Length of geodesics on a two-dimensional sphere.

Alexander Nabutovsky and Regina Rotman

Department of Mathematics, University of Toronto, Toronto, Ontario, M5S2E4, CANADA; and Department of Mathematics, Mc Allister Bldg., The Pennsylvania State University, University Park, PA 16802, USA

Abstract. Let M be an arbitrary Riemannian manifold diffeomorphic to S^2. Let x, y be two arbitrary points of M. We prove that for every $k = 1, 2, 3, \ldots$ there exist k distinct geodesics between x and y of length less than or equal to $(4k^2 - 2k - 1)d$, where d denotes the diameter of M.

1. Main results. Here are the main results of the present paper.

Theorem 1. Let M be an arbitrary Riemannian manifold diffeomorphic to S^2, and x, y be two arbitrary points of M. Denote the diameter of M by d. (Recall that the diameter of a compact Riemannian manifold is, by definition, the maximal distance between two points on the manifold.) For every positive integer k there exist at least k distinct geodesics starting at x and ending at y of length not exceeding $(4k^2 - 2k - 1)d$.

Theorem 1.1. For every point x on a Riemannian manifold M diffeomorphic to S^2 and any $k = 1, 2, 3, \ldots$ there exist at least k non-trivial geodesic loops based at x of length $(4k^2 + 2k)d$.

Theorem 1.1 deals with a particular case of Theorem 1 when $y = x$. In this case the shortest geodesic between x and y is trivial, and all other geodesics starting and ending at x are non-trivial geodesic loops based at x. Thus, we can apply Theorem 1 for $k+1$ and obtain an upper bound for the length of the kth non-trivial geodesic loop based at x. However, the upper bound provided by Theorem 1.1 is somewhat better.

Theorem 1 is a result in the direction of our conjecture made in [NR1]. Recall, that for every two points x in a closed Riemannian manifold M there exists an infinite set of distinct geodesics connecting x and y. (This is a well-known theorem of J. P. Serre, [Se]). In [NR1] we observed that if M is a non-simply connected closed Riemannian manifold with a torsion-free fundamental group, then for every $x, y \in M$ and every k there exist k distinct geodesics between x and y of length $\leq kd$. Of course, the same fact is true also for round spheres of all dimensions. And, if $k = 1$, then the existence of one geodesic between x and y of length $\leq d$ is a trivial corollary of the definition of diameter. All these facts led us to a conjecture that there exists a universal upper bound of the form $f(n, k)d$ for the length of k distinct geodesics between points of a closed n-dimensional Riemannian manifold of diameter d. The main point here is that this upper bound does not involve any information about the metric invariants of the Riemannian manifold other than its diameter.

Previously we established this conjecture for $k = 2$ and an arbitrary closed Riemannian manifold in [NR2]. More precisely, we proved that for any two points $x, y \in M^n$ there exist two distinct geodesics starting at x and ending at y of length $\leq 2nd$ (and even $\leq 2qd$, where $q = \min_i\{\pi_i(M^n) \neq 0\}$). The proof of this result in [NR2] was heavily based on methods developed in [R], where it had been proven that for any n, M^n, $x \in M^n$ the length of the shortest non-trivial geodesic loop based at x does not exceed $2qd(\leq 2nd)$, where d denotes the diameter of M^n. Since in our situation $n = 2$, we obtain the upper bound $4d$ which is better than the estimate asserted in Theorem 1 for $k = 2$.

Note that one can make even a bolder conjecture: For every closed Riemannian manifold M, every pair of points $x, y \in M$ and every positive integer k there exist k distinct geodesics between x and y of length $\leq kd$, where d denotes the diameter of M. If this stronger conjecture is true, then the estimate kd cannot be improved, as one can see by considering the case of round spheres. But one does not know if this stronger conjecture is true even in the simplest non-trivial case when $k = 2, x = y$ and M is a convex surface (diffeomorphic to S^2) in R^3.

Also note that A. Schwarz ([S]) noticed that a modification of the proof of J.P. Serre implies that for any closed Riemannian manifold M^n, any two points $x, y \in M^n$ and any k there exist at least k geodesics on M^n connecting x and y of length $\leq kC(M^n)$, where $C(M^n)$ depends on the ambient Riemannian manifold M^n. Of course, this result immediately implies that there exists a *scale-invariant* constant $c(M^n)$ such that there are k distinct geodesics connecting x and y of length $\leq c(M^n)kd$. Whenever the stronger version of our conjecture asserts that $c(M^n) = 1$, we do not know if there exists a uniform bound for $c(M^n)$ that depends only on n or even only on the diffeomorphism class of M^n. In fact, it is quite possible that such an upper bound does not exist, and the quadratic dependence on k in Theorem 1 is optimal.

Yet we are able to prove some specific linear estimates in k for the length of k distinct geodesics between arbitrary points. Define *the geodesic complexity* of M as follows:

Definition. For every pair of points $x, y \in M$ denote by $g_1(x, y)$ the number of distinct geodesics between x and y of length $\leq 2d + dist(x, y)$ such that each of these geodesics provides the local minimum of the length functional on the space of paths connecting x and y. Denote by $g_2(x, y)$ the number of distinct geodesics between x and y of length $\leq 2d$ such that each of these geodesics provides the local minimum of the length functional on the space of paths connecting x and y. The geodesic complexities T_i of M, ($i = 1$ or 2), are defined as $\min_{x,y\in M} g_i(x, y)$.

In particular, this definition implies that every two points x and y of M can be connected by at least T_2 geodesics of length $\leq 2d$ and T_1 geodesics of length $\leq 2d + dist(x, y)$. Observe, that obviously $T_1 \geq T_2 \geq 1$. The following theorem is non-trivial only for $k > T_i$.

Theorem 1.A. Let M be a Riemannian manifold diffeomorphic to S^2 with diameter d and geodesic complexities T_i, $i = 1, 2$. Then for every pair of points x, y of M and every k there exist at least k distinct geodesics between x and y of length $\leq (2(k-1)(2T_1+3)+2)d \leq kd(4T_1 + 6)$. Also, there exist at least k distinct geodesics between x and y of length $\leq (2(k-1)(T_2^2 - T_2 + 5) + 2)d \leq kd(2T_2^2 - 2T_2 + 10)$.

The methods of the present paper are quite different from the methods of [R] and [NR2]. To explain our methods let us start from the following definition:

Definition. Let M be a Riemannian manifold diffeomorphic to S^2, and L be a positive real number. An *L-slicing* of M is a non-zero degree map from S^2 to M such that the length of the image of every meridian of S^2 in M does not exceed L.

The importance of this definition for our purposes is due to the following lemma:

Lemma 3. Let M be a Riemannian manifold of diameter d, which is diffeomorphic to S^2 and admits an L-slicing for some positive L. Then for every two points $x, y \in M$ and every positive integer k there exist at least k distinct geodesics starting at x and ending at y of length $2(k-1)L + 2d$. If the L-slicing maps the South pole of S^2 into either x or y, then the upper bound for the length can be improved to $2(k-1)L + d$. If, in addition, $x = y$ then the upper bound can be improved to $2(k-1)L$.

This lemma will be proven in the next section. Its proof is based on Morse theory and is mostly a compilation of known facts and ideas.

As a corollary, one might be tempted to look for an L-slicing of M where $L \leq cd$ for an appropriate constant c. Yet examples of metrics on D^2 constructed by S. Frankel and M. Katz ([FK]) can be used to show that such a slicing does not always exist: Although we did not check all the details, it seems not difficult to prove that if one takes smoothed out doubles of Riemannian D^2 constructed by Katz and Frenkel, then one obtains a sequence of metrics on S^2 such that for no c all of them admit a cd-slicing.

So, instead we establish the following dichotomy: Either there exists an cd-slicing of M for a controlled not very large c, or for every two points $x, y \in M$ there exist many distinct geodesics between x and y of length $\leq 2d$ which are local minima of the distance function on the space of all paths between x and y. Moreover, if there is no cd-slicing of M for a not very large c there must be many short geodesics between x and y which are "deep" local minima for the length functional. We need the following definition to state a precise form of this result:

Definition: Let x, y be two points of M, and S be a non-negative real number. Let γ_1, γ_2 be two geodesics from x to y providing local minima for the length functional on the space of paths from x to y. We say that γ_1 and γ_2 are *S-distinct* if every path homotopy (i.e. a homotopy with fixed endpoints) between γ_1 and γ_2 must pass through a path of length $\geq \max\{\text{ length}(\gamma_1), \text{length}(\gamma_2)\} + S$.

Theorem 2. (Dichotomy Theorem) Let M be a Riemannian manifold diffeomorphic to S^2, x, y be two points of M, $k \geq 2$ and $S \geq 0$. Then:
I. One of the following two assertions is true:
(A_1) There exist at least k pairwise S-distinct geodesics between x and y of length $\leq 2d + dist(x, y)$ that are local minima for the length functional on the space of paths between x and y.
(B_1) There exists an L-slicing of M with $L = (2k-1)d + 2dist(x, y) + S$. Moreover, this L-slicing maps the South pole of S^2 into x.
II. Also one of the following two assertions is true:
(A_2) There exist at least k pairwise S-distinct geodesics between x and y of length $\leq 2d$ that are local minima for the length functional on the space of paths between x and y.
(B_2) There exists an L-slicing of M with $L = (k^2 - 3k + 7)d + S$. Moreover, this L-slicing maps the South pole of S^2 into x.

Observe that Theorem 2 immediately implies Theorem 1. Indeed, if (A_1) is true, then the theorem is true. But, if (B_1) is true, the theorem immediately follows from Lemma 3. Similarly, it implies Theorem 1.1. (We need to apply Theorem 2 and Lemma 3 in the case when $x = y$ for $k+1$ instead of k. Indeed, one of the $k+1$ geodesic loops based at x will be trivial.)

Theorem 2 also easily implies Theorem 1.A. Indeed, observe that the definition of the geodesic complexity T_i of M implies the existence of a pair of points $x, y \in M$ for which the alternative (A_i) in Theorem 2 does not hold for $k = T_i + 1$. (Here $i = 1$ or 2, of course.) Therefore there exists an L-slicing with $L = (2T_1 + 3)d$, if $i = 1$, or $L = (T_2^2 - T_2 + 5)d$, if $i = 2$. Now Lemma 3 immediately implies Theorem 1.A.

The above argument can also be combined with the following observation: Assume that M admits an L slicing. Let γ be a closed curve of length l. We can use the images of meridians under the L-slicing to contract γ to a point. The length of the trajectory of every point of γ during a contracting homotopy will not exceed L, and the homotopy will pass only through closed curves of length $\leq l + 2L$. Thus, we arrive to the following corollary of Theorem 2:

Corollary: Let M be a Riemannian manifold of diameter d, which is diffeomorphic to S^2. Let T_1, T_2 be geodesic complexities of M. Denote $\min\{2T_1 + 3, T_2^2 - T_2 + 5\}$ by τ. Then every closed curve of length l can be contracted to a point by a homotopy passing through closed curves of length not exceeding $l + 2\tau d$ such that the length of the trajectory of every point does not exceed τd.

So, we need only to prove Lemma 3 and Theorem 2 in order to establish the rest of the results of this paper. We are going to prove Lemma 3 in the next section. Theorem 2 will be proven in section 3.

Section 4 contains some strengthenings of the Dichotomy Theorem. In particular, one of the results that can be found there asserts that if some two points x, y of M are connected by only one geodesic of length $\leq 2d$, then there exists a $3d$-slicing of M that maps the South pole of S^2 into x. In section 4 we present also another dichotomy theorem (Theorem 2.A). One of its corollaries is that for every $k \geq 2$ and every pair of points $x, y \in M$ either the space of paths of length $\leq 2d$ connecting x and y has at least k connected components, or M admits a $(k^2 - 3k + 7)d$-slicing.

2. Proof of Lemma 3.

We are going to deduce Lemma 3 from the following Lemma 4 that will be proven at the end of this section, To state this lemma we need the following notation: For any two points x, y in a manifold X let $\Omega_{x,y}X$ denote the space of all paths starting at x and ending at y. Of course, all these spaces are homotopy equivalent to the loop space ΩX.

Lemma 4. Let S and N be the South and the North poles of S^2 with the standard round metric. There exists a generator of $H_2(\Omega_{S,S}S^2)$, which can be represented as the image of the fundamental class of S^2 under the homomorphism induced by a map $\lambda : S^2 \longrightarrow \Omega_{S,S}(S^2)$, such that every loop in the image of λ is either an arc of a meridian traversed twice in opposite directions, or consists of a meridian that goes from S to N, and a meridian that returns from N to S.

Proof of Lemma 3. To see that Lemma 4 implies Lemma 3 we will follow the exposition in [S]. Denote the homology class $\lambda_*([S^2])$ introduced in Lemma 4 by h and a two-dimensional cohomology class dual to h by c. (In other words, c must satisfy $< c, x >= 1$.) For every $m = 1, 2, \ldots$ define h^m as $(\lambda^m)_*([(S^2)^m])$, where $\lambda^m : (S^2)^m \longrightarrow \Omega_{S,S}(S^2)$ is the map defined by the formula $\lambda^m(s_1, \ldots, s_m) = (\lambda(s_1), \ldots, \lambda(s_m))$. It had been demonstrated in [S] that for every $m < h^m, c^m >= m!$, where c^m denotes the mth cup power of c.

Alternatively, we can use the results on rational homology and cohomology algebras of ΩS^n that easily follow from the rational homotopy theory (cf. [FHT], ch. 16). (The multiplication on homology groups of loop spaces is induced by the composition of paths regarded as a map $\Omega S^n \times \Omega S^n \longrightarrow \Omega S^n$.) In particular, the loop space of S^2 is rationally homotopy equivalent to the product of S^1 and CP^∞. The class c introduced above corresponds to a generator of the algebra $H^*(CP^\infty, Q)$. The rational homology and cohomology groups of ΩS^2 in every dimension are isomorphic to Q. The rational homology algebra of S^2 is isomorphic to the algebra of polynoimials $Q[t]$ of one variable (of degree one) ([FHT], p.234-235). The class h introduced above corresponds to the square of a generator of this algebra.

Let ϕ be a L-slicing of M. It induces a map $\Omega_{S,S}(S^2) \longrightarrow \Omega_{\phi(S),\phi(S)}(M)$ that will be denoted by $\tilde{\phi}$. Note that the homomorphisms in homology groups (in all positive dimensions) induced by $\tilde{\phi}$ are also non-trivial. Consider classes $\tilde{h}, \tilde{c}$ corresponding to h, c in $\Omega_{x,y}(M)$. More precisely, consider some minimizing geodesics τ_1 from x to $\phi(S)$ and τ_2 from $\phi(S)$ to y, and for every $s \in S^2$ define $\tilde{\lambda}(s)$ as the join of τ_1, $\tilde{\phi}(\lambda(s))$ and τ_2. The map $\tilde{\lambda}$ induces isomorphisms between all rational homology and cohomology groups of $\Omega_{S,S}(S^2)$ and $\Omega_{x,y}(M)$. Denote classes of $\Omega_{x,y}(M)$ corresponding to homology and cohomology classes h^m and c^m by $\tilde{h^m}$ and $\tilde{c^m}$. We have exhibited above a specific cycle representing h^m. The definition of L-slicings and the definition of $\tilde{\lambda}$ immediately imply that, when we apply $\tilde{\lambda}$ to this cycle we obtain a $2m$-dimensional cycle in $\Omega_{x,y}(M)$ made of paths of length $\leq 2mL + 2d$ ($\leq 2mL + d$, if either $x = \phi(S)$, or $y = \phi(S)$; $\leq 2mL$, if $x = y = \phi(S)$).

Note that when one pulls down $\tilde{h^m}$ as far as possible, it gets stuck at a critical point of the length (or the energy) functional, which is a geodesic between x and y. Here one can just use the gradient flow of the energy functional on $\Omega_{x,y}(M)$, and to use a modification of the work of N. Koiso [Ko] or M. Grayson [G] where the space of closed curves is replaced by the space of paths with fixed endpoints to prove the local and global existence of solutions of the corresponding parabolic PDE. But, of course, this is not needed. Classically one circumvents the technical difficulties related to the appearance of non-linear parabolic PDEs by using a finite-dimensional approximation of $\Omega_{x,y}(M)$ and a gradient flow of the energy functional on this finite-dimensional approximation (cf. [B] or [Kl]). One can define the same critical values of the energy functional in terms of cohomology classes $\tilde{c^m}$ as the minimal sublevel set of the energy functional which can be a support of a cochain representing $\tilde{c^m}$.

In general critical points corresponding to homology (or cohomology) classes of different dimension need not be different, but Lyusternik and Schnirelman observed that if one of the cohomology classes is a cup product of the other and a third cohomology class of dimension l, and these two classes correspond to the same critical level, then the dimension of the critical point set at this level is at least l. See, for example, [Kl] for a proof of this fact. (In terms of homology classes one needs to require that one of these classes is a cap product of the other and a cohomology class of dimension l. In this case one says that the first homology class is *subordinate* to the second.)

Since the difference of dimensions of classes $\tilde{c^m}$ is at least two, and the dimension of any critical level (=the set of geodesics between two fixed points of the same length on M) is at most one, one concludes that $\tilde{h^m}$ correspond to geodesics between x and y of diffferent length for different values of $m = 0, 1, 2, \ldots$, and the lemma follows. QED

Proof of Lemma 4. First, recall that $\pi_1(\Omega S^2) = H_1(\Omega S^2) = Z, H_2(\Omega S^2) = Z, \pi_2(\Omega S^2) = \pi_3(S^2) = Z$ and is induced by the Hopf fibration $H : S^3 \longrightarrow S^2$. Since $H_2(\pi_1(\Omega S^2)) = 0$ the Hurewicz homomorphism $\pi_2(\Omega S^2) \longrightarrow H_2(\Omega S^2)$ is surjective by virtue of the Hopf theorem. Therefore we can construct a generator of $H_2(\Omega_{S,S} S^2)$ by merely regarding the Hopf fibration as a cycle in $H_2(\Omega_{S,S} S^2)$. Let $x_0 \in H^{-1}(S) \subset S^3$ be a point in the inverse image of the South pole of S^2. Consider a slicing of S^3 by loops based at x_0 and transversally intersecting a big $S^2 \subset S^3$ passing through x_0 at one point (pairwise intersecting only at x_0). (The resulting picture will be a three-dimensional analog of slicing of S^2 into loops depicted on Fig. 1.) Then the images of these loops under H will be elements of $\Omega_{S,S} S^2$. Together they will form a generator of $H_2(\Omega_{S,S} S^2)$.

Here is an explicit description of the generator. Big circles passing through the South pole S on S^2 constitute a circle in $\Omega_{S,S} S^2$. Take a big circle B passing through S. Take the perpendicular big circle b passing through S. It consists of two meridians m_1 and m_2. Contract B to S along circles passing through S and tranversally intersecting m_1 (correspondingly, m_2) at one point. These will be two halves of the homologically non-trivial circle in ΩS^2 hanging at B (see Fig. 1).

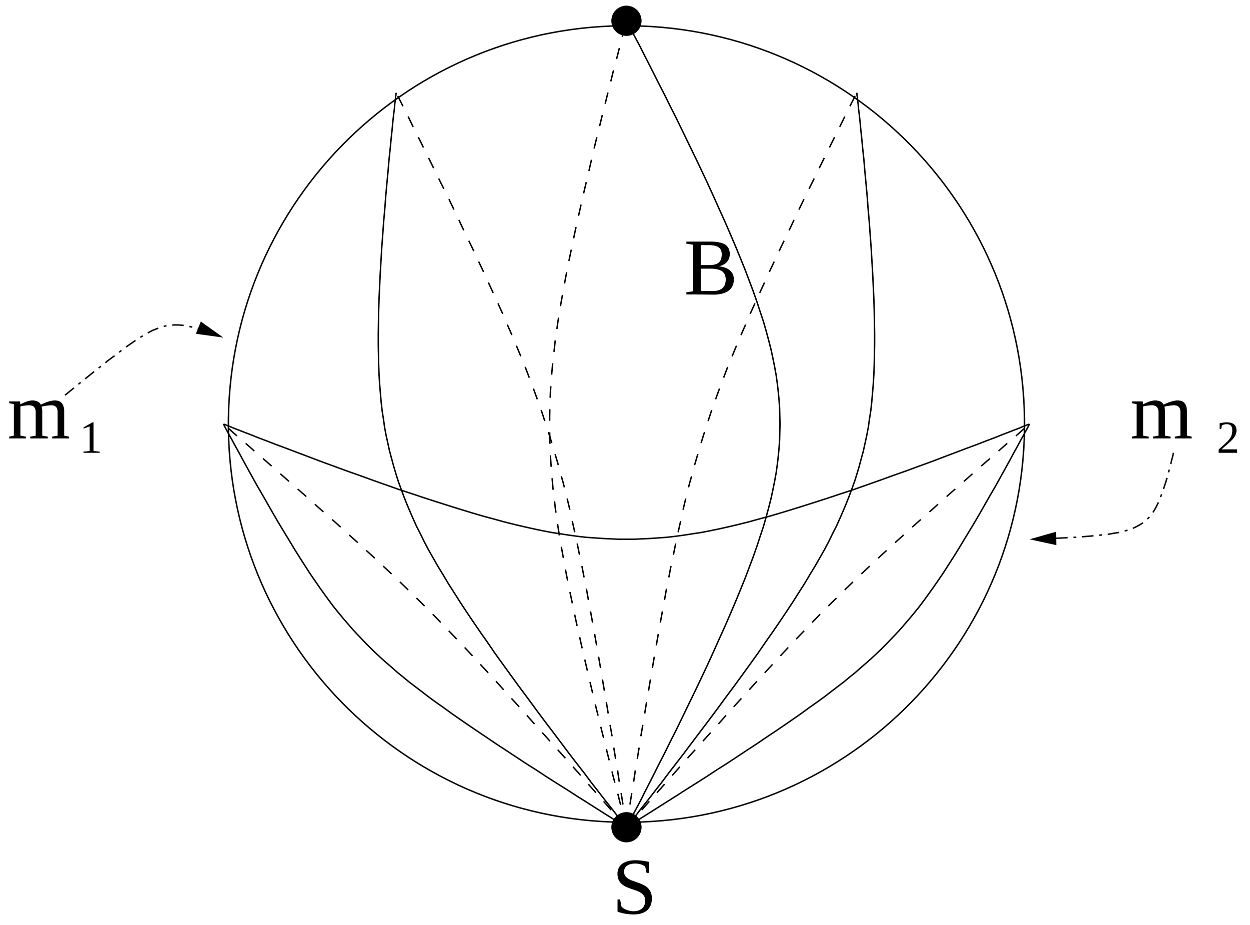

Figure 1.

We prefer to replace this circle by the following homotopic circle: Start from S regarded as a trivial loop. Continuously extend it along arcs of m_1 passed twice with opposite orientations until we obtain $m_1 * m_1^{-1}$. Now go along pairs of meridians forming with m_1 angles α and $-\alpha$, where α varies from 0 to π. At the end of this stage we obtain $m_2 * m_2^{-1}$. Contract the doubled m_2 to S along itself. (See Fig. 2.)

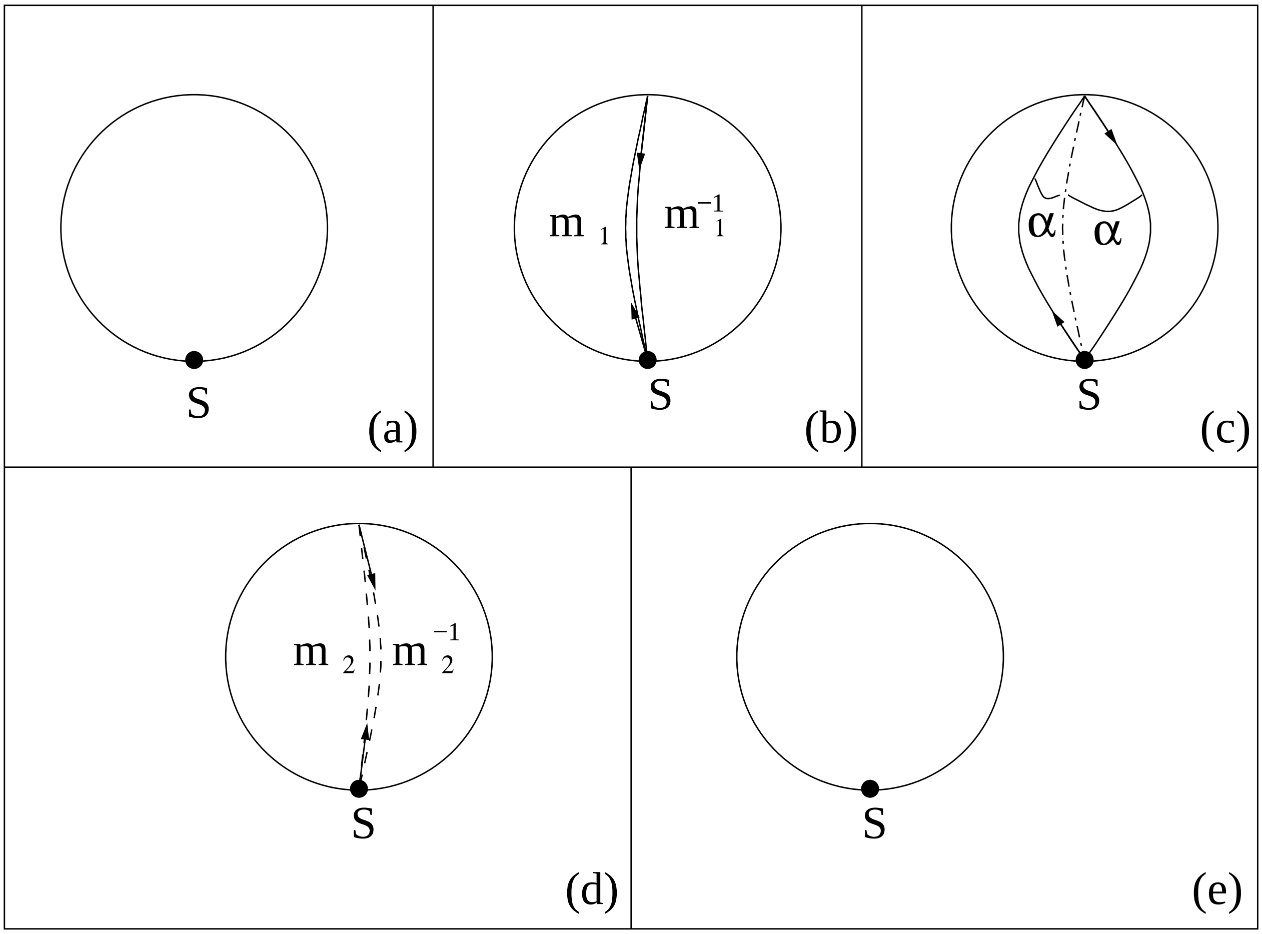

Figure 2.

Consider this circle as a map of a meridian μ of S^2 into $\Omega_{S,S}S^2$. (Here both poles of S^2 are mapped to the trivial loop.)

To obtain the desired two-dimensional homology class in $\Omega_{S,S}(S^2)$ we need to rotate B (and thereby b and the whole circle in $\Omega_{S,S}S^2$ obtained by the two different contractions of B to S). We obtain a desired map of S^2 to $\Omega_{S,S}S^2$. QED.

3. Proof of Theorem 2.

First, we are going to prove Theorem 2 in the case when M is analytic. In section 3.4 we will prove the theorem in the general (smooth) case. In sections 3.1-3.3 we will assume that M is a real analytic Riemannian manifold. Throughout this section we are going to use the following notation: The cut locus of a point $x \in M$ will be denote C_x.

3.1. Some facts about cut loci. Here we collected some facts about cut loci of points on analytic Riemannian manifolds diffeomorphic to S^2 that we will need for our proof. Most of these facts are known and were apparently first discovered by S.B. Myers ([M]). See also [K], [Be] and references there for more information.

3.1.1. A cut locus of a point on a closed analytic n-dimensional manifold is a finite CW complex of dimension $\leq n-1$ ([Bu]). (This assertion does not hold if we assume that the manifold is only C^∞, [GS]). In dimension 2 it had been first proven by S.B. Myers ([M]). In fact, Buchner proved that the cut locus is subanalytic and used Hironaka's results on triangulability of subanalytic sets. We refer the reader to a survey [BM] for properties of subanalytic sets.

3.1.2. The cut-locus of every point x of M is a finite tree ([M]).

Proof: Indeed, its homotopy type is the homotopy type of S^2 minus a 2-cell, i.e. of a point. But a finite-dimensional 1-complex contractible to a point is a tree. If it is 0-dimensional, it is a point, which we consider as a degenerate tree.

3.1.3. Assume that τ_1, τ_2 are two distinct minimizing geodesics between x and a point v in the cut locus of x, C_x. Consider one of two open domains, D, bounded by the digon $\tau_1 * \tau_2$. Either $D \bigcap C_x$ is not empty, or $C_x = \{v\}$, and M can be partitioned into minimizing geodesics from x to v.

Proof: Consider the plane tangent to M at x. Consider the angle formed by vectors tangent to M at x pointing inwards D and bounded by the rays generated by the tangent vectors to τ_1 and τ_2 at x. Consider the geodesic rays from x in directions of all vectors in this angle. If $D \bigcap C_x$ is empty, then all these geodesics must be minimizing until they leave D. None of them can intersect τ_1 or τ_2 at points different from x, v, since otherwise τ_1 (or τ_2) will stop being minimizing. Therefore each of these geodesic rays must pass through v. Let $l = dist(x, v)$, $T_l M_x$ denote the set of all tangent vectors at x of length equal to l, and $A_l \subset T_l M_x$ denote the subset of $T_l M_x$ formed by tangent vectors to geodesics from the considered set. (A_l is an arc of the circle $T_l M_x$ formed by all vectors in the angle between tangent vectors to τ_1 and τ_2). Consider the restriction of the exponential map at x on $T_l M_x$. This is an analytic map, which is constant on A_l. (Its value is equal to v.) The analytic continuation principle implies that it is constant on $T_l M_x$, and is equal to v. Thus, every geodesic of length $\leq dist(x, v)$ issued from x must be minimizing, and therefore does not contain any points of C_x other than v. On the other hand, it stops being minimizing after passing through v. Since every point of M can be connected with x by a minimizing geodesic, we see that M is partitioned into these minimizing geodesics from x to v (intersecting only at x and v).

3.1.4. Let v be a point inside an edge of the cut locus of x, C_x. Then the minimizing geodesics from x come to v from both sides of the cut locus. There exists exactly one minimizing geodesic from x to v coming from each side.

Proof: Consider the boundary B_ϵ of ϵ-heighborhood of C_x. Consider its intersection with the 2ϵ-neighborhood of v. For small positive ϵ it will consist of two arcs A_1^ϵ, A_2^ϵ on both sides of C_x. Each point of A_i^ϵ can be connected with x by exactly one minimizing geodesic. As $\epsilon \longrightarrow 0$, these minimizing geodesics converge to a limit set of minimizing geodesics between x and v. All the geodesics in these limit sets are minimizing.

We need to prove that these two limit sets contain exactly one geodesic. Let γ_1, γ_2 be two distinct minimizing geodesics from x to v coming from one side of C_x. They can intersect C_x only at v. C_x is connected. Therefore one of two open digons bounded by γ_1, γ_2 has the empty intersection with C_x. Now we can apply 3.1.3 and conclude that $C_x = \{v\}$, which contradicts our assumption that v is a point on an edge of C_x.

3.1.5. Let v be a vertex of the tree C_x which is not a point. Denote the number of edges of the cut locus meeting at v by k_v. These k_v edges divide a small open neighborhood U of v into k_v connected components that we will denote K_i, $i = 1, \ldots, k_v$. The set of all minimizing geodesics from x to v consists of k_v geodesics. There exists exactly one minimizing geodesic from x to v approaching v from within K_i for every i.
Proof: We can use 3.1.3 to conclude that there is at most one minimizing geodesic from x to v approaching v from within K_i exactly as it has been done in the proof of 3.1.4. We can obtain one minimizing geodesic from x to v approaching v from within K_i proceeding as follows: Take a point v_ϵ in the intersection of the boundary of the ϵ-neighborhood of v with K_i. Consider the minimizing geodesic from x to v_ϵ. As $\epsilon \longrightarrow 0$, a subsequence of the sequence of these minimizing geodesics will converge to a minimizing geodesic from x to v approaching v from within the closure of K_i. But minimizing geodesics from x to v cannot intersect C_x at points other than x and v, and therefore each of these minimizing geodesics must approach v from within of one of the sets K_i.
3.1.6. (Sliding of minimizing geodesics along edges) Let $[v_1, v_2]$ be an edge 3.1.6. Let γ be a minimizing geodesic from x to v_1. Then there exists a continuous family of minimizing geodesics connecting x with all points of $[v_1, v_2]$.
Proof: Consider a very small open neighborhood U of $[v_1, v_2]$ and the connected component K of $U \setminus C_x$ that contains $\gamma \setminus \{v\}$. According to 3.1.4, 3.1.5 for every $v \in [v_1, v_2]$ there exists exactly one minimizing geodesic from x to v that approaches v from within K. We claim that this family of minimizing geodesics continuously depends on v. To see that observe that if this family is not continuous at a point v_*, then there must be at least two distinct minimizing geodesics between x and v_* approaching v_* from K providing a contradiction.

Definition. The minimizing geodesic between x and v_2 obtained as in the proof above will be called the result of a sliding of γ along the edge $[v_1, v_2]$.

3.2. Curve-shortening processes. Let x, y be two points, on M and $\Omega_{x,y}(M)$ denotes the space of paths between x and y. Below we will need a curve-shortening process on $\Omega_{x,y}(M)$. This means that we would like to choose for every path $\rho \in \Omega_{x,y}$ a length-nonincreasing homotopy connecting ρ with a geodesic between x and y providing a local minimum for the length functional on $\Omega_{x,y}(M)$. In fact many such processes are known. For example, one can use the obvious modification of the Birkhoff curve shortening process for $\Omega_{x,y}(M)$ instead of the space closed curves ΛM on M. (We refer a reader to [C] for a detailed description of this process for ΛM. In order to modify it for $\Omega_{x,y}$ we consider only broken geodesics that start at $x_0 = x$ and end at $x_N = y$. The endpoints x and y remain fixed during the whole process.)

Yet another option is a class of flows suggested by J. Hass and P. Scott in [HS]. Let $conv(M)$ denotes the convexity radius of M. Consider a covering of M by discs $D_0, \ldots, D_{n-1}$. We demand that the radii of these discs are less than $conv(M)/2$, that the discs with the same centers of half the radius cover M, the discs meet transversely, and no three boundaries of these discs intersect. Define for every $i > n$ D_i as $D_{i(mod n)}$. Then for every $i = 1, 2, 3, \ldots$, one constructs an obvious homotopy that replaces every arc of the curve inside D_i by the minimal geodesic with the same end points. The desired curve shortening process on the space of all closed curves on M is the composition of all these homotopies. In order to modify this flow to make it work for $\Omega_{x,y}(M)$ instead of the space of all closed curves on M we demand that x, y must be inside some of the discs D_i. Also, when we modify the initial and the last arc of the curve, we find path homotopies connecting these arcs with the minimizing geodesics between x (or, correspondingly, y) and the other endpoint of the arc (on the boundary of the considered disc).

Either of these flows can be used to prove Theorems 1, 1.1 and 2 in the case of an analytic M. We can assume that "a curve shortening process" mentioned below in sections 3.1-3.3 is one of these processes. If M is analytic, we do not have any restrictions on the curve shortening flow at all. In particular, we do not need any kind of continuity.

However, in order to prove our results in the smooth case we need to impose one restriction on the flow. This restriction will be used only in the last section to extend the main results of the paper from the analytic case to the smooth case. Namely, let γ be a piecewise smooth path on S^2 connecting x and y. Consider a surface generated by a path homotopy obtained using a chosen curve shortening flow. Consider this surface as a map from the standard two-dimensional disc into $M = (S^2, g)$. Consider the Lipschitz constant of this map, λ, as a function of the Riemannian metric g on M. We need to define a curve shortening flow for all smooth Riemannian metrics on S^2 so that for every compact set K of Riemannian metrics on S^2 $\sup_{g\in K} \lambda(g) < \infty$. Here we can assume that for every $g \in K$ the length of γ is either $\leq D_0 = 2d + \delta$ (for the purposes of proving Theorem 2.II) or $\leq D_0 = 2d + dist(x, y) + \delta$ (for the purposes of proving Theorem 2.I). Here is the simplest way to do this. Let $D = D_0 + S$, where S is the same as in Theorem 2. Let γ be a piecewise-smooth path in $M = (S^2, g)$ connecting x and y of length $\leq D$. Let $\Omega_{x,y}(M)^D$ denotes the space of all paths of length $\leq D$ between x and y. Consider the connected component Ω_γ of $\Omega_{x,y}(M)^D$ that contains γ. Choose a path p in Ω_γ providing the global minimum of the length functional on Ω_γ. Define a metric on $\Omega_{x,y}(M)$ as follows: The distance between two paths γ_1 and γ_2 will be the infimum over all homotopies between these two paths of the maximal length of a trajectory of a point of γ_1 during this homotopy. Of course, homotopies between some γ_1, γ_2 connecting x and y can be regarded as paths between γ_1 and γ_2 in $\Omega_{x,y}(M)$. Choose an (almost) optimal homotopy between γ and p. That is, we choose a homotopy such that the maximal length of a trajectory is almost equal to the distance between γ and p. This homotopy will be the curve shortening process for γ. It is obvious that the Lipschitz constant of an (optimal) map of $[0, 1] \times [0, 1]$ into M generated by this homotopy can be majorized in terms of the length of γ and the width of the homotopy (=the maximal length of the trajectory of a point of γ during this homotopy). The width of this homotopy can, in turn, be majorized by the diameter of Ω_γ. (The last assertion trivially follows from the definition of the metric on $\Omega_{x,y}(M)$.)

Now we are going consider γ as a fixed path on S^2 between x and y, where S^2 is endowed with Riemannain metrics from a compact set K of Riemannian metrics. The length of γ and D continuously depend on the Riemannian metric and, are therefore bounded on K.

So, everything boils down to proving the existence of a uniform upper bound for the diameter of every connected component of $\Omega_{x,y}(M)^D$ in the considered metric for all Riemannian metrics $g \in K$ on S^2. The existence of such a bound is well-known. Here we sketch how it can be obtained. We refer the reader to a much more detailed exposition in [R0], [NR0]. (In these papers such an estimate was obtained for spaces of closed curves instead of spaces of paths with fixed endpoints, but the argument is completely similar.) The sectional curvature, volume and diameter of (S^2, g) also depend continuously on g. The classical lower bound for the injectivity radius proven by J. Cheeger easily implies that the convexity radius of M is bounded from below in terms of an upper bound for the absolute value of the sectional curvature of (S^2, g), a positive lower bound for the volume and an upper bound for diameter (cf. [Ch]). Therefore the convexity radius of M is uniformly bounded from below on K by a constant $\nu > 0$. Denote the upper bound for the volume of (S^2, g), where $g \in K$ by V, and an upper bound for D by D_K. The classical inequality proven by C. Croke implies that the volume of any ball of radius $\nu/10$ is bounded from below by $const\nu^2$. This gives us a uniform upper bound $M = [V/(const\nu^2)] + 1$ for the number of points in a minimal $\nu/4$-net Net_g on (S^2, g) for all $g \in K$. Consider all broken geodesics from x to y of length $\leq 5D$ such that the length of every segment is $\leq \nu$ and all ends of geodesic segments other than the endpoints x, y are in Net_g. Denote the set of these broken geodesics by B_g. B_g is a ν-net in the set of paths in $\Omega_{x,y}(M)^D$ made of all paths parametrized by the arclength (in the sense of the metric on $\Omega_{x,y}(M)$ that we are considering). The cardinality of B_g can be obviously uniformly bounded on K. Moreover, every path between $\gamma_1, \gamma_2 \in \Omega_{x,y}(M)^D$ can be ν-approximated by a path with short segments connecting points from B_g. Therefore, if γ_1 and γ_2 are in the same connected component of $\Omega_{x,y}(M)^D$ they can be connected by the following homotopy: It goes from γ_1 to a nearest point of B_g, then it goes via short segments connecting pairs of close points of B_g, and at the end it goes from a point of B_g that is ν-close to γ_2 to γ_2. Now note that if this homotopy visits a point of B_g twice (thus, forming a loop), then we can just eliminate this loop shortening the homotopy. So, the homotopy can be chosen so that it visits every point of B_g at most once, and it length is bounded in terms of the product of the cardinality of B_g and ν. These observations provide us with an upper bound for the diameter of every connected component of $\Omega_{x,y}(M)^D$.

Note that although this flow will eventually shorten any path between x and y that was not a geodesic, it can increase the length of the path in the process. (It differs in this aspect from the Birkhoff or Hass-Scott flows.) But the lengths of the intermediate paths will be bounded by D.

3.3. Proof of Theorem 2.

3.3.1. An outline of the proof. Let $f : S^2 \longrightarrow M$ be a diffeomorphism. Endow S^2 with a very fine triangulation, so that the images of all simplices of this triangulation under f are δ-small for a very small positive δ. We are going to attempt to extend f to the disc D^3 triangulated as the cone over the chosen triangulation of S^2. (Such extension is obviously impossible, but in the process we will construct a desired L-slicing of M unless we will be prevented from doing so by the appearance of k distinct geodesics connecting x and y.) We start from mapping the only new vertex (=the center of D^3) into x and all new 1-dimensional simplices (that connect the center of D^3 with points on its boundary, S^2) into (arbitrary) minimizing geodesics between x and the images of the corresponding points on S^2 under f.

Consider now the boundaries of the new 2-simplices. Each of them is formed by two new 1-dimensional simplices mapped into minimizing geodesics between x and some δ-close points $a, b \in M$ and the (very short) minimal geodesic $[a, b]$. We will think of the boundaries of the new 2-dimensional singular simplices in M as being formed by two paths of length $\leq d + \delta/2$ connecting x with the midpoint of the short side, and going along the sides of the triangle from x along two different directions. We will try to fill each of these triangles in M by finding a path homotopy between these two paths.

We will try to find the desired path homotopies using some geometric ideas explained below. The upshot will be that either there exists a path homotopy via paths of controlled length, or we will get the desired k distinct or S-distinct short geodesics between x and y. (This will be the main part of the proof.)

To finish the proof we need to consider only the case when all boundaries of new 2-dimensional simplices can be contracted by means of a path homotopy passing via paths of controlled length.

Since the extension of f to D^3 is impossible, we cannot extend f from the boundary of at least one three-dimensional simplex of the chosen triangulation of D^3. This boundary consists of three new 2-dimensional simplices and a small 2-dimensional simplex coming for the original triangulation of S^2. First, assume that this simplex is mapped into a point. In particular, all new 1-dimensional simplices in the 1-skeleton of the 3-dimensional simplex were mapped to minimal geodesics between x and this point. On the previous steps of the proof we concluded that for each of the three resulting geodesic digons one its side can be connected with the other by a path homotopy passing through paths of controlled length. Now we can combine these three path homotopies in the most obvious way into a L-slicing of M. (Each of these three path homotopies constitutes a third of the L-slicing. The boundary of the considered 3-simplex is mapped into M with a non-zero degree. Here L is an upper bound for the length of the paths in the three path homotopies.)

The situation when the small 2-dimensional 1-simplex is non-trivial can be easily reduced to the situation when it is trivial. We contract the small 2-simplex over itself to a point q. Let x_1, x_2, x_3 denote its vertices. We extend the minimal geodesics from x to x_1, x_2, x_3 by joining them correspondingly with $[x_1, q], [x_2, q], [x_3, q]$ to obtain three paths, p_1, p_2, p_3 connecting x with q. For every vertex x_i denote the midpoint of the opposite side of the triangle $x_1x_2x_3$ by m_i. Assume that for every $i = 1, 2, 3$ there exists a path homotopy connecting two paths between x and m_i along two different sides of the triangle xx_jx_k, where x_j and x_k denote the vertices of the side of the triangle $x_1x_2x_3$ opposite to x_i. (Note that $m_i \in (x_jx_k)$.) Assume that this homotopy goes via paths of length $\leq L$ (for some L). Then is is easy to find a path homotopy between p_i and p_j (for all pairs of i and j) that goes via paths of length $\leq L + o(1)$. (Here is a description of one such homotopy. Denote the midpoint of $[x_ix_j]$ by m_k.

$$p_i = [xx_i] * [x_iq] \longrightarrow [xx_i] * [x_im_k] * [m_kx_i] * [x_iq] \longrightarrow [xx_j] * [x_jm_k] * [m_kx_i] * [x_iq] =$$

$$= [xx_j] * [x_jx_i] * [x_iq] \longrightarrow [xx_j] * [x_jq] = p_j.$$

Here the first path homotopy is the insertion of $[x_im_k]$ traversed twice in the opposite direction, the second path homotopy passes only though paths of length $\leq L$, and the third path homotopy is, in fact, between $[x_jx_i] * [x_iq]$ and $[x_jq]$ via (a part of) the very small triangle $x_1x_2x_3$.

Thus, we demonstrated that proving of Theorem 2 boils down to finding path homotopies between the halves of the boundaries of new 2-dimensional singular simplices in M that pass via paths of controlled length (or obtaining distinct minimizing geodesics between x and y as obstructions if the desired path homotopies do not exist).

We will need the following easy lemma that can be used to reduce finding path homotopies between two paths γ_1, γ_2 connecting the same endpoints to contracting the loop $\gamma_1 * \gamma_2^{-1}$ formed by these paths:

Lemma 5. Consider two paths τ_1, τ_2 from x to y. Assume that the loop $\tau_1 * \tau_2^{-1}$ based at x can be contracted to x via loops based at x of length $\leq K$. Then there exists a path homotopy between τ_1 and τ_2 that passes only through paths of length $\leq \text{length}(\tau_2) + K$.

Proof: We start from τ_2 and use the contraction of $\tau_1 * \tau_2^{-1}$ in reverse time to create $\tau_1 * \tau_2^{-1}$ out of the point x. As the result we obtain $\tau_1 * \tau_2^{-1} * \tau_2$. Then we contract $\tau_2^{-1} * \tau_2$ over itself. QED.

Now let us describe how we construct path homotopies between halves of the boundaries of the new 2-dimensional simplices. Denote two vertices of a considered 2-simplex by x_1 and x_2; the third vertex is, of course, x. One of two paths in question is the join of the minimal geodesic from x to x_1 and the part of $[x_1x_2]$ between x_1 and the midpoint m of $[x_1x_2]$, the other is the join of $[xx_2]$ and $[x_2m]$. Denote these two paths by π_1, π_2. The intersection of $\pi_1 \bigcup \pi_2$ with the cut locus C_x of x is entirely contained in the short simplex $[x_1x_2]$. By perturbing the short geodesic segment between x_1 and x_2 as a smooth path with fixed endpoints we can ensure that the number of its points of intersection with C_x is finite. (Here we use the subanaliticity of the cut-locus, cf. [B], [BM]). Denote these points by $P_1, P_2, \ldots P_L$ (see Fig. 3). These points divide $\pi_1 \bigcup \pi_2$ into segments $[P_iP_{i+1}]$, where $P_0 = P_{L+1} = x$.

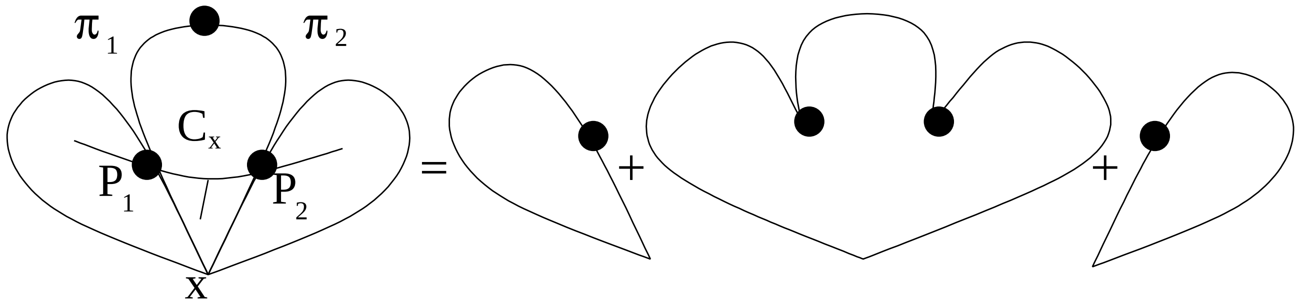

Figure 3.

For every i consider the set of all minimizing geodesics connecting x with P_i. It is easy to prove using induction with respect to the number L of points of intersection that one can find a path homotopy between π_1 and π_2 as the composition of path homotopies 1) between pairs of these minimizing geodesics connecting x with P_i and cobounding a domain that does not contain other minimizing geodesics between x and P_i; 2) between one side of a triangle xP_iP_{i+1} and the join of two other sides for the following triangles xP_iP_{i+1}: The sides xP_i and xP_{i+1} are some of the considered geodesics, P_iP_{i+1} is a part of $\pi_1 * \pi_2^{-1}$, and the intersection of the closure of the triangle xP_iP_{i+1} with C_x is the set $\{P_i, P_{i+1}\}$. Finding path homotopies of the second type does not constitute any problem: Since the intersection of the cut-locus of x with the triangle consists only of its two vertices, one can easily find such a path homotopy using a continuous family of minimizing geodesics emanating from x to all points of the side P_iP_{i+1} of the triangle. (Here we are using 3.1.6 to obtain this continuous family of geodesics emanating from x.) The desired path homotopy passes through the joins of the minimizing geodesics from x to $P_t \in [P_iP_{i+1}]$ and $[P_tP_{i+1}] \subset [P_iP_{i+1}]$. The length of paths in such path homotopies will clearly be less than or equal to $3d$.

Thus, everything boils down to finding path homotopies between pairs of minimizing geodesics between x and P_i.

Our next observation is that one can reduce finding controlled path homotopies between pairs of minimizing geodesics between x and P_i to finding controlled contractions of loops formed by these pairs. Indeed, Lemma 5 implies that if a loop formed by a pair of these minimizing geodesics can be contracted via loops of length $\leq L_*$ then there is a path homotopy between these minimizing geodesics passing through paths of length $\leq L_* + d$. Combining these observation together we arrive at the following proposition:

Proposition A. Let $\tilde{L}$ be an upper bound for the length of loops in optimal path homotopies contracting loops formed by all pairs of minimizing geodesics between x and P_i (for all i. Of course, here we consider only pairs of minimizing geodesics with the same endpoints.) Then the length of paths in a path homotopy connecting π_1 to π_2 does not exceed $\max\{\tilde{L} + d, 3d\}$.

Combining this proposition with the previous discussion we obtain:

Proposition B. Assume that for every $z \in C_x$ and every two minimizing geodesics γ_1, γ_2 between x and z there exists a path homotopy contracting the loop $\gamma_1 * \gamma_2^{-1}$ that passes through loops of length $\leq \tilde{L}$, where $\tilde{L} \geq 2d$. Then there exists an $(\tilde{L} + d)$-slicing of M.

Thus, everything boils down to finding an upper bound L for the maximal length of loops in optimal path homotopies contracting loops formed by pairs of minimizing geodesics connecting x with a point on the cut locus C_x of x. Below we will see how one can get such an estimate if there exists less than k short geodesics between x and y. Our estimates will be $L = (k^2 - 3k + 6)d$, if there exist less than k distinct geodesics of length $\leq 2d$ (see Corollary 7 below), or $L = (2k - 2)d + 2dist(x, y)$, if there exist less than k distinct geodesics of length $\leq 2d + 2dist(x, y)$ (Corollary 9). Let S be a non-negative number. Then one can also obtain similar upper bounds, if one assumes the existence of at most $k - 1$ S-distinct geodesics (instead of assuming the existence of $k - 1$ distinct geodesics.) In the first case our estimate becomes $L = (k^2 - 3k + 6)d + S$ (if there exist less than k S-distinct geodesics of length $\leq 2d$), in the second case it becomes $L = (2k - 2)d + 2dist(x, y) + S$ (if there exist less than k distinct S-geodesics of length $\leq 2d + 2dist(x, y)$). Combining this estimates with Proposition B, we immediately complete the proof of Theorem 2 (in the considered now analytic case). Thus, it remains to prove the above mentioned Corollaries 7 and 9. This will be done in section 3.3.2 below (and is the crucial part of the proof of all main results of the present paper).

3.3.2. Path homotopies contracting loops formed by two minimizing geodesics.

3.3.2.1 Let γ_1, γ_2 be two minimizing geodesics from x to a point $p \in C_x$. We would like to find a bound $L \geq d$ such that there exists a path homotopy contracting $\gamma_1 * \gamma_2^{-1}$ passing through loops based at x of length $\leq L$. We are going to assume either that there exist at most $k - 1$ geodesics between x and y of length $\leq 2d$ (or $2d + dist(x, y)$) or at most $k - 1$ S-distinct geodesics between x and y of length $\leq 2d$ (or $\leq 2d + dist(x, y)$), where S is a non-negative number.

3.3.2.2. Obstructing pairs. Let us first try to find the desired path homotopy as follows. Connect p with y by a minimizing geodesic σ. Apply a curve shortening process in the space of paths starting at x and ending at y to $\gamma_1 * \sigma$ and $\gamma_2 * \sigma$. We end at geodesics between x and y of length $\leq 2d$ providing local minimuma for the length functional. Denote these geodesics by ω_1, ω_2. If $\omega_1 = \omega_2$ then we can combine these two homotopies into a path homotopy

$$\gamma_1 * \gamma_2^{-1} \longrightarrow \gamma_1 * \sigma * \sigma^{-1} * \gamma_2^{-1} \longrightarrow \omega_1 * \sigma^{-1} * \gamma_2^{-1} \longrightarrow \omega_1 * \omega_2^{-1} = \omega_1 * \omega_1^{-1} \longrightarrow \{x\},$$

and hence obtain a path homotopy contracting $\gamma_1 * \gamma_2^{-1}$ via loops based at x of length $\leq 4d$. The difficult case is when $\omega_1 \neq \omega_2$. In this case we will call the (unordered) pair of geodesics ω_1, ω_2 *an obstructing pair.* If ω_1 and ω_2 are S-distinct, we will call them *an S-obstructing pair.* If they are not S-distinct, then we can contract $\omega_1 * \omega_2^{-1}$ via loops of length $S + 4d$, thereby obtaining a homotopy that contracts $\gamma_1 * \gamma_2^{-1}$ via loops of length $\leq S + 4d$.

3.3.2.3. Obstructing geodesics. Here is another approach to contracting $\gamma_1 * \gamma_2^{-1}$. Fix one of the minimizing geodesics between x and y and denote it γ. Consider the following path between x and y : $\gamma_1 * \gamma_2^{-1} * \gamma$. If we apply a curve-shortening process to this path we will end at a geodesic τ of length $\leq 2d + dist(x, y)$ between x and y providing a local minimum for the length functional on the space of paths between x and y. If this geodesic is not γ, we will call it an *obstructing geodesic* corresponding to the pair γ_1, γ_2. If S is a non-negative number, and this geodesic is S-distinct from γ, we call it an S-obstructing geodesic. Note that $\gamma_1 * \gamma_2^{-1}$ can be connected by a path homotopy first with $\gamma_1 * \gamma_2^{-1} * \gamma * \gamma^{-1}$, and then with $\tau * \gamma^{-1}$. If $\tau = \gamma$, we can contract $\tau * \gamma^{-1}$ over itself, obtaining a path homotopy contracting $\gamma_1 * \gamma_2^{-1}$ via loops of length $\leq 2d + 2dist(x, y)$. If τ and γ are not S-distinct, then $\tau * \gamma^{-1}$ can be contracted via loops of length $\leq 2d + 2dist(x, y) + S$. Therefore, in the last case, $\gamma_1 * \gamma_2^{-1}$ can be contracted via loops of length $\leq 2d + 2dist(x, y) + S$.

3.3.2.4. These considerations provide the base of inductive proofs of the following two lemmae (Lemma 6, Lemma 8):

Lemma 6. Let γ_1 and γ_2 be minimizing geodesics between x and a point $p \in M$, D be a domain in M bounded by $\gamma_1 \bigcup \gamma_2$, and S be a non-negative number. Denote the union of the set of all vertices of C_x inside D and the one point set $\{p\}$ by Q. For every $q \in Q$ and every pair of minimizing geodesics between x and q consider the corresponding obstructing pair (respectively, S-obstructing pair), if it exists. Let N be the cardinality of the resulting set of obstructing pairs (respectively, S-obstructing pairs). Then there exists a path homotopy contracting loop $\gamma_1 * \gamma_2^{-1}$ via loops of length $\leq (2N + 4)d$ (respectively, $(2N + 4)d + S$).

Corollary 7. Let γ_1 and γ_2 be minimizing geodesics between x and a point $p \in M$. Assume that there exists at most $k-1$ geodesics (respectively, S-distinct geodesics) between x and y of length $\leq 2d$ providing local minima for the length functional on the space of curves connecting x and y. Then the loop $\gamma_1 * \gamma_2^{-1}$ can be contracted via loops of length $\leq (k^2 - 3k + 6)d$ (respectively, $(k^2 - 3k + 6)d + S$).

Indeed, it is obvious that N in Lemma 8 is majorized by by the number $(k-1)(k-2)/2$ of all possible pairs of geodesics between x and y that provide a local minimum of the length functional on $\Omega_{x,y}(M)$.

Proof of Lemma 6: The proof will be by induction with respect to N. The base of induction $N = 0$ was provided by an argument that followed the definition of obstructing and S-obstructing pairs in 3.3.2.2. To prove the induction step assume that the lemma is true for $N - 1$. We will prove the induction step only in the case, when N is the number of the obstructing pairs. The proof in the case, when N measures the number of S-obstructing pairs is nearly identical.

The main idea is to use the following obvious observation: If $\gamma_1 * \gamma_2^{-1}$ and another geodesic loop λ have the same obstructing pair ω_1, ω_2, then $\gamma_1 * \gamma_2^{-1}$ can be connected with either λ or λ^{-1} by a path homotopy passing through $\omega_1 * \omega_2^{-1}$ and involving only loops of length $\leq \max\{length(\lambda), length(\gamma_1 * \gamma_2)\}$. This observation will enable us to reduce finding of a path homotopy contracting $\gamma_1 * \gamma_2^{-1}$ to finding a path homotopy contracting the loop λ formed by the innermost in D pair of minimizing geodesics between x and a vertex of C_x that has the same obstructing pair as γ_1, γ_2. Then, we will reduce finding a path homotopy contracting λ to finding path homotopies between pairs of minimizing geodesics (with the same endpoints) that are even deeper inside D. Therefore the obstructing pair ω_1, ω_2 cannot appear anymore, and the induction assumption will apply. Here are the details.

First, we are going to introduce a partial order on the set of all pairs of minimizing geodesics connecting x with points of Q. For brevity we will call such pairs MGD (an abbreviation of "minimal geodesic digons"). Note that no pair of such MGD can intersect at points other than their common endpoints. (Otherwise, the intersecting geodesics will stop being minimizing.) In particular, they can intersect $\gamma_1 \bigcup \gamma_2$ only at x and, possibly at p (if p is an endpoint of the considered MGD). Therefore each of the considered MGD bounds a unique domain contained in D. We say that one such MGD is *less* than the other if the domain inside D bounded by the first MGD is properly contained in the domain bounded by the second. It is clear that γ_1, γ_2 is the maximal element of this (finite) poset. Denote the resulting poset $P(\gamma_1, \gamma_2, D)$.

Note that N is the cardinality of the set of obstructing pairs of all MGDs from $P(\gamma_1, \gamma_2, D)$. If γ_1, γ_2 do not have an obstructing pair, then the assertion of Lemma 6 immediately follows from 3.3.2.2. So, we can assume that γ_1, γ_2 have an obstructing pair. Among all considered MGD that have the same obstructing pair as γ_1, γ_2 choose an MGD μ that has the following property: Consider the set of MGD that are less than μ . None of these MGD has the same obstructing pair as γ_1, γ_2. (In other words, μ is a minimal element in the set of all MGDs in $P(\gamma_1, \gamma_2, D)$ that have the same obstructing pair as γ_1, γ_2.) The existence of μ follows from the finiteness of $P(\gamma_1, \gamma_2, D)$. μ consists of two minimizing geodesics with common endpoints. Denote these minimizing geodesics by δ_1 and δ_2. Let D_1 be the domain bounded by δ_1, δ_2 and contained in D. Denote the common endpoint of δ_1, δ_2 different from x by p_*.

As we already observed, $\gamma_1 * \gamma_2^{-1}$ can be connected with $\delta_1 * \delta_2^{-1}$ via loops of length $\leq 4d$. Therefore it is sufficient to contract $\delta_1 * \delta_2^{-1}$.

Consider the intersection of C_x with D_1. There are three cases:

Case 1: This intersection is empty. Then according to 3.1 there is a length non-increasing path homotopy from δ_1 to δ_2, that can be used to find a length non-increasing contraction of $\delta_1 \bigcup \delta_2$.

Case 2. There is exactly one edge of C_x approaching p_* from inside of D_1. We can slide δ_1, δ_2 along this edge to the nearest vertex (see 3.1.6). This sliding will provide us with a homotopy of the loop $\delta_1 * \delta_2^{-1}$ via loops of length $\leq 2d$. At the end of this process we will obtain a MGD that bounds a proper subdomain D_2 of D_1. We can define a set Q_2 and a number N_2 for D_2 exactly as Q and N were defined for D in the text of Lemma 6. The definition of μ implies that $N_2 < N$. Now we can apply the induction assumption to

contract the loop formed by this new MGD that bounds D_2.
Case 3: There exist more than one edge of C_x approaching p_* from D_1. In this case there is a sequence of minimizing geodesics between x and p_* $\beta_1 = \delta_1, \beta_2, \ldots, \beta_{q-1}, \beta_q = \delta_2$, $q > 2$ such that for every i β_i and β_{i+1} bound a proper subdomain D_i^* of D_1 such that no geodesic β_j passes inside D_i^*. One can construct a path homotopy between δ_1 and δ_2 by joining path homotopies between β_i and β_{i+1} for all i. Because of the definition of μ every pair β_i, β_{i+1} satisfies conditions of Lemma 6 with some number $N_i < N$ of potential obstructing pairs instead of N. Therefore, the induction assumption implies that loops $\beta_i * \beta_{i+1}^{-1}$ can be contracted via loops of length $\leq 2(N-1)+4 = 2N+2$. Therefore Lemma 5 implies the existence of path homotopies between β_i and β_{i+1} that pass through paths of length $\leq 2N+3$. These path homotopies together form a path homotopy $\delta_t, t \in [1,2]$ between δ_1 and δ_2. We can contract $\delta_1 * \delta_2^{-1}$ by, first, going through loops $\delta_t * \delta_2^{-1}$, and then cancelling $\delta_2 * \delta_2^{-1}$ along itself. QED.
We can prove analogues of Lemma 6 and Corollary 7, where the notion of obstructing pair is replaced by the notion of obstructing geodesic.

Lemma 8. Let γ_1 and γ_2 be distinct minimizing geodesics between x and a point $p \in M$, D be a domain in M bounded by $\gamma_1 \bigcup \gamma_2$, and S be a non-negative number. Let Q be the set of points formed by p and all vertices of C_x inside D. For every vertex $q \in Q$ and every pair of minimizing geodesics between x and q consider the corresponding obstructing geodesic (respectively, S-obstructing geodesic), if it exists. Let N be the cardinality of the resulting set of obstructing geodesics (respectively, S-obstructing geodesics). Then there exists a path homotopy contracting $\gamma_1 * \gamma_2^{-1}$ and passing only through loops of length $\leq (2N+2)d + 2dist(x,y)$ (respectively, $\leq (2N+2)d + 2dist(x,y) + S$).

Corollary 9. Let γ_1 and γ_2 be minimizing geodesics between x and a point $p \in M$. Assume that there exists at most $k-1$ geodesics (respectively, S-obstructing geodesics) between x and y of length $\leq 2d + dist(x,y)$ providing local minima for the length functional on the space of curves connecting x and y. Then $\gamma_1 * \gamma_2^{-1}$ can be connected by a path homotopy via loops of length $\leq (2k-2)d + 2dist(x,y) \leq 2kd$ (respectively, $\leq (2k-2)d + 2dist(x,y) + S \leq 2kd + S$).

Indeed, if there exist at most $k-1$ distinct geodesics between x and y, then there exist at most $(k-1)-1 = k-2$ distinct obstructing geodesics. So, N in Lemma 8 cannot exceed $k-2$.

Proof of Lemma 8. The proof of Lemma 8 uses induction with respect to N. It is completely parallel to the proof of Lemma 6. The difference is that the notion of obstructing geodesic replaces the notion of obstructing pair. Correspondingly, the construction of a path homotopy contracting $\gamma_1 * \gamma_2^{-1}$, when it has no obstructing geodesic (correspondingly, S-obstructing geodesic) replaces a similar argument for obstructing pairs. (Recall that this argument was given right after the definition of obstructing and S-obstructing geodesics in 3.3.2.3.) QED.

3.4. The smooth case. In the previous sections we proved Theorem 2 (and thereby the rest of the results of this paper) in the case when M is a real analytic Riemannian manifold. Here we demonstrate that the analiticity is not necessary.

First, approximate a given smooth Riemannian metric g on S^2 by a converging sequence of analytic Riemannian metrics g_n, $\lim_{n\longrightarrow\infty} g_n = g$ in C^3-topology on the space of smooth Riemannian metrics on S^2. Denote (S^2, g_n) by M_n and (S^2, g) is M.

We are going to use the curve shortening flow introduced at the end of section 3.2 in order to construct path homotopies between a path from x and y and a geodesic between x and y.

Let $i = 1$ or 2. The already proven analytic case of Theorem 2 implies that for every n either A_i or B_i hold for M_n. Consider two cases. In the first case A_i holds not only for the considered value, S_0, of S but also for every value of S in the interval $[S_0, S_0 + \zeta]$ for some small positive ζ and for all but finitely many g_n. In the second case there exists a decreasing sequence $S_m < S_0 + \zeta$ converging to S and a strictly increasing sequence n_m such that B_i holds for S_n and M_{n_m}. We would like to prove that in the first case A_i holds for M, and in the second case B_i holds for M thereby establishing Theorem 2 for M.

Case 1. Passing to a subsequence and changing the notations we can assume that A_i holds for all Riemannian metrics g_n for $S = S_0+\zeta$. We claim that it holds also for M for $S = S_0$. Otherwise, there will be less than k S_0-distinct geodesics between x and y. Note that an easy compactness argument implies that for every $\epsilon > 0$ for all sufficiently large n every geodesic of length $\leq 2d + dist(x, y)$ (correspondingly, $2d$) between x and y on M_n will be ϵ-close to a geodesic of length $\leq 2d + dist(x, y)$ (correspondingly, $\leq 2d$) between x and y on M. Therefore for all sufficiently large n there exist two $(S_0 + \zeta)$-distinct geodesics on M_n that are ϵ-close to geodesics between x and y on M that are not S_0-distinct. Choosing ϵ sufficiently small we ensure that these geodesics on M_n are not $(S_0 + \zeta)$-distinct thereby obtaining the desired contradiction.

In the second case we would like to obtain an L-slicing of M for the postulated value of L as the limit of $(L+\epsilon_m)$-slicings that exist for M_{n_m} (possibly passing to a subsequence of n_m). The only thing that we need in order to ensure the existence of such a limit is a uniform bound for the Lipschitz constant for all slicings of M_{n_m} (possibly after a suitable reparametrization of S^2 that is being mapped to M_{n_m}). Since these slicings were constructed by gluing several path homotopies between paths between x and y and geodesics between x and y we need a uniform bound for Lipschitz constant for these path homotopies. These homotopies were just applications of a chosen curve shortening process. We arrive to the question that had been posed and resolved in 3.2: We demonstrated that one can choose a curve shortening process so that such a bound always exists (for all paths of length $\leq 2d + dist(x, y) + S + \delta$ or $\leq 2d + S + \delta$ and all Riemannian metrics from a compact set K. In our situation $K = \{g\} \bigcup \{g_1, g_2, \ldots\}$.) So, for the chosen curve shortening processes a limit L-slicing of M exists. QED.

4. Concluding remarks.

4.1. One can find path homotopies between pairs of minimizing geodesics connecting x and P_i (see 3.2) not using a path homotopy contracting the geodesic loops formed by such pairs. An alternative approach enabled us to improve the Dichotomy Theorem (Theorem 2) for $k = 2$. We obtained the following result:

Theorem. Let M be a Riemannain manifold diffeomorphic to S^2 of diameter d. Let $S \geq 0$. Let x, y be two points of M such that there exists only one geodesic between x and

y of length $\leq 2d$ (respectively, every two geodesics between x and y of length $\leq 2d$ can be connected by a path homotopy that increases the length by a summand not exceeding S). Then there exists a $3d$-slicing (respectively, a $(3d+S)$-slicing) of M that maps the South pole of S^2 into x.

We omit the details of the proof of this improvement.

4.2. As it had been noted the main purpose of chosing a curve shortening process as it had been done in 3.2 was to extend the proof of main results from the analytic case to the smooth case. Yet, such a choice of the curve shortening process can be used also to prove some strenghenings of Theorem 2. Indeed, note that the curve shortening process introduced in 3.2 applied to a closed curve can end only at a geodesic between x and y that is a global minimum of the length on a connected component of $\Omega_{x,y}(M)^{2d+S+\delta}$ (or, correspondingly, $\Omega_{x,y}(M)^{2d+dist(x,y)+S+\delta}$.) (Here $\Omega_{x,y}(M)^R$ denotes the space of all paths of length $\leq R$ between x and y.) Therefore, we can strengthen A_1, A_2 by demanding that the k S-distinct geodesics minimize the length in different connected components of $\Omega_{x,y}(M)^{2d+dist(x,y)+S}$ (or, correspondingly, $\Omega_{x,y}(M)^{2d+S}$). As a corollary we obtain the following dichotomy theorem:

Theorem 2.A. (Dichotomy theorem II.) Let M be a Riemannian manifold diffeomorphic to S^2, d be the diameter of M and x and y be two arbitrary points of M.
A. For every $k>1$ either the space of all paths between x and y of length $\leq 2d$ has at least k connected components, or M admits a $(k^2-3k+7)d$-slicing that maps the South pole of S^2 into x.
B. For every $k>1$ either the space of all paths between x and y of length $\leq 2d+dist(x,y)$ has at least k connected components or M admits a $(2k-1)d+2dist(x,y)$-slicing that maps the South pole of S^2 into x.

Interpreting the non-existence of a L-slicing that maps the South pole of S^2 into x (for a controlled L) as a roughness of the Morse landscape of the distance from x function, we can interprete this theorem as a manifestation (for S^2) of the following empirical principle:

Principle. If the Morse landscape of the distance function from a point on a closed Riemannian manifold is rough, then the Morse landscape of the length functional on the space of loops based at this point is rough. The same will be true for all spaces of paths with fixed endpoints starting at this point.

We expect that this principle has manifestations that are true for all closed Riemannian manifolds. It seems very plausible that various modifications of this principle involving spaces of closed curves, cycles etc. are also true.

Acknowledgments: Both authors would like to acknowledge a partial support of this research by their NSF and NSERC grants. Regina Rotman acknowledges a partial support of this research by the NSERC University Faculty award. The authors are grateful to the Max-Planck-Institut für Mathematik, where this paper had been partially written for its kind hospitality during their visit in July-August, 2006. A.N. would like to thank Shmuel Weinberger for useful discussions of the topology of the space of loops on S^2.

References.

[Be] M. Berger, “A panoramic view of Riemannain geometry”, Springer-Verlag, 2003.
[BM] E. Bierstone, P. Milman, “Semianalytic and subanalytic sets”, Publ. IHES, 67(1988), 5-42.
[B] R. Bott, “Lectures on Morse theory, old and new”, Bull. of the AMS, 7(1982), 331-358.
[Bu] M. Buchner, “Simplicial structure of the real analytic cut locus”, Proc. of the AMS, 64(1977), 118-121,
[Ch] I. Chavel, “Riemannian Geometry - a modern introduction”, Cambridge University Press, 1993.
[C] C. Croke, “Area and the length of the shortest closed geodesic”, J. Diff. Geom. 27(1988), 1-21.
[FHT] Y. Felix, S. Halperin, J.-C. Thomas, “Rational homotopy theory”, Springer, 2001.
[FK] S. Frankel, M. Katz, “The Morse landscape of a Riemannian disc”, Annales de l’Institut Fourier, 43(2)(1993),503-507.
[GS] H. Gluck, D. Singer, “Scattering of geodesic fields I, II”, Ann. Math. 108(1978), 347-372 and 110(1979), 205-225.
[G] M. Grayson, “Shortening embedded curves”, Ann. Math., 129(1989), 71-111.
[HS] J. Hass, P. Scott, “Shortening curves on surfaces”, Topology, 33(1994), 25-43.
[Kl] W. Klingenberg, “Lectures on closed geodesics”, Springer-Verlag, 1978.
[K] S. Kobayashi, “On conjugate and cut loci”, in “Global Differential Geometry”, ed. by S.S. Chern, MAA Studies in Mathematics 27, 96-122, Mathematical Association of America, Washington, DC, 1989.
[Ko] N. Koiso, “Convergence to a geodesic”, Osaka J. Math., 30(1993), 559-565.
[M] S. B. Myers, “Connections between differential geometry and topology. I. Simply-connected surfaces”, Duke Math. J. 1(1935), 376-391.
[NR0] A. Nabutovsky, R. Rotman, “Upper bounds for the length of a shortest closed geodesic and quantitaive Hurewicz theorem”, J. of Eur. Math. Soc. (JEMS), 5(2003), 203-244.
[NR1] A. Nabutovsky, R. Rotman, “Lengths of geodesics between two points on a Riemannian manifold”, preprint, available at http://www.math.toronto.edu/alex/
[NR2] A. Nabutovsky, R. Rotman, “The length of the second shortest geodesic”, preprint, available at http://www.math.toronto.edu/alex/
[R0] R. Rotman, “Upper bounds on the length of the shortest closed geodesic on a simply connected manifold”, Math. Z., 233(2000), 365-398.
[R] R. Rotman, “The length of a shortest geodesic loop at a point”, preprint, available at http://comet.lehman.cuny.edu/sormani/others/rotman.html
[S] A.S. Schwarz, “Geodesic arcs on Riemannian manifolds”, Uspekhi Math. Nauk (translated from Russian as “Russian Math. Surveys”), 13(6)(1958), 181-184.
[Se] J.P. Serre, “Homologie singulière des espaces fibrés. Applications”, Ann. Math. 54(1951), 425-505.