THE THREE GAP THEOREM AND RIEMANNIAN GEOMETRY

IAN BIRINGER & BENJAMIN SCHMIDT

The classical Three Gap Theorem asserts that for $n \in \mathbb{N}$ and $p \in \mathbb{R}$, there are at most three distinct distances between consecutive elements in the subset of $[0, 1)$ consisting of the reductions modulo 1 of the first n multiples of p (see e.g. [Só58], [Św58]). There are several interesting generalizations of this theorem in [FrSó92], [Vi08], and the references therein. Regarding it as a statement about rotations of the circle, we find results in a similar spirit pertaining to isometries of compact Riemannian manifolds and the distribution of points along their geodesics.

Let M^k denote a complete Riemannian k-manifold with associated distance function $\mathrm{d} : M \times M \to \mathbb{R}$, and let X be a finite subset of M with $|X| \geq 2$. For $x \in X$, define the distance from x to its nearest neighbor in X to be $\mathrm{nnd}(x, X) = \min_{y \in X \setminus \{x\}} \mathrm{d}(x, y)$. We denote the set of all nearest neighbor distances in X by

$$\mathrm{NND}(X) = \{\mathrm{nnd}(x, X) \mid x \in X\}.$$

The following variant of the Three Gap Theorem is the starting point for our work:

Three Gap Theorem (Geometric Version)**.** *Let S^1 denote the unit circle. For any rotation R, point $p \in S^1$, and $n \in \mathbb{N}$,*

$$|\,\mathrm{NND}(\{R^i(p) \,|\, i = 0, \ldots, n\})| \leq 3.$$

Note that this result is slightly weaker than the classical Three Gap Theorem as stated above. Specifically, the classical result considers the distance from a point to its nearest neighbor on the left separately from the distance to its nearest neighbor on the right, stating that even then the number of distinct distances is at most 3. However, this left/right distinction has no obvious analogue in higher dimensions, so for us the geometric version will be the natural one to generalize.

Our first result shows that the phenomenon exhibited in the theorem above is common to all isometries of compact Riemannian manifolds with bounds depending only on dimension, sectional curvatures, and diameter.

Theorem 1. *For each $k \in \mathbb{N}$, $\kappa \in \mathbb{R}$ and $D > 0$, there is a constant $K(k, \kappa, D) \in \mathbb{N}$ such that for any complete Riemannian k-manifold M^k with $\sec \geq \kappa$ and $\operatorname{diam}(M) \leq D$, and for any $I \in \operatorname{Isom}(M)$, $p \in M$, and $n \in \mathbb{N}$,*

$$|\operatorname{NND}(\{I^i(p) \,|\, i = 0, \ldots, n\})| \leq K.$$

A cleaner statement can be made for manifolds with nonnegative sectional curvatures.

Corollary 1. *Let M be a complete Riemannian k-manifold with non-negative sectional curvatures. Then for any $I \in \operatorname{Isom}(M)$, $p \in M$, and $n \in \mathbb{N}$,*

$$|\operatorname{NND}(\{I^i(p) \,|\, i = 0, \ldots, n\})| \leq 3^k + 1.$$

Working in a different direction, we also consider Riemannian metrics with uniform bounds for the number of nearest neighbor distances appearing in finite subsets of equally spaced points along geodesics. Unless stated differently, geodesics are assumed throughout to be unit-speed paramaterized curves $\gamma : \mathbb{R} \to M$ satisfying the geodesic differential equation.

Definition (Bounded Geodesic Combinatorics)**.** *Let $K \in \mathbb{N}$. A complete Riemannian manifold M is defined to have K-bounded geodesic combinatorics if for every (unit-speed) geodesic $\gamma : \mathbb{R} \to M$, $T \in \mathbb{R}$, and $n \in \mathbb{N}$,*

$$|\operatorname{NND}(\{\gamma(iT) \,|\, i = 0, \ldots, n\})| \leq K.$$

When the particular $K \in \mathbb{N}$ is irrelevant, we shall simply say that M has bounded geodesic combinatorics. By the Three Gap Theorem, the circle has bounded geodesic combinatorics. We will show below that all compact symmetric spaces have bounded geodesic combinatorics. Other examples include SC-Riemannian manifolds, manifolds all of whose geodesics are simple closed and of common shortest period. The Zoll type metrics on spheres provide explicit non-symmetric examples of these. Finally, finite products of compact symmetric spaces and SC-metrics have bounded geodesic combinatorics. It is possible that these are the only examples. Our final theorem shows that this is indeed the case in dimension 2.

Theorem 2 (Classification in Dimension 2)**.** *Assume that M is a closed Riemannian surface. Then M has bounded geodesic combinatorics if and only if it is isometric to $\mathbb{T}^2$ with a flat metric, $\mathbb{RP}^2$ with a round metric, or $\mathbb{S}^2$ with an SC-metric.*

Acknowledgements Thank are due to Sujith Vijay for e-mail correspondence concerning ([Vi08]). His paper provides an important idea used in the proof of Theorem 1. The referee has helped us enormously by contributing many useful comments, and in particular by suggesting that we prove Lemma 9, which allowed us to remove an unnecessary assumption from an earlier version of Theorem 2. We are also grateful to Juan Souto for the suggestion to consider asymptotic geodesics in section 3 and to Jean-François Lafont and Krastio Lilov for bringing the three gap theorem to the second author's attention.

1. Nearest Neighbors and Orbits of Isometries

The goal of this section is to understand nearest neighbor distances in orbit segments of isometries on closed Riemannian manifolds. We show that the number of distinct distances is limited by a certain packing number, and then use this to prove theorem 1. Finally, we construct orbit segments in k-dimensional flat tori with at least k distinct nearest neighbor distances.

Throughout the following, let (M, d) be a complete Riemannian k-manifold. Define $P(M, r)$ to be the maximum number of points that can be packed pairwise r-apart into some open r-ball in M, and let $P(M) = \sup_r P(M, r)$.

Lemma 1. *For all $I \in \mathrm{Isom}(M)$, $p \in M$, and $n \in \mathbb{N}$, we have*

$$|\,\mathrm{NND}(\{I^i(p)\,|\,i = 0, \ldots, n\})| \leq P(M) + 1.$$

Proof. We can assume that the points defining $O = \{p, I(p), \ldots, I^n(p)\}$ are distinct, for otherwise we can reduce to this case by considering a smaller orbit segment. Let $\mathrm{nnd}(i) = \mathrm{nnd}(I^i(p), O)$ for $0 \leq i \leq n$. Then

$$\begin{aligned} \mathrm{nnd}(i) &= \min\{\mathrm{d}(I^j(p), I^i(p))\,|\,j = 0, \ldots, n,\ j \neq i\} \\ &= \min\{\mathrm{d}(p, I^{|i-j|}(p))\,|\,j = 0, \ldots, n,\ j \neq i\} \\ &= \min\{\mathrm{d}(p, I^k(p))\,|\,k = 1, \ldots, \max\{i, n-i\}\}. \end{aligned}$$

Set $m = \lceil n/2 \rceil$ and observe that $m = \min_{0 \leq i \leq n} \max(i, n-i)$. If $r = \min_{1 \leq j \leq m} \mathrm{d}(p, I^j(p))$, we have $\mathrm{nnd}(i) \geq r$ for $0 \leq i \leq n$. Furthermore, if $\mathrm{nnd}(i) < r$, then $\mathrm{nnd}(i) = \mathrm{d}(p, I^k(p))$ for an index k satisfying $m < k \leq n$ and $\mathrm{d}(p, I^k(p)) < r$. Note that if $m < k < l \leq n$ are two such indices, then $\mathrm{d}(I^k(p), I^l(p)) = \mathrm{d}(p, I^{l-k}(p)) \geq r$ since $l - k < m$. Hence, the number of distinct values of $\mathrm{nnd}(i)$ less than r is at most the maximal number of r-separated points that can be packed into the open ball of radius r about the point p, concluding the proof. □

Observe that only two points can be packed pairwise r-apart into an open interval of length $2r$, so Lemma 1 provides a proof of the Geometric Three Gap Theorem. It is easy to see in general that if M is closed, then $P(M)$ is finite. Therefore, any closed Riemannian manifold has an upper bound for the number of nearest neighbor distances occurring in finite orbit segments of isometries. To prove theorem 1, however, we must bound $P(M)$ in terms of geometric data.

Lemma 2 (Bounding the Packing Constant)**.** *Assume that M is a Riemannian k-manifold with sectional curvatures bounded below by κ. Let M^k_κ be the k-dimensional model space of constant curvature κ. Then for each $r > 0$,*

$$P(M, r) \leq P(M^k_\kappa, r).$$

Proof. Consider n points $\{p_1, \ldots, p_n\} \subset M$ packed pairwise r-apart into the open r-ball around a point $p \in M$. We claim that we can also pack n points pairwise r-apart into an open r-ball in M^k_κ.

Fix a basepoint $\overline{p} \in M_\kappa$ and a linear isometry $L : T_pM \to T_{\overline{p}}M_\kappa$. For $i = 1, \ldots, n$, let $\gamma_i : [0, 1] \to M$ be a minimizing geodesic in M joining $\gamma_i(0) = p$ to $\gamma_i(1) = p_i$, and

$$\overline{\gamma}_i : [0, 1] \to M_\kappa,\ \overline{\gamma}_i(t) := \exp_{\overline{p}}(t \cdot \gamma_i'(0))$$

the comparison geodesic in M^k_κ. Note that $\mathrm{d}(\overline{p}, \overline{\gamma}_i(1)) = \mathrm{d}(p, \gamma_i(1)) \leq r$. Furthermore, by Toponogov's Comparison Theorem (see e.g. [ChEb75, Theorem 2.2]), we have for $i \neq j$ that

$$d(\overline{\gamma}_i(1), \overline{\gamma}_j(1)) \geq d(\gamma_i(1), \gamma_j(1)) \geq r.$$

Therefore the points $\{\overline{\gamma}_i(1) \mid i = 1, \ldots, n\}$ are packed pairwise r-apart into the r-ball around $\overline{p} \in M^k_\kappa$. □

We are now ready to prove the following theorem, stated previously in the introduction.

Theorem 1. *For each $k \in \mathbb{N}$, $\kappa \in \mathbb{R}$ and $D > 0$, there is a constant $K(k, \kappa, D) \in \mathbb{N}$ such that for any complete Riemannian k-manifold M^k with $\sec \geq \kappa$ and $\mathrm{diam}(M) \leq D$, and for any $I \in \mathrm{Isom}(M)$, $p \in M$, and $n \in \mathbb{N}$,*

$$|\,\mathrm{NND}(\{I^i(p)\,|\,i = 0, \ldots, n\})| \leq K.$$

Proof. It suffices by Lemma 1 to provide a bound for $P(M)$. If $\kappa < 0$, it follows easily from the law of cosines that $P(M^k_\kappa, r)$ is nondecreasing with r. Therefore $P(M) \leq P(M^k_\kappa, D)$, a constant depending only on k, κ and D. If $\kappa \geq 0$, we have $P(M, r) \leq P(\mathbb{E}^k, r)$, which is independent of r and therefore determined by k. □

Bounds on the packing constant of k-dimensional Euclidean space can be calculated explicitly. Recall that a ball of radius r in $\mathbb{E}^k$ has volume $C(k)r^k$, where $C(k)$ is a dimensional constant. If n points are packed pairwise r-apart into the r ball around $p \in \mathbb{R}^k$, then the $r/2$ balls centered at these n points are disjoint and are all contained in the $3r/2$ ball around p. Hence,

$$n \cdot C(k) \cdot (r/2)^k \leq C(k) \cdot (3r/2)^k,$$

so $P(\mathbb{E}^k) \leq 3^k$. This produces the following:

Corollary 1. *Let M be a complete Riemannian k-manifold with non-negative sectional curvatures. Then for any $I \in \mathrm{Isom}(M)$, $p \in M$, and $n \in \mathbb{N}$,*

$$|\,\mathrm{NND}(\{I^i(p) \,|\, i = 0, \ldots, n\})| \leq 3^k + 1.$$

1.1. **Tori with many distinct nearest neighbor distances.** We now construct orbit segments of translations on flat k-dimensional tori with at least k distinct nearest neighbor distances. The idea is fairly simple. First, we construct a flat k-torus M by gluing together sides of a k-dimensional box with side-lengths equal to distinct primes. This allows us to find a translational isometry of M with an orbit consisting of the entire integer lattice in M. A point p in this orbit is the nearest neighbor to the $2k$ lattice points from which it is distance 1. We can then perturb our isometry slightly so that k of these points now have distinct distances to p, while preserving the fact that p is their nearest neighbor.

Fix an orthonormal basis $\{e_1, \ldots, e_k\} \subset \mathbb{R}^k$. Given $v \in \mathbb{R}^k$, let $T_v : \mathbb{R}^k \to \mathbb{R}^k$ be the translation by v. Let $p_1, \ldots, p_k$ be distinct odd primes, and set $N = \prod_{i=1}^k p_i$.

Consider the flat torus

$$M = \mathbb{R}^k/\Gamma, \ \Gamma = \langle T_{p_1 \cdot e_1}, \ldots, T_{p_k \cdot e_k} \rangle \,.$$

The translations T_v descend to isometries $t_v : M \to M$, and we denote the projection to M a of a point $p \in \mathbb{R}^k$ by $\overline{p} \in M$.

Set $\pi_j \in \{1, \ldots, p_j - 1\}$ to be the mod-p_j inverse of $\prod_{i \neq j} p_i$, and let

$$a_j = \max(\pi_j, p_j - \pi_j) \cdot \prod_{i \neq j} p_i.$$

Then $\frac{N+1}{2} \leq a_j < N$, $a_j \equiv \pm 1 \,(mod\ p_j)$, and $a_j \equiv 0 \,(mod\ p_i)$ for $i \neq j$. Thus for some $\delta_j \in \{-1, +1\}$, we have $\overline{(a_j, \ldots, a_j)} = \overline{\delta_j \cdot e_j}$.

Pick some small $s > 0$, and let $v = (1 - s\delta_1, \ldots, 1 - s\delta_k)$. Define $x_i = \overline{i \cdot v}$, and let $X = \{x_0, \ldots, x_{N-1}\}$. Note that X consists of the first N elements of the orbit of $\overline{0}$ under the isometry $t_v : M \to M$. We claim that for $j = 1, \ldots, k$, the nearest neighbor in X to x_{a_j} is x_0 and the distances $\mathrm{d}(x_{a_j}, x_0)$ are all distinct.

First, observe that $x_{a_j} = \overline{\delta_j \cdot e_j - a_j s(\delta_1, \ldots, \delta_k)}$, so

$$\mathrm{d}(x_{a_j}, x_0) = \sqrt{1 - 2a_j s + k\,(a_j s)^2}.$$

Thus if s is small, the distances $\mathrm{d}(x_{a_j}, x_0)$ are less than 1 and are monotonically decreasing with a_j. In particular, they are distinct, so we need only show that x_0 is the nearest neighbor of x_{a_j} for each j.

Assume on the contrary that $\mathrm{d}(x_{a_j}, x_i) < \mathrm{d}(x_{a_j}, x_0)$ for some i. Then $\mathrm{d}(x_{|a_j - i|}, x_0) < \mathrm{d}(x_{a_j}, x_0)$ as well. Since s is small, $x_{|a_j-i|}$ is very close to the projection of an element of the integer lattice in $\mathbb{R}^k$. But $\mathrm{d}(x_{|a_j-i|}, x_0) < 1$, so in fact $x_{|a_j-i|}$ must be close to $\overline{\pm e_l}$ for some $1 \le l \le k$. So either $|a_j - i| = a_l$ or $|a_j - i| = N - a_l$. The latter case is impossible since $\mathrm{d}(x_{N-a_l}, x_0) = \sqrt{1 + 2a_j s + k\,(a_j s)^2}$, and is therefore greater than 1. If $|a_j - i| = a_l$, then since $\frac{N-1}{2} \le a_j < N$ and $1 \le i < N$ we must have $a_l < a_j$. This cannot be, because as mentioned earlier $\mathrm{d}(x_{a_j}, x_0)$ decreases monotonically with a_j.

2. Bounded Geodesic Combinatorics

Recall our definition given in the introduction.

Definition (Bounded Geodesic Combinatorics)**.** *Let $K \in \mathbb{N}$. A complete Riemannian manifold M is defined to have K-bounded geodesic combinatorics if for every unit-speed geodesic $\gamma : \mathbb{R} \to M$, $T \in \mathbb{R}$, and $n \in \mathbb{N}$,*

$$|\operatorname{NND}(\{\gamma(iT) \mid i = 0, \ldots, n\})| \le K.$$

We begin with some examples. First, if M is a flat torus then any sequence of equally spaced points along a geodesic in M can be constructed as a segment of an orbit of a translational isometry. Therefore, Theorem 1 implies that flat tori have bounded geodesic combinatorics. This principle can be applied more generally:

Proposition 1. *Compact symmetric spaces have bounded geodesic combinatorics.*

Proof. Assume that M is a compact Riemannian symmetric space. Let $\gamma : \mathbb{R} \to M$ be a unit speed paramaterization of a geodesic, and let $T \in \mathbb{R}$. By definintion, the geodesic involution $s_m : M \to M$ at each point $m \in M$ defined by $\exp_m(w) \mapsto \exp_m(-w)$ for each $w \in TM$ is a well-defined isometry of M. Let $p = \gamma(0) \in M$ and $q = \exp_p(-\frac{T}{2}\dot\gamma(0)) \in M$. Define the isometry $I = s_p \circ s_q \in \operatorname{Isom}(M)$. Then for each natural number $n \in \mathbb{N}$, $\gamma(nT) = I^n(p)$, and hence Theorem 1 implies M has bounded geodesic combinatorics. □

Round spheres are symmetric spaces, and therefore have bounded geodesic combinatorics. To see this more directly, note that image of any geodesic $\mathbb{S}^k$ is an embedded circle γ for which the distance function satisfies $\mathrm{d}_\gamma(x,y) = \mathrm{d}_{\mathbb{S}^k}(x,y)$ for all $x, y \in \gamma$. Hence, the Three Gap Theorem implies that $\mathbb{S}^k$ has 3-bounded geodesic combinatorics. This proof motivates considering the following class of spaces.

Definition. *A closed Riemannian manifold (M, h) is said to be an SC-manifold if all its geodesics are simple, periodic, and of common least period.*

The Zoll metrics on spheres give explicit non-symmetric examples of SC-metrics on spheres of each dimension (see e.g. [Be78]). Note that the definition above does not imply that the image γ of a closed geodesic satisfies $\mathrm{d}_\gamma = \mathrm{d}_M$ as is the case for round spheres. However, the following lemma allows us to work around this to show that SC-manifolds have bounded geodesic combinatorics. Note that the conclusion of the lemma is equivalent to the existence of a uniform lower bound for the radii of tubular neighborhoods of geodesics in M.

Lemma 3. *Let M be a closed Riemannian SC-manifold. Then there exists a constant $r_M > 0$ such that for every geodesic $\gamma \subset M$ and any points $p, q \in \gamma$ satisfying $\mathrm{d}_M(p,q) < r_M$, we have $\mathrm{d}_M(p,q) = \mathrm{d}_\gamma(p,q)$.*

Proof. Assume this is not the case, and that the common least period of the geodesics in M is L. Then we can find a sequence of unit speed geodesics $\gamma_i : \mathbb{R} \to M$ and times $0 < d_i \le \frac{L}{2}$ such that $\mathrm{d}(\gamma_i(0), \gamma_i(d_i)) \to 0$ but $\gamma_i([0, d_i])$ is not minimizing. Since d_i cannot be less than the injectivity radius of M, we can pass to an appropriate subsequence so that the γ_n converge pointwise to a geodesic $\gamma : \mathbb{R} \to M$ and $d_i \to d \in [\mathrm{inj}(M), \frac{L}{2}]$. By continuity, $\gamma(0) = \gamma(d)$, contradicting the fact that γ is simple with least period L. □

Proposition 2. *Suppose that M is a closed Riemannian SC-manifold. Then M has bounded geodesic combinatorics.*

Proof. Let $\gamma : \mathbb{R} \to M$ be a geodesic, $T > 0$, and $n \in \mathbb{N}$. As M is compact, there is a number $K_M \in \mathbb{N}$ such that at most K_M points in M can be pairwise r_M apart, where r_M is as in the previous lemma. Therefore, at most K_M of the nearest-neighbor distances $\mathrm{NND}(\{\gamma(iT) \,|\, i \in \{0, \dots, n\}\})$ are not realized as distances measured in γ. By the Three Gap Theorem,

$$|\,\mathrm{NND}(\{\gamma(iT) \,|\, i \in \{0, \dots, n\}\})| \le K_M + 3,$$

concluding the proof. □

The argument above can be applied more generally to spaces in which all geodesics are contained in nicely embedded flat tori:

Lemma 4. *Let M be a closed Riemannian manifold. Assume that there exists a constant $r_M > 0$ such that every geodesic $\gamma \subset M$ is contained in an embedded totally geodesic flat torus $F_\gamma \subset M$ with the property that if two points $p, q \in F_\gamma$ satisfy $\mathrm{d}_M(p,q) < r_M$, then $\mathrm{d}_M(p,q) = \mathrm{d}_{F_\gamma}(p,q)$. Then M has bounded geodesic combinatorics.*

Proof. The proof is identical to that of Proposition 2, except that it uses the bounded geodesic combinatorics in tori instead of the Three Gap Theorem. □

A useful quality of Lemma 4 is that if two manifolds satisfy its hypotheses, then their Riemannian product does as well. It follows from Lemma 3 that SC-manifolds satisfy the hypotheses of Lemma 4. Any compact symmetric space does as well: its maximal flats are embedded tori and any two of these differ by an isometry of the ambient manifold, [He01, Theorem 6.2], so any r_M that works for a single flat works for all flats simultaneously. This proves the following result.

Corollary 2. *Finite products of SC-manifolds and compact symmetric spaces have bounded geodesic combinatorics.*

To appreciate our approach to Corollary 2, note that it is not obvious, although probably true, that the set of manifolds with bounded geodesic combinatorics is closed under finite products.

3. Global Consequences of Bounded Geodesic Combinatorics

This section is devoted to understanding the global behavior of geodesics in manifolds with bounded geodesic combinatorics.

Lemma 5. *Let M be a closed manifold with bounded geodesic combinatorics. Then for every constant speed geodesic $\gamma : \mathbb{R} \to M$, the set of points $\gamma(\mathbb{N})$ has finitely many isolated points.*

Proof. Let $A \subset \mathbb{N}$ be the set of indices i for which $\gamma(i)$ is isolated, and assume that A is infinite. Set $X = \gamma(\mathbb{N})$, and for each $i \in A$, let $L_i = \mathrm{d}(\gamma(i), X \setminus \gamma(i))$. Since $\gamma(i)$ is isolated in X, $L_i > 0$. Note that the $\frac{L_i}{2}$-balls around $\gamma(i)$ are all disjoint, so the compactness of M implies that $L_i \to 0$. In particular, if M has K-bounded geodesic combinatorics, we can pick $I \subset A$ with $|I| > K$ such that L_i, $i \in I$ are all distinct. Set $\epsilon = \min\{|L_i - L_j| \mid i, j \in I, i \neq j\}$. Then there

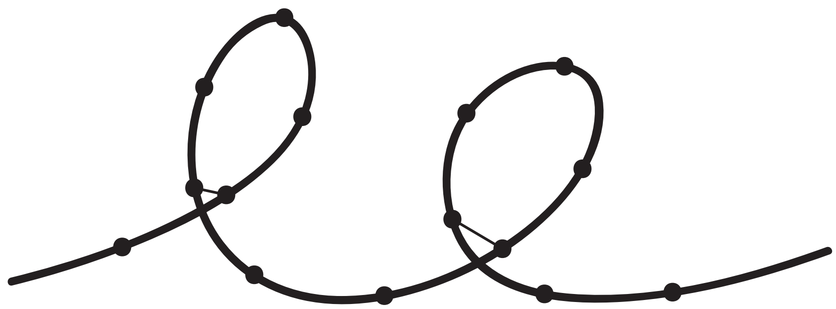

FIGURE 1. If a geodesic intersects itself many times, we can lay down points along it at regular intervals so that near to every intersection there is a cross-branch pair of nearest neighbors. By perturbing the initial point and increment, we can ensure that the distances corresponding to the cross-branch pairs are all distinct.

are indices $n(i),\ i \in I$ with $L_i \leq \mathrm{d}(\gamma(i), \gamma(n(i))) < L_i + \frac{\epsilon}{2}$. To finish the proof, set $N = \max\big(I \cup \{n(i) \mid i \in I\}\big)$ and $X_N = \gamma(\{0, \ldots, N\})$. For each $i \in I$, we have that $L_i \leq \mathrm{nnd}(\gamma(i), X_N) < L_i + \frac{\epsilon}{2}$. Thus $\mathrm{nnd}(\gamma(i), X_N)$ are all distinct, contradicting the assumption that M has K-bounded geodesic combinatorics. □

Definition (Asymptotic ray)**.** *Let $\alpha \subset M$ be a closed geodesic in a complete Riemannian manifold M. A geodesic ray $\gamma : [0, \infty) \to M$ is defined to be asymptotic to α if*

$$\lim_{t \to \infty} \mathrm{d}(\gamma(t), \alpha) = 0.$$

We say γ is non-trivially asymptotic to α if the image of γ is not equal to α.

Lemma 6. *Let M be a closed Riemannian manifold with bounded geodesic combinatorics. Then there are no non-trivial asymptotic rays.*

Proof. Let $\gamma : [0, \infty) \to M$ be a unit-speed geodesic ray non-trivially asymptotic to a closed geodesic $\alpha \subset M$. Since γ and α can intersect only countably many times (see e.g. [Be78, Lemma 7.10]), we can reparameterize γ at a different speed so that $\gamma(i) \notin \alpha,\ \forall i \in \mathbb{N}$. Then since γ is asymptotic to α, each $\gamma(i)$ is an isolated point of $\gamma(\mathbb{N})$. By Lemma 5, M cannot have bounded geodesic combinatorics. □

Remark. *The above lemma can be used to show that for a closed manifold M of dimension at least two, there is a C^2 open and dense set in the space of C^∞ Riemannian metrics on M which do not have bounded* *geodesic combinatorics. Indeed, according to* [Co08, page 3, paragraph 3]*, there is a C^2 open and dense set of metrics containing a periodic orbit with a transversal homoclinic point, and in particular, a non-trivial asymptotic ray.*

Lemma 7 (Geodesic Self Intersections Are Bounded)**.** *Assume that M is a Riemannian manifold that has K-bounded geodesic combinatorics. Let $\gamma : (0, a) \to M$ be a geodesic segment and assume that the image of γ has n transverse self intersections of multiplicities $m_1, \dots, m_n$. Then*

$$\sum_{i=1}^{n} m_i \leq 2K.$$

Proof. Pick some small $\epsilon, \delta > 0$, and let $X(\epsilon, \delta) = \{\gamma(i\delta + \epsilon) | i = 0, \dots, \lfloor \frac{a-\epsilon}{\delta} \rfloor .\}$ Observe that if δ is small then the distance between a pair of nearest neighbors $\gamma(i\delta+\epsilon), \gamma(j\delta+\epsilon) \in X(\epsilon, \delta)$ is exactly δ unless $(i\delta + \epsilon, j\delta + \epsilon)$ lies very close to a pair of times (s_i, t_j) describing a self intersection of γ. Furthermore, near a multiplicity m self intersection point of γ there are at least m points in $X(\epsilon, \delta)$ that lie at a distance less than δ from some other point of $X(\epsilon, \delta)$: this follows easily from the fact that on each of the m geodesic subsegments having a common intersection point, there is a point from $X(\epsilon, \delta)$ whose distance to the intersection is at most $\frac{\delta}{2}$.

Let $N \subset \{0, \dots, \lfloor \frac{a-\epsilon}{\delta} \rfloor\}^2$ be a maximal subset such that if $(i, j) \in N$ then

- $\mathrm{d}(\gamma(i\delta + \epsilon), \gamma(j\delta + \epsilon)) < \delta$
- $\mathrm{d}(\gamma(i\delta + \epsilon), \gamma(j\delta + \epsilon)) \in \mathrm{NND}(X(\epsilon, \delta))$
- if $(k, l) \in N$ and $(k, l) \neq (i, j)$, then $\{s_i, t_j\} \neq \{s_k, t_l\}$.

The last condition means that no two pairs in N lie on the same (unordered) pair of branches of γ at the same self intersection. It is not hard to see that $|N| \geq \frac{1}{2} \sum_{i=1}^{n} m_i$.

Without loss of generality, we may assume that ϵ and δ have been chosen so that each element of N determines a point of $X(\epsilon, \delta)$ which has a *unique* nearest neighbor. Therefore, elements of N still determine pairs of nearest-neighbors in $X(\epsilon + x, \delta(1 + y))$ for all choices of x and y sufficiently close to zero. Hence, the lemma will be proved if we show that ϵ and δ can be perturbed by making appropriate choices of x and y so that the distances $\mathrm{d}(\gamma(i\delta(1 + y) + (\epsilon + x)), \gamma(j\delta(1 + y) + (\epsilon + x)))$ for $(i, j) \in N$ are all distinct.

Assume that $s, t \in (0, a)$ and consider for small $c > 0$ the function $f_{s,t} : (-c, c) \times (-c, c) \to M$ defined by

$$f_{s,t}(x, y) = \mathrm{d}(\gamma(s + x + sy), \gamma(t + x + ty)).$$

Note that the functions $(f_{s,t})^2$ vary smoothly with both s and t. If $s = i\delta + \frac{\epsilon}{1+y}$ and $t = j\delta + \frac{\epsilon}{1+y}$, then $f_{s,t}(x,y)$ can be interpreted geometrically as the nearest-neighbor distance realized between the points $\gamma(i\delta(1+y) + (\epsilon + x))$ and $\gamma(j\delta(1+y) + (\epsilon + x)))$ from $X(\epsilon + x, \delta(1+y))$.

Lemma 8. *Assume that $\gamma(s) = \gamma(t) = p \in M$. Let α be the angle between $\gamma'(s)$ and $\gamma'(t)$ in TM_p. Then*

$$\begin{aligned} \frac{\partial^2}{\partial \xi^2}\big(f_{s,t}(\xi,0)\big)^2\big|_{\xi=0} &= 4 - 4\cos(\alpha) \\ \frac{\partial^2}{\partial \xi^2}\big(f_{s,t}(0,\xi)\big)^2\big|_{\xi=0} &= 2s^2 + 2t^2 - 4st\cos(\alpha) \\ \frac{\partial^2}{\partial \xi^2}\big(f_{s,t}(\xi,\xi)\big)^2\big|_{\xi=0} &= 2s^2 + 2t^2 - 4st\cos(\alpha) + \\ &\quad + (4 - 4\cos(\alpha))\,(1 + s + t)\,. \end{aligned}$$

Proof. We will prove the first equality, since the others are proven similarly. Observe that $f_{s,t}(0,0)$ and $\frac{\partial}{\partial \xi}\big(f_{s,t}(\xi,0)\big)^2|_{\xi=0}$ both vanish — the latter does because $f_{s,t}(\xi,0) \le 2\xi$ by the triangle inequality. It follows that

$$\begin{aligned} \frac{\partial^2}{\partial \xi^2}\big(f_{s,t}(\xi,0)\big)^2\big|_{\xi=0} &= 2\lim_{\xi\to 0}\frac{(f_{s,t}(\xi))^2}{\xi^2} \\ &= 2\big(\frac{\partial}{\partial \xi}f_{s,t}(\xi,0)\big|_{\xi=0^+}\big)^2. \end{aligned}$$

Now for any $v, w \in TM_p$ we have

$$\mathrm{d}(\exp_p(v), \exp_p(w)) = |v - w| + O(\max\{|v|,|w|\}^2).$$

Therefore,

$$\frac{\partial}{\partial \xi}f_{s,t}(\xi,0)\big|_{\xi=0^+} = \frac{\partial}{\partial \xi}\Big|\exp^{-1}(\gamma(s+\xi)) - \exp^{-1}(\gamma(t+\xi))\Big|_{\xi=0^+}.$$

Since $\exp^{-1}(\gamma(s+\cdot))$ and $\exp^{-1}(\gamma(t+\cdot))$ are simply two lines through the origin in TM_p intersecting with angle α, this can be calculated directly using the law of cosines. □

The point of Lemma 8 is that the triple

$$D(s,t) = \frac{\partial^2}{\partial \xi^2}\Big(f_{s,t}(\xi,0)^2, f_{s,t}(0,\xi)^2, f_{s,t}(\xi,\xi)^2\Big)\Big|_{\xi=0}$$

determines the set $\{s,t\}$ when $\gamma(s) = \gamma(t)$.

We now return to our discussion of N. Recall that by choosing δ small, we can ensure that if $(i,j) \in N$ then $(i\delta + \epsilon, j\delta + \epsilon)$ is arbitrarily close to some pair of times (s_i, t_j) with $\gamma(s_i) = \gamma(t_j)$. Also, no two

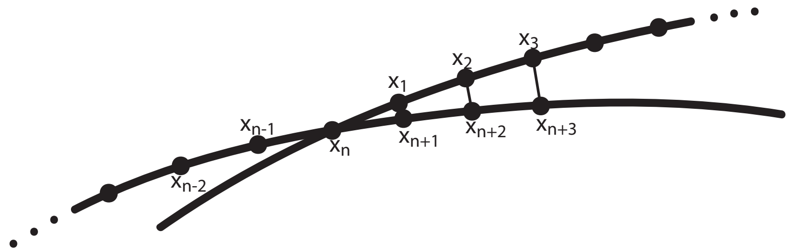

FIGURE 2. If a geodesic intersects itself with a shallow angle, a large number of nearest neighbor distances can be realized near to the intersection. The distances will increase as the points move farther away from the intersection, and therefore will be distinct.

elements of N give the same unordered pair (s_i, t_j). So, by Lemma 8, the triples $D(s_i, t_j)$ are distinct for $(i, j) \in N$. If δ is chosen small enough, the triples $D(i\delta + \epsilon, j\delta + \epsilon)$ will approximate $D(s_i, t_j)$ very closely, and therefore will be distinct as well.

The terms of $D(i\delta + \epsilon, j\delta + \epsilon)$ are second (directional) derivatives of the function $(\epsilon, \delta) \mapsto \mathrm{d}(\gamma(i\delta+\epsilon), \gamma(j\delta+\epsilon))$. If two real valued functions differ on some (even higher-order) partial derivative at a point, there is a neighborhood of that point on which the two functions disagree almost everywhere. Therefore, for almost every small perturbation of ϵ and δ, the distances $\mathrm{d}(\gamma(i\delta+\epsilon), \gamma(j\delta+\epsilon))$ will be distinct for $(i, j) \in N$. These constitute at least $\frac{1}{2}\sum_{i=1}^{n} m_i$ nearest neighbor distances in $X(\epsilon, \delta)$, so the lemma follows. □

Here is another result in a similar vein.

Lemma 9 (Self Intersections Have Large Angles). *Let M be a closed Riemannian manifold with K-bounded geodesic combinatorics. Then there is some $\theta(K) > 0$ such that no geodesic in M intersects itself with an angle less than $\theta(K)$.*

Proof. The idea of the proof is fairly clear from Figure 2, so we will try to present a rigorous proof as concisely as possible. Choose $\theta(K) > 0$ small enough so that if $0 \leq \theta < \theta(K)$ and $i \in \{1, \ldots, K+1\}$, then

(1) $i(2 - 2\cos(\theta)) < 1$,
(2) $(1 + \frac{1}{i})^2 > \frac{2}{1+\cos(\theta)}$.

This will be the angle described in the statement of the lemma.

We argue by contradiction, assuming that $\gamma : \mathbb{R} \to M$ is a geodesic in M that intersects itself at $p \in M$ with an angle θ less than $\theta(K)$. We will assume that the angle θ is actually made by $\gamma'(s)$ and $\gamma'(t)$

for some $s < t \in \mathbb{R}$. The other case, in which $\gamma'(s)$ and $-\gamma'(t)$ make a small angle, is handled similarly. Furthermore, note that by passing to a subinterval of $[s,t]$ if necessary, we may assume that θ is the smallest angle of self-intersection at p in the interval $[s,t]$.

Let us now define the points shown in Figure 2. For $n \in \mathbb{N}$, let $x_i = \gamma(s + i\big(\frac{t-s}{n}\big))$ and

$$X(n) = \{x_i \,|\, i = 0, \dots, n + K + 1\}.$$

To begin with, we estimate the distance between x_i and x_{i+n}, where $1 \le i \le K+1$. If n is large then these points are very close to p, so $\mathrm{d}(x_i, x_{i+n})$ is closely approximated by the distance between the vectors $v_i = i\big(\frac{t-s}{n}\big)\gamma'(s)$ and $v_{i+n} = i\big(\frac{t-s}{n}\big)\gamma'(t)$ in TM_p which exponentiate to x_i and x_{i+n}:

$$\begin{aligned} \mathrm{d}(x_i, x_{i+n}) &= ||v_i - v_{i+n}||_{TM_p} + O(\max(||v_i||, ||v_{i+n}||)^2) \\ &= \frac{t-s}{n}\, ||i\gamma'(s) - i\gamma'(t)||_{TM_p} + O(\frac{1}{n^2}) \\ &= i\frac{t-s}{n}\big(2 - 2\cos(\theta)\big)^{\frac{1}{2}} + O(\frac{1}{n^2}). \end{aligned} \tag{3.1}$$

Let $\mathrm{nnd}(i) = \mathrm{nnd}(x_i, X(n))$, and note that by (1),

$$\mathrm{nnd}(i) \le \mathrm{d}(x_i, x_{i+n}) < \frac{t-s}{n},$$

for sufficiently large $n \in \mathbb{N}$. Hence, for n sufficiently large the nearest-neighbor of x_i in $X(n)$ is in a branch of γ passing through p distinct from the branch containing x_i. In particular, $\mathrm{nnd}(i)$ is at least as large as the minimum distance of x_i to another branch and since θ was chosen to be minimal, this distance is bounded below by $i\frac{t-s}{n}\sin(\theta) + O(\frac{1}{n^2})$. Combining these estimates, yields that

$$\mathrm{nnd}(i) \in (i\frac{t-s}{n}\sin(\theta) + O(\frac{1}{n^2}), i\frac{t-s}{n}\big(2 - 2\cos(\theta)\big)^{\frac{1}{2}} + O(\frac{1}{n^2})).$$

A simple calculation using (2) shows that these intervals are disjoint for $i = 1, 2, \dots, K+1$ when n is sufficiently large, concluding the proof. □

The final lemma applies only to surfaces, utilizing that geodesics locally separate space:

Lemma 10. *Let M be a closed surface with bounded geodesic combinatorics. Then every non-closed geodesic in M is simple and accumulates on itself in TM.*

Proof. Assume that $\gamma \subset M$ is a non-closed geodesic. As M has bounded geodesic combinatorics, Lemma 5 provides a sequence of times $t_n \to \infty$ such that $\gamma(t_n)$ converges to $\gamma(t)$ for some $t \in \mathbb{R}$. Moreover, γ cannot intersect itself infinitely many times by Lemma 7, so $\gamma'(t_n)$ converges to $\gamma'(t)$ in TM. Thus γ accumulates on itself in TM. Finally, by a continuity argument we see that for arbitrary $s \in \mathbb{R}$, $\gamma'(t_n + (s-t)) \to \gamma'(s)$ in TM. This implies that γ must be simple, for otherwise any self intersection of γ will be accompanied by infinitely a many other self intersections, which again contradicts Lemma 7. □

Perhaps surprisingly, there are metrics on the two sphere with non-closed geodesics $\gamma : \mathbb{R} \to \mathbb{S}^2$ that are simple and accumulate on themselves in the tangent bundle [CoHi06]. Thus, Lemma 10 is not sufficient to prove that metrics on the two sphere with bounded geodesic combinatorics have all geodesics closed.

4. Surfaces with Bounded Geodesic Combinatorics

In this section, we prove Theorem 2. Before giving the proof, we record the following two results about surfaces:

Theorem 3 (Gaĭdukov, [Ga66]). *Let M be a closed, oriented surface of positive genus, $p \in M$, and Γ a non-trivial free homotopy class of closed cuves in M. Then there are a closed geodesic $\gamma \in \Gamma$ and a ray $r : [0,\infty) \to M$ starting at p that is asymptotic to γ. Moreover, any cover of r and any cover of γ in the universal cover $\tilde{M}$ is a globally minimizing geodesic.*

Theorem 4 (Innami, [In86]). *Let M be a closed, oriented surface of positive genus. Suppose that for each non-trivial free homotopy class of closed curves Γ in M, there is a foliation of M by geodesics all belonging to the class Γ. Then M is a flat torus.*

We now begin the proof of Theorem 2.

Theorem 2 (Classification in Dimension 2). *Assume that M is a closed Riemannian surface. Then M has bounded geodesic combinatorics if and only if it is isometric to $\mathbb{T}^2$ with a flat metric, $\mathbb{RP}^2$ with a round metric or $\mathbb{S}^2$ with an SC-metric.*

Proof. It was shown in Section 2 that the three candidates above all have bounded geodesic combinatorics. So, fix a closed Riemannian surface M that has bounded geodesic combinatorics; we will show that it is isometric to one of the indicated manifolds.

Consider first the case when M is not homeomorphic to $\mathbb{S}^2$ or $\mathbb{RP}^2$; we will prove that M is isometric to a flat torus. To begin, let M_O denote a connected component of the oriented double cover of M. Fix a simple closed curve $\gamma_0 \in M_O$, and let $p \in M_O$. As in Theorem 3, pick a geodesic ray passing through p asymptotic to a closed geodesic γ_p in the homotopy class of γ_0. Since M has bounded geodesic combinatorics, Lemma 6 implies that the ray's projection in M cannot be nontrivially asymptotic, so its image in M_O is exactly γ_p. Therefore, we have a closed geodesic γ_p containing p in the homotopy class of γ_0. Recall from Theorem 3 that γ_p lifts to a minimizing geodesic in the universal cover.

We claim that the collection $\{\gamma_p\}$ is a foliation of M_O by circles. First, each γ_p must be simple. For otherwise, since it can be homotoped to be simple, there must be a pair of arcs on γ_p that bound a bigon in M_O. We can then replace one of these arcs by the other and smooth out the corners to create a new curve (based) homotopic to γ_p of shorter length, violating the condition that γ_p lifts to a distance minimizing geodesic in the universal cover. Next, assume that γ_p and γ_q intersect but are not equal. Since they can be homotoped to be disjoint, there is an arc on each which when combined bound a bigon in M_O. This is a contradiction as before.

Since M_O is orientable and can be foliated by circles, it is topologically a torus. Each free homotopy class of closed curves on the torus admits a simple representative. Therefore, repeating the argument above proves that for each homotopy class there is a foliation of M_O by closed geodesics in that class. Theorem 4 now implies that M_O is flat. Therefore, M is either a flat torus or a flat Klein bottle. It is easy, however, to find geodesics on a flat Klein bottle that intersect themselves in arbitrarily many distinct points, so since M has bounded geodesic combinatorics, Lemma 7 implies that M is a flat torus.

Assume now that M is homeomorphic to either $\mathbb{S}^2$ or $\mathbb{RP}^2$. It suffices to show that all geodesics in M are closed. Indeed, Gromoll and Grove [GrGr81] have proven that any Riemannian metric on $\mathbb{S}^2$ with all geodesics closed is in fact an SC-metric, and Pries [Pr07] showed that a Riemannian metric on $\mathbb{RP}^2$ has all its geodesics closed if and only if it has constant curvature.

As in the previous case of the proof, we will never explicitly reference the bounded geodesic combinatorics property. In fact, by Lemmas 7, 9, and 10, any metric on $\mathbb{S}^2$ or lift of a metric on $\mathbb{RP}^2$ with bounded geodesic combinatorics satisfies the hypotheses of the following claim:

Claim. *Let M be a Riemannian manifold homeomorphic to $\mathbb{S}^2$. If*

- *every non-closed geodesic in M is simple and accumulates on itself, and*
- *there is an upper bound for the number of self intersections of any geodesic and a lower bound for the angle of any self-intersection,*

then all geodesics in M are closed.

To prove the claim, we fix a point $p \in M$ and show that all geodesics through p are closed. For $v \in SM_p$, let $\gamma_v : \mathbb{R} \to M$ denote the unit speed paramaterization of the geodesic with $\dot{\gamma_v}(0) = v$. Define

$$U = \{v \in SM_p \mid \exists\, T > 0 \text{ such that } \gamma_v(T) \text{ is conjugate to } p\}.$$

The following lemma shows that there can be no conjugate points along any geodesic in a Riemannian surface that is simple and accumulates on itself. It follows that γ_v is closed for all $v \in U$.

Lemma 11. *Assume that M is a Riemannian surface, $\gamma : \mathbb{R} \to M$ is a geodesic, and $\gamma(a)$ and $\gamma(b)$ are conjugate along γ. Then given $\epsilon > 0$, there exists $\delta > 0$ such that if $\alpha : \mathbb{R} \to M$ is a geodesic with*

$$|\alpha(t) - \gamma(t)| < \delta, \ \forall t \in [a - \epsilon, b + \epsilon]$$

then $\alpha([a - \epsilon, b + \epsilon])$ must intersect $\gamma([a - \epsilon, b + \epsilon])$.

Proof. Assume without loss of generality that there is no point on γ before $\gamma(b)$ that is conjugate to $\gamma(a)$. Extend γ to a geodesic variation $\gamma_t : [a - \epsilon, b + \epsilon] \to M$ with $\gamma_0 = \gamma$ such that the Jacobi field $\vec{v}(s) = \frac{\partial}{\partial t}\gamma_t(s)$ is nontrivial, orthogonal to $\gamma'(s)$ and vanishes at $s = a$ and $s = b$, but not between. Because $\vec{v}$ is not identically zero and determined by the covariant derivative $\frac{\nabla}{ds}\vec{v}(s)|_{s=a}$, this derivative must be nonzero. Thus $\vec{v}(s)$ must lie for $s > a$ on the side of $\gamma'(s)$ opposite to that on which it lies for $s < a$. A similar statement holds for s near b. Therefore, by considering only t sufficiently close to 0, we can ensure that each γ_t intersects γ very close to $\gamma(a)$ and $\gamma(b)$. Moreover, γ_t cannot intersect γ far away from $\gamma(a)$ and $\gamma(b)$ because $\vec{v}(s)$ vanishes only when $s = a$ or $s = b$. Finally, since the intersections between two geodesics cannot be too close together we see that γ and γ_t intersect exactly twice. Thus the middle portion of each γ_t lies on one side of γ and the ends lie on the other side.

Choose δ small enough so that if α is as in the statement of the Lemma then it intersects the middle portion of some geodesic in the variation γ_t. Assume that $\alpha([a-\epsilon, b+\epsilon])$ does not intersect $\gamma([a-\epsilon, b+\epsilon])$ and let γ_{t_0} be the first geodesic in the variation that $\alpha([a - \epsilon, b + \epsilon])$ does intersects. The intersection of α and γ_{t_0} must lie in the interior of both segments, since close to their endpoints the segments lie on opposite sides of γ. Because α does not intersect γ_t for $t < t_0$, its intersection with γ_{t_0} cannot be transverse. This is a contradiction, because geodesics always intersect transversely. □

Our goal, then, is to show that $U = SM_p$. Since $\tilde{M}$ is compact, U must be nonempty. Moreover, a result of Warner [Wa65, Theorem 3.1] implies that each component of the tangential conjugate locus of M at p is a properly embedded 1-submanifold of TM_p transverse to the radial direction. Since U is the image of the tangential conjugate locus under the radial projection $TM_p \backslash \{0\} \to SM_p$, U is an open non-empty subset of SM_p. Seeking a contradiction, assume that $U \neq SM_p$ and let $I \subset U$ be a maximal open interval.

Lemma 12. *The geodesics $\{\gamma_v \mid v \in I\}$ have a common period.*

The proof of Lemma 12 is fairly long, so let us first describe how to finish the proof of the above claim assuming this lemma. The lemma implies that if v is an endpoint of I, then γ_v is closed and L is a period for γ_v. Furthermore, γ_v is part of a (one-sided) geodesic variation all of whose geodesics start at p and close up at time L. Therefore γ_v contains a point conjugate to p. This is a contradiction, by definition of I. Thus, the proof of the claim and therefore the proof of Theorem 2 is concluded with the following proof.

Proof of Lemma 12. The proof proceeds in several steps. We will first show that for any compact subset $C \subset I$, there is an upper bound for the least period of $\gamma_v, v \in C$. Next, we prove that there is a continuous function $l : C \to \mathbb{R}$ such that $l(v)$ is a period for γ_v, and then prove that the function is smooth. We then use l to parameterize the family γ_v as a proper geodesic variation and deduce from the first variational formula that in fact l is constant. In the final step, we use Lemma 9 to extend this common period over all of I.

To begin with, for each $v \in C$ let $lp(v)$ be the least period of γ_v. We first show that $lp(v)$ is bounded on C. Assume on the contrary that there is a sequence $v_n \in C$ with the property that $lp(v) \to \infty$. Without loss of generality, we may assume that $v_n \to v \in C$. Define $\hat{\gamma}_v, \hat{\gamma}_{v_n} : \mathbb{R} \to SM$ to be lifts of γ_v and γ_{v_n} to the unit tangent bundle of M, and let N be a small regular neighborhood of $\hat{\gamma}_v(\mathbb{R}) \subset SM$.

Since $v \in C$, there is some $T > 0$ such that $\gamma_v(0)$ and $\gamma_v(T)$ are conjugate along γ_v. Assume that $\beta \subset M$ is a geodesic segment that lifts into N and has length at least $K = T + lp(v)$. If N is small, β has a subsegment that closely tracks $\gamma_v([0, T])$. Lemma 11 implies that if this subsegment stays close enough to $\gamma_v([0, T])$, then β must in fact intersect γ_v. Therefore we can choose N small enough so that any such geodesic segment β must intersect γ_v.

Assume that n is very large. Then $\hat{\gamma}_{v_n}$ spends at least a duration of K inside of N. Moreover, $\hat{\gamma}_{v_n}$ must eventually exit N. For otherwise, γ_{v_n} will be a closed curve homotopic within an annular neighborhood of γ_v to a large power of γ_v, and therefore will intersect itself more often than is allowed by Lemma 7. Set t_n to be the first time at which $\hat{\gamma}_{v_n}$ exits N, and let $g_n = \gamma_{v_n}([t_n - K, t_n])$. Since g_n lifts into N and has length K, it must intersect γ_v. If the angle of intersection is very small, then g_n will be forced to track γ_v for a long time, and therefore cannot exit N before a duration of K had elapsed. Thus there is some $\alpha > 0$ such that each g_n intersects γ_v with angle at least α. But if n is very large, $\gamma_{v_n}([0, t_n - K])$ winds many times around γ_v while staying close enough to pass through any geodesic segment intersecting γ_v with angle at least α. Therefore, $\gamma_{v_n}([0, t_n - K])$ must intersect $g_n = \gamma_{v_n}([t_n - K, t_n])$ many times as well. This is impossible, since the number of self intersections of γ_{v_n} is limited by Lemma 7.

We now know that $lp(v)$ has an upper bound on C, which we set to be L. One geometric consequence of this is that a convergent sequence $v_n \to v \in C$ gives a sequence of geodesics γ_{v_n} whose images converge to the image of γ_v in the Hausdorff topology on closed subsets of M. Note that for large n, this implies that $lp(v_n)$ is very close to a period of γ_v.

Our aim is now to find a continuous function $l : C \to \mathbb{R}$ with $l(v)$ a period for γ_v. We will start by defining it locally. So, fix a vector $w \in C$; our goal is to produce a neighborhood $N_w \subset C$ of w and a continuous function $l_w : N_w \to \mathbb{R}$ that gives periods for $\gamma_v,\ v \in N_w$. To begin with, set

$$l_w(w) = \left\lfloor \frac{L}{lp(w)} \right\rfloor \; lp(w).$$

Define N_w to be a neighborhood of w small enough so that if $v \in N_w$, then $lp(v)$ is within $\frac{lp(w)}{8\left\lfloor \frac{L}{lp(w)} \right\rfloor}$ of a period of γ_w. Then for $v \in N_w$, set

$$l_w(v) = \text{the period of } \gamma_v \text{ that is within } \tfrac{lp(w)}{8} \text{ of } l_w(w).$$

We claim that $l_w : N_w \to \mathbb{R}$ is continuous. Assume that $v_n \to v \in N_w$. By definition, $l_w(v)$ and each $l_w(v_n)$ are within $\frac{lp(w)}{8}$ of $l_w(w)$. Therefore $|l_w(v) - l_w(v_n)| < \frac{lp(w)}{4}$. Since

$$lp(v) > lp(w) - \frac{lp(w)}{8\left\lfloor \frac{L}{lp(w)} \right\rfloor} > \frac{lp(w)}{2},$$

this implies

$$|l_w(v) - l_w(v_n)| < \frac{lp(v)}{2}.$$

So, $l_w(v_n)$ is closer to $l_w(v)$ than to any other period of γ_v. As in the previous paragraph, the fact that least periods are bounded in C implies that when n is large, $l_w(v_n)$ is very close to a period of γ_v. This period must then be $l_w(v)$. Therefore, $l_w(v_n) \to l_w(v)$.

We now have for each $w \in C$, a continuous function $l_w : N_w \to \mathbb{R}$ such that $l_w(v)$ is a period for γ_v. Pick a finite set $\{w_1, \ldots, w_n\} \subset C$ so that $N_{w_i} \cap N_{w_{i+1}} \neq \emptyset$ and $\cup_{i=1}^n N_{w_i}$ covers C. Since l_{w_1} and l_{w_2} are continuous and their quotient is rational valued, they can be multiplied by appropriate positive integers so that they coincide on the intersection of their domains. Note that this does not change the property that they pick out periods for $\gamma_v,\ v \in C$. A similar trick applies to make l_{w_3} agree with the previous two. Continuing inductively, we can piece our locally defined functions together to create a continuous function $l : C \to \mathbb{R}$.

To show that l is smooth, pick a vector $w \in C$ and choose a coordinate chart $\phi : O \to \mathbb{R}^2$ with $\phi(p) = (0,0)$ and $\mathrm{D}\,\phi_p(\gamma_w'(l(w)w)) = (1,0)$. Let $\pi : \mathbb{R}^2 \to \mathbb{R}$ be the projection onto the first coordinate. If $V \subset TM_p$ is a small neighborhood of $l(w)w$, then $\pi \circ \phi \circ \exp_p : V \to \mathbb{R}$ is defined and a submersion. By the Implicit Function Theorem, $(\pi \circ \phi \circ \exp_p)^{-1}(0)$ is a smooth 1-submanifold of V. Since $\{l(v)v \mid v \in C\} \cap V$ is a continuous 1-manifold contained in that preimage, the connected components of both that contain $l(w)w$ must coincide. This shows that l is smooth in a neighborhood of $w \in C$.

To show that l is constant, define a smooth geodesic variation $G : C \times [0,1] \to M$ by $G(v,t) = \gamma_v(l(v)t)$. Then $G(v,0) = G(v,1) = p$ for all $v \in C$, and the length of each segment $G(v,[0,1])$ is $l(v)$. By the First Variational Formula, $\frac{\partial}{\partial v} l(v) = 0$. So l is constant.

Finally, we must extend the common period L of $\{\gamma_v \mid v \in C\}$ over all of I. Assuming that this is not possible, we can enlarge C as much as possible so that any open neighborhood of it in SM_p contains some v for which γ_v does not have period L. There is then some endpoint a of C which is the limit of a sequence of points $v_n \in SM_p$ for which γ_{v_n} does not have period L. However, the v_n and C are all contained in some larger compact subinterval of I, over which there must be a common period by our earlier work. So, γ_{v_n} all have a common period kL for some $k \in \mathbb{N}$. Therefore, if n is large then γ_{v_n} closely tracks some nontrivial power of γ_a.

As $M \cong \mathbb{S}^2$, γ_a is the image of the core curve of some immersed annulus $A \looparrowright M$. When n is large, γ_{v_n} lifts into A and closely tracks some nontrivial multiple of the core curve. It therefore must intersect itself, and since it is always nearly tangent to the core curve the intersection must happen with small angle. This violates the assumptions of the claim, and finishes the proof.

□

References

[Be78] A. Besse. Manifolds all of whose geodesics are closed. *A Series of Modern Surveys in Mathematics*, **93**, Springer-Verlag, 1978.

[ChEb75] J. Cheeger and D.G. Ebin, *Comparison theorems in Riemannian geometry*, Amsterdam, Oxford: North-Holland 1975.

[CoHi06] T. Colding and N. Hingston, *Geodesic laminations with closed ends on surfaces and Morse index; Kupka-Smale metrics.* Comment. Math. Helv. 81 (2006), no.3, 495-522.

[Co08] G. Contreras, *Geodesic flows with positive topological entropy, twist maps and hyperbolicity.* To appear in Annals of Mathematics.

[do Ca92] M. do Carmo. Riemannian Geometry, Birkhauser, 1992.

[FrSó92] E. Fried and V. Sós, *A generalization of the three-distance theorem for groups*, Algebra Universalis, 29 (1992) 136-149.

[Ga66] E.V. Gaĭdukov, *Asymptotic geodesics on a Riemannian manifold nonhomeomorphic to the sphere*, Soviet Math. Dokl., Vol. 7 (1966), No.4.

[GrGr81] D. Gromoll and K. Grove, *On metrics on S^2 all of whose geodesics are closed*, Invent. math. 65 (1981), 175-177.

[He01] S. Helgason. *Differential Geometry, Lie Groups, and Symmetric Spaces.* Graduate Studies in Mathematics, 34. American Mathematical Society, Providence RI (2001).

[In86] N. Innami, *Families of geodesics which distinguish flat tori*, Math. J. Okayama Univ. 28 (1986), 207-217.

[Pr07] C. Pries, *If all geodesics are closed on the projective plane*, Archive preprint 0710.0951v1[math.DG], October 2007.

[Só58] V.T. Sós, *On the distribution mod 1 of the sequence* $\{n\alpha\}$, Ann. Univ. Sci. Budapest, Eotvos Sect. Math. 1 (1958), 127-134.

[Św58] S. Świerczkowski, *On successive settings of an arc on the circumference of a circle*, Fund. Math. 46 (1958), 187-189.

[Vi08] S. Vijay, *Eleven Euclidean Distances are Enough*, Journal of Number Theory, 128 (2008), 1655-1661.

[Wa65] F. Warner., *The conjugate locus of a Riemannian manifold*, American Journal of Mathematics, Vol. 87, No. 3 (Jul. 1965), 575-604.

Ian Biringer, University of Chicago, biringer@math.uchicago.edu

Benjamin Schmidt, University of Chicago, schmidt@math.uchicago.edu