℓ and forming an angle $\theta \in [\frac{2\pi}{3}, \pi]$. Without loss of generality we can then assume that

$$\begin{aligned}\ell &= \{(0,0,s) : s \in \mathbb{R}\}\\ \tau &= \{(r,0,s) : s \in \mathbb{R}, r \geq 0\}\\ \tau' &= \{(r\cos\theta, r\sin\theta, s) : s \in \mathbb{R}, r \geq 0\}\,.\end{aligned}$$

Assume that E is not Almgren minimizing with respect to ℓ and fix a map Φ which contradicts (1.1) but satisfies the requirements of Definition 1.2. By assumptions there are $R > \delta > 0$ such that

$$\operatorname{spt}(\Phi - \mathrm{Id}) \subset (B_R(\ell) \setminus \overline{B}_\delta(\ell)) \cap \{|x_3| < R\}\,.$$

In particular the "slice"

$$\Gamma(s) := \{x_3 = s\} \cap \Phi\big(\overline{B}_R(\ell) \cap \{|x_3| \leq R\} \cap E\big)$$

is a subset of the horizontal plane $\alpha(s) := \{x_3 = s\}$ which, for every $s \in (-R, R)$, contains

- The two points $p_0 = (0,0,s)$ and $q = (R,0,s)$ in case (a);
- The three points p_0, q, and $q' = (R\cos\theta, R\sin\theta, s)$ in case (b).

From Lemma 19.2 below we infer that these points belong to the same connected component of $\Gamma(s)$, which we denote by $\Gamma^c(s)$: in case (a) this is obvious, while in case (b) we apply the lemma separately to each of the two rectangles forming $E \cap \overline{B}_R(\ell) \cap \{|x_3| \leq R\}$, which are the intersections of the latter sets with the two halfplanes τ' and τ. We next apply the coarea formula to get

$$\begin{aligned}\int_{-R}^{R} \mathcal{H}^1(\Gamma(s))\, ds &\leq \mathcal{H}^2\big(\Phi(E \cap \overline{B}_R(\ell) \cap \{|x_3| \leq R\})\big)\\ &< \mathcal{H}^2\big(E \cap \overline{B}_R(\ell) \cap \{|x_3| \leq R\}\big) = 2c_0R^2\,,\end{aligned}$$

where $c_0 = 1$ in case (a) and $c_0 = 2$ in case (b). In particular we infer the existence of an s such that

$$\mathcal{H}^1(\Gamma^c(s)) \leq \mathcal{H}^1(\Gamma(s)) < c_0R\,.$$

In both cases $\mathcal{H}^1(\Gamma^c(s))$ (which is, respectively, strictly less than R in case (a) and strictly less than $2R$ in case (b)) gives an upper bound for the length of the solution of the corresponding Steiner problem: the minimum of the length over connected sets containing, respectively, p_0 and q in case (a) and p_0, q, and q' in case (b). It is however elementary and well known to that in case (a) the minimizer is a straight segment with endpoints p_0 and q while in case (b) it is the union of two straight segments with one common endpoint p_0 and the other endpoints being q and q'. With this we reach a contradiction. □

Lemma 19.2. *Let $R > 0$ and let*

$$Q := \big\{(x_1, x_2, x_3) \in \mathbb{R}^3 : x_2 = 0,\ 0 \leq x_1 \leq R,\ |x_3| \leq R\big\}.$$

Suppose that $\Phi\colon Q \to \mathbb{R}^3$ is continuous and agrees with the identity in a neighborhood of ∂Q in Q. For each $s \in (-R, R)$, set $\Gamma(s) := \Phi(Q) \cap \{x_3 = s\}$. Then the points $a_s = (0,0,s)$ and $b_s = (R,0,s)$ belong to the same connected component of $\Gamma(s)$.

Proof. We first suppose that Φ is smooth. Let $\pi_3\colon \mathbb{R}^3 \to \mathbb{R}$ denote the projection onto the third coordinate, and set $f := \pi_3 \circ \Phi$. Since Φ agrees with the identity near ∂Q, we have $f(x_1, 0, x_3) = x_3$ on a neighborhood of ∂Q. By Sard's theorem, there is a sequence of regular values $t_k \in (-R, R)$ of f such that $t_k \to s$. For every k, the level set $L_k := f^{-1}(t_k)$ is a compact one-dimensional manifold. Moreover, since $f = x_3$ near ∂Q, $\partial L_k = L_k \cap \partial Q = \{a_{t_k}, b_{t_k}\}$. Every compact one-dimensional manifold is a disjoint union of circles and compact intervals. It follows that the connected component A_k of L_k containing a_{t_k} is (topologically) a compact interval whose other endpoint must be b_{t_k}.

Since Q is compact, after passing to a subsequence the sets A_k converge in the Hausdorff topology to a nonempty compact set $A \subset Q$. As the Hausdorff limit of compact connected sets, A is connected, and moreover $a_s, b_s \in A$. Furthermore, if $x \in A$, we can choose $x_k \in A_k$ such that $x_k \to x$. Consequently,

$$f(x) = \lim_{k\to\infty} f(x_k) = \lim_{k\to\infty} t_k = s.$$

Thus $A \subset f^{-1}(s)$. It follows that $\Phi(A)$ is a connected subset of $\Gamma(s)$ containing a_s and b_s. This proves the claim when Φ is smooth.

We now consider a merely continuous map Φ and approximate it uniformly with a sequence of smooth maps Φ_j which are the identity in a fixed neighborhood of the boundary of the rectangle Q.

By the smooth case, for every j there is a compact connected set $C_j \subset Q$ such that

$$a_s, b_s \in C_j \qquad \text{and} \qquad \pi_3\big(\Phi_j(x)\big) = s \quad \text{for every } x \in C_j.$$

After passing to a subsequence, C_j converges in the Hausdorff topology to a compact connected set $C \subset Q$. Since every C_j contains a_s and b_s, so does C. On the other hand the uniform convergence of Φ_j to Φ immediately gives $\Phi(C) \subset \Phi(Q) \cap \{x_3 = s\} = \Gamma(s)$. The set $\Phi(C)$ is connected and, because Φ fixes ∂Q, contains a_s and b_s. Hence these two points belong to the same connected component of $\Gamma(s)$. □

19.2. **The cube family.** In case (e) we know the following.

Proposition 19.3. *There is a positive $\varepsilon > 0$ with the following property. Let E be the cone over the 1-skeleton of a cube with the eight vertices placed at the points $(i\frac{\sqrt{3}}{3}, j\frac{\sqrt{3}}{3}, k\frac{\sqrt{3}}{3})$ with $(i, j, k) \in \{-1, 1\}^3$. Let ℓ be the line going through zero and through the point*

$$q(\lambda) = \lambda(\tfrac{\sqrt{3}}{3}, \tfrac{\sqrt{3}}{3}, \tfrac{\sqrt{3}}{3}) + (1-\lambda)(-\tfrac{\sqrt{3}}{3}, \tfrac{\sqrt{3}}{3}, \tfrac{\sqrt{3}}{3})$$

with $\lambda \in [\frac{1}{2}, 1]$. If $|\lambda - \frac{1}{2}| < \varepsilon$, then E is not Almgren minimizing with respect to ℓ.

Proof. Observe that, when $\lambda = \frac{1}{2}$, then the link of the cone E, together with the two points $\{q(\lambda), -q(\lambda)\}$, forms a stationary geodesic net with antipodal pins which is already listed as an element of the goggles family, cf. Definition 15.2. We have already shown in the proof of Theorem 1.3 that in this case E is not Almgren minimizing with respect to ℓ. But in fact our argument provides (with the assistance of a computer) a simplicial complex Σ and a piecewise linear map Φ which satisfy the assumptions of Proposition 17.4. But then we can apply the latter proposition to conclude the proof of the claim that E is not Almgren minimizing with respect to ℓ when λ is sufficiently close to $\frac{1}{2}$. □

Remark 19.4. The interested reader can easily avoid the assistance of a computer, because one can use the explicit competitor already furnished by Taylor in [25] to exclude the Almgren minimality of E when ℓ is replaced by the empty set. For the reader's convenience we

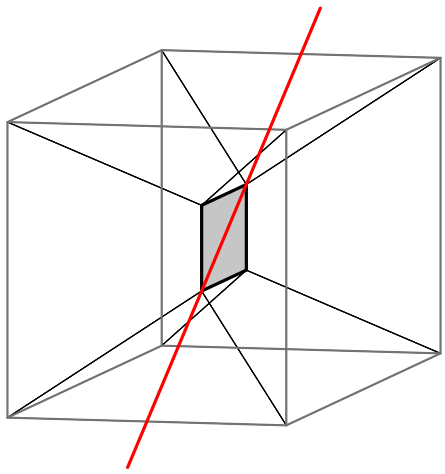


FIGURE 32. Taylor's competitor for the cubical cone when ℓ bisects two edges of the 1-skeleton.

describe it quickly, a visual illustration of the competitor is moreover given in Figure 32. We let $V(i,j,k)$ the eight vertices of the cube and introduce further the four vertices $(0,\delta,\delta)$, $(0,-\delta,\delta)$, $(0,-\delta,-\delta)$, and $(0,\delta,-\delta)$, where δ is a small positive number. The competitor consists of the "central square" (the one which has the latter four points as vertices), four triangles, and eight trapezoids. For each σ which is an edge of the central square there are two trapezoids which have σ as shortest basis and one of the two parallel edges of the cube closest to σ as longest basis. To construct the four triangles take the four remaining edges of the cube and for each add the closest vertex of the central square. An obvious triangulation of this competitor gives an embedded simplicial complex satisfying the assumptions of Proposition 17.4 (as convex set K we take the convex hull of the vertices $V(i,j,k)$). Computing the area of the competitor as a function of δ, one sees that the competitor has strictly less area than the cone over the skeleton of the cube if δ is positive and sufficiently small: we omit the computations and rather refer to [25].

Brakke's surface evolver suggests that E is not area minimizing if the angle formed by ℓ and the plane $\{x_1 = 0\}$ is smaller than 0.11 radians. The tools developed in this paper certainly allow to give an effective bound for the number ε in Proposition 19.3 and with some effort one might be able to even prove the latter statement rigorously. However we do not pursue this issue here. Note that for $\lambda = 1$ the angle between ℓ and the same plane is $\frac{1}{2}\arccos(\frac{1}{3}) \sim 0.61$ radians.

APPENDIX A. ALLARD'S THEORY AND MINIMIZING SETS

In this section we give the proofs of Proposition 4.4, Corollary 4.5, and Lemma A.2. The cornerstone of all the arguments is the following simple lemma.

Lemma A.1. *Let $U \subset \mathbb{R}^N$ be an open set, $\Gamma \subset U$ a C^2 $(k-1)$-dimensional submanifold of U, and E a minimizing set in U with respect to Γ. Then the varifold induced by E (which is rectifiable by assumption) is stationary in the sense of Allard in $U \setminus \Gamma$.*

Proof. The stationarity in the sense of Allard is defined as follows (see e.g. [10, Definition 1.3]). Having fixed a vector field $X \in C^\infty_c(U \setminus \Gamma)$ we consider the one-parameter family of

diffeomorphisms generated by it, namely the map $\Phi : U \times \mathbb{R} \to U$ such that $\Phi(x,0) = x$ for every x and $\frac{d\Phi}{dt}(x,t) = X(\Phi(x,t))$. The stationarity in $U \setminus \Gamma$ means that, for every such X, after choosing $U' \subset\subset U$ containing its support, the first variation

$$\delta E(X) := \frac{d}{dt}\bigg|_{t=0} \mathcal{H}^k(U' \cap \Phi(E,t))$$

vanishes. Note however that, because of the minimizing property, the point $t = 0$ must be a minimum of the function $f(t) = \mathcal{H}^k(U' \cap \Phi(E,t))$: indeed $\Phi(\cdot, t)$ is Lipschitz and it is the identity on Γ and on the complement of U', therefore the minimizing property must imply that $\mathcal{H}^k(U' \cap \Phi(E,t)) \geq \mathcal{H}^k(U' \cap E)$. Finally, we recall that the differentiability of the function f is a consequence of the area formula. □

Thanks to the lemma we can apply Allard's regularity theory in our setting. An important byproduct of Allard's theory is that, if $p \in E$ (irrespectively of whether $p \in \Gamma$ or not), then we have a lower bound on $\mathcal{H}^k(E \cap B_r(p))$ which depends only upon the curvature of Γ. This is a consequence of the monotonicity formulae and the following lower bound for the density.

Lemma A.2. *Assume U, Γ, and E are as in Lemma A.1. Let ω_k be the Lebesgue measure of the unit ball in $\mathbb{R}^k$. Then:*

$$\lim_{r \downarrow 0} \frac{\mathcal{H}^k(E \cap B_r(p))}{\omega_k r^k} \geq \begin{cases} 1 & \text{if } p \in E \setminus \Gamma; \\ \frac{1}{2} & \text{if } p \in E \cap \Gamma. \end{cases} \tag{A.1}$$

Proof. Given that the varifold induced by E is stationary, the interior case ($p \notin \Gamma$) can be found e.g. in [10, Proposition 2.2(ii)]. The boundary case ($p \in \Gamma$) is in [2, Theorem 3.5, conclusion (2)]. □

But in fact Allard's theory is much more powerful: his papers [1] and [2] imply the existence of a geometric constant $\delta > 0$ such that:

- If $p \in E$, $B_{2r}(p) \subset U \setminus \Gamma$ and $\mathcal{H}^k(E \cap B_{2r}(p)) < (1+\delta)\omega_k (2r)^k$, then $E \cap B_r(p)$ is a $C^{1,\alpha}$ submanifold of $B_r(p)$ without boundary [1, Section 8, Regularity Theorem];
- If $p \in E \cap \Gamma$, $B_{2r}(p) \subset U$ and $\mathcal{H}^k(E \cap B_{2r}(p)) < (\frac{1}{2}+\delta)\omega_k (2r)^k$, then $E \cap B_r(p)$ is a $C^{1,\alpha}$ submanifold of $B_r(p)$ with boundary $\Gamma \cap B_r(p)$ see [2, Section 4, Regularity Theorem].

Clearly, the lower bounds of Lemma A.2 can also be inferred as immediate consequences of the latter regularity results.

We can now tackle the proof of Corollary 4.5.

Proof of Corollary 4.5. We will use here that the varifold induced by E is stationary. Therefore, (4.1) is a direct consequence of Allard's interior monotonicity formula, see [1, Theorem 5.1], [24, Equation 4.15], or [10, Theorem 2.1]. (4.2) is a direct consequence of his boundary monotonicity formula, see [2, Theorem (2) of Section 3.4] or [5, Equation (31)]. Concerning the second part of the corollary, the case $p \notin \Gamma$ is contained in [1, Theorem (2) of Section 5.2] and can also be found in the proof of [24, Theorem 4.5.3]. The second case, namely $p \in \Gamma$ and $A_\Gamma \equiv 0$ can be inferred from the first case using the Reflection Principle of [2, Section 3.2]. In fact, since $A_\Gamma \equiv 0$, Γ must be a $(k-1)$-dimensional plane. The reflection principle of [2, Section 3.2] implies that the varifold W induced by E and its reflection along Γ is stationary in some ball $B_{r_0+\varepsilon}(p)$: for W p is then an interior point and we can apply the first case of

the statement to conclude that the support of W is a cone with vertex p. This easily implies that E itself is a cone. □

We finally come to Proposition 4.4. Part of the proof will be based on the following lemma.

Lemma A.3. *Let $U \subseteq \mathbb{R}^N$ be an open set and $\Sigma \subset U$ a k-dimensional C^2 submanifold without boundary (in U). For every open $U' \subset\subset U$ there is a number $\bar\delta(U', \Sigma) > 0$ with the following property. For any $0 < \delta < \bar\delta$, there is a Lipschitz map $\Phi : U \to U$ such that:*

- *The Lipschitz constant of Φ is at most a geometric constant C_0;*
- *Φ is the identity outside some open $V \subset\subset U$;*
- *Φ maps $B_\delta(\Sigma \cap U')$ in $\Sigma \cap U'$.*

Proof. Using the C^2 regularity of Σ and the existence of a normal neighborhood in some open set $U'' \subset\subset U$ which contains U', the lemma is a simple exercise. We fix two parameters $1 >> \eta >> \delta > 0$. For a point $p \in B_\eta(\Sigma) \cap U'$. If the distance $d(p)$ of p to Σ is smaller than δ we define the value of a map Ψ to be $\Psi(p) = q(p)$, the orthogonal projection of p onto Σ. If $d(p) = \eta$ we define it as $\Psi(p) = p$. For intermediate values of the distance d we interpolate and map p in

$$\Psi(p) := \frac{\eta - d(p)}{\eta - \delta} q(p) + \frac{d(p) - \delta}{\eta - \delta} p\,.$$

We next fix a smooth cut-off function which is identically equal to 1 on U', compactly supported in U'' and ranging in the interval $[0, 1]$. We define Φ in U'' as

$$\Phi(p) = \varphi(p)\Psi(p) + (1 - \varphi(p))p\,.$$

We will show below that $|D\Psi|$ can be bounded by a geometric constant C_0 choosing the ratio $\frac{\delta}{\eta}$ and the number η sufficiently small. On the other hand we can estimate

$$|D\Phi(p)| \leq |\nabla\varphi(p)||\Psi(p) - p| + \varphi(p)|D\Psi(p)| + (1 - \varphi(p))\,.$$

Since $|\Psi(p) - p| \leq \eta$, it suffices to choose η and $\frac{\delta}{\eta}$ sufficiently small to estimate the Lipschitz constant of Φ.

It remains to estimate $D\Psi$. We compute

$$|D\Psi(p)| \leq |\nabla(p)|\frac{|p - q(p)|}{\eta - \delta} + \max\{|Dq(p)|, |\mathrm{Id}|\}$$

in the range where $D\Psi$ is not the identity, which is contained in an η tubular neighborhood of some compact subset K of Σ. In particular, we can assume that, if A is the second fundamental form of Σ, $|A|$ is bounded over K by some fixed number M. We then use the well-known estimates

$$|\nabla d(p)| \leq 1 + C_1 M|p - q(p)| \leq 1 + C_1 M\eta$$

and

$$|\nabla q(p)| \leq |\mathrm{Id}| + C_1 M|p - q(p)| \leq C_1(1 + M\eta)\,,$$

where C_1 is a geometric constant. The claimed estimate for $D\Psi$ follows easily. □

Proof of Proposition 4.4. We consider the Radon measures $\mu_j := \mathcal{H}^k \llcorner E_j$ and, up to extraction of a subsequence, not relabeled, we assume that $\mu_j \rightharpoonup^* \mu$ for some Radon measure μ on U. Using Lemma A.1 and Allard's compactness theorem for stationary integral varifolds, cf. [1, Theorem 6.4], we know that, in every open set $U' \subset\subset U \setminus \Gamma$, we have a rectifiable set R and an integer-valued Borel function Θ such that $\mu \llcorner U' = \Theta\mathcal{H}^k \llcorner R$. Clearly the same conclusion

holds on the whole set $U \setminus \Gamma$ for a single integer rectifiable R with (locally) finite Hausdorff measure and an integer-valued Borel function. Since V is stationary in $U \setminus \Gamma$, we can assume that R is the intersection of the support of μ with $U \setminus \Gamma$ (hence a relatively closed set) and that Θ is the upper semicontinuous function

$$\Theta(p) := \lim_{r \downarrow 0} \frac{\mu(B_r(p))}{\omega_k r^k} .$$

Moreover, thanks to the lower density bounds, we know that $E_j \cap U'$ converge locally in the Hausdorff topology to $R \cap U'$. Indeed:

- If $p \in R \cap U'$ we must have that $\mu(B_r(p)) > 0$ for evert $r > 0$, in particular there must be at least one sequence of points $\{p_j\}$ with $p_j \in E_j$ and $p_j \in B_r(p)$. Otherwise there is a subsequence of E_j, not relabeled, with $E_j \cap B_r(p) = \emptyset$, leading to $\mu(B_r(p)) \leq \liminf_j \mathcal{H}^k(E_j \cap B_r(p)) = 0$.
- We now take $p \in U'$ which is the limit of some (subsequence of a) sequence $\{p_j\}$ with $p_j \in E_j$ and we wish to show that $p \in \operatorname{spt}(\mu)$, namely that $\mu(B_\rho(p)) > 0$ for every $\rho < \operatorname{dist}(p, U')$. We distinguish two cases.

  In the first case we assume $p \in U' \setminus \Gamma$. Then, for any $\rho > 0$ sufficiently small so that $\operatorname{dist}(p, \Gamma) \geq 2\rho$, we must have

  $$\mu(\overline{B}_\rho(p)) \geq \limsup_{j \to \infty} \mathcal{H}^k(E \cap \overline{B}_\rho(p)) \geq \limsup_{j \to \infty} \mathcal{H}^k(E_j \cap \overline{B}_{\rho/2}(p_j)) \geq C_0^{-1} \omega_k \rho^k ,$$

  by Lemma A.2. This implies $\mu(B_\rho(p)) > 0$ for every ρ, namely that p belongs to the support of μ.

  In the second case we assume that $p \in \Gamma \cap U'$ and fix ρ such that $2\rho < \operatorname{dist}(p, \partial U)$ .Then we let $q_j \in \Gamma_j$ be the closest point to p_j in Γ_j. Set $\varepsilon_j := |p_j - q_j|$, which clearly is converging to 0. Moreover $q_j \to p$, so $B_\rho(q_j) \subset U'$ for j sufficiently large. If $\varepsilon_j = 0$ then Lemma A.2 and the monotonicity formula implies immediately that

  $$\mathcal{H}^k(E_j \cap B_\rho(q_j)) \geq c_0 \rho^k .$$

  If $\varepsilon_j > 0$, note however that

  $$\mathcal{H}^k(E_j \cap B_{2\varepsilon_j}(q_j)) \geq \mathcal{H}^k(E_j \cap B_{\varepsilon_j}(p_j)) \geq C_0 \varepsilon_j^k$$

  because we can use the interior case of Lemma A.2 and the monotonicity formula for radii up to ε_j, given that $B_{\varepsilon_j}(p_j) \subset U' \setminus \Gamma$. But now we can use the boundary case of the monotonicity formula and $2\varepsilon_j \leq \rho/2$ (which holds for j large enough) to bound

  $$\mathcal{H}^k(E_j \cap B_{\rho/2}(q_j)) \geq C_1 \rho^k . \tag{A.2}$$

  We now have established the validity of (A.2) irrespectively of whether $q_j = p_j$ or not. Since $q_j \to q$ we can now argue as in the interior case and conclude that $\mu(\overline{B}_\rho(p)) > 0$ passing into the limit in j and using (A.2) together with $q_j \to q$ to bound

  $$\mathcal{H}^k(E_j \cap \overline{B}_\rho(p)) \geq C_1 \rho^k .$$

In the next step we show that the density Θ equals 1 almost everywhere. It suffices to show that the density is equal to one at every point where the rectifiable set R has a weak tangent plane π. Fix such a point and assume, without loss of generality, that it is the origin. Moreover define the rescaling maps $\lambda_r(x) := \frac{x}{r}$ We consider the rescaled integral varifolds with rectifiable sets given by $\lambda_r(R)$ and densities $\Theta \circ \lambda_r^{-1}$. These varifolds converge to the varifold induced by the plane π with constant density equal to $\Theta(0)$. Again by the

compactness of integral varifold we conclude that in fact $\Theta(0)$ is an integer, while the density lower bound implies that it is positive. Our aim is thus to exclude that $\Theta(0) \geq 2$. Using a simple diagonal argument, the measure given by $\nu = \Theta(0)\mathcal{H}^k \llcorner \pi$ is in fact the weak* limit of $\nu_j := \mathcal{H}^k \llcorner (\lambda_{r_j}(E_j))$ for an appropriately chosen sequence $r_j \downarrow 0$. Moreover, arguing as in the previous step, we can see that the sets $\lambda_{r_j}(E_j)$ must in fact be converging locally in the sense of Hausdorff to the plane π. We can now apply Lemma A.3 and, for every fixed $\varepsilon > 0$, find a Lipschitz map Φ with the following properties.

- $\mathrm{Lip}\,(\Phi) \leq C_0$;
- Φ is the identity outside B_1;
- Φ maps $B_\eta(\pi) \cap B_{1-\varepsilon}$ onto $\pi \cap B_{1-\varepsilon}$.

We now wish to estimate $\mathcal{H}^k(B_1 \cap \Phi(\lambda_{r_j}(E_j)))$. Given the properties of the map Φ we can easily come up with the upper bound

$$\mathcal{H}^k(B_1 \cap \Phi(\lambda_{r_j}(E_j))) \leq \omega_k(1-\varepsilon)^k + C_0^k \mathcal{H}^k((B_1 \setminus B_{1-\varepsilon}) \cap \lambda_{r_j}(E_j))$$

and using the properties of the weak* convergence of measures we infer

$$\limsup_{j\to\infty} \mathcal{H}^k(B_1 \cap \Phi(\lambda_{r_j}(E_j))) \leq \omega_k(1-\varepsilon)^k + C_0^k \Theta(0)(1-(1-\varepsilon)^k)\,.$$

On the other hand we also know that

$$\lim_{j\to\infty} \mathcal{H}^k(B_1 \cap \lambda_{r_j}(E_j)) = \Theta(0)\omega_k\,.$$

Given the minimizing property of E_j, which is invariant under dilations, we must have

$$\Theta(0) \leq (1-\varepsilon)^k + C_0^k \Theta(0)(1-(1-\varepsilon)^k)\,.$$

We now let $\varepsilon \downarrow 0$ to conclude that $\Theta(0) \leq 1$.

We can at this point set $E := \mathrm{spt}\,(\mu)$. In order to complete the proof of the proposition it suffices to show the following two things:

(a) $\mu(\Gamma) = 0$;
(b) E is minimizing with respect to Γ.

In order to show both properties we will use the following fact, a consequence of the C^2 convergence of Γ_j to Γ. There is a sequence of C^1 maps Ψ_j and a sequence of sets $\Omega_j \uparrow U$ with the following property:

- Ψ_j is a diffeomorphism of Ω_j with itself;
- $\Psi_j - \mathrm{Id}$ and $\Psi_j^{-1} - \mathrm{Id}$ converge to 0 in the C^1 norm;
- Ψ_j is a diffeomorphism of $\Omega_j \cap \Gamma_j$ with $\Omega_j \cap \Gamma$.

It is then easy to see that:

(i) $\mathcal{H}^k \llcorner \Psi_j(E_j \cap \Omega_j)$ converges weakly* to μ;
(ii) for every fixed Lipschitz map Φ which is the identity on $\Gamma \cup (U \setminus U')$ for some $U' \subset\subset U$ with $\mu(\partial U') = 0$, we must have

$$\liminf_{j\to\infty} (\mathcal{H}^k(U' \cap \Phi(\Psi_j(E_j \cap \Omega_k)) \geq \mu(U')\,. \tag{A.3}$$

Assume therefore that there is a $U' \subset\subset U$ such that $\mu(U' \cap \Gamma) > 0$. By varying U' slightly we can assume that $\mu(\partial U') = 0$. For every fixed $U'' \subset\subset U'$ and for every fixed $\varepsilon > 0$ we can find a $\delta > 0$ with the following property. For any $\eta < \delta$ there is a Lipschitz map Φ with the following properties:

- $\operatorname{Lip}\Phi \leq C_0$;
- Φ is the identity on the complement of U', on the complement of $B_\varepsilon(\Gamma \cap U')$ and on $B_{\kappa^2\delta}(\Gamma)$;
- Φ shrinks the normal neighborhood of $\Gamma \cap U'$ by a factor $C_0\kappa$ in $B_\delta(\Gamma \cap U') \setminus B_{\kappa^2\delta}(\Gamma)$.

Since Γ has dimension $k-1$, the tangent space to any k-dimensional set at a point $x \in B_\delta(\Gamma \cap U') \setminus B_{\kappa^2\delta}(\Gamma)$ contains at least one vector which is shrinked by the factor κ. Moreover as long as $\kappa > 0$ we can arrange that Φ satisfies the condition in Definition 1.2 to be an admissible competitor we can thus estimate

$$\begin{aligned}
0 \leq & \mathcal{H}^k(U' \cap \Phi(\Psi_j(E_j \cap \Omega_j))) - \mathcal{H}^k(U' \cap \Psi_j(E_j \cap \Omega_j)) \\
\leq & C_0^k \mathcal{H}^k(U'' \cap (\overline{B}_\eta(\Gamma) \setminus B_\delta(\Gamma)) \cap \Psi_j(E_j \cap \Omega_j)) + C_0^k \mathcal{H}^k((U' \setminus U'') \cap \Psi_j(E_j \cap \Omega_j)) \\
& - (1 - C_0\kappa)\mathcal{H}^k((B_\delta(\Gamma) \setminus B_{\kappa^2\delta}(\Gamma)) \cap U'' \cap \Psi_j(E_j \cap \Omega_j)) \,.
\end{aligned}$$

Next, by slightly varying the sets we can assume:

$$\mu(\partial U'') = \mu(\partial(B_\eta(\Gamma))) = \mu(\partial(B_\delta(\Gamma))) = 0 \,.$$

We now first let $\kappa \downarrow 0$ and then let $j \uparrow \infty$ to get

$$0 \leq C_0^k \mu(U'' \cap (\overline{B}_\eta(\Gamma) \setminus \Gamma)) + C_0^k \mu(U' \setminus U'') - \mu(\Gamma \cap U'') \,.$$

Now, letting $\eta \downarrow 0$ and $U'' \uparrow U'$, we get in the limit

$$0 \leq -\mu(\Gamma \cap U'') \,,$$

which clearly is a contradiction.

In order to prove (b) we will take advantage of Allard's regularity theory. Given that the density is 1 a.e. on the rectifiable set E, we actually know that there is a set $U' \subset U$, containing Γ with the property that $\mathcal{H}^k(E \cap (U \setminus U')) = 0$ and that $E \cap U'$ is a classical embedded minimal surface in U', i.e. an analytic submanifold.

Assume now, by contradiction, that E is not minimizing with respect to Γ and therefore that there is a map Ψ satisfying the condition of Definition 1.2 and with the property that

$$\mathcal{H}^k(\Psi(E) \cap W) < \mathcal{H}^k(E \cap W) \,.$$

Without loss of generality we can assume that $\mathcal{H}^k(E \cap \partial W) = 0$. We fix an open set $U'' \subset\subset W \cap U'$ and use Lemma A.3 to find a Lipschiz map Ψ which maps some δ-neighborhood of $E \cap U''$ onto E, is the identity outside U', and has Lipschitz constant at most C_0. We can then estimate that

$$\mathcal{H}^k(\Psi \circ \Phi(\Phi_j(E_j \cap \Omega_j)) \cap W) \leq C_0^k \operatorname{Lip}(\Phi) \mathcal{H}^k(E_j \cap (W \setminus U'')) + \mathcal{H}^k(\Psi(E) \cap W) \,.$$

Letting $j \uparrow \infty$ and using the minimizing condition we infer

$$0 \leq C_0^k \operatorname{Lip}(\Phi) \mathcal{H}^k(E \cap (W \setminus U'')) + \mathcal{H}^k(\Psi(E) \cap W) - \mathcal{H}^k(E \cap W) \,.$$

On the other hand we can let $U'' \uparrow U'$ and conclude that

$$0 \leq C_0^k \operatorname{Lip}(\Phi) \mathcal{H}^k(E \cap (W \setminus U')) + \mathcal{H}^k(\Psi(E) \cap W) - \mathcal{H}^k(E \cap W) \,.$$

Since $W \setminus U' \subset U \setminus U'$ and $\mathcal{H}^k(E \cap (U \setminus U')) = 0$, we thus conclude

$$0 \leq \mathcal{H}^k(\Psi(E) \cap W) - \mathcal{H}^k(E \cap W) \,.$$

This gives a contradiction and hence completes the proof. □

Appendix B. Convex spherical open sets and positively curved arcs

First of all, given an open set $\Omega \subset \mathbb{S}^2$ we say that it is convex if and only if the open cone with link Ω and vertex the origin is a convex subset of $\mathbb{R}^3$. As it can be easily seen using separating planes, this is equivalent to being the intersection of a collection of open hemispheres of $\mathbb{S}^2$. For open sets which is not the whole sphere this is also equivalent to not containing any pair of antipodal points and always containing the length-minimizing geodesic arc connecting any two points.

The following criterion for recognizing open convex sets plays a crucial role in several places and will be proved in this section (the statement is folklore but we have not been able to find it in the literature; in the Euclidean space it is a classical result proved independently by Nakajima and Tietze, see e.g. [26]).

Lemma B.1. *Let $\Omega \subseteq \mathbb{S}^2$ be a connected open domain, which is* weakly locally convex*, i.e.:*

(WLC) For every $p \in \partial\Omega$ there is a neighborhood U such that $U \cap \Omega$ has finitely many connected components which are all convex.

Then Ω is convex. In particular $\overline{\Omega}$ is convex and, if $\Omega \neq \mathbb{S}^2$, then $\partial\Omega$ is a Jordan curve.

Closely connected to the notion of convexity is the following one of positively curved oriented arcs, cf. Figure 33.

Definition B.2. Let $\gamma : [a,b] \to \mathbb{S}^2$ be a continuous locally injective map. We will say that γ is a positively curved oriented arc if for every $T_0 \in [a,b]$ there are an equator $\alpha \subset \mathbb{S}^2$ and a positive number ε with the following properties:

(a) α goes through $\gamma(t_0)$;
(b) Upon setting $I := [t_0 - \varepsilon, t_0 + \varepsilon] \cap [a,b]$ the arc $\gamma(I)$ is contained in a closed hemisphere H bounded by α;
(c) $\gamma(I)$ is a normal graph over α, namely the nearest point projection $\mathbf{p} : \gamma(I) \to \alpha$ is well defined and injective;
(d) If $\eta : (-\pi, \pi)$ is the arclength parametrization of α with $\dot\eta \times \ddot\eta$ pointing towards H and $\eta(0) = \gamma(t_0)$, then $\gamma^{-1} \circ \mathbf{p}^{-1} \circ \eta$ is a strictly increasing function from the interval $\eta^{-1}(\mathbf{p}(\gamma(I)))$ to I.

Whether an oriented arc is positively curved is invariant under parametrizations of the arc which are orientation preserving. The property can be easily extended to oriented loops. The following lemma plays an important role in our arguments.

Lemma B.3. *Let γ be a positively curved oriented arc which is embedded and whose endpoints are antipodal. Then γ is necessarily half of an equator.*

The rest of this section is devoted to the proofs of Lemma B.1 and Lemma B.3.

B.1. Elementary remarks.

First of all we record the following remarks, which are simple consequences of standard facts in convex geometry

Proposition B.4.

(i) A connected open set $\Omega \subsetneq \mathbb{S}^2$ is convex if and only it is contained in at least one open hemisphere and its boundary (oriented so that the unit normal points inwards) is a positively curved oriented loop.
(ii) If $\Omega \subsetneq \mathbb{S}^2$ is open and convex than its boundary is a Jordan curve.

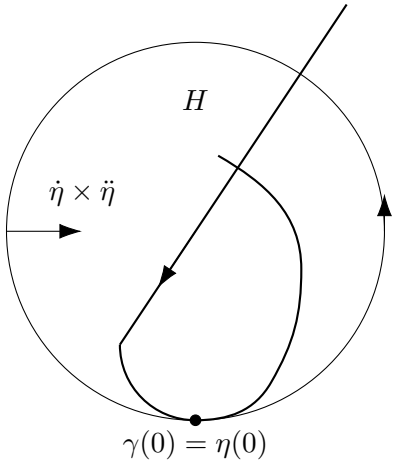


FIGURE 33. An example of a positively curved oriented arc γ, in a stereographic projection of the sphere. Note that the arc is not necessarily embedded (even though it is locally embedded) and not necessarily differentiable. η parametrizes an equator with the property that, locally at $\gamma(0)$, the curve γ is locally a graph over the equator and all "on the left side" of it (assuming γ and η are compatibly oriented). The property is only local, e.g. portions of the curve away from $\gamma(0)$ are allowed to escape the hemisphere H.

(iii) If γ is a positively curved oriented arc, then it has finite length and the normal graph in Definition B.2(c) is Lipschitz.

(iv) Consider a positively curved oriented arc $\gamma : [0, L] \to \mathbb{S}^2$ that is parametrized by arc-length. For any $t_0 \in [0, L]$, the following statements hold:
- *γ is left and right differentiable at t_0.*
- *The properties (c) and (d) in Definition B.2 are valid for a sufficiently small ε and any equator passing through $\gamma(t_0)$ and tangent to $\dot\gamma(t_0^-)$ or $\dot\gamma(t_0^+)$.*
- *When $t_0 \in (0, L)$ the oriented angle between $v^- = \dot\gamma(t_0^-)$ and $v^+ = \dot\gamma(t_0^+)$ is strictly smaller than π and the properties (c) and (d) in Definition B.2 are valid for a sufficiently small ε and any equator passing through $\gamma(t_0)$ and tangent to any convex combination of v^- and v^+.*

(v) If A and B are two disjoint open convex sets, then $\overline{A} \cap \overline{B}$ is either an equator, or a geodesic arc long at most π, or two antipodal points, or a singleton, or the empty set.

Proof. **Proof of (i).** The statement follows immediately from a general theorem of Karcher, cf. [15, Proposition 1] on open connected sets of Riemannian manifolds: the latter states that an open connected set is convex if and only if it does not intersect its own cut locus and has a support element at every boundary point. The notion of support element is equivalent, in our context, to the existence of the equator α as in Definition B.2.

Proof of (ii). The statement follows immediately from the fact that the open cone with link Ω and vertex the origin is a convex set: this implies that Ω is topologically a disk and that its boundary is locally Lipschitz, in particular $\partial\Omega$ is an embedded loop.

Proof of (iii). By compactness, without loss of generality we can assume that the arc is embedded. Consider the cone E with vertex the origin and link γ. Fix any t_0. Then we

see immediately that the latter is, in a neighborhood of $\gamma(t_0)$ the graph of a convex function φ over the plane π containing the equator α. It is therefore a Lipschitz graph in a slightly smaller neighborhood and the claim follows at once.

Proof of (iv). Consider φ as in (iii). Let D be the subdifferential of φ at $\gamma(t_0)$, or equivalently consider the supporting planes of the graph at the point. It follows immediately from the invariance under dilation that any supporting plane must contain the vector $\gamma(t_0)$. This constrains immediately the subdifferential D to be either a point (in which case φ is differential at $\gamma(t_0)$) or a closed segment and the claim follows easily.

Proof of (v). We denote by A' and B' the open convex cones with links A and B and the origin as vertex. We note that $\overline{A} \cap \overline{B} = \overline{A}' \cap \overline{B}' \cap \mathbb{S}^2$. Since $\overline{A}'$ and $\overline{B}'$ are convex, so is $D = \overline{A}' \cap \overline{B}'$. On other hand D has also empty interior, so it is either empty, or of dimension 1 (and hence contained in an affine line l) or of dimension 2 and contained in some 2-dimensional affine plane α. Since it is also a cone, namely $\lambda D = D$ for every $\lambda > 0$, the plane α and the line l must is in fact be linear. If D is not empty we have therefore the following four possibilities:

- D coincides with the plane α and thus $D \cap \mathbb{S}^2$ is an equator.
- D is 2-dimensional, but a strict subset of the plane; then $D \cap \mathbb{S}^2$ is a geodesic arc of length at most π.
- D is the 1-dimensional line l, and hence $D \cap \mathbb{S}^2$ consists of two antipodal poins.
- D is 1-dimensional, but a strict subset of the line l; then D is necessarily a half-line starting at the origin, hence $D \cap \mathbb{S}^2$ is a singleton.

□

B.2. **Proof of Lemma B.1. Step 1. Locally convex case.** We first prove the lemma under the following stronger assumption:

(LC) for every $q \in \partial\Omega$ there is a neighborhood U of q and an open hemisphere H such that $q \in \partial H$ and $U \cap \Omega \subset H$.

Fix indeed a point $q \in \partial\Omega$ and an hemisphere H as in (LC). Let $\Omega' = \Omega \cap H$. Observe that the property (LC) still holds for Ω' in place of Ω and moreover Ω' does not contain any pair of antipodal points. We can therefore apply [15, Proposition 1] and conclude that Ω' is convex. Let now $H'' := \mathbb{S}^2 \setminus \overline{H}$. If $\Omega'' := H'' \cap \Omega$ is empty, then we are done. Assume otherwise that Ω'' is nonempty. Then H'' is also convex. Consider that $\alpha := \overline{\Omega}' \cap \overline{\Omega}''$ is a subset of the equator $\gamma := \overline{H} \cap \overline{H}''$. Since Ω is connected, it must contain a point of γ and thus α is necessarily one-dimensional. On the other hand $\overline{\Omega}''$ cannot contain the point q because $U \cap \Omega \subset H$ for some neighborhood U of q. From Proposition B.4(v) we conclude that α is an arc of γ, with length at most π. By convexity $\beta := \overline{\Omega}' \cap \gamma$ is also an arc and it must contain α strictly. So there is an endpoint $\bar{q}$ of α which is in the interior of β By (LC) there must be one hemisphere H''' which contains $B_r(\bar{q}) \cap \Omega$. On the other hand, given that β contains $\bar{q}$ in the interior and Ω' is convex, we easily see that there is $\rho > 0$ such that $B_\rho(\bar{q}) \cap \Omega' = H \cap \Omega'$. So H''' must necessarily coincide with H. But by definition this would imply $q \notin \overline{\Omega}''$, which is a contradiction.

Step 2. Regularization. In the general case we will first approximate Ω with a new connected open set which satisfies the local convexity property (LC). First of all consider the

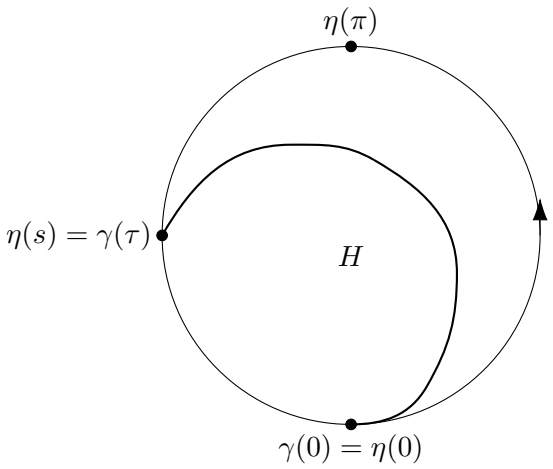


FIGURE 34. A picture illustrating the Exit Lemma B.5, using a stereographic projection of the half-sphere H. The boundary of ∂H is oriented in such a way that the normal points outward. The "exit point" $\gamma(\tau)$ of the curve γ is "past" the antipodal point $\eta(\pi)$ of the "entry point" $\eta(0) = \gamma(0)$.

function $f : \mathbb{S}^2 \to \mathbb{R}$ given by $f(p) = \operatorname{dist}(p, \partial\Omega)$, where the distance is the geodesic distance on $\mathbb{S}^2$. Observe that, for every $p \in \partial\Omega$ there is a neighborhood U of p with the property that:

- f is concave on any connected component of $U \cap \Omega$.

In particular, by a simple compactness argument, there is a positive number $\varepsilon_0 > 0$ such that, for every positive $\varepsilon < \varepsilon_0$, the set $\Omega_\varepsilon := \{f > \varepsilon\}$ satisfies (LC). We can now argue as in Step 1 and conclude that any connected component of Ω_ε is an open convex set. It is elementary to construct a sequence $\Lambda_k \uparrow \Omega$ of connected open sets with the property that $\Lambda_k \subset \Omega_{\varepsilon_k}$ for some sequence $\varepsilon_k \downarrow 0$. In particular, if Ω'_{ε_k} is the connected component containing Λ_k, we then conclude $\Omega'_{\varepsilon_k} \uparrow \Omega$. We now show that Ω is convex. For ever point $q \notin \Omega$ we clearly have $q \notin \Omega'_{\varepsilon_k}$. Thus, we can find an open hemisphere H_k which contains Ω_{ε_k} but does not contain q. Up to subsequences we can assume that the hemispheres H_k converge to some hemisphere H^q. We see immediately that $\Omega \subset H^q$ and $q \notin H^q$. The intersection of all H^q's must then be Ω, this showing that the latter must be convex.

B.3. **Exit lemma.** The main tool to prove Lemma B.3 is the following.

Lemma B.5 (Exit lemma)**.** *Let $H \subset \mathbb{R}^2$ be an open hemisphere and let $\gamma : [0, L] \to \mathbb{S}^2$ be an arc-length parametrization of positively curved oriented arc of $\mathbb{S}^2$. Assume that $\gamma(0) \in \partial H$, $\gamma((0, \delta)) \subset H$ for some positive δ, but $\gamma((0, L])$ is not contained in H. Let $\eta : [0, 2\pi) \to \partial H$ be a parametrization by arclength of ∂H such that $\eta(0) = \gamma(0)$ and $\dot\eta(0) \times \ddot\eta(0)$ points towards H. Let $\tau > 0$ be the smallest positive real such that $\gamma(\tau) \in \partial H$. Then $\gamma(\tau) = \eta(s)$ for some $s \geq \pi$. Moreover, if $s = \pi$ then $\gamma([0, \tau])$ is half of an equator passing through the antipodal points $\gamma(0)$ and $\gamma(\tau)$. See Figure 34.*

Proof. Fix γ, η, τ, and s as in the statement. Assume moreover that $s \leq \pi$. First of all observe that $\dot\gamma(0^+) \neq -\dot\eta(0)$ and $\dot\gamma(\tau^-) \neq -\dot\eta(s)$ (recalling Proposition B.4 γ is left and right differentiable at every point). We choose coordinates in $\mathbb{R}^3$ such that $H = \{v \in \mathbb{S}^2 :$

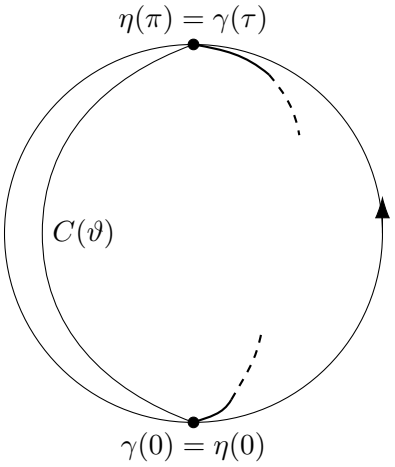


FIGURE 35. A stereographic projection of the hemisphere H. We are assuming that $\gamma(\tau) = \eta(s)$ for $s \leq \pi$. Given the angle formed by $\dot\gamma(0)$ and $\dot\eta(0)$ and, in case $s = \pi$, the one formed by $\dot\gamma(s)$ and $\dot\eta(\pi)$, for a sufficiently small ϑ the half-equator $C(\vartheta)$ does not intersect $\gamma((0, s))$.

$\langle v, e_3\rangle > 0\}$ and $\gamma(0) = e_1$. In particular $\eta(0) = e_1$ while $\eta(\frac{\pi}{2}) = e_2$ For $\vartheta > 0$ we let $e(\vartheta) = \cos\vartheta e_3 + \sin\vartheta e_2$ and $C(\vartheta) = \{v \in \mathbb{S}^2 : \langle v, e(\vartheta)\rangle = 0 \text{ and } \langle v, e_3\rangle \geq 0\}$. For small enough ϑ the intersection of $\gamma([0,\tau])$ with $C(\vartheta)$ is contained in $\{\gamma(0), \gamma(\tau)\}$, cf. Figure 35.

Let ϑ_0 be the largest ϑ with the property that $\gamma([0,\tau])$ is contained in the closure of the open convex set Ω bounded by $C(\vartheta) \cup \eta([0,\pi])$. Let $\tilde\eta$ be a parametrization by arc-length of $\partial\Omega$ which makes it a positively curved oriented loop as in Definition B.2 such that $\eta(0) = \tilde\eta(0)$ and $\eta(t) = \tilde\eta(t)$ for all $t \in [0,\pi]$.

Note that ϑ must be strictly smaller than π. Hence we cannot have $\dot\gamma(0^+) = \dot{\tilde\eta}(2\pi^-)$, given that the latter vector points towards $\mathbb{S}^2 \setminus H$. On the other hand, if $\dot\gamma(0^+) = -\dot{\tilde\eta}(2\pi^-)$, then we would in fact conclude that $\gamma([0,\tau])$ lies in $C(\vartheta)$. Likewise, if it is the case that $s = \pi$, $\dot\gamma(\tau^-)$ must be different from $\dot{\tilde\eta}(\pi^+)$, while it can be $-\dot{\tilde\eta}(\pi^+)$ only if $\gamma([0,\tau])$ coincides with $C(\vartheta)$. In all the remaining hypothetical cases ($s < \pi$ and $s = \pi$ but $\gamma([0,\tau])$ is not half of an equator) we conclude that there must be at least one third point $t \in (0,\tau)$ such that $\gamma(t) \in C(\vartheta)$, namely $\gamma(t) = \tilde\eta(t')$ for some $t' \in (\pi, 2\pi)$. This would however force γ to have a self-intersection, which is a contradiction, cf. Figure 36. □

B.4. **Proof of Lemma B.3.** Fix an arclength parametrization γ over $[0, L]$ of an arc as in the statement. Recalling Proposition B.4 γ is left and right differentiable at every point and in particular it is differentiable at the two endpoints. Let n be the tangent unit vector to $\mathbb{S}^2$ which is the counterclockwise rotation of $\dot\gamma(0)$ by $\frac{\pi}{2}$ and let H be the open hemisphere $\{p \in \mathbb{S}^2 : \langle p, n\rangle < 0\}$. If γ is entirely contained in $\overline{H}$ the conclusion follows immediately from the Exit Lemma B.5. Otherwise let $[0, t_1]$ be a maximal interval with the property that $\gamma([0,t_1]) \subset \overline{H}$. Parametrize ∂H with η as in the Exit Lemma and observe that we must necessarily have $\gamma(t_1) = \eta(s_1)$ for some $s_1 \in (\pi, 2\pi)$. Consider now a maximal interval $[t_1, t_2]$ with the property that $\gamma([t_1,t_2]) \subset \mathbb{S}^2 \setminus H$ and again it must be that $\gamma(t_2) = \eta(s_2)$ for some $s_2 \in (0, 2\pi]$. We can again apply the exit lemma to the open hemisphere $\mathbb{S}^2 \setminus \overline{H}$. Consider

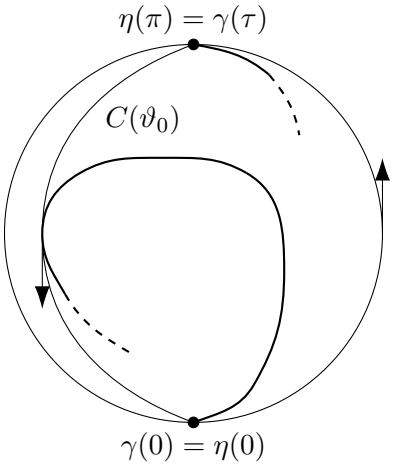


FIGURE 36. A picture of $C(\vartheta_0)$ in the hemisphere H, under the assumption that $\gamma(\tau) = \eta(s)$ for $s \leq \pi$. The existence of the point $\gamma(t_0) = \tilde{\eta}(s_0) \in C(\vartheta_0)$ forces a self-intersection.

however that a parametrization as in Definition B.2 for the latter is now $-\eta$. In particular we conclude that $s_2 \in (s_1, 2\pi) \cup (0, s_1 - \pi]$. We now distinguish two cases.

Case 1 $s_2 \in (s_1, 2\pi)$. Clearly then $\gamma([s_1, L])$ cannot be contained in $\overline{H}$ because, since $\gamma(L) = \eta(\pi)$, this would violate the injectivity of γ. Moreover, for the same reason, if we let $[t_2, t_3]$ be the maximal interval such that $\gamma([t_2, t_3]) \subset \overline{H}$, we must have $\gamma(t_3) = \eta(s_3)$ for some $s_3 \in (s_1, s_2)$. But then it cannot be that $\gamma([t_3, L])$ is entirely contained in $\overline{H}$ and arguing in the same way we would conclude the existence of a maximal interval $[t_3, t_4]$ with $\gamma([t_3, t_4]) \subset \mathbb{S}^2 \setminus H$ and $\gamma(t_4) = \eta(s_4)$ for some $s_4 \in (s_3, s_2)$. Cf. the left picture in Figure 37. This process can be iterated indefinitely leading to the conclusion that γ cannot reach the antipodal of $\gamma(0)$.

Case 2 $s_2 \in (0, s_1 - \pi]$. In this case we conclude from the Exit Lemma that, if $[t_2, t_3]$ is the the maximal interval such that $\gamma([t_2, t_3]) \subset \overline{H}$, then $\gamma(t_3) = \eta(s_3)$ for some $s_3 \in [s_2 + \pi, 2\pi)$. On the other hand, since γ is injective, we must also have $s_3 < s_1$ and thus $s_3 \in [s_2 + \pi, s_1)$. We can now keep iterating as above and prove inductively that:

- When we are looking at the a maximal intervals $[t_{2k}, t_{2k+1}]$ with $\gamma([t_{2k}, t_{2k+1}]) \subset \overline{H}$ and $[t_{2k+1}, t_{2k+2}]$ with $\gamma([t_{2k+1}, t_{2k+2}] \subset \mathbb{S}^2 \setminus H$, we have $s_{2k} \geq s_{2k-2} > 0$ and $s_{2k+1} > s_{2k-1}$ by injectivity of γ.
- The Exit Lemma then implies that $s_{2k+1} \geq s_{2k} + \pi > s_0 + \pi = \pi$, while $s_{2k+2} \leq s_{2k+1} - \pi \leq s_1 - \pi < 2\pi - \pi = \pi$.

Cf. the right picture in Figure 37. In particular we conclude again that γ cannot reach the point $\eta(\pi)$, namely the antipodal of $\gamma(0)$.

APPENDIX C. PROOF OF PROPOSITION 6.8

The two formulae (6.2) and (6.3) will be proven at the end and the argument for (6.2) will show first that opposite sides of the quadrilateral are congruent.

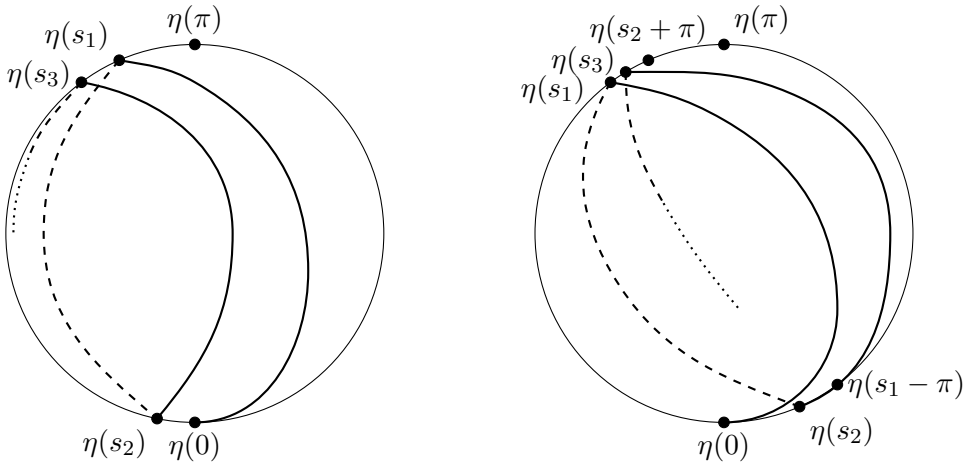


FIGURE 37. This picture illustrates the argument by contradiction that γ is contained in the closure of the hemisphere H in the proof of Lemma B.3. On the left we see an illustration of Case 1, while on the right we see an illustration of Case 2. The circle serves as a projection of both $\overline{H}$ and $\mathbb{S}^2 \setminus H$. The continuous arcs are those contained in $\overline{H}$, while the dashed arcs are those contained in $\mathbb{S}^2 \setminus H$.

The statement about equiangular triangles is an immediate consequences of the formulae in Theorem 6.4. Next, if two adjacent sides of a convex quadrilateral with angles equal $\frac{2\pi}{3}$ are equal, it follows from the congruence between opposite dies that all four sides are congruent. Draw therefore a diagonal of the quadrilateral. Such diagonal is then also a bisector of the angles at the vertices that it connects. If we label the vertices consecutively as $ABCD$ and assume that the diagonal connects A and C, we then have that ABC is a convex triangle with angles at A and C congruent and equal to $\frac{\pi}{3}$. It then follows immediately from Theorem 6.4 that the diagonal AC measures λ_T.

The fact that the lengths of two adjacent sides of a convex pentagon with angles $\frac{2\pi}{3}$ determine the pentagon is a very particular case of the more general Theorem 6.7. If we know the length of two sides which are not adjacent, we can use formula (6.3) to determine the length of a third side (cf. the function `pentagon2Constraint` in `metrization_constraints.cpp` in [13]). This side must then be adjacent to one of the other two and hence we have determined the pentagon thanks to Theorem 6.7.

We now come to (6.2) and we will prove at the same time that opposite sides in the quadrilateral are congruent. Let $ABCD$ be the quadrilateral with vertices labeled clockwise. We let E and E' be the two antipodal points which are the intersections of the two equators containing AD and BC (note that the convexity excludes that this equator could be the same). We label these two intersections in such a way that the interiors of the triangles EAB and $E'CD$ both are disjoint from the quadrilateral. Note that, in these triangles, the (interior) angles at E and E' are congruent and that the angles at A, B, C, and D all measure $\frac{\pi}{3}$. By Theorem 6.4 the two triangles are congruent (which shows the congruence of the sides AB and CD) and isosceles. Now let γ be the equator bisecting the angles in E and E' and let H be its intersection of the equator with the side AB. Since ABE is isosceles, the triangle

EHA is rectangle in H. Moreover, if we denote by a the length of AB and by b the length of AD we have:

$$|AH| = \frac{a}{2}$$
$$|EA| = \frac{\pi}{2} - \frac{b}{2}\,.$$

Let α be the angle at E of the triangle EHA. We first apply the law of sines to get

$$\frac{\sin\frac{a}{2}}{\sin\alpha} = \frac{\sin(\frac{\pi}{2} - \frac{b}{2})}{\sin\frac{\pi}{2}}\,,$$

which we rewrite as

$$\frac{\sin\frac{a}{2}}{\cos\frac{b}{2}} = \sin\alpha\,. \tag{C.1}$$

We next use Theorem 6.4 a first time to get

$$\cos|EH| = \cos\tfrac{a}{2} = \frac{\cos\alpha + \cos\frac{\pi}{3}\cos\frac{\pi}{2}}{\sin\frac{\pi}{3}\sin\frac{\pi}{2}}\,,$$

which we rewrite as

$$\cos\tfrac{a}{2} = \frac{\cos\alpha}{\sin\frac{\pi}{3}}\,, \tag{C.2}$$

and a second time to get

$$\cos|EA| = \cos(\tfrac{\pi}{2} - \tfrac{b}{2}) = \frac{\cos\frac{\pi}{2} + \cos\frac{\pi}{3}\cos\alpha}{\sin\frac{\pi}{3}\sin\alpha}\,,$$

which we rewrite as

$$\sin\tfrac{b}{2} = \frac{\cos\frac{\pi}{3}\cos\alpha}{\sin\frac{\pi}{3}\sin\alpha}\,. \tag{C.3}$$

We now first divide (C.3) by (C.2) in order to get

$$\frac{\sin\frac{b}{2}}{\cos\frac{a}{2}} = \frac{\cos\frac{\pi}{3}}{\sin\alpha} \tag{C.4}$$

and then multiply the latter by (C.1) to get

$$\tan\tfrac{a}{2}\tan\tfrac{b}{2} = \cos\tfrac{\pi}{3} = \tfrac{1}{2}\,.$$

We then come to (6.3). We label the vertices of the pentagon $P = ABCDE$ counterclockwise and we set $b = |AB|$, $a = |BC|$, and $c = |DE|$. We then first prolong the edges AB and DC "after" B and C until they meet in a point F, forming a (convex) quadrilateral $Q = AFDE$. Hence we prolong AE and DF on both sides until they meet in the (antipodal) points G (on the side of E and D) and G'. Clearly, in this way we find a convex region Ω delimited by two half-equators joining the antipodal points G and G' and containing Q (which in turn contains the initial pentagon P). We denote by

- γ the internal angles at G and G' of Ω;
- α the internal angle at F of Q.

Note moreover that the two triangles GDE and $G'FA$ have, respectively:

- Internal angles at E, D, and A measuring $\frac{\pi}{3}$;
- Internal angle at G and G' measuring γ;
- Internal angle at F measuring $\pi - \alpha$.

Applying Theorem 6.4 to the triangle GDE we can compute

$$\cos\gamma = \cos|ED|\sin^2\frac{\pi}{3} - \cos^2\frac{\pi}{3} = \tfrac{3}{4}\cos c - \tfrac{1}{4}\,. \tag{C.5}$$

If we denote by $\bar b$ the length of the side AF, we can do the same to the triangle $G'FA$ to derive

$$\begin{aligned}\cos\gamma &= \cos|AF|\cos\tfrac{\pi}{3}\sin(\pi-\alpha) - \cos\tfrac{\pi}{3}\cos(\pi-\alpha)\\ &= \tfrac{\sqrt3}{2}\cos\bar b\sin\alpha + \tfrac{1}{2}\cos\alpha\,.\end{aligned} \tag{C.6}$$

In particular we can use (C.5) and (C.6) to eliminate γ and derive

$$\tfrac{3}{4}\cos c - \tfrac{1}{4} = \tfrac{\sqrt3}{2}\cos\bar b\sin\alpha + \tfrac{1}{2}\cos\alpha\,. \tag{C.7}$$

We now use Theorem 6.4 on the triangle BCF, which is isosceles and has angles at B and C measuring $\frac{\pi}{3}$. We then can get

$$\cos\alpha = \cos a\sin^2\tfrac{\pi}{3} - \cos^2\tfrac{\pi}{3} = \tfrac{4}{3}\cos a - \tfrac{1}{4}\,. \tag{C.8}$$

Substituting the latter in (C.7) we then get

$$\tfrac{3}{4}\cos c - \tfrac{3}{8}\cos a - \tfrac{1}{8} = \tfrac{\sqrt3}{2}\cos\bar b\sin\alpha\,. \tag{C.9}$$

we next wish to compute $\cos\bar b$. We note that $\bar b = |AB| + |BF| = b + |BF|$. Setting $b' = |BF|$ we can use the formula for the sum of cosines and substitute it in (C.9) to find

$$\tfrac{3}{4}\cos c - \tfrac{3}{8}\cos a - \tfrac{1}{8} = \tfrac{\sqrt3}{2}\cos b\cos b'\sin\alpha - \tfrac{\sqrt3}{2}\sin b\sin b'\sin\alpha\,. \tag{C.10}$$

Applying the law of sines to the triangle CBF we can first find

$$\sin b'\sin\alpha = \sin a\sin\tfrac{\pi}{3} = \tfrac{\sqrt3}{2}\sin a\,,$$

which we can insert in (C.10) to get rid of $\sin b\sin\alpha$:

$$\tfrac{3}{4}\cos c - \tfrac{3}{8}\cos a - \tfrac{1}{8} + \tfrac{3}{4}\sin b\sin a = \tfrac{\sqrt3}{2}\cos b\cos b'\sin\alpha\,. \tag{C.11}$$

Now we can again apply Theorem 6.4 to deduce

$$\cos b'\sin\alpha\sin\tfrac{\pi}{3} = \cos\tfrac{\pi}{3} + \cos\alpha\cos\tfrac{\pi}{3}\,.$$

This can be rewritten, using (C.8), as

$$\cos b'\sin\alpha = \tfrac{1}{\sqrt3}(1+\cos\alpha) = \tfrac{1}{\sqrt3}(\tfrac{3}{4}\cos a + \tfrac{3}{4})\,.$$

We then complete the proof using the latter identity to get rid of the $\cos b'\sin\alpha$ term in (C.11) getting to

$$\tfrac{3}{4}\cos c - \tfrac{3}{8}\cos a - \tfrac{1}{8} + \tfrac{3}{4}\sin b\sin a = \tfrac{3}{8}\cos a\cos b + \tfrac{3}{8}\,.$$

This is easily seen to be equivalent to (6.3).

Finally, it remains to show that all sides of a pentagon are strictly shorter than λ_T. Thanks to (6.3) we know

$$2\cos(DE) - \frac{1}{3} = f(AB, BC)$$

and thus $DE < \lambda_T = \arccos(-\frac{1}{3})$ is equivalent to

$$f(AB, BC) + 1 > 0. \tag{C.12}$$

Thanks to the trigonometric identity

$$f(x,y)+1=4\cos^2\frac{x}{2}\cos^2\frac{y}{2}\Big(1-2\tan\frac{x}{2}\tan\frac{y}{2}\Big),$$

we deduce that Equation (C.12) is equivalent to

$$\tan\frac{AB}{2}\tan\frac{BC}{2}<\frac{1}{2}\,.$$

To prove the latter, we use what we learned for quadrilaterals. Assume, by contradiction, that $\tan\frac{AB}{2}\tan\frac{BC}{2}\geq\frac{1}{2}$; then find A',C' on AB, BC respectively such that $\tan\frac{A'B}{2}\tan\frac{BC'}{2}=\frac{1}{2}$. By convexity, the triangle $A'BC'$ is strictly contained in the pentagon. Moreover, recalling Equation (6.2), we know that $A'BC'$ is exactly half of a convex spherical quadrilateral with angles equal to $\frac{2\pi}{3}$. We find the contradiction by looking at the areas:

$$\frac{1}{2}\frac{2\pi}{3}=\mathrm{Area}(A'BC')<\mathrm{Area}(ABCDE)=\frac{\pi}{3}.$$

Appendix D. Proof of Theorem 6.7

Observe that we can assume $N\geq 4$, because the case $N=3$ is already covered by Theorem 6.4. Clearly, it suffices to show that the lengths of the three remaining arcs are determined. Consider now

- the equator α containing V_{N-2} and V_{N-1} (which is uniquely determined because the length of $[V_{N-2},V_{N-1}]$ is strictly shorter than π);
- the equator β containing V_1 and V_N.

The two equators α and β must be distinct, otherwise the four points would be all on the same equator and hence the three arcs $[V_{N-2},V_{N-1}]$, $[V_{N-1},V_N]$, and $[V_N,V_1]$ would also be contained in that equator (given that they are all shorter than π). This would mean that the angles at V_{N-1} and V_N are equal to 0 or to π, which is excluded by assumption (the angles at the vertices of a closed polygonal path are supposed to be distinct from 0 and π). For the same reason it cannot be that V_{N-1} belongs to β or that V_N belongs to α. In particular:

- $\mathbb{S}^2\setminus(\alpha\cup\beta)$ consists of four convex open connected sets;
- one of them, which we call Ω, is determined by the property that both V_{N-1} and V_{N-2} are in its boundary;
- the arc $[V_{N-1},V_N]$ is contained in $\overline{\Omega}$ (because the latter is convex and the arc is length-minimizing).

Observe next that another connected component, which we call Ω', is determined by the following property:

- $\partial\Omega'$ contains V_1, V_{N-2}, a portion of $[V_1,V_N]$ and a portion of $[V_{N-2},V_{N-1}]$.

Let A and B be the two antipodal points at the intersection of α and β. In particular we mark A so that if we travel from V_1 towards V_N along the arc $[V_1,V_N]$, then we get closer to A. We now distinguish eight cases depending on:

(i) whether traveling along the arc $[V_{N-2},V_{N-1}]$ from V_{N-2} to V_{N-1} we move towards A (we will then say that the arcs point in the same direction) or B;
(ii) the relative position of Ω and Ω'.

First of all we see that (i) is in fact determined by the rest of the polygonal path. More precisely, consider the original situation in which we had another polygonal path $W_1,\dots,W_N$ with the same angles and the first $N-3$ segments which are congruent and move the second

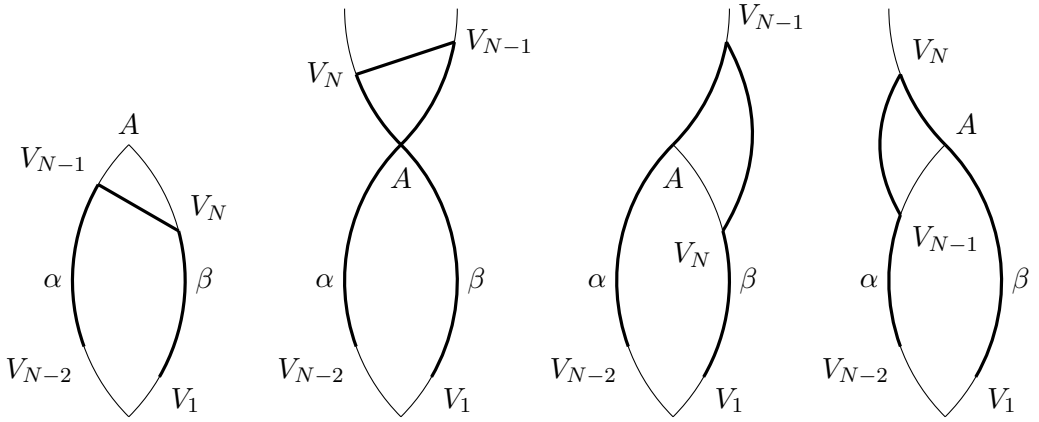


FIGURE 38. Given the relative orientation of the arcs $[V_{N-2}, V_{N-1}]$ and $[V_1, V_N]$ the relative position of Ω and Ω' is determined by the concavity of the angles in V_{N-2} and V_{N-1}. The picture illustrates the situation when the arcs $[V_{N-2}, V_{N-1}]$ and $[V_1, V_N]$ point in the same direction. From the far left to the right we see: $\Omega = \Omega'$, when both angles are convex; $\Omega \cap \Omega' = \{A, B\}$ when both angles are concave; $\Omega \cap \Omega' \supset [V_1, V_N]$ when $\hat{V}_{N-1}$ is convex and $\hat{V}_N$ is concave; $\Omega \cap \Omega' \supset [V_{N-2}, V_{N-1}]$ when $\hat{V}_{N-1}$ is concave and $\hat{V}_N$ is convex.

path with an orientation preserving isometry S so that we make the first $N-2$ vertices coincide with $O(V_i) = W_i$ for all $i \in \{1, \dots, N-2\}$. Then we easily see that, starting from V_1 and traveling towards $O(W_N)$, we in fact move in the same direction as traveling towards V_N, while starting from V_{N-2} and traveling towards $O(W_{N-1})$ we move towards the same direction as traveling towards V_{N-1}. This holds because the angles at V_1 and V_{N-2} are the same as those at W_1 and W_{N-2} and because we are assuming that $N \geq 4$.

The two cases of (i) are entirely analogous and here we just focus on one of them, namely when the arcs $[V_{N-2}, V_{N-1}]$ and $[V_1, V_N]$ point in the same direction. We then observe that the relative position of Ω and Ω' is completely determined by the concavity/convexity of the two angles at the vertices V_{N-1} and V_{N-2}. This is illustrated in Figure 38.

At this point, having established in which of the four situations we are, we consider the convex triangle with vertices A, V_{N-1}, and V_N. We immediately see that the angles of this triangle are completely determined by the angles of the polygonal path at V_{N-1} and V_N. In particular its sides are also completely determined. This in turn determines the length of the arc $[V_{N-1}, V_N]$ and the relative position with respect to A of the points V_{N-1} and V_N on the equators α and β. In particular we have completely determined the path and hence proved the claim of the theorem.

APPENDIX E. RECOGNIZING ISOMORPHIC PLANAR MAPS

Let us present an algorithm which detects whether two combinatorial maps are isomorphic (cf. [18, 21]). The algorithm will be used in order to avoid "multiple counts" in both the list of twice-derived and once-derived candidates.

We first introduce the darts' graph of a graph.

Definition E.1 (Darts' graph)**.** Let (G, next) be a combinatorial map. Its associated Darts' graph is the oriented graph with $\text{Darts}(G)$ as vertices and such that there is a directed edge from d_1 to d_2 if either $\text{rev}(d_1) = d_2$ or $\text{next}(d_1) = d_2$. The edges of the Darts' graph keep track of whether they are rev-edges or next-edges.

Remark E.2 (Isomorphism test)**.** We are given two connected combinatorial maps (G, next) and (G', next') with at least one edge. Let us describe an algorithm, that runs in polynomial time, which determines whether they are isomorphic. Let us fix $d_0 \in \text{Darts}(G)$ and let us search for an isomorphism $f : \text{Darts}(G) \to \text{Darts}(G')$ such that $f(d_0) = d'_0$, where $d'_0 \in \text{Darts}(G')$ is an arbitrarily chosen dart. We will repeat this process for each $d'_0 \in \text{Darts}(G')$.

Now, run a depth-first search on the darts' graph associated to G, starting from d_0, which gives priority to rev-edges over next-edges. Let $d_0, d_1, \dots, d_{k-1}$ be the elements of $\text{Darts}(G)$ in the order they were visited by this search. Do the same for $\text{Darts}(G')$ starting from d'_0, obtaining $d'_0, d'_1, \dots, d'_{k-1}$. Then, the isomorphism we are looking for shall satisfy $f(d_i) = d'_i$ for all $i = 0, \dots, k-1$. In particular, we have identified the candidate isomorphism and we need only to check whether it works.

Note that the isomorphism test works for embeddings in $\mathbb{S}^2$, as the isomorphism of the embedding is equivalent to the isomorphism of the associated combinatorial maps, but it does not work for planar maps because of the special unbounded face. To make up for this issue, it is sufficient to ensure that the unbounded face of (G, next) is mapped into the unbounded face (G', next).

For our purposes (see Section 8.1), we need to perform the following operation: given a family of planar maps, we want to select a subset containing exactly one representative of each isomorphism class. If we do this naïvely, the time complexity would be quadratic in the size of the family.

To speed up the process, we come up with a *signature* for combinatorial maps and for planar maps. This is a $\mathbb{Z}$ invariant, namely a function which to each planar map X associate an integer, with the property that it takes the same value on any other planar map X' isomorphic to X. Having a signature allows us to solve the above-mentioned problem as follows: first we split the family in buckets indexed by the value of the signature, then we perform the isomorphism test on each pair in the same bucket.

We next describe the signature functions we employ for combinatorial maps. We choose two arbitrary integers p_{rev} and p_{next}. For a combinatorial map $X = (G, \text{next})$, let M_X be the matrix[12] such that

$$M_X(d, d') = \begin{cases} p_{\text{rev}} + p_{\text{next}} & \text{if } d' = \text{next}(d) = \text{rev}(d) \\ p_{\text{rev}} & \text{if } d' = \text{rev}(d) \\ p_{\text{next}} & \text{if } d' = \text{next}(d) \\ 0 & \text{otherwise} \end{cases} \tag{E.1}$$

Assume that X is isomorphic to $\bar{X}$ via the isomorphism $\sigma : \text{Darts}(X) \to \text{Darts}(\bar{X})$. Then $M_X(d, d') = M_{\bar{X}}(\sigma(d), \sigma(d'))$. It is then immediate to check that the following function Sign, taking values in $\mathbb{Z}$, is a signature:

$$\text{Sign}\,(X) := \sum_{d \in \text{Darts}(X)} M_X^{256}(d, \text{rev}(d)) \oplus M_X^{256}(d, \text{next}(d)) \tag{E.2}$$

[12] This is the adjacency matrix of the darts' graph, up to the choice of weights.

where M_X^k denotes the k-th power of the matrix M_X, and $\oplus$ denotes the bitwise exclusive OR[13].

Let us now present the signature for planar maps. Observe that the one given above would be a valid signature, but it would not distinguish two planar maps which are equal as combinatorial maps but have different unbounded faces. The signature for planar maps is the sum of the signature in (E.2) and

$$\sum_{d\in \text{unbounded}(X)} \bigoplus_{e\in \text{Darts}(X)} M_X^{256}(d,e)\,.$$

We leave to the reader to check that two isomorphic planar maps have the same signature.

The algorithms described above are used verbatim to avoid multiple count while producing the list of Proposition 8.2. It must be slightly modified in order to account for the special features of once-derived marked spherical graphs of Proposition 8.5: we just need to add the requirement that, when the algorithm checks whether a candidate is in fact an isomorphism, it also checks the following additional condition: if a degree two vertex A is mapped into A', then the special edge of Definition 8.3 adjacent to A must be mapped into the special edge adjacent to A'.

Remark E.3. Finally, even if we avoid isomorphic planar combinatorial maps we are, for our purposes, still double-counting the spherical graphs in question, because we do not take into account mirror symmetry, cf. Remark 5.10. We in fact ignore this issue in the list of Proposition 8.2 because it is not as computationally costly as ignoring *all* isomorphisms: keeping a copy of the mirror symmetric in each list just doubles the cardinality of the $\mathcal{C}_k$ of Section 8.4. Eventually this redundancy is instead accounted for when listing the *once-derived* candidate topologies in Proposition 8.5: the list produced by the latter proposition does not contain mirror symmetric copies.

Appendix F. Computations for the proofs of Propositions 11.2, 11.5, and 11.6

F.1. **Computations for Proposition 11.2.** The aim of this section is to complete the proof of Proposition 11.2 giving the argument for the claim (a).

First of all we recall the function

$$f(\delta,\alpha) := \mathcal{H}^2(\Sigma_0) - \mathcal{H}^2(\Sigma)\,,$$

which we have defined when the parameters (δ,α) belong to $(0, 0.011] \times [\frac{\pi}{4}, \frac{\pi}{2}]$. However, $\mathcal{H}^2(\Sigma_0)$ is the sum of the areas of two triangles and $\mathcal{H}^2(\Sigma)$ is the sum of the areas of 10 triangles. While D' is defined as the Torricelli point of A, Q, and B, for the other vertices we know the coordinates as explicit smooth functions of δ and α which extend smoothly to $\delta = 0$. We will moreover show below that the Torricelli point is a Lipschitz function of the three other points. We thus have a Lipschitz extension of f to $\mathcal{D} = [0, 0.011] \times [\frac{\pi}{4}, \frac{\pi}{2}]$.

We next give an upper bound for

$$L_\alpha := \max_{\mathcal{D}} \left|\tfrac{\partial f}{\partial \alpha}\right|$$
$$L_\delta := \max_{\mathcal{D}} \left|\tfrac{\partial f}{\partial \delta}\right|\,.$$

[13]Any commutative operation that is sufficiently "distinct" from sum and product of integers would work. The choice of bitwise exclusive OR is standard.

We first observe that, given two triangles XYZ and XYZ' we have the obvious estimate

$$|\mathcal{H}^2(XYZ) - \mathcal{H}^2(XYZ')| \le \frac{1}{2}|(X-Y)\times(Z-Z')| \le \frac{1}{2}|XY||Z-Z'|\,. \tag{F.1}$$

Hence the area of a triangle is Lipschitz in each vertex, with a Lipschitz constant bounded by the diameter of the triangle divided by 2.

In the case of α, only the areas of the 10 triangles forming Σ depend on it. Moreover only the points A, A', M, M', and D' depend on α, while the others are fixed. The Lipschitz constants of A, A', and M' as functions of α can be bounded by 0.2, $0.2\cdot 0.8$, $\delta \le 0.011$, and $0.8\delta \le 0.008$ respectively. The point D' is rD, where D is the Torricelli point of A, Q, and B, where only A depends on α. We will show below that the Torricelli point has a Lipschitz dependence on each of the three vertices, with Lipschitz constant bounded by $\frac{2}{\sqrt{3}}$, cf. Lemma F.1 below. Therefore the Lipschitz constant of D' with respect to α is bounded by $0.8\cdot 0.2\cdot 2/\sqrt{3}$. The diamaters of all 10 triangles are bounded by $\sqrt{1+0.2^2}$. Therefore it is easy to see that the Lipschitz bound for the dependence in α is

$$L_\alpha \le \tfrac{\sqrt{1+0.2^2}}{2}\left(0.008\cdot 2 + 0.011\cdot 6 + 0.16\cdot 2 + 0.2\cdot 3 + \tfrac{2}{\sqrt{3}}\cdot 0.16\cdot 8\right) \le 1.3\,. \tag{F.2}$$

Now, consider that only the points M and M' depend on δ. So, of the 12 triangles which enter in the computations of f, seven have vertices which depend on δ and are AMO, CMO, $M'OD'$, BCM, AMD', BMD', and $MM'D'$. In the first six only one vertex, M, or M', depends δ, with Lipschitz constant respectively 1 and $r = 0.8$. Moreover their diameters can all be bounded by $\sqrt{1+0.2^2}$. For the final triangle two vertices depend on δ, while again the diameter can be bounded by the same number. Therefore we have the bound

$$L_\delta \le \frac{6 + 0.8\cdot 2}{2}\cdot\sqrt{1+0.2^2} \le 4\,. \tag{F.3}$$

In order to prove that the function f is positive we will then run a machine calculation of the values of f on the centers of the rectangles of a grid with size $s_\alpha = 0.001$ in α and size $s_\delta = 0.0001$ in δ. Given the Lipschitz bound it suffices to show that the minimum value is sufficiently (accounting for precision in floating point computations) over the threshold

$$\frac{L_\delta\cdot s_\delta}{2} + \frac{L_\alpha\cdot s_\alpha}{2} = 0.00085\,.$$

The code, which can be found in the file `pin-lower-bound.cpp` in [13], computes the coordinates of the various points for δ and α varying in the centers of the grid and then uses the usual vector algebra formula for evaluating the areas of the various triangle in the expression of f. The code returns the minimum value of f found in the search.

All the coordinates have been explicitly given except for the ones of the Torricelli point D. For the latter, one possibility would be to use an efficient general routine like the Weiszfeld algorithm. In this case we would need to certify error bounds. But given the structure of our problem we actually have explicit trigonometric expressions. In fact elementary geometric considerations show that D lies on the bisector of the angle AQB. Therefore its coordinates take the form $(|DQ|\cos\frac{\alpha}{2}, |DQ|\sin\frac{\alpha}{2}, 1)$. Moreover, again by elementary geometric considerations, $|DQ| = L(\cos\frac{\alpha}{2} - \sin\frac{\alpha}{2}\tan\frac{\pi}{6})$. After elementary manipulations we then get

$$D(\alpha) = \left(L\left[\tfrac{1}{2} + \tfrac{1}{\sqrt{3}}\sin\left(\tfrac{\pi}{3}-\alpha\right)\right], L\left[-\tfrac{1}{2\sqrt{3}} + \tfrac{1}{\sqrt{3}}\cos\left(\tfrac{\pi}{3}-\alpha\right)\right], 1\right)\,. \tag{F.4}$$

The maximum value found in the grid search is 0.00095565. The difference with the threshold is $0.00010565 > 10^{-4}$, well above the accuracy of the floating point computations. In fact

the maximum is found at (an approximation of) $\delta = 0.011$ and $\alpha = \frac{\pi}{2}$, which should be the maximum point of the function on the domain $\mathcal{D}$, a claim which can be likely proved with careful analytical considerations.

We conclude the argument with a geometric proof that the Torricelli point has the claimed Lipschitz dependence. Note that in our specific case the Lipschitz bound in α could also be proved directly by simply differentiating the explicit formula (F.4).

Lemma F.1. *Fix two points P and R on the Euclidean plane, and vary a third point S. Consider the Torricelli point T of the triangle PRS. Then $Lip\,(T) \leq \frac{2}{\sqrt{3}}$.*

Proof. When one of the angles of the triangle is $\frac{2\pi}{3}$, the Torricelli point is the obtuse angle of the triangle and so $T(S)$ is constant. Otherwise, the Torricelli point forms a triangle of degree $\frac{2\pi}{3}$ at T with the vertices P and R, so it lies on a circle β which contains the points P and R and is centered at a point S' which forms an equilateral triangle with P and R, which is opposite to S. And it is also well known from the Fermat-Torricelli construction that T lies on the segment connecting S with S'. Hence $T(S)$ is the unique intersection of the circle β with the segment SS' and the Lipschitz constant is bounded by an elemetary computation using this fact. □

F.2. **Computations for Proposition 11.5.** In this section we complete the proof of Proposition 11.5 giving the arguments for the claims (a) and (b).

For claim (a) consider first the equators containing the edges α and γ. They subdivide the sphere in four convex open sets and there is exactly one of them, say Λ, which contains the convex triangles $A2DA$ and BCB_2. The geodesic segment A_2B_2 is contained in Λ and for this reason the quadrilateral AA_2B_2B has internal angles at A_2 and B_2 which are strictly smaller than π. Given that its internal angles at A and B are $\frac{2\pi}{3}$, we conclude that AA_2B_2B is convex. Since D and C are internal to this quadrilateral, the convex triangles A_2DA and B_2CB are obviously conained in AA_2B_2B. Observe also that the length-minimizing geodesic segment CD connecting C and D is contained in AA_2B_2B. The four geodesic arcs AB, BC, CD, and DA bound a quadrilateral $ABCD$ contained in AA_2B_2B. The distance between an interior point of CD and the edge AB is strictly larger than that of C and D to AB. Since the points of the arc η are equidistant from AB, the arc is contained in the quadrilateral $ABCD$. Therefore the region Q is contained in the quadrilateral AA_2B_2B. This completes the proof of claim (a).

It remains to show (b). We set $x = \mathrm{len}(\beta)$ and recall that $|AA_2| = |BB_2| = 10x$. Observe moreover that $\mathrm{len}(\eta) \leq x$. Using the fomulae in Theorem 6.4 (it is convenient to apply them to the right-angled triangle which is half of AA_2D) we get the identity

$$\tan(5x) = \tfrac{\sqrt{3}}{2} \tan \tfrac{|AD|}{2}\,. \tag{F.5}$$

In particular $5x < \frac{|AD}{2}$. Since $\frac{\tan x}{x}$ is monotone increasing, we then get

$$\frac{\tan 5x}{5x} < \frac{\tan \frac{|AD|}{2}}{\frac{|AD|}{2}}$$

and dividing the two sides of the latter inequality by the two sides of (F.5) we immediately conclude $y < \frac{10x}{\sqrt{3}}$. We can then bound

$$|AA2| + |BB2| - (|A_2D| + |DA| + \mathrm{len}(\eta) + |CB2|) > (19 - 10\sqrt{3})x\,. \tag{F.6}$$

We next observe that the region Q is foliated by a one parameter family of arcs $\{\eta_s\}_{s\in[0,|AD|]}$ which are at distance s from AB. Since each of these arcs is shorter than AB, the coarea formula immediately gives

$$\mathcal{H}^2(Q) \leq x|AD| \leq \frac{10}{\sqrt{3}}x^2\,. \tag{F.7}$$

We next wish to bound the area of the triangle AA_2D (which is the same as the area of the triangle BB_2C). We split the triangle AA_2D into two equal right angled triangles by drawing the altitude from D to the edge AB. Take the one which has A among its vertices and denote it by T. T has internal angles $\frac{\pi}{2}$ (at the foot of the altitude) and $\frac{\pi}{6}$ at the vertex A. If we denote by θ the angle at the vertex D, then Gauss-Bonnet gives the formula

$$\mathcal{H}^2(T) = \theta - \frac{\pi}{3}\,.$$

On the other hand

$$\cos\theta = \cos\tfrac{|AA_2|}{2}\sin\tfrac{\pi}{6}\,.$$

We now claim that

$$\mathcal{H}^2(T) \leq \tfrac{|AD|^2}{8} \leq \tfrac{25x^2}{2}\,. \tag{F.8}$$

This is implied equivalent to $\theta \leq \frac{\pi}{3} + \frac{100x^2}{8}$, which by the monotonicity of the cosine, is equivalent to

$$\cos\left(\tfrac{\pi}{3} + \tfrac{|AD|^2}{8}\right) \leq \cos\tfrac{|AA_2|}{2}\cos\tfrac{\pi}{6}\,.$$

From the cosine duplication formula, the latter is in fact equivalent to

$$\tfrac{1}{2}\cos\tfrac{|AD|^2}{8} \leq \tfrac{1}{2}\cos\tfrac{|AD|}{2} + \tfrac{\sqrt{3}}{2}\sin\tfrac{|AD|^2}{8}$$

Using the known inequalities $1 - \frac{t^2}{2} \leq \cos t \leq 1$ and $\sin t \geq t - \frac{t^3}{3!}$, we immediately see that, since $|AD| = 10\mathrm{len}(\beta) \leq 0.1$, the inequality above holds.

Using (F.6), (F.7), and (F.8), we see that (11.6) is implied by

$$(19 - 10\sqrt{3})x \geq \left(\tfrac{20}{\sqrt{3}} + 100\right)x^2\,,$$

namely by $x \leq (19 - 10\sqrt{3})(100 + \frac{20}{\sqrt{3}})$. Since $x \leq 0.01$, the latter is equivalent to $0.01 \leq (19 - 10\sqrt{3})(100 + \frac{20}{\sqrt{3}})$, which can be easily verified by a calculator.

F.3. **Computations for Proposition 11.6.** To complete the proof of the proposition we give here the arguments for the claims (a), (b), and (c). It is convenient to mark the points X and Y closest to B where the equator λ intersects the two equators $\bar\alpha$ and $\bar\delta$ containing, respectively, the edges α and δ, cf. Figure 39. Moreover, before proving (a), we can simply mark the points D and E as the closest ones to M where $\bar\lambda$ intersects $\bar\alpha$ and $\bar\delta$. The claim (b) is then obvious, since by construction all internal angles of the pentagon $ABCDE$ are smaller than π.

In order to prove (a) we just need to check that the segments $|AE|$ and $|CD|$ are shorter, respectively, than $\mathrm{len}(\alpha)$ and $\mathrm{len}(\delta)$. Because the latter are strictly larger than $10(|AB| + |BC|) = 2|BM|$, it suffices to show that $|AE| \leq 2|BM|$ and $|CD| \leq 2|BM|$. The argument is the same and we focus therefore on the first inequality. Observe that, since sum of the angles

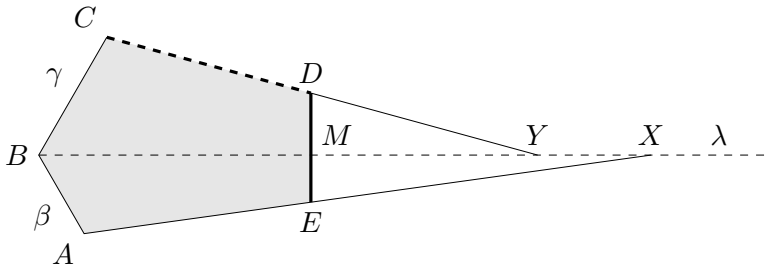


FIGURE 39. The competitor used to prove Proposition 11.6. The dashed segment is removed from the net and replaced with the thick segment. The gray area is the region spanned by the homotopy from the original net to the competitor. The points X and Y are intersections of the equator λ with the equators containing the segments α and δ.

$X\hat{A}B$ and $A\hat{B}X$ measure $\frac{2\pi}{3}$ and $\frac{\pi}{3}$ (and in particular they sum to π), elementary geometry implies that

$$|AX| + |BX| = \pi \tag{F.9}$$

$$|AB| > B\hat{X}A \tag{F.10}$$

$$|BX| > |AX|\,. \tag{F.11}$$

We moreover observe that

$$|EM| = \operatorname{dist}(E, XB) \leq \operatorname{dist}(A, XB) = B\hat{X}A \leq |AB|$$

(the second inequality follows because $|EX| < |AX| < \pi$ by (F.9) and (F.11)). In particular

$$|AE| \leq |BM| + |BA| + |ME| \leq |BM| + 2|AB| \leq 2|BM|\,.$$

We are thus left with proving (c). Observe that the argument just given above also implies

$$\begin{aligned} |CD| - |DE| &\geq |BM| - |CB| - |DM| - |DE| \geq |BM| - 2|CB| - |DE| \\ &\geq |BM| - 3|BC| - |AB| \geq \tfrac{2}{5}|BM|\,. \end{aligned} \tag{F.12}$$

We will prove below that

$$\mathcal{H}^2(BMEA) \leq |BM| \cdot |BA|\,. \tag{F.13}$$

Assuming the latter fact for a moment, the same inequality applied the quadrilateral $BCDM$ then gives

$$\mathcal{H}^2(ABCDE) \leq |BM| \cdot (|AB| + |BC|) \leq \tfrac{1}{5}|BM|^2\,. \tag{F.14}$$

In particular, taking into account (F.12) and (F.14), (c) holds if $|BM| < 1$, which is certainly the case because $|BM| = 5(|AB| + |BC|) < \frac{1}{10}$.

In order to show (F.13), consider the equator λ and the equator $\bar{\alpha}$ containing the edge α and let X and X' be their two intersections. The quadrilateral $BMEA$ is contained in the convex region delimited by two halves of these equators connecting X and X', cf. Figure 40, and we denote by η the half-equator (of the two) which is contained in $\bar{\alpha}$. Denote further by H the hemisphere bounded by λ which contains η, by N the center of H, and by $r < \frac{\pi}{2}$ the radius such that the disk $D_r(N)$ centered at N of radius r is tangent to η. By convexity η touches the disk from the outside. In particular we conclude that

$$BMEA \subset \{r < \operatorname{dist}(\cdot, N) < \tfrac{\pi}{2}\}\,.$$

The point of tangency between γ and $\partial D_r(N)$ is given by the midpoint of η and it follows therefore that

$$r = \tfrac{\pi}{2} - A\hat{X}B\,. \tag{F.15}$$

We now let NB and NM be the shortest geodesic segments connecting N with, respectively, B and M: these segments meet the equator λ perpendicularly at B and M. Because the internal angle of $ABME$ at B is $\frac{\pi}{3} < \frac{\pi}{2}$ and the one at M is $\frac{\pi}{2}$, if we denote by T the convex triangle with vertices N, B, and M we have

$$BMEA \subset T \cap \{r < \mathrm{dist}(\cdot, N) < \tfrac{\pi}{2}\} =: \Lambda\,.$$

Using that the function $p \mapsto \mathrm{dist}\,(p, N)$ is smooth in Λ and its gradient has modulus 1, we can use the coarea formula to compute

$$\mathcal{H}^2(BMEA) \leq \mathcal{H}^2(\Lambda) = \int_r^{\frac{\pi}{2}} \mathcal{H}^1(T \cap \{\mathrm{dist}(\cdot, N) = \rho\})\, d\rho \leq |BM|(\tfrac{\pi}{2} - r)\,,$$

where we have used that the function $\rho \mapsto \mathcal{H}^1(T \cap \{\mathrm{dist}(\cdot, N) = \rho\})$ reaches its maximum at $\rho = \frac{\pi}{2}$. Using (F.15) we then get $\mathcal{H}^2(BMEA) \leq |BM| \cdot A\hat{X}B$, which then implies (F.13) by (F.10).

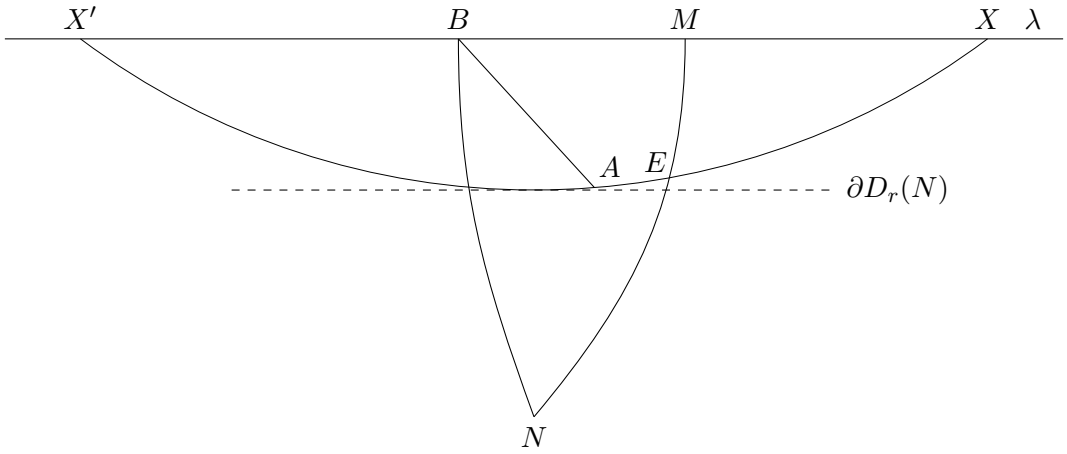


FIGURE 40. The argument for the inequality (F.13).

APPENDIX G. PROOF OF LEMMA 15.5

Consider a quadrilateral $ABCD$ as in the lemma, with vertices labeled clockwise. Assume that the equators γ_A and γ_C bisect the angles at A and C. These equators coincide if and only if the diagonal AC bisects both angles, which happens if and only if the quadrilateral is also equilateral. We can therefore draw the diagonal connecting A and C and assume, without loss of generality, that $\theta := A\hat{C}B$ is strictly smaller than $\frac{\pi}{3}$. Of course the angle $B\hat{A}C$ measures $\frac{2\pi}{3} - \theta$. We can explicitly compute the cosine of the length λ of AC as

$$\cos\lambda = \frac{-\frac{1}{2} + \cos(\theta)\cos\left(\frac{2\pi}{3} - \theta\right)}{\sin(\theta)\sin\left(\frac{2\pi}{3} - \theta\right)}$$

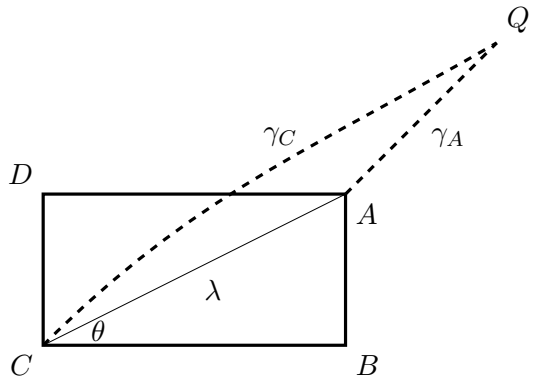


FIGURE 41. The quadrilateral $ABCD$ and the triangle ACQ.

using the three angles of the convex triangle ABC via the spherical law of cosines, cf. Figure 41 We apply the product-to-sum identities to the numerator and denominator. For the cosines:

$$\cos(\theta)\cos\left(\tfrac{2\pi}{3}-\theta\right)=\tfrac{1}{2}\left[\cos\left(2\theta-\tfrac{2\pi}{3}\right)+\cos\left(\tfrac{2\pi}{3}\right)\right]=\tfrac{1}{2}\left[\cos\left(2\theta-\tfrac{2\pi}{3}\right)-\tfrac{1}{2}\right]$$

For the sines:

$$\sin(\theta)\sin\left(\tfrac{2\pi}{3}-\theta\right)=\tfrac{1}{2}\left[\cos\left(2\theta-\tfrac{2\pi}{3}\right)-\cos\left(\tfrac{2\pi}{3}\right)\right]=\tfrac{1}{2}\left[\cos\left(2\theta-\tfrac{2\pi}{3}\right)+\tfrac{1}{2}\right]$$

Substituting these expansions back into the equation for $\cos\lambda$ and multiplying the numerator and denominator by 4 yields the simplified closed-form expression:

$$\cos\lambda=\frac{2\cos\left(2\theta-\frac{2\pi}{3}\right)-3}{2\cos\left(2\theta-\frac{2\pi}{3}\right)+1}$$

Consider now one of the two points Q of intersections of γ_A and γ_C, more precisely the first which we encounter on the equator γ_A if we start at A and exit the quadrilateral $ABCD$. Q, A, and C form a convex triangle in which the angles adjacent to AC measure $\frac{2\pi}{3}$ and $\frac{\pi}{3}-\theta$, cf. Figure 41. The cosine of the angle $\alpha=C\hat{Q}A$ can thus be computed applying again the spherical law of cosines to the triangle CQA:

$$\cos\alpha=-\cos\left(\tfrac{2\pi}{3}+\theta\right)\cos\left(\tfrac{\pi}{3}-\theta\right)+\sin\left(\tfrac{2\pi}{3}+\theta\right)\sin\left(\tfrac{\pi}{3}-\theta\right)\cos\lambda\,.$$

As above we evaluate the products using the sum and difference of the angles. The sum is π and the difference is $2\theta+\frac{\pi}{3}$:

$$\cos\left(\tfrac{2\pi}{3}+\theta\right)\cos\left(\tfrac{\pi}{3}-\theta\right)=\tfrac{1}{2}\left[\cos\left(2\theta+\tfrac{\pi}{3}\right)+\cos(\pi)\right]=\tfrac{1}{2}\left[\cos\left(2\theta+\tfrac{\pi}{3}\right)-1\right]$$

$$\sin\left(\tfrac{2\pi}{3}+\theta\right)\sin\left(\tfrac{\pi}{3}-\theta\right)=\tfrac{1}{2}\left[\cos\left(2\theta+\tfrac{\pi}{3}\right)-\cos(\pi)\right]=\tfrac{1}{2}\left[\cos\left(2\theta+\tfrac{\pi}{3}\right)+1\right]$$

Substituting these into the equation for $\cos\alpha$ and using $\cos\left(2\theta+\frac{\pi}{3}\right)=-\cos(\frac{2\pi}{3}-2\theta)$, we get

$$\cos\alpha=\frac{1+\cos\left(\frac{2\pi}{3}-2\theta\right)}{2}+\frac{1-\cos\left(\frac{2\pi}{3}-2\theta\right)}{2}\left(\frac{2\cos\left(\frac{2\pi}{3}-2\theta\right)-3}{2\cos\left(\frac{2\pi}{3}-2\theta\right)+1}\right).$$

After a few simple algebraic manipulations we find

$$\cos\alpha=\frac{8\cos\left(\frac{2\pi}{3}-2\theta\right)-2}{2\left(2\cos\left(\frac{2\pi}{3}-2\theta\right)+1\right)}=\frac{4\cos\left(\frac{2\pi}{3}-2\theta\right)-1}{2\cos\left(\frac{2\pi}{3}-2\theta\right)+1}\,.$$

The claim of the lemma is then that α uniquely determines θ. Set then $x = \cos(\frac{2\pi}{3} - 2\theta) \in (-\frac{1}{2}, 1)$. Given that cos on $[0, \pi]$ is an homeomorphism onto $[-1, 1]$ and the fact that $0 < \alpha < \pi$, our claim is equivalent to the injectivity of the function

$$x \mapsto \frac{4x-1}{2x+1} = 2 - \frac{3}{2x+1}$$

on the interval $(-\frac{1}{2}, 1)$. This is however an obvious fact.

References

[1] W. K. Allard. On the first variation of a varifold. *Ann. of Math. (2)*, 95:417–491, 1972. 7, 13, 115, 116

[2] W. K. Allard. On the first variation of a varifold: boundary behavior. *Ann. of Math. (2)*, 101:418–446, 1975. 7, 13, 115

[3] F. J. Almgren, Jr. Existence and regularity almost everywhere of solutions to elliptic variational problems with constraints. *Mem. Amer. Math. Soc.*, 4(165):viii+199, 1976. doi:10.1090/memo/0165. 4, 13

[4] Frederick J. Almgren and Jean E. Taylor. The geometry of soap films and soap bubbles. *Scientific American*, 235(1):82–93, 1976. URL: http://www.jstor.org/stable/24950395. 5

[5] Theodora Bourni. Allard-type boundary regularity for $C^{1,\alpha}$ boundaries. *Adv. Calc. Var.*, 9(2):143–161, 2016. doi:10.1515/acv-2014-0032. 115

[6] Kenneth A. Brakke. The surface evolver. *Experiment. Math.*, 1(2):141–165, 1992. URL: http://projecteuclid.org/euclid.em/1048709050. 5, 7, 110

[7] Guy David. A local description of 2-dimensional almost minimal sets bounded by a curve. *Annales de la Faculté des sciences de Toulouse : Mathématiques*, Ser. 6, 31(1):1–382, 2022. URL: https://www.numdam.org/articles/10.5802/afst.1697/, doi:10.5802/afst.1697. 7

[8] Guy David. Personal communication. 2026. 3

[9] C. De Lellis, F. Ghiraldin, and F. Maggi. A direct approach to Plateau's problem. *J. Eur. Math. Soc. (JEMS)*, 19(8):2219–2240, 2017. doi:10.4171/JEMS/716. 5, 6, 7, 13

[10] Camillo De Lellis. Allard's interior regularity theorem: an invitation to stationary varifolds. In *Nonlinear analysis in geometry and applied mathematics. Part 2*, volume 2 of *Harv. Univ. Cent. Math. Sci. Appl. Ser. Math.*, pages 23–49. Int. Press, Somerville, MA, 2018. 114, 115

[11] Camillo De Lellis and Federico Glaudo. Boundary plateau laws ii. *arXiv preprint*, 2026. 3, 5, 6

[12] Federico Glaudo. Boundary plateau laws, 2026. URL: https://glaudo.github.io/boundary-plateau-laws/index.html. 10, 30, 60, 108

[13] Federico Glaudo. Boundary regularity for minimal sets, 2026. URL: https://github.com/dario2994/boundary-plateau-laws. 10, 30, 60, 62, 64, 68, 79, 82, 83, 104, 108, 126, 133

[14] Aladár Heppes. Isogonale sphärische Netze. *Ann. Univ. Sci. Budapest Eötvös Sect. Math*, 7(41-48), 1964. 16, 20, 26, 27, 37, 76, 84

[15] Hermann Karcher. Schnittort und konvexe Mengen in vollständigen Riemannschen Mannigfaltigkeiten. *Math. Ann.*, 177:105–121, 1968. doi:10.1007/BF01350787. 121, 122

[16] Darren King, Francesco Maggi, and Salvatore Stuvard. Plateau's problem as a singular limit of capillarity problems. *Comm. Pure Appl. Math.*, 75(3):541–609, 2022. doi:10.1002/cpa.22019. 5

[17] Ernest Lamarle. Sur la stabilité des systèmes liquides en lames minces. *Mémoires de l'Académie Royale des Sciences, des Lettres et des Beaux-Arts de Belgique*, 35:1–104, 1865. URL: http://eudml.org/doc/179074. 16

[18] Sergei K. Lando and Alexander K. Zvonkin. *Graphs on surfaces and their applications*, volume 141 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin, 2004. With an appendix by Don B. Zagier, Low-Dimensional Topology, II. doi:10.1007/978-3-540-38361-1. 20, 21, 31, 32, 130

[19] Gary Lawlor and Frank Morgan. Curvy slicing proves that triple junctions locally minimize area. *J. Differential Geom.*, 44(3):514–528, 1996. URL: http://projecteuclid.org/euclid.jdg/1214459219. 7

[20] Frank Morgan. *Geometric measure theory*. Elsevier/Academic Press, Amsterdam, fifth edition, 2016. A beginner's guide, Illustrated by James F. Bredt. 7

[21] T. Nishizeki and N. Chiba. *Planar Graphs: Theory and Algorithms*. Dover Books on Mathematics Series. Dover Publications, 2008. URL: https://books.google.com/books?id=1Nl4BpacvpwC. 130

[22] OEIS Foundation Inc. Entry A000109 in the On-Line Encyclopedia of Integer Sequences. https://oeis.org/A000109, 2026. Accessed: 2026-02-08. 35

[23] Joseph Antoine Ferdinand Plateau. *Statique Expérimentale Et Théorique Des Liquides Soumis Aux Seules Forces Moléculaires*, volume 1 of *Classic Reprint*. Gauthier-Villars, Paris, 1873. URL: https://books.google.com/books/about/Statique_exprimentale_et_thorique_de.html?id=t13TAkMklz8C. 1

[24] Leon Simon. *Lectures on geometric measure theory*, volume 3 of *Proceedings of the Centre for Mathematical Analysis, Australian National University*. Australian National University, Centre for Mathematical Analysis, Canberra, 1983. 115

[25] Jean E Taylor. The structure of singularities in soap-bubble-like and soap-film-like minimal surfaces. *Annals of Mathematics*, 103(3):489–539, 1976. 4, 5, 7, 26, 111, 113, 114

[26] Heinrich Tietze. über Konvexheit im kleinen und im großen und über gewisse den Punkten einer Menge zugeordnete Dimensionszahlen. *Math. Z.*, 28(1):697–707, 1928. doi:10.1007/BF01181191. 120

[27] Isaac Todhunter. *Spherical Trigonometry: For the Use of Colleges and Schools*. Macmillan, London, 5th edition, 1886. 24

[28] J. H. C. Whitehead. Simplicial spaces, nuclei and m-groups. *Proceedings of the London Mathematical Society*, 45(1):243–327, 1939. doi:10.1112/plms/s2-45.1.243. 87

School of Mathematics, Institute for Advanced Study, 1 Einstein Dr., Princeton NJ 08540, USA

Gran Sasso Science Institute, viale Francesco Crispi, 7, 67100 L'Aquila, Italy
Email address: camillo.delellis@ias.edu

Email address: federico.glaudo@gmail.com