Algebraic & Geometric Topology 6 (2006) 459–492

Regular homotopy and total curvature I: circle immersions into surfaces

TOBIAS EKHOLM

We consider properties of the total absolute geodesic curvature functional on circle immersions into a Riemann surface. In particular, we study its behavior under regular homotopies, its infima in regular homotopy classes, and the homotopy types of spaces of its local minima.

53C42; 53A04, 57R42

1 Introduction

An *immersion* of manifolds is a map with everywhere injective differential. Two immersions are *regularly homotopic* if there exists a continuous 1–parameter family of immersions connecting one to the other. The Smale–Hirsch h–principle [8; 4] says that the space of immersions $M \to N$, $\dim(M) < \dim(N)$ is homotopy equivalent to the space of injective bundle maps $TM \to TN$. In contrast to differential topological properties, differential geometric properties of immersions do not in general satisfy h–principles, see [3, (A) on page 62]. In this paper and the sequel [2], we study some aspects of the differential geometry of immersions and regular homotopies in the most basic cases of codimension one immersions. We investigate whether or not it is possible to perform topological constructions while keeping control of certain geometric quantities.

Let Σ be a Riemann surface, ie, an orientable 2–manifold with a Riemannian metric, and let $c\colon S^1 \to \Sigma$ be an immersion of the circle parameterized by arc length. If $\vec{c}\colon S^1 \to U\Sigma$, where $U\Sigma$ is the unit tangent bundle of Σ, denotes the natural lift of c, then the h–principle mentioned above implies that the map $c \mapsto \vec{c}$ induces a weak homotopy equivalence between the space of circle immersions into Σ and the space of continuous circle maps into $U\Sigma$. In particular, regular homotopy classes of circle immersions into Σ are in one to one correspondence with the homotopy classes of (free) loops in $U\Sigma$.

Published: 23 March 2006 DOI: 10.2140/agt.2006.6.459

The *total absolute geodesic curvature* κ of a circle immersion c into a Riemann surface is given by the integral

$$\kappa(c) = \int_c |k_g|\, ds,$$

where k_g is the geodesic curvature of c, and where ds denotes the arc length element along c. We study properties of the functional κ on the space of circle immersions, starting with the following question. *If c_0 and c_1 are regularly homotopic circle immersions into a Riemann surface, what is the infimum, over all regular homotopies c_t, $0 \le t \le 1$, connecting c_0 to c_1, of $\max_{0\le t\le 1} \kappa(c_t)$?*

Theorem 1.1 answers this question for the simplest Riemann surfaces of constant curvature. We use the following notational conventions: All Riemann surfaces are assumed to be complete unless otherwise explicitly stated. If Σ is a Riemann surface then $K\colon \Sigma \to \mathbb{R}$ denotes its Gaussian curvature function. For topological spaces X and Y, we write $X \approx Y$ to indicate that X is homeomorphic to Y.

Theorem 1.1 *Let Σ be a Riemann surface and let $c_0, c_1\colon S^1 \to \Sigma$ be regularly homotopic.*

(a) *If Σ has constant curvature $K = 0$ (the completeness assumption then implies $\Sigma \approx \mathbb{R}^2$, $\Sigma \approx S^1 \times \mathbb{R}$, or $\Sigma \approx T^2$), or if $\Sigma \approx \mathbb{R}^2$ and has constant curvature $K < 0$, then there exists a regular homotopy c_t, $0 \le t \le 1$, connecting c_0 to c_1 with*

$$\kappa(c_t) \le \max\{\kappa(c_0), \kappa(c_1)\}, \quad 0 \le t \le 1.$$

(b) *If Σ is the 2–sphere with a constant curvature metric (with $K > 0$) then, for any $\epsilon > 0$, there exists a regular homotopy c_t, $0 \le t \le 1$, connecting c_0 to c_1 such that*

$$\kappa(c_t) \le \max\{\kappa(c_0), \kappa(c_1), 2\pi + \epsilon\}, \quad 0 \le t \le 1.$$

Moreover, if $c_0\colon S^1 \to \Sigma$ runs m times around a geodesic and $c_1\colon S^1 \to \Sigma$ runs $m+2$ times around a geodesic then any regular homotopy c_t, $0 \le t \le 1$, connecting c_0 to c_1 has an instant c_τ, $0 < \tau < 1$, with

$$\kappa(c_\tau) > 2\pi.$$

Theorem 1.1 is proved in Section 5.3. The proof of (b) uses Arnol'd's J^-–invariant for immersed curves on the sphere, see Arnol'd [1], Inshakov [5] and Tchernov [9]. In Remark 5.1 we present a metric on $\mathbb{R}^2$ with $K \le 0$ for which the conclusion in (a) does not hold.

The proof of Theorem 1.1 also gives information about infima of κ. To state these results we first introduce some notation. If Σ is a flat Riemann surface then parallel translation gives a trivialization of $U\Sigma$ and the free homotopy classes of curves in $U\Sigma$ are in natural one to one correspondence with $\pi_1(\Sigma)\times\mathbb{Z}$, where $\pi_1(\Sigma)$ encodes the homotopy class of a circle immersion and $\mathbb{Z}$ its tangential degree. Thus, if $\Sigma\approx\mathbb{R}^2$, then we denote a regular homotopy class of circle immersions by the integer m which equals the tangential degree of any of its representatives, and, similarly, if $\Sigma\approx S^1\times\mathbb{R}$ or $\Sigma\approx T^2$, then we denote a regular homotopy class by $(\xi,m)\in\pi_1(\Sigma)\times\mathbb{Z}$, where ξ and m is the homotopy class in Σ and the tangential degree, respectively, of any of its representatives. If Σ is the 2–sphere then there are exactly two regular homotopy classes: one represented by a simple closed curve, the other by such a curve traversed twice. Finally, if α is a regular homotopy class of curves in a Riemann surface then let $\hat{\kappa}(\alpha)=\inf_{c\in\alpha}\kappa(c)$.

Theorem 1.2

(a) *Let Σ be a Riemann surface with $K(p)<0$ for all $p\in\Sigma$ and assume that either Σ is closed or $\Sigma\approx\mathbb{R}^2$. Then the infimum $\hat{\kappa}(\alpha)$ is attained at some curve in the regular homotopy class α if and only if α is representable by (a multiple of) a closed geodesic. Moreover, if $K(p)=K<0$ is constant and $\Sigma\approx\mathbb{R}^2$ then $\hat{\kappa}(m)$, $m\in\mathbb{Z}$, satisfies*

$$\hat{\kappa}(m)=\begin{cases}2\pi & \text{for } m=0,\\ \pi(|m|+1) & \text{for } m\neq 0.\end{cases}$$

(b) *Let Σ be a Riemann surface of constant curvature $K=0$ (the completeness assumption then implies $\Sigma\approx\mathbb{R}^2$, $\Sigma\approx S^1\times\mathbb{R}$, or $\Sigma\approx T^2$). Then the infimum $\hat{\kappa}((\xi,m))$ is attained at some curve in the regular homotopy class $(\xi,m)\in\pi_1(\Sigma)\times\mathbb{Z}$ if and only if $\xi\neq *$ or $m\neq 0$, where $*$ denotes the homotopy class of the constant loop. Moreover, $\hat{\kappa}((\xi,m))$, $(\xi,m)\in\pi_1(\Sigma)\times\mathbb{Z}$, satisfies*

$$\hat{\kappa}((\xi,m))=\begin{cases}2\pi & \text{for } (\xi,m)=(*,0),\\ 2\pi(|m|) & \text{otherwise.}\end{cases}$$

(c) *Let Σ be the 2–sphere with any metric and let α be a regular homotopy class of circle immersions into Σ. Then the infimum $\hat{\kappa}(\alpha)$ equals 0 and is attained at some curve in α.*

Theorem 1.2 is proved in Section 5.1.

A curve in a Riemann surface Σ with $K(p)\neq 0$ for all $p\in\Sigma$, which is a local minimum of κ is in fact a closed geodesic, see Proposition 2.2 (a). For flat Riemann surfaces this is not the case. Here any local minimum of κ is a locally convex curve, see Proposition 2.2 (b). We say that a curve c is *locally convex* if $k_g\geq 0$ everywhere for some orientation of c. If $k_g > 0$ everywhere, we say that c is *strictly locally convex*. In the terminology of Gromov [3, page 8], strictly locally convex curves are called free curves.

The following result describes the homotopy types of the spaces of local minima of κ for a flat Riemann surface. (Here we think of circle immersions as oriented unit speed curves parameterized by arc length.)

Theorem 1.3 *On a flat Riemann surface Σ (the completeness assumption implies $\Sigma \approx \mathbb{R}^2$, $\Sigma \approx S^1 \times \mathbb{R}$, or $\Sigma \approx T^2$), the space $\Omega_{(\xi,m)}$ of (strictly) locally convex curves of regular homotopy class $(\xi, m) \in \pi_1(\Sigma) \times \mathbb{Z}$ satisfies*

$$\Omega_{(\xi,m)} \simeq \begin{cases} \varnothing & \text{if } (\xi, m) = (*, 0), \\ \Sigma & \text{if } \xi \neq * \text{ and } m = 0, \\ U\Sigma & \text{if } m \neq 0, \end{cases}$$

where $\simeq$ denotes weak homotopy equivalence.

Theorem 1.3 is proved in Section 5.2.

Acknowledgements

During the preparation of this paper, the author was a research fellow of the Royal Swedish Academy of Sciences sponsored by the Knut and Alice Wallenberg foundation. He also acknowledges support from the Alfred P Sloan Foundation and from NSF–grant DMS-0505076. The author thanks B Pausader for reading and commenting on earlier versions of the paper.

2 First variation and local minima

In this section we compute the first variation of total absolute geodesic curvature. We use the result to classify local minima. Since the absolute value function is not differentiable at zero, the first variation is expressed as a statement about differences rather than as a statement about derivatives.

2.1 First variation of total absolute geodesic curvature

Let $c\colon [0, L] \to \Sigma$ be an immersion into a Riemann surface, parameterized by arc length. Let e_1 be the unit tangent vector field of c and let e_2 be a unit vector field along c everywhere orthogonal to e_1. We consider variations $\omega\colon [0, L] \times (-\delta, \delta) \to \Sigma$

of c with the following three properties: $\omega(s,0) = c(s)$, $\partial_\epsilon\omega(s,0) = \alpha(s)e_2(s)$, where ∂_ϵ denotes differentiation with respect to the second variable, for some function $\alpha\colon [0,L] \to \mathbb{R}$, and the curves $\omega_\epsilon(s) = \omega(s,\epsilon)$ (ϵ fixed) are immersions for $\epsilon \in (-\delta,\delta)$. We also introduce the sign function $\sigma\colon \mathbb{R} \to \mathbb{R}$ as follows

$$\sigma(x) = \begin{cases} 1 & \text{for } x > 0, \\ 0 & \text{for } x = 0, \\ -1 & \text{for } x < 0. \end{cases}$$

Lemma 2.1 *For $\epsilon \in (-\delta,\delta)$,*

$$\begin{aligned}\kappa(\omega_\epsilon) - \kappa(c) =& \epsilon\left(\int_{\{k_g \neq 0\}} \sigma(k_g)\,(\ddot\alpha + \alpha K)\;ds\right) \\ &+ |\epsilon|\left(\int_{\{k_g=0\}} |\ddot\alpha + \alpha K|\,ds\right) + \mathcal{O}(\epsilon^2),\end{aligned}$$

where $\ddot\alpha = \frac{d^2\alpha}{ds^2}$, and where $\mathcal{O}(\epsilon^2)$ denotes a function such that $\epsilon^a\mathcal{O}(\epsilon^2) \to 0$ as $\epsilon \to 0$ for all $a > -2$.

Proof To simplify notation, let $\frac{\partial\omega}{\partial s} = \dot\omega$ and $\frac{\partial\omega}{\partial\epsilon} = \omega'$. Let ∇ denote the Levi–Civita connection and let $\nabla_s = \nabla_{\dot\omega}$ and $\nabla_\epsilon = \nabla_{\omega'}$. If $d\tau$ denotes the arc length element of the curve ω_ϵ and $k_g(s,\epsilon)$ denotes the geodesic curvature of ω_ϵ at s then

$$|k_g|\,d\tau = \left|k_g(s,\epsilon)\right|\,|\dot\omega|\;ds = \frac{\left|\langle\nabla_s\,\dot\omega, \iota\,\dot\omega\rangle\right|}{|\dot\omega|^2}\,ds,$$

where ι denotes rotation by $\frac{\pi}{2}$. Assuming $k_g(s,0) \neq 0$ and remembering that c is parameterized by arc length, we compute

$$\begin{aligned}\partial_\epsilon\big(|k_g(s,0)||\dot\omega|\big) =& \Big(\partial_\epsilon|\dot\omega|^{-2}\Big)|k_g| + \sigma(k_g)\Big(\partial_\epsilon\,\langle\nabla_s\,\dot\omega, \iota\,\dot\omega\rangle\Big) \\ =& -2\langle\nabla_\epsilon\,\dot\omega, \dot\omega\rangle|k_g| + \sigma(k_g)\Big(\langle\nabla_\epsilon\nabla_s\,\dot\omega, \iota\,\dot\omega\rangle + \langle\nabla_s\,\dot\omega, \iota\nabla_\epsilon\,\dot\omega\rangle\Big) \\ =& -2\langle\nabla_s\,\omega', \dot\omega\rangle|k_g| \\ &+ \sigma(k_g)\Big(\big\langle\nabla_s\nabla_s\,\omega', \iota\,\dot\omega\big\rangle + \big\langle R(\omega',\dot\omega)\dot\omega, \iota\,\dot\omega\big\rangle + \big\langle\nabla_s\,\dot\omega, \iota\nabla_s\,\omega'\big\rangle\Big),\end{aligned}$$

where R is the curvature tensor. Noting that $\omega' = \alpha e_2$, $\nabla_s e_1 = k_g e_2$, and $\nabla_s e_2 = -k_g e_1$, we conclude

$$\begin{aligned} \partial_\epsilon |k_g(s,0)| &= 2\alpha k_g |k_g| + \sigma(k_g)\Big(\ddot\alpha - \alpha k_g^2 + K\alpha - \alpha k_g^2\Big) \\ &= \sigma(k_g)\Big(\ddot\alpha + K\alpha\Big). \end{aligned} \tag{2–1}$$

A similar calculation at s where $k_g(s,0) = 0$ gives

$$\partial_\epsilon\big(k_g(s,0)|\dot\omega|\big) = \ddot\alpha + K\alpha.$$

Hence for such s,

$$|k_g(s,\epsilon)||\dot\omega| = |\epsilon||\ddot\alpha + K\alpha| + \mathcal{O}(\epsilon^2). \tag{2–2}$$

The result follows by integration of (2–1) and (2–2). □

2.2 Local minima of κ

Proposition 2.2 *Let Σ be a Riemann surface.*

(a) *If $K(p) \neq 0$ for all $p \in \Sigma$, then an immersion $c\colon S^1 \to \Sigma$ is a local minimum of κ if and only if it is a geodesic.*

(b) *If Σ is flat then an immersion $c\colon S^1 \to \Sigma$ is a local minimum of κ if and only if it is a locally convex curve.*

In particular, in both cases (a) *and* (b)*, any local minimum of κ is a global minimum in its regular homotopy class.*

Proof Consider case (a). A curve $c\colon S^1 \to \Sigma$ is a geodesic if and only if $k_g(s) = 0$ for all $s \in S^1$ and geodesics are global minima of κ. Let $c\colon S^1 \to \Sigma$ be a local minimum of κ. Note that $U = \{s \in S^1 \colon k_g(s) \neq 0\}$ is open. Assume U is nonempty. Then there exists a nonempty open subinterval $J \subset U$. For any variation $\omega(s,\epsilon)$ of c with $\partial_\epsilon\omega = \alpha e_2$ where $\alpha\colon S^1 \to \mathbb{R}$ is supported in J we have

$$\int_J \sigma(k_g)\ddot\alpha\, ds = 0.$$

Thus, since c is a local minimum, we conclude from Lemma 2.1 that

$$\int_J \sigma(k_g)\alpha K\, ds = 0.$$

This contradicts $K(p) \neq 0$ for all $p \in \Sigma$. It follows that U is empty and thus c is a geodesic.

Consider case (b). Let $c\colon S^1 \to \Sigma$ be a local minimum. We show that k_g cannot change sign along c. Assume it does, then there exist two disjoint open subintervals J_+ and J_- of S^1 such that $k_g > 0$ on J_+ and $k_g < 0$ on J_-. Let A be a subinterval of S^1 containing both J_+ and J_-. Let $\alpha\colon A \to \mathbb{R}$ be a function such that $\ddot{\alpha}$ is supported in small subintervals of $J_+ \cup J_-$ and such that $\dot{\alpha} = r \neq 0$, where r is a non-zero constant, between J_+ and J_-. For a variation ω of c with $\partial_\epsilon \omega = \alpha e_2$, Lemma 2.1 implies that

$$\kappa(\omega_\epsilon) - \kappa(c) = \pm 2r\epsilon + \mathcal{O}(\epsilon^2).$$

This contradicts c being a local minimum. Consequently, k_g does not change sign along c and c is locally convex.

It remains to show that c is a global minimum. Fix a unit speed parametrization of c so that $k_g(s) \geq 0$ for all $s \in S^1$. As in Section 1, we construct an orthonormal trivialization of $T\Sigma$ by parallel translation with respect to the flat metric. This identifies the unit tangent bundle $U\Sigma$ of Σ with $\Sigma \times S^1$ and the regular homotopy class of c is determined by its homotopy class in Σ and the degree of $\pi_2 \circ \vec{c}\colon S^1 \to S^1$, where $\pi_2\colon \Sigma \times S^1 \to S^1$. (Recall that $\vec{c}\colon S^1 \to U\Sigma$ denotes the natural lift of the unit speed curve c.) Moreover, $\kappa(c)$ is simply the length of the curve $\pi_2 \circ \vec{c}$. Now, $k_g(s) \geq 0$ for all $s \in S^1$ implies that the length of $\pi_2 \circ \vec{c}$ equals 2π times the degree of $\pi_2 \circ \vec{c}$ and it follows that local minima are global minima also in this case. □

Remark 2.3 Proposition 2.2 does not hold for arbitrary Riemann surfaces. Consider for example the boundary of a convex body in $\mathbb{R}^3$ which agrees with the standard 2–sphere except that it has a flat region near the north pole. Any locally convex curve in this flat region is a local minimum of κ but it is certainly not a global minimum in its regular homotopy class.

3 Curvature concentrations and approximations

In this section we define piecewise geodesic curves with curvature concentrations and show that circle immersions can be approximated by such curves without increasing the total absolute geodesic curvature.

3.1 Piecewise geodesic curves with curvature concentrations

Let Σ be a Riemann surface. A *piecewise geodesic curve* in Σ is a continuous curve $c\colon S^1 \to \Sigma$ which is a finite union of geodesic segments. More formally such a curve c can be described as follows. Consider a finite collection of geodesics $c_j\colon [0,1] \to \Sigma$, $j = 1, \dots, m$, with $c_j(1) = c_{j+1}(0)$ for each j (here $c_{m+1} = c_1$). Let I_j be $[0,1]$

thought of as the domain of c_j and let $I_{m+1} = I_1$. Then the space obtained by identifying $1 \in I_j$ with $0 \in I_{j+1}$ is a circle S^1 which can be considered as the domain of a continuous map $c\colon S^1 \to \Sigma$ such that if $p \in S^1$ is the image of $p' \in I_j$ under the quotient projection then $c(p) = c_j(p')$. We say that the points $p \in S^1$ with two preimages under the quotient projection are the *vertices* of the piecewise geodesic curve c. We will often deal with images of vertices and we call also these image points vertices of c, when no confusion can arise.

Let $U\Sigma$ denote the unit tangent bundle of Σ. Note that at each vertex $c(p)$ of a piecewise geodesic curve c as above, there is an *incoming* unit tangent vector $\dot c_j(1)/|\dot c_j(1)| \in U_{c(p)}\Sigma$ and an *outgoing* unit tangent vector $\dot c_{j+1}(0)/|\dot c_{j+1}(0)| \in U_{c(p)}\Sigma$. A piecewise geodesic curve *with curvature concentrations* is a piecewise geodesic curve c together with a *vertex curve* $\gamma\colon [0,1] \to U_{c(p)}\Sigma$ for each vertex p which connects the incoming– to the outgoing unit tangent of c at p and which satisfies the following condition: γ is a continuous piecewise geodesic curve (with finitely many geodesic arcs) in the fiber circle $U_{c(p)}\Sigma$ equipped with the metric induced by the Riemannian metric on Σ. We use the abbreviations PGC–curve to denote piecewise geodesic curves with curvature concentrations, and we often write $(c_1, \gamma_1, \dots, c_m, \gamma_m)$ for a PGC–curve with geodesic segments c_j and vertex curves γ_j. We say that the length $l(\gamma_j)$ of the vertex curve γ_j is the *curvature concentration* of the PGC–curve c at the vertex $c(p) = c_j(1) = c_{j+1}(0)$, and that the piecewise geodesic curve with geodesic arcs $c_1, \dots, c_m$ is the *underlying curve* of c.

We note that any PGC–curve $c = (c_1, \gamma_1, \dots, c_m, \gamma_m)$ in Σ has a natural continuous lift $\vec c\colon S^1 \to U\Sigma$ which consists of the usual lifts $\vec c_j$ of c_j, $j = 1, \dots, m$, connected by the curves γ_j, $j = 1, \dots, m$, in the fibers of $U\Sigma \to \Sigma$ over vertices of c. The lift $\vec c$ of c is thus a piecewise smooth curve. In particular, its derivative is smooth except for finitely many jump discontinuities where the curve has left and right derivatives. If c is a PGC–curve we consider $\vec c\colon S^1 \to U\Sigma$ as a parameterized curve with its natural arc length parametrization scaled by a suitable factor so that its domain becomes the unit circle.

Definition 3.1 The *total absolute geodesic curvature* of a PGC–curve

$$c = (c_1, \gamma_1, \dots, c_m, \gamma_m)$$

is
$$\kappa(c) = \sum_{j=1}^{m} l(\gamma_j),$$

where $l(\gamma_j)$ is the length of the vertex curve γ_j.

Let $\mathrm{Pgc}(S^1, \Sigma)$ denote the set of all PGC–curves in Σ. We define the distance between two elements b and c in $\mathrm{Pgc}(S^1, \Sigma)$ to be the C^0–distance (with respect to the metric on $U\Sigma$ induced from the metric on Σ) between their lifts $\vec{b}$ and $\vec{c}$ endowed with parameterizations proportional to arc length, as discussed above. A PGC*–homotopy* is a continuous 1–parameter family of PGC–curves or equivalently a continuous map from the interval to $\mathrm{Pgc}(S^1, \Sigma)$.

3.2 Approximation

If Σ is a Riemann surface then let $\mathrm{Imm}(S^1, \Sigma)$ denote the space of circle immersions into Σ with the C^2–topology. Let $c\colon S^1 \to \Sigma$ be a circle immersion, let $\pi = (p_0, \dots, p_m)$ be a partition of S^1 and let $|\pi| = \max_j d(p_j, p_{j+1})$ (d is the distance function on S^1 and we use the convention $p_{m+1} = p_0$). If $|\pi|$ is sufficiently small then we associate a PGC–curve c^π to c, as follows: c^π is the PGC–curve with underlying piecewise geodesic curve consisting of the shortest geodesic segments between $c(p_j)$ and $c(p_{j+1})$, and with vertex curves γ_j which are the shortest arcs in $U_{c(p_j)}\Sigma$ connecting the incoming– to the outgoing unit tangent of the underlying piecewise geodesic curve. We note that $\max_j l(\gamma_j) \to 0$ and $|\kappa(c) - \kappa(c^\pi)| \to 0$ as $|\pi| \to 0$.

If $f\colon \Lambda \to \mathrm{Imm}(S^1, \Sigma)$ is a continuous family of circle immersions parameterized by a compact space Λ and if $\epsilon > 0$ is arbitrary, then there exists $\delta > 0$ such that for all partitions π with $|\pi| < \delta$, $f^\pi\colon \Lambda \to \mathrm{Pgc}(S^1, \Sigma)$, defined by $f^\pi(\lambda) = (f(\lambda))^\pi$, is a continuous family of PGC–curves, and there exists an ϵ–small homotopy connecting the family of continuous curves $\vec{f}\colon \Lambda \to U\Sigma$ to the family of continuous curves $\vec{f}^\pi\colon \Lambda \to U\Sigma$.

Lemma 3.2 *Let $c_0\colon S^1 \to \Sigma$ be a circle immersion into a Riemann surface with $K < 0$, $K > 0$, or $K = 0$ everywhere. Then there exists a regular homotopy c_t, $0 \le t \le 1$, of c_0 such that*

$$\kappa(c_t) \le \kappa(c_0), \quad 0 < t \le 1, \tag{3–1}$$

and such that the PGC*–curve c_1^π (defined using any sufficiently fine partition π), satisfies*

$$\kappa(c_1^\pi) \le \kappa(c_0). \tag{3–2}$$

Moreover, if c is not a local minimum of κ then the non-strict inequalities in (3–1) and (3–2) can be replaced by strict inequalities.

Proof Assume $K \neq 0$ everywhere. Then Proposition 2.2 implies that c_0 is a local minimum if and only if c_0 is a geodesic. Hence, if c_0 is not a geodesic then there exists a κ–decreasing regular homotopy connecting c_0 to some curve c_1. For sufficiently fine partition π, the curve c_1^π then satisfies (3–2).

Assume $K = 0$. If c is not locally convex then the above argument can be repeated. Recall from Section 1 that parallel translation in the flat metric gives $U\Sigma = \Sigma \times S^1$ and let $\pi_2\colon U\Sigma \to S^1$ denote the projection. If c is locally convex it is elementary to see that for sufficiently fine partitions π, $\pi_2 \circ \vec{c^\pi}$ is monotone and thus $\kappa(c^\pi) = \kappa(c)$. $\square$

Remark 3.3 Lemma 3.2 does not hold for general metrics. Consider for example $\mathbb{R}^2$ with coordinates (x, y) and a metric given by

$$ds^2 = \exp\Big(2\left(\sqrt{x^2+y^2}-1\right)^3\Big)(dx^2+dy^2)$$

in a small neighborhood of $c = \{x^2+y^2 = 1\}$. Then $K(x,y) < 0$ for $x^2+y^2 > 1$ and $K(x,y) > 0$ for $x^2+y^2 < 1$ and the geodesic curvature of c is identically equal to 1. Let b be any curve which is a C^1–small perturbation of c. Assume that b meets c transversely in $2m$ points. These intersection points subdivide b and c into unions of arcs $c = \gamma_{\rm out} \cup \gamma_{\rm in}$ and $b = \beta_{\rm out} \cup \beta_{\rm in}$, where $\beta_{\rm out} \subset \{x^2+y^2 \geq 1\}$ and $\beta_{\rm in} \subset \{x^2+y^2 \leq 1\}$ and where the endpoints of an arc in $\gamma_{\rm in}$ ($\gamma_{\rm out}$) agree with the endpoints of some arc in $\beta_{\rm out}$ ($\beta_{\rm in}$). Let $\Omega^j_{\rm out}$, $j = 1, \dots, m$, be the m regions bounded by an arc in $\beta_{\rm out}$ and an arc in $\gamma_{\rm in}$ and let $\Omega^j_{\rm in}$, $j = 1, \dots, m$ be the m region bounded by an arc in $\gamma_{\rm out}$ and an arc in $\beta_{\rm in}$. Then the Gauss–Bonnet theorem implies that

$$\int_{\beta_{\rm out}} k_g\, ds - \int_{\gamma_{\rm in}} k_g\, ds + \sum_{j=1}^{m} \int_{\Omega^j_{\rm out}} K\, dA + \sum_{j=1}^{2m} \alpha_j = 2\pi m,$$

$$\int_{\gamma_{\rm out}} k_g\, ds - \int_{\beta_{\rm in}} k_g\, ds + \sum_{j=1}^{m} \int_{\Omega^j_{\rm in}} K\, dA + \sum_{j=1}^{2m} \alpha_j = 2\pi m,$$

where α_j is the exterior angle at the $j^{\rm th}$ intersection point between b and c. Thus

$$\kappa(b) \geq \int_b k_g\, ds > \int_c k_g\, ds = \kappa(c).$$

By approximation we conclude that for any PGC–curve e which lies in a sufficiently small tubular neighborhood of c, $\kappa(e) > \kappa(c)$.

4 Curvature non-increasing homotopies, smoothing, and locally convex curves

In this section we construct special PGC–homotopies of PGC–curves on flat Riemann surfaces and on the hyperbolic plane which decrease the total absolute geodesic curvature of a given initial curve and which ends at a curve of certain standard shape. We construct similar special PGC–homotopies of curves on the 2–sphere with a constant curvature metric. We show that these special PGC–homotopies can be smoothed to regular homotopies, increasing the total curvature arbitrarily little. We also study the space of locally convex curves on flat surfaces.

4.1 Flat surfaces

Let Σ be a Riemann surface. Let $\Pi\colon \widetilde{\Sigma} \to \Sigma$ be a smooth covering map and endow $\widetilde{\Sigma}$ with the pull-back metric. Define a *lift* of a PGC–curve c in Σ to be a PGC–curve b in $\widetilde{\Sigma}$ such that $\vec{b}$ is a lift of $\vec{c}$ with respect to the induced covering $\Pi_U\colon U\widetilde{\Sigma} \to U\Sigma$.

Let c be a PGC–curve in Σ with lift b in $\widetilde{\Sigma}$ and let b_t, $0 \le t \le 1$, be a PGC–homotopy of b with the following properties: the start-point $b_t(\alpha)$ and endpoint $b_t(\omega)$ of b_t satisfy $\Pi(b_t(\alpha)) = \Pi(b_t(\omega))$ and $\Pi_U(\dot{b}_t(\alpha)) = \Pi_U(\dot{b}_t(\omega))$, for all t, where $\dot{b}_t(\alpha)$ and $\dot{b}_t(\omega)$ are the outgoing– and incoming tangent vectors of b_t at $b_t(\alpha)$ and $b_t(\omega)$, respectively. Then b_t induces a PGC–homotopy c_t of c by transporting geodesic segments of b_t in $\widetilde{\Sigma}$ to geodesic segments of c_t in Σ with the projection $\Pi\colon \widetilde{\Sigma} \to \Sigma$, and by transporting vertex curves of b_t to vertex curves of c_t with the induced projection $\Pi_U\colon U\widetilde{\Sigma} \to U\Sigma$. (Note that the conditions on the PGC–homotopy b_t holds if $\vec{b}_t$ is a closed curve for each t.)

Lemma 4.1 *Let Σ be a Riemann surface with constant curvature $K = 0$ and let c be a PGC–curve in Σ. Then there exists a PGC–homotopy c_t, $0 \le t \le 1$, with $c_0 = c$, with $\kappa(c_t) \le \kappa(c_0)$, for all t, and with the following property: the underlying curve of c_1 is a geodesic and any curvature concentration of c_1 equals π.*

Proof Consider the case $\Sigma = \mathbb{R}^2$. We claim that any PGC–curve which does not have underlying curve a line segment admits a PGC–homotopy which does not increase κ and which decreases the number of vertices with curvature concentration not a multiple of π. Together with an obvious inductive argument this shows that any PGC–curve is PGC–homotopic through an homotopy with properties as above to a curve with all curvature concentrations integral multiples of π. Noting that any curvature concentration at a vertex of magnitude $m\pi$, m a positive integer, can be split up into

m vertices each with curvature concentration π by a PGC–homotopy preserving κ we conclude that the claim implies the lemma when $\Sigma = \mathbb{R}^2$.

Consider the claim. If the underlying curve of a PGC–curve is not a line segment then it has three consecutive line segments connected by two vertices with curvature concentrations which are not integral multiples of π. At such a vertex the curvature concentration is either the sum of the exterior angle of the underlying curve and a multiple of 2π or the sum of the interior angle of the curve and an odd multiple of π. We call the former type of vertices *exterior* and the latter *interior*. We also need to distinguish two types of configurations of the three consecutive segments. We say that the configuration is *convex* if all three segments are contained in the closure of one of the half planes determined by the line containing the middle segment, otherwise we say it is *non-convex*. To establish the claim, consider three consecutive segments e_1, e_2, e_3 as above connected at vertices v_1 and v_2. We separate the cases:

Case 1 If both v_1 and v_2 are exterior, then move v_2 along e_3 until it reaches the next vertex following it. It is straightforward to check that this PGC–homotopy preserves κ in the convex case and decreases κ in the non-convex case.

Case 2 If v_1 is exterior and v_2 interior, then move v_1 backwards along e_1 until it reaches the vertex preceding it. This PGC–homotopy preserves κ in the non-convex case and decreases it in the convex case.

Case 3 Assume that both v_1 and v_2 are interior. If the configuration is convex, then move v_2 along e_3 until it reaches the vertex following it. This is a κ–preserving PGC–homotopy. If the configuration is non-convex, consider the lines l_j containing e_j. Assume first that l_1 and l_3 are not parallel. Note that the segments e_1 and e_3 lie in different components of $\mathbb{R}^2 - l_2$. If the point $l_1 \cap l_3$ lies in the component of $\mathbb{R}^2 - l_2$ which contains e_3 then move v_1 along l_1 to $l_1 \cap l_3$. Otherwise, move v_2 along l_3 to the intersection point. This PGC–homotopy strictly decreases κ. In the case that l_1 and l_3 are parallel, start by moving v_1 as described above. Note that this decreases κ. Thus we may change e_3 slightly without increasing κ past $\kappa(c_0)$ so that l_1 and l_3 intersect. We then apply the above.

Note that in either case, the PGC–homotopy described reduces the number of vertices with curvature concentration not an integral multiple of π. The claim follows.

Consider the non-simply connected case. Since Σ is a flat Riemann surface, there is a covering map $\Pi\colon \mathbb{R}^2 \to \Sigma$ with deck transformations which are translations of $\mathbb{R}^2$. Let b be a lift of c. If b is closed, we apply the above result to construct a PGC–homotopy b_t, $0 \le t \le 1$, with properties as in the formulation of the lemma. The induced PGC–homotopy c_t, $0 \le t \le 1$, of c then satisfies the lemma. If b is non-closed we lift it in the middle of some segment and apply the construction above to the lift. If at some stage of the PGC–homotopy a moving vertex passes the endpoint of the lifted curve, we stop the PGC–homotopy and choose the midpoint of the next segment of the curve in the direction the vertex is moving and then proceed. Note that this construction applies as long as the underlying curve has at least three distinct geodesic segments. Thus, to complete the proof it remains only to study curves with two line segments and two vertices.

Let b be the lift of a curve c with two vertices. Let v_0 and v_1 be the endpoints of the lift b and let l be the straight line segment in $\mathbb{R}^2$ connecting v_0 to v_1. Let e_0 and e_1 be the straight line segments of b with one endpoint at v_0 and at v_1, respectively. Let e be the segment of b connecting the endpoint of e_0 which is not equal to v_0 to the endpoint of e_1 not equal to v_1. Since we lift c at a midpoint of a geodesic segment it follows that e passes the midpoint p of l.

Let w_0 and w_1 be the vertices where e_0 and e meet and where e_1 and e meet, respectively. Choose an orienting basis $(\widehat{e}, \widehat{f})$ of the plane where $\widehat{e}$ is a unit vector in direction of e oriented from v_0 to v_1 and where $\widehat{f}$ is a vector perpendicular to $\widehat{e}$ oriented so that the direction vector of e_0 has positive $\widehat{f}$ component. If $\Delta(r,s,t)$ is a triangle with corners at $r,s,t \in \mathbb{R}^2$ we write $\alpha(r;r,s,t)$ and $\beta(r;r,s,t)$ for the exterior– respectively interior angle of $\Delta(r,s,t)$ at the corner r.

We consider separate cases. Consider first the case when both w_0 and w_1 are exterior vertices. Rotate the line containing e in the positive directions around p, and rotate the lines containing e_0 and e_1 in the negative direction around v_0 and v_1, respectively, so that angles change at linear speed. More precisely if w_0^t and w_1^t denotes the intersection points between the rotated line containing e and the rotated line containing e_0 and e_1, respectively. Then $\beta(p;p,w_j^t,v_j) = (1-t)\beta(p;p,w_j,v_j)$, $j=0,1$. Note that the segments in the triangles $\Delta(p,w_j^t,v_j)$, $j=0,1$, which are not parallel to l give a PGC–homotopy b_t of b which via the projection Π induces a PGC–homotopy c_t of c. Moreover,

$$\frac{d}{dt}\Big(\kappa(c)-\kappa(c_t)\Big) = 2\big(\beta(p;p,v_0,w_0)+\beta(v_0;p,v_0,w_0)\big) > 0,$$

and the curvature concentrations of c_1 are integral multiples of 2π.

Consider second the case when both w_0 and w_1 are interior angles. In this case we repeat the construction above with the important difference that we rotate also the lines containing e_0 and e_1 in the positive direction. Again the projection gives a PGC–homotopy c_t of c and we have

$$\frac{d}{dt}\Big(\kappa(c_t)-\kappa(c)\Big) = 2\big(-\alpha(v_0;p,v_0,w_0)+\beta(p;p,v_0,w_0)\big),$$

but $\alpha(v_0; p, v_0, w_0) = \beta(p; p, v_0, w_0) + \beta(w_0; p, v_0, w_0)$ and hence $\kappa(c_t) < \kappa(c)$. Again c_1 has the desired form.

Finally, if one of w_0 and w_1 is an exterior vertex and the other one is an interior vertex, then any one of the above procedures may be used. The result is a PGC–homotopy which does not change κ: the sum of the interior angle and the exterior angle at w_0 and w_1 is constantly equal to π. This finishes the proof in the non-simply connected case. □

Lemma 4.2 *If c_0 and c_1 are two PGC–homotopic PGC–curves on a Riemann surface as in Lemma 4.1 then there exists a PGC–homotopy c_t, $0 \le t \le 1$, connecting them such that*

$$\kappa(c_t) \le \max\{\kappa(c_0), \kappa(c_1)\},$$

for all t.

Proof After Lemma 4.1 it is sufficient to consider the case when c_0 and c_1 both have underlying curves geodesics and all curvature concentrations equal to π. Any vertex curve is thus either a positive or a negative π–rotation. It is easy to see that neighboring vertex curves of different orientations cancel. The lemma follows. □

4.2 The hyperbolic plane

The counterpart of Lemma 4.1 for curves in the hyperbolic plane differs from the flat case in an essential way: the limit curve which arises as the end of a κ–decreasing PGC–homotopy, is not a PGC–curve *in* the hyperbolic plane, in fact it often has infinite length. To deal with this phenomenon we define a generalized PGC–curve in the hyperbolic plane as a PGC–curve which is allowed to have vertices at infinity. More concretely, consider the disk model of the hyperbolic plane,

$$D = \{x = (x_1, x_2) \in \mathbb{R}^2 \colon |x| < 1\}, \quad ds^2 = \frac{4(dx_1^2 + dx_2^2)}{(1 - |x|^2)^2}, \tag{4–1}$$

and add to it the circle at infinity ∂D. Define a *generalized* PGC*–curve* as a piecewise smooth curve in $\bar{D} = D \cup \partial D$ which consists of geodesic segments (ie, arcs of circles meeting ∂D at right angles), which is allowed to have vertices on ∂D, and which have vertex curves at all vertices connecting the incoming– to the outgoing unit tangent. Note that the length of any vertex curve at a vertex on ∂D is an odd multiple of π. We extend κ to generalized PGC–curves by defining it as the sum of the lengths of all vertex curves (the sum of all curvature concentrations).

To connect generalized PGC–curves to PGC–curves we measure the distance between generalized PGC–curves as the C^0–distance between their lifts in $U\bar{D}$ with respect to the metric on $U\bar{D}$ induced by the Euclidean metric on the plane. Using this metric we define a generalized PGC–homotopy as a continuous 1–parameter family of generalized PGC–curves. Moreover, it is clear that the following approximation result holds: if b_λ, $\lambda \in \Lambda$ is any continuous family of generalized PGC–curves parameterized by a compact space Λ and if $\epsilon > 0$ is arbitrary then there exists a family of PGC–curves c_λ, $\lambda \in \Lambda$, such that the distance between b_λ and c_λ is less than ϵ and such that $|\kappa(b_\lambda) - \kappa(c_\lambda)| < \epsilon$ for every $\lambda \in \Lambda$.

Lemma 4.3 *Let $\Sigma \approx \mathbb{R}^2$ be a Riemann surface of constant curvature $K < 0$ and let c be a PGC–curve in Σ. Then there exists a generalized PGC–homotopy c_t, $0 \le t \le 1$, with $c_0 = c$, with $\kappa(c_t)$ a non-increasing function of t, and with the following property: c_1 is a generalized PGC–curve with all its vertices on the circle at infinity, and with all curvature concentrations equal to π.*

Proof Note that κ is invariant under scaling. Therefore we may assume that $K = -1$. Then $\Sigma = D$, where D is as in (4–1). Let c be a PGC–curve in Σ. Consider a vertex curve γ of c at the vertex $p \in D$. Let e_1 be the incoming geodesic arc of c at p. We consider two cases separately.

Case 1 Assume that $l(\gamma) = a + n2\pi$ where $0 < a < \pi$ and $n \ge 0$. In this case the vertex curve is, modulo 2π–rotations, the exterior angle of the curve. We push the endpoint of e_1 backwards along e_1 until we reach its start-point. At this moment we have reduced the number of non-infinite vertices of c with curvature concentrations not a multiple of π by one. To see that this generalized PGC–homotopy does not increase κ we calculate with notation as in Figure 1,

$$\alpha' + \beta' = \alpha + \int_\Omega K\, dA \le \alpha,$$

and hence κ does not increase.

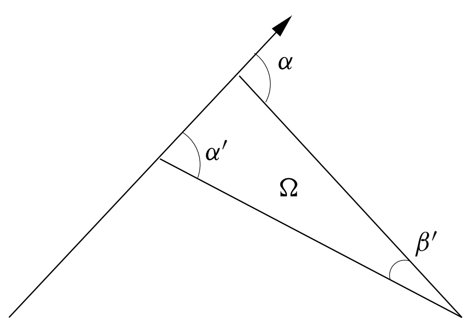

Figure 1: Removing an exterior angle

Case 2 Assume that $l(\gamma) = a + n2\pi$ where $\pi \le a < 2\pi$ and $n \ge 0$. In this case the vertex curve is, modulo 2π–rotations, the complementary angle to the exterior angle. We push the vertex in the positive direction along e_1 until it hits ∂D, thereby reducing the number of non-infinite vertices by one. To see that this PGC–homotopy does not increase κ we calculate with notation as in Figure 2,

$$\alpha' + \beta' = \alpha + \int_\Omega K\, dA \le \alpha,$$

and hence κ does not increase.

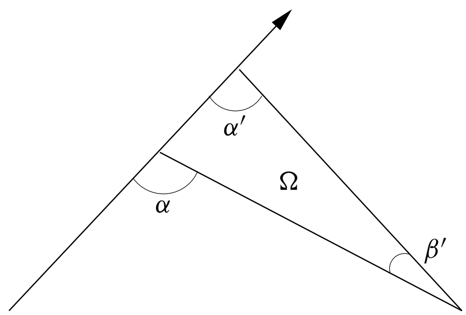

Figure 2: Removing an interior angle

Repeating this argument a finite number of times we remove all non-infinite vertices of c. To finish the proof we note that any vertex curve at infinity has length $\pi + 2m\pi$, for some integer $m \ge 0$, and that the $2\pi m$–concentration can be pushed along one of the geodesics which ends at the infinite vertex so that it lies in D (not in ∂D). Finally, such a curvature concentration of magnitude $2\pi m$ in the finite part of the disk can be split up and pushed to $2m$ curvature concentrations at infinity, each of length π. In Figure 3 this is illustrated for $m = 1$. (If the multiplicity of the geodesic on the left hand picture in Figure 3 is k, then the multiplicity between the two curvature concentrations on the right hand side is $k + 2$.) $\square$

Our next goal is to deform any generalized PGC–curve with its vertices at infinity to a standard form. To this end, we distinguish two different vertices at infinity: let x be a vertex on ∂D of a generalized PGC–curve with incoming geodesic segment along the geodesic c_i and outgoing geodesic segment along the geodesic c_o. Note that c_i subdivides D into two components. Let D_+ be the component with inward normal ν along c_i such that if $n(x)$ is the outward normal of ∂D at x then $n(x), \nu(x)$ is a positively oriented basis of $\mathbb{R}^2$. Let D_- be the other component. If $c_i = c_o$ then we say x is a *degenerate vertex*. Assume that x is a non-degenerate vertex and that the vertex

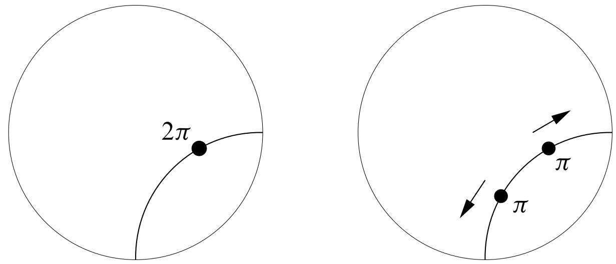

Figure 3: Splitting a curvature concentration

curve at x is a positive (negative) π–rotation. Then we say that x is an *over-rotated vertex* if c_o lies in D_- (D_+), otherwise we say it is an *under-rotated vertex*.

We say that a generalized PGC–curve of tangential degree (Whitney index) m in the hyperbolic plane is in *standard position* if it has the form of the curve in Figure 4, where if, $|m| \neq 0$, all vertex curves have length π and have the same orientation. Clearly any two PGC–homotopic generalized PGC–curves in standard position are PGC–homotopic through such curves.

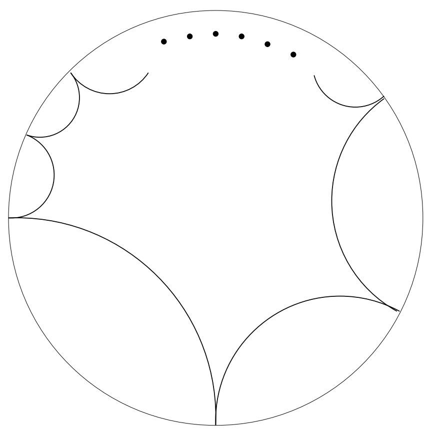

Figure 4: A curve in standard position

Lemma 4.4 *If c_0 and c_1 are two PGC–homotopic PGC–curves in the hyperbolic plane each with at least one curvature concentration not an integral multiple of π, then there exists a PGC–homotopy c_t, $0 \leq t \leq 1$, from c_0 to c_1 such that*

$$\kappa(c_t) \leq \max\{\kappa(c_0), \kappa(c_1)\},$$

for all t.

Proof Using Lemma 4.3, we deform any given PGC–curve to a generalized PGC–curve with all its vertices at infinity and all curvature concentrations equal to π without increasing κ. In fact, the condition that some curvature concentration is not an integral multiple of π implies that κ is decreased by this deformation. It is thus sufficient to show that any generalized PGC–curve with all its vertices at infinity and all curvature concentrations equal to π may be deformed to a curve in standard position by a generalized PGC–homotopy which increases κ arbitrarily little.

Note that any generalized PGC–curve with all its vertices at infinity which has only two infinite vertices is automatically in standard position. Assume inductively that any such curve with $< m$ vertices can be brought to standard position by a PGC–homotopy which increases κ arbitrarily little and consider a curve with m vertices.

If the curve has an under-rotated vertex then this vertex can be removed using the method of Case 1 in the proof of Lemma 4.3. The inductive assumption finishes the proof in that case. We thus assume that the curve does not have any under-rotated vertices.

Consider a curve which satisfies this assumption and which has two curvature concentrations of opposite signs. Such a curve must have two vertices a and b with curvature concentrations of opposite signs which are connected by a geodesic arc C. Let A and B be the other geodesics with endpoints at a and b, respectively. Since every vertex is over-rotated, it follows that C separates A from B. Push the second endpoint of B which is not equal to b until it is very close to a and the endpoint of A which is not equal to a until it is very close to b. With this done, push the curvature concentrations inwards along C and cancel them increasing κ arbitrarily little. The initial phase of the latter homotopy is shown in Figure 5.

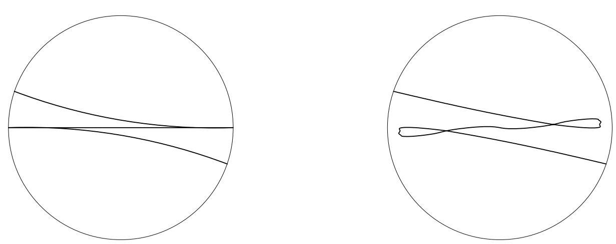

Figure 5: Canceling opposite concentrations

Finally, assume that there are no under-rotated vertices and that all of the vertices have curvature concentrations of the same sign. In this case we pick a first vertex and move the second toward it, eventually creating a small geodesic. We then move the third vertex toward the second and so on. We claim that either this process creates a standard curve or it creates an under-rotated vertex. To see this, let p_j be the first vertex which has not been moved. There are two cases, either p_j can be moved without passing p_{j+2}. In this case the construction continues. Or, p_j must pass p_{j+2} in which case an under-rotated vertex at p_{j+1} is created. The claim and the lemma follow. □

4.3 Constant curvature spheres

Lemma 4.5 *Let Σ be the 2–sphere with a constant curvature metric and let c be a* PGC*–curve in Σ. Then there exists a* PGC*–homotopy c_t, $0 \le t \le 1$, with $c_0 = c$, with $\kappa(c_t)$ a non-increasing function of t, and with the following property: the underlying curve of c_1 is a geodesic and any curvature concentration of c_1 equals π.*

Proof As in Lemma 4.3, scaling invariance of κ implies we may assume $K = 1$. We claim that any PGC–curve c which does not have underlying curve a geodesic admits a PGC–homotopy which does not increase κ and which decreases the number of vertices with curvature concentration not a multiple of π. As in the proof of Lemma 4.1 this finishes the proof.

Let p_0 denote the start and p_1 denote the endpoint of a geodesic e_1 of c. Let e_0 and e_2 be the other geodesic segments connecting to p_0 and p_1, respectively, and let p_2 denote the other vertex of e_2. Let G_j be the great circle in which e_j has its image. If g is a geodesic on Σ we let $l(g)$ denote its length. We must consider two separate cases. In the first case $l(e_1) \neq m\pi$ for all integers $m > 0$. In this case we deform the curve by moving p_1 along G_2 in such a way that the length $l(\gamma_0)$ of the vertex curve γ_0 at p_0 *decreases*. We stop this deformation the first time p_1 hits p_2 or one of the points in $G_0 \cap G_2$, or, when $l(\gamma_0) = n\pi$, for some integer n. At this instant we obtain a curve with the number of vertices with curvature concentration not an integral multiple of π one smaller than the corresponding number for c.

A straightforward case by case check shows that this deformation does not increase κ. More precisely, there are 16 subcases to check. They arise as follows. First $l(e_1) = a + 2\pi m$, $0 < a < \pi$ where $m \ge 0$ is even or odd, second the tangent vectors of e_0 and e_2 at p_0 and p_1, respectively, points into different components of $\Sigma - G_1$ or into the same component, third and fourth the lengths of the vertex curve γ_j at p_j satisfies $l(\gamma_j) = a + 2\pi m$, $0 < a < \pi$, where $m \ge 0$ is even or odd, $j = 0, 1$. However, the fact that κ does not increase follows in all of these subcases from one of the following two computations.

First, with notation as in Figure 6, we calculate

$$\beta' + (\pi - (\alpha - \alpha')) + \pi - \beta + \int_{\Omega} K \, dA = 2\pi.$$

Thus,

$$\Delta\kappa = (\alpha' + \beta') - (\alpha + \beta) = -\int_{\Omega} K \, dA < 0.$$

Second, with notation as in Figure 7, we calculate

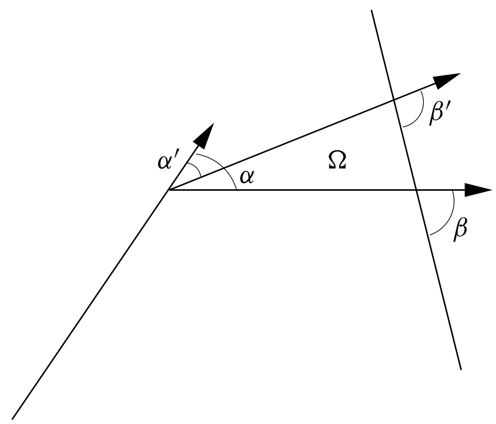

Figure 6: Removing a vertex I

$$\beta + (\pi - \beta') + (\pi - (\alpha - \alpha')) + \int_{\Omega} K \, dA = 2\pi.$$

Thus,

$$\Delta\kappa = (\alpha' + \beta') - (\alpha + \beta) = \int_{\Omega} K \, dA - 2(\alpha - \alpha') < 0,$$

where the last inequality follows since $2(\alpha - \alpha') = \int_{\Gamma} K \, dA$ where Γ is the angular region between the geodesics connecting antipodal points and intersecting at an angle $\alpha - \alpha'$, and since $\Omega \subset \Gamma$.

In the second case $l(e_1) = n\pi$ for some integer $n > 0$. In this case we rotate e_1 in such a way that the length of at least one of the vertex curves at the endpoints of e_1 decreases. The process stops when one of the vertex curves becomes an integral multiple of π. This finishes the proof. □

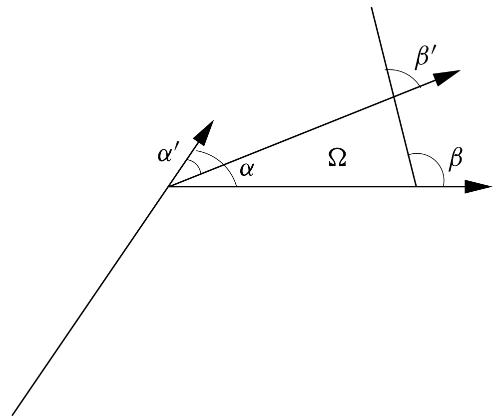

Figure 7: Removing a vertex II

Lemma 4.6 *Let c_0 and c_1 be any two* PGC*–homotopic curves on a Riemann surface Σ as in Lemma 4.5. Then there exists a* PGC*–homotopy c_t, $0 \leq t \leq 1$, connecting c_0 to c_1 with*

$$\kappa(c_t) \leq \min\{\kappa(c_0), \kappa(c_1), 2\pi\},$$

for all t.

Proof Lemma 4.5 shows that it is enough to consider two curves with the properties of c_1 there. We first show how to deform such a curve to a multiple of a closed geodesic. Fix a first vertex p and the orientation of the great circle of its incoming geodesic. Move the second vertex q to the antipodal point of the first and rotate the arc which connects p and q an angle π. Note that after the rotation, the orientation of the arc agrees with the fixed one. If the orientation of the vertex curves at p and q are the same then this rotation removes two vertices and decreases κ. If the two vertex curves have opposite orientations this rotation does not change κ and one vertex with vertex curve of length 2π is created. Splitting this new born vertex into two and repeating the above argument removes them. In this way, we eventually remove all vertices.

To finish the proof we need only show how to increase the number of times a curve encircles a geodesic by 2 not increasing κ by more than 2π. We use the following procedure. Create two curvature concentrations of length π and of opposite orientations keeping κ not larger than 2π. Applying the procedure described above, we arrive at a curve which goes two more times around the geodesic. □

4.4 Curves with self-tangencies on the sphere

A generic circle immersion into a surface has only transverse double points. In generic regular homotopies there appear isolated instances of triple points and self-tangencies. The self-tangencies are of two kinds *direct* when the tangent vectors at the tangency point agree and *opposite* when they do not. Let Σ denote the 2–sphere with a constant curvature metric throughout this subsection.

Lemma 4.7 *Any circle immersion $c\colon S^1 \to \Sigma$ with an opposite self-tangency satisfies $\kappa(c) > 2\pi$.*

Proof Since c is not a geodesic we may decrease $\kappa(c)$ by a small deformation, keeping the self-tangency, see the proof of Proposition 2.2. Let b be the curve resulting from such a deformation. Fix a partition π of the circle such that the PGC–approximation b^π of b satisfies $\kappa(b^\pi) < \kappa(c)$, see Lemma 3.2. Let $\widetilde{b}$ be a PGC–curve close to b^π which contains two segments of the same geodesic close to the self-tangency point of c and such that $\kappa(\widetilde{b}) < \kappa(c)$.

Apply the deformation in the proof of Lemma 4.5 to $\widetilde{b}$ with the endpoint of one of the self-tangency segments as p_1. Note that as this process reaches the endpoint of the other self-tangency segment, the PGC–curve constructed must contain a geodesic segment with at least one vertex curve of length π. Since the process does not increase κ and since the lift of a PGC–curve is closed, it follows that $\kappa(\widetilde{b}) \geq 2\pi$ and therefore $\kappa(c) > 2\pi$. □

In order to prove Theorem 1.1 (b) we are going to apply Arnol'd's J^-–invariant of immersed curves on S^2. Arnol'd introduced his invariant for circle immersions in the plane, see [1]. Its existence for curves on more general surfaces was established by Inshakov [5] and Tchernov [9]. The existence of the J^-–invariant stems from the following fact: if c_0 and c_1 are two self-transverse regularly homotopic curves on S^2 then the algebraic numbers of opposite self-tangencies in any two generic regular homotopies c_t, $0 \leq t \leq 1$, connecting them are equal. This number is called the relative J^-–invariant of c_0 and c_1, we will denote it $\Delta J^-(c_0, c_1)$. To compute the algebraic number of self-tangencies each self-tangency moment is equipped with a sign as follows: it is a *positive* moment if it increases the number of double points of the curve, otherwise it is *negative*.

Lemma 4.8 *Let $k > 0$ be an integer and consider the circle immersions $S^1 \to \Sigma$ which go k and $k+2$ times respectively around a geodesic (a great circle). Let c_0 and c_1 be any small perturbations of these curves. Then $|\Delta J^-(c_0, c_1)| = 2$.*

Proof Any sufficiently small perturbation of a curve going around a geodesic can be thought of as a multi-graph over the 0–section in a tubular neighborhood of a great circle. This implies that there exist deformations without opposite self-tangencies connecting any two such perturbations. To compute $\Delta J^-(c_0, c_1)$ it is thus sufficient to pick two perturbations and count the algebraic number of opposite self-tangencies in one regular homotopy connecting them. In Figure 8, the first picture illustrates a curve in a neighborhood of the north pole which is obtained by shrinking a perturbed multiple of the equator and in the first deformation one kink is pulled over the south pole. The lemma follows from Figure 8. □

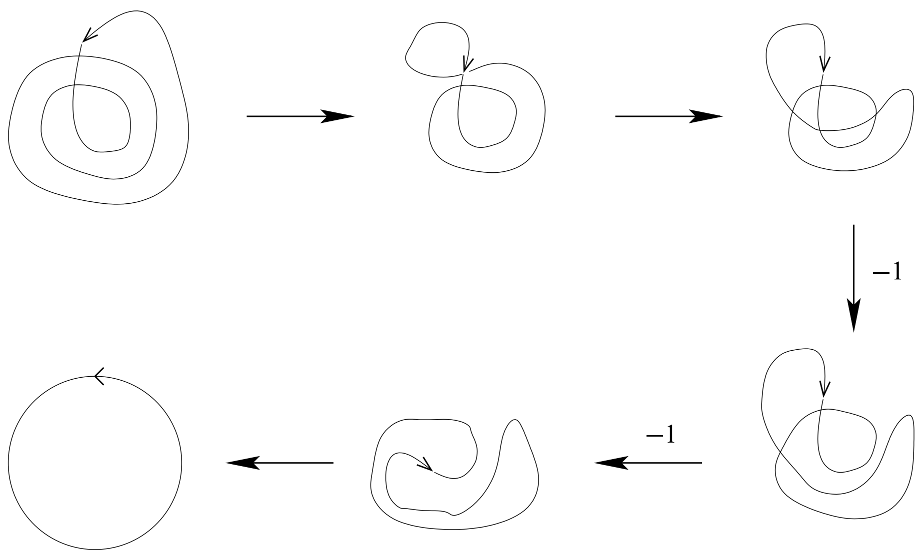

Figure 8: Computation of ΔJ^-

4.5 Smoothing PGC–homotopies

The PGC–homotopies in Lemma 4.1, Lemma 4.3 and Lemma 4.5 have very special forms. We show that such PGC–homotopies can be made smooth in a standard way, increasing the total curvature arbitrarily little.

Consider first two unit vectors $v_{\mathrm{in}} \in \mathbb{R}^2$ and $v_{\mathrm{out}} \in \mathbb{R}^2$ and let γ be a geodesic in S^1 connecting v_{in} to v_{out} and let $\delta > 0$ be given. Fix a family of reference curves $b\big(v_{\mathrm{in}}, v_{\mathrm{out}}, \gamma\big)$ inside the unit disk such that b agrees with the straight line in direction v_{in} (v_{out}) near the endpoints. The tangent map of b is homotopic to γ with endpoints fixed and the total curvature of b exceeds the length of γ by at most δ. Clearly there exists such families which depend continuously on the data for all $\delta > 0$.

We first discuss smoothing of a fixed PGC–curve. Let $c = (c_1, \gamma_1, \dots, c_m, \gamma_m)$ be a PGC–curve such that all vertex curves are immersions. For sufficiently small $\epsilon > 0$, we define an ϵ–smoothing of c as follows. Fix a disk of radius $\epsilon > 0$ around each vertex of c. Under the inverse of the exponential map at a vertex the curve c looks like the model discussed above. More precisely, in the tangent space of the surface Σ at a vertex p we have an incoming unit tangent vector v_{in} and an outgoing one v_{out}, and we glue in the curve $b\big(v_{\mathrm{in}}, v_{\mathrm{out}}, \gamma\big)$ in this tangent space. We then scale the glued in curve by ϵ and map it back into Σ with the exponential map at p. Applying this procedure at each vertex we get a smoothing $\widetilde{c}$ of c. Note that $\kappa(\widetilde{c})$ can be made arbitrarily close to $\kappa(c)$ by choosing δ and ϵ sufficiently small. (The deviation from the flat case is measured by a curvature integral over a region with area going to 0 with ϵ.) We illustrate this smoothing process in Figures 9 and 10.

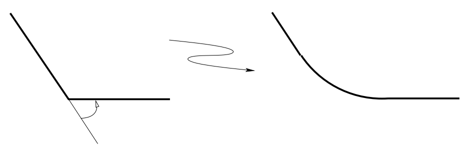

Figure 9: Smoothing of exterior angle

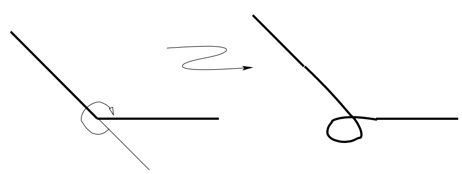

Figure 10: Smoothing of interior angle

We next note that all PGC–homotopies in the proofs of Lemma 4.1, Lemma 4.3 and Lemma 4.5 have one of the following forms

- One geodesic segment of a PGC–curve moves along a consecutive segment.
- Two curvature concentrations of magnitude π and with opposite orientations are created or annihilated somewhere along a PGC–curve.
- A homotopy of the form presented in the last part of Lemma 4.1.

We call such PGC–homotopies *simple*.

Lemma 4.9 *Let c_t, $0 \le t \le 1$, be a simple* PGC*–homotopy between two* PGC*–curves c_0 and c_1 with immersed vertex curves. Then, for any $\epsilon > 0$ there exists a regular homotopy $\widetilde{c}_t$ connecting the ϵ–smoothings $\widetilde{c}_0$ and $\widetilde{c}_1$ with*

$$\max_{0\le t\le 1} \kappa(\widetilde{c}_t) \le \max_{0\le t\le 1} \kappa(c_t) + 10\epsilon.$$

Proof The proof is straightforward. Consider a PGC–homotopy of the first type. We define $\widetilde{c}_t$ as the ϵ–smoothing of c_t for t outside a neighborhood of 1. Inside a neighborhood of 1 we may again use local flat models and define $\widetilde{c}_t$ by composing with the exponential map. A picture of such a local model is shown in Figure 11. For

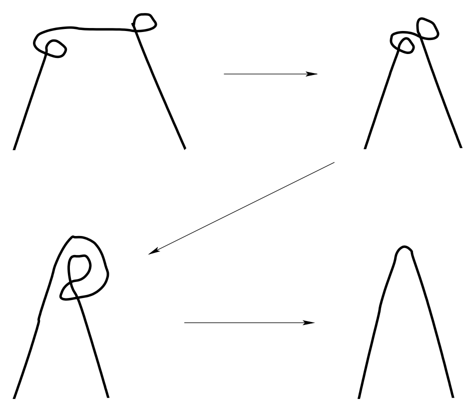

Figure 11: A regular homotopy near two meeting vertices

the local model of a simple homotopy of the second type, see Figure 12.

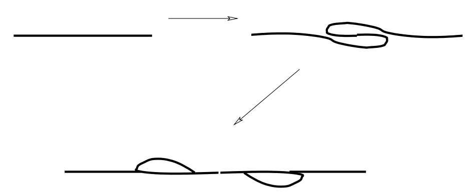

Figure 12: Creation of curvature concentrations

For simple homotopies of the third kind one may simply take $\widetilde{c}_t$ to equal the ϵ–smoothing of c_t for all t. □

4.6 Locally convex curves on flat surfaces

Let Σ be a flat Riemann surface ($\Sigma \approx T^2$, $\Sigma \approx S^1 \times \mathbb{R}$, or $\Sigma \approx \mathbb{R}^2$). For $p \in \Sigma$, $v \in U_p\Sigma$, and $(\xi, m) \in \pi_1(\Sigma) \times \mathbb{Z}$ let $\Omega_{(\xi,m)}(p,v)$ ($\widehat{\Omega}_{(\xi,m)}(p,v)$) denote the space of all (strictly) locally convex circle immersions $c\colon S^1 \to \Sigma$ of regular homotopy class (ξ, m), see Section 1, such that $c(1) = p$ and $\dot{c}(1) = v$. (As usual we use the C^2–topology on the space of circle immersions.)

Lemma 4.10 *The spaces $\Omega_{(\xi,m)}(p,v)$ and $\widehat{\Omega}_{(\xi,m)}(p,v)$ are weakly contractible.*

Proof We start in the simply connected case, $\Sigma = \mathbb{R}^2$. Let F denote $\Omega_m(p,v)$ or $\widehat{\Omega}_m(p,v)$. Let $\Gamma\colon S^n \to F$ be a continuous map from the n–sphere, $\Gamma(x) = c_x\colon S^1 \to \mathbb{R}^2$. We think of S^1 as an interval $[0, L]$ with endpoints identified. Thus $c_x(0) = p$ and $\dot{c}_x(0) = v$, for some fixed point p, some unit vector v, and all $x \in S^n$. To prove the lemma we must extend Γ continuously to the $(n+1)$–ball B^{n+1}.

Fix a small $\epsilon > 0$ and a unit vector w such that $\langle v, w\rangle = 0$ and such that in the orientation of the plane induced by the basis w, v the tangential degree of the curves c_x are positive. We claim that there exist continuous maps $t_j\colon S^n \to \mathbb{R}$, $j = 1, 2$, with $0 < t_1(x) < t_2(x) < L$, and with the following properties: $\dot{c}_x(t_1(x))$ lies in the short sub-arc A_ϵ of S^1 between $(\cos\epsilon)v + (\sin\epsilon)w$ and $(\cos 2\epsilon)v + (\sin 2\epsilon)w$, $\dot{c}_x(t_2(x))$ lies in the short sub-arc B_ϵ between $(\cos\epsilon)v - (\sin\epsilon)w$ and $(\cos 2\epsilon)v - (\sin 2\epsilon)w$, $t_1(x)$ lies in the component of $\dot{c}_x^{-1}(A_\epsilon)$ closest to 0, and $t_2(x)$ lies in the component of $\dot{c}_x^{-1}(B_\epsilon)$ closest L.

In the strictly locally convex case this claim is obviously true: consider preimages under $\dot{c}_x$ of fixed points in A_ϵ and B_ϵ to define $t_1(x)$ and $t_2(x)$, respectively. To see that it holds also in the non-strictly locally convex case we argue as follows. By continuity, the subset $\Lambda \subset S^n \times S^1$,

$$\Lambda = \{(x,t)\colon \dot{c}_x(t) \in \operatorname{int}(A_\epsilon)\},$$

where $\operatorname{int}(X)$ denotes the interior of X, is open. In particular, for each $x \in S^n$ there exists $r_x > 0$ such that

$$\bigcap_{y \in B(x, r_x)} \dot{c}_y^{-1}(\operatorname{int}(A_\epsilon))$$

contains an interval I_x. Cover S^n by balls $B(x, r_x)$ with this property. This cover has a Lebesgue number $\delta > 0$. Triangulate S^n by simplices which are so small that for every vertex in the triangulation, the union of all simplices in which this vertex lies is a subset of diameter less than δ. It is then straightforward to construct $t_1\colon S^n \to \mathbb{R}$: if v is a vertex take $t_1(v)$ as any point in I_x where x is some point such that the union of all simplices containing v lies in $B(x, r_x)$. By contractibility of the interval we can now inductively extend this function over higher dimensional skeleta of the triangulation. In the final step we get the desired function on S^n.

Let q_x be the intersection point of the tangent lines to c_x at $c_x(t_1(x))$ and $c_x(t_2(x))$. Define an initial deformation of the curves c_x in the family which pushes c_x toward the piecewise linear curve obtained by replacing $c_x([0, t_1(x)]) \cup c_x([t_2(x), L])$ with the curve $\overline{c_x(t_1(x))q_x} \cup \overline{c_x(t_2(x))q_x}$, where $\overline{ab}$ denotes the line segment between a and b. To stay in the space of pointed curves we also translate and slightly rotate the curves to ensure that they pass through p with the right tangent, see Figure 13. Clearly, this

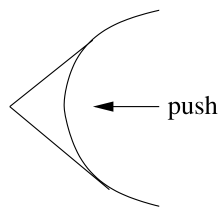

Figure 13: An initial deformation

deformation can be made continuous in x and chosen in such a way that the resulting curves, still denoted c_x, are strictly locally convex at $c_x(0) = p$ for each $x \in S^n$.

The next step is to deform the curves so that the marked point is a global minimum of the height function in direction w. To this end, let $T_j\colon S^n \to \mathbb{R}$, $j = 1, 2$, $0 < T_1(x) < T_2(x) < L$ be continuous functions on S^n such that $\dot{c}_x(T_1(x))$ ($\dot{c}_x(T_2(x))$) lies in an η–arc C_η ending at w (beginning at $-w$) in the orientation of S^1 determined by $\dot{c}_x$ and such that $T_1(x)$ ($T_2(x)$) lies in the component of $\dot{c}_x^{-1}(C_\eta)$ closest to 0 (to L). In the strictly locally convex case such functions are easily constructed using suitable points in the preimage $\dot{c}_x^{-1}(\pm w)$. In the non-strictly locally convex case, such functions can be constructed using the same arguments that were used in the construction of the functions t_j, $j = 1, 2$, just given. Let $w_1(x) = \dot{c}_x(T_1(x))$ and let $w_2(x) = -\dot{c}_x(T_2(x))$.

Pick $M > 0$ such that the minimum point $c_x(q)$ in the w–direction of any curve c_x, $x \in S^n$, satisfies $\langle w, c_x(q) - p\rangle > -\frac{1}{2}M$. Pick $\eta > 0$ sufficiently small so that the intersection point r of the lines in direction $w_1(x)$ and $w_2(x)$ passing through $c_x(T_1(x))$ and $c_x(T_2(x))$ respectively satisfies $\langle w, r - c_x(q)\rangle < -20M$. Consider the subdivision of the source circle of c_x into two arcs $D_1(x)$ and $D_2(x)$ as follows: the endpoints of the arcs are $T_1(x)$ and $T_2(x)$, $D_1(x)$ contains the marked point, and $D_2(x)$ does not contain it.

For $0 \le s \le \frac{1}{2}$, define c_{sx} as the curve which consists of the following four pieces: the curve $c_x(D_2(x))$, the line segment l_s^1 in direction $w_1(x)$ starting at $c_x(T_1(x))$ and such that the projection of this segment to a line in the w–direction has length $10sM$, a suitably scaled version of the translate of the curve $c_x(D_1(x))$ along l_s^1, which is tangent to the line through $c_x(T_2(x))$ in direction $w_2(x)$, and finally the line segment l_s^2 in direction $-w_2(x)$ connecting the scaled and translated $c_x(D_1(x))$ to $c_x(D_2(x))$ at $c_x(T_2(x))$, see Figure 14. In the case when F is a space of strictly convex curves

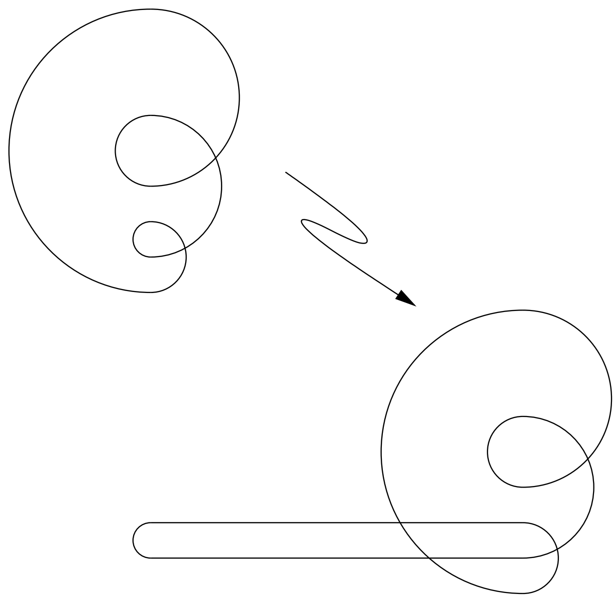

Figure 14: Making the marked point a global extremum

we replace the straight line segments in $c_{sx}(t)$ above with slightly curved circular arcs. (In fact, since S^n is compact we can take these arcs to have curvature smaller than $\min_{x\in S^n}\min_{t\in[0,L]}(|k_g(t)|)$.) Finally, to have the curves mapping the marked point to p we also compose with a suitable translation.

The second step in the deformation takes all the curves in the family to curves with image in a large circle. Let C be a circle through p with tangent v at p. Let D be the bounded component of the plane with boundary C. If the radius of C is sufficiently big then we may choose C so that all curves $c_{\frac{1}{2}x}$ has image in D and $C \cap c_{\frac{1}{2}x}([0,L]) = c_{\frac{1}{2}x}(0)$. For $0 \le s \le 1$ let $c_{(-\frac{s}{2}+\frac{1}{2})x}$ be the continuous family of convex curves which is the union of a curve in C starting at p and ending at the intersection point of the negative tangent half-line of $c_{\frac{1}{2}x}$ at $c_{\frac{1}{2}x}(sL)$, follows this half-line, and then goes along $c_{\frac{1}{2}x}$. For $s = 1$, we get a curve wrapping around C, see Figure 15. The construction is continuous in x and therefore gives a continuous

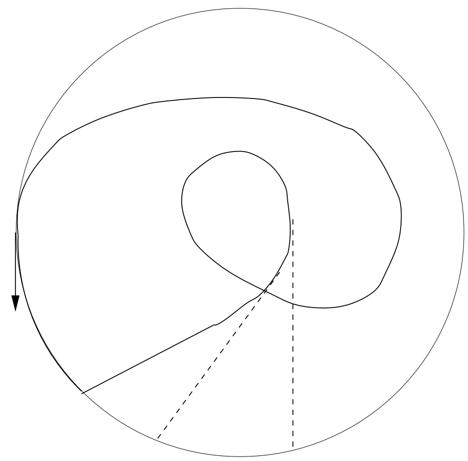

Figure 15: Making a locally convex curve circular

extension of the map $\Gamma\colon S^n \to F$ over B^{n+1}, as desired. (Again, in the case of strictly locally convex curves we replace the tangent half-line with a very slightly curved circular arc.)

We next consider the non-simply connected case. In this case we have a covering $\Pi\colon \mathbb{R}^2 \to \Sigma$ where the deck transformations are translations. We lift the curves c_x at the marked point. In the case when F is the space of non strictly locally convex curves a similar local deformation as the initial deformation above makes all curves strictly locally convex at the marked point. Let b_x denote the lifted curves. We have $b_x(0) = 0$, $b_x(L) = q$, $\dot{b}_x(0) = v = \dot{b}_x(L)$.

Let l_0 and l_q be the straight lines through 0 and q respectively with tangent vector v. By strict local convexity, for one of the unit vectors w orthogonal to v, we have $\langle b_x(t), w\rangle > 0$ for t in a (punctured) neighborhood of 0 and for t in a (punctured) neighborhood of L. For convenience, assume that $\langle q, w\rangle > 0$ (otherwise change coordinates in $\mathbb{R}^2$ so that $q = 0$). We construct $T_j\colon S^n \to \mathbb{R}$, $j = 1, 2$, analogous to the functions with the same names above, in a similar way as above so that $\dot{b}_x(T_1(x)) \approx w$ and so that $\dot{b}_x(T_2(x)) \approx -w$. As above, we add straight line segments to b_x so that $b_x(0)$ (or $b_x(q)$) is the global minimum of the height function in direction w and so that b_x intersects the region between the two lines l_0 and l_q in an arc.

Pick two circles C_0 and C_q through 0 and q, respectively, both with tangents v at these points. Let D_0 and D_q be the bounded components of the plane which are bounded by C_0 and C_q, respectively. If the circles have sufficiently large radii then $b_x([0,L]) \subset D_0$ for each x and b_x intersects $D_0 - D_q$ in an arc. Let C be the boundary of the convex hull of C_0 and C_q. As above, we deform b_x using its negative tangent half-line. For $0 \le s \le \tau$ where τ is the smallest number such that the negative tangent half-line at $b_x(\tau)$ intersects C in a point in C_q. We let b_{sx} be a part of C followed by the negative tangent half-line, in turn followed by the rest of b_x. For $s \ge \tau$ we let b_{sx} be the curve with initial part as above followed by a curve in C_q, in turn followed by the tangent half-line, and finally the rest of b_x. At $s = L$ we find that every curve in the family is a curve which is a segment in C followed by a curve with image in C_q, see Figure 16. In the case of strictly locally convex curves it is easy to

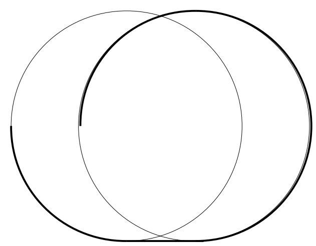

Figure 16: The final stage of a deformation of the lift of a locally convex curve

modify the deformation described so that it keeps all curves strictly locally convex. This finishes the proof. □

Lemma 4.11 *The inclusion $\hat{\Omega}_{(\xi,m)} \subset \Omega_{(\xi,m)}$ induces surjections*

$$\pi_r(\hat{\Omega}_{(\xi,m)}) \to \pi_r(\Omega_{(\xi,m)}),$$

on homotopy groups, for all r.

Proof The proof of Lemma 4.10 homotopes an arbitrary family of (based) locally convex curves to a family of strictly locally convex curves. Hence we need only consider what happens to the base point. Let $\Gamma\colon S^n \to \Omega_{(\xi,m)}$ be a family of curves. First use the initial deformation of the proof above to make all curves strictly locally convex in a neighborhood of $1 \in S^1$. With this accomplished we lift all curves in the family to $\mathbb{R}^2$, by lifting at $c(1)$. (If $n \ge 2$ then we can lift the whole family in this way however when $n = 1$ the start- and endpoints of our lifts may differ by a translation.) Now apply the procedure of the proof of Lemma 4.10. Note that (in the case $n = 1$) the procedure behaves well with respect to translations. Hence, any $\Gamma\colon S^n \to \Omega_{(\xi,m)}$ can be homotoped to a map $\Gamma'\colon S^n \to \hat{\Omega}_{(\xi,m)}$. This finishes the proof. □

5 Proofs

In this section we prove the theorems stated in Section 1.

5.1 Infima

Proof of Theorem 1.2 Consider first case (c). A well-known theorem of Lyusternik and Fet [7] says that the 2–sphere with any metric (actually any closed Riemannian manifold) has a non-constant simple closed geodesic, see also Jost [6, Section 5.5]. This geodesic traversed once and twice respectively gives representatives with $\kappa = 0$ for both regular homotopy classes on the 2–sphere.

In case (b) it is easy to construct a locally convex curve in any regular homotopy class (ξ, m) with $\xi \neq *$ or $m \neq 0$. Note that the class $(*, 0)$ can neither be represented by a closed geodesic nor by a locally convex curve. Hence $\kappa(c)$ can be decreased for each c of regular homotopy class $(*, 0)$. Moreover, the curvature decreasing procedure in the proof of Lemma 4.1 gives in this case a PGC–curve with underlying curve a geodesic segment. Since the lift of that PGC–curve is closed it must have at least two curvature concentrations of magnitude π. Hence $\widehat{\kappa}(*, 0) \geq 2\pi$. Approximating a segment traversed twice with two vertex curves which are rotations π and $-\pi$ we find that $\widehat{\kappa}(*, 0) = 2\pi$.

Finally, in case (a), it follows as above that the infimum is not attained in classes not representable by geodesics. To find the infima in case $\Sigma \approx \mathbb{R}^2$ we apply Lemma 4.3 to conclude that it is enough to find the minimal κ of a generalized PGC–curve in standard form in a given regular homotopy class. It is straightforward to check that a curve in standard form representing the regular homotopy class m has two vertex curves of length π if $|m| = 0$ and $|m| + 1$ vertex curves of length π otherwise. □

5.2 Locally convex curves

Proof of Theorem 1.3 Let v be a (covariantly) constant unit vector field on Σ. Let $\widehat{\Omega}_{(\xi,m)}(v)$ and $\Omega_{(\xi,m)}(v)$ denote the space of strictly locally convex–, respectively, locally convex curves c with $\dot{c}(1) = v$. Consider the evaluation maps $e(c) = c(1)$,

$$e\colon \widehat{\Omega}_{(\xi,m)}(v) \to \Sigma \quad \text{and } e\colon \Omega_{(\xi,m)}(v) \to \Sigma.$$

These maps are clearly Serre fibrations: using translations we can lift any map from an n–disk into Σ. The fibers of these fibrations are $\widehat{\Omega}_{(\xi,m)}(p, v)$ and $\Omega_{(\xi,m)}(p, v)$, respectively, which are both weakly contractible by Lemma 4.10. Hence

$$\pi_r(\widehat{\Omega}_{(\xi,m)}(v)) \approx \pi_r(\Sigma) \approx \pi_r(\Omega_{(\xi,m)}(v))$$

with isomorphisms induced by evaluation.

Assume $m \neq 0$ and consider the fibration

$$e'\colon \widehat{\Omega}_{(\xi,m)} \to S^1,$$

$e'(c) = \dot{c}(0)$. This is a Serre fibration since $\dot{c}\colon S^1 \to S^1$ is a covering map (here it is essential that the curves are *strictly* locally convex). Since the fiber of e' is $\widehat{\Omega}_{(\xi,m)}(v)$ we find that

$$\pi_r(\widehat{\Omega}_{(\xi,m)}) = \pi_r(U\Sigma),$$

with isomorphism induced by the evaluation map. Since $\pi_r(U\Sigma) = 0$ if $r > 1$ we conclude from Lemma 4.11 that $\pi_k(\Omega_{(\xi,m)}) = \pi_k(\widehat{\Omega}_{(\xi,m)})$ for all k.

We finally consider $m = 0$. In this case the space under consideration is the space of closed geodesics in a fixed homotopy class. If $\Sigma \approx S^1 \times \mathbb{R}$ any element in such a space is uniquely determined by its intersection with $\mathbb{R} \times \{1\}$ and $\pi_{S^1}(c(1))$, where $\pi_{S^1}\colon S^1 \times \mathbb{R} \to S^1$ is the natural projection. Thus $\Omega_{(\xi,0)} \simeq \Sigma$. Similar arguments give the same result when $\Sigma \approx T^2$. □

5.3 Regular homotopies

Proof of Theorem 1.1 Consider first case (a). If Σ is flat and both c_0 and c_1 are closed geodesics or locally convex curves then the theorem follows from Theorem 1.3. In all other cases we may first decrease the total curvature a little by Proposition 2.2 and then approximate by a PGC–curve as in Lemma 3.2. In the flat case the theorem then follows from Lemma 4.1, which allows us to deform both curves to a standard form without increasing κ, Lemma 4.2 which allows us to connect these curves, and Lemma 4.9 which allows us to smooth the entire homotopy keeping control of κ. In the case of negative curvature the theorem follows in a similar way from Lemmas 4.3, 4.4 and 4.9.

In case (b) we argue in the same way to prove the first statement using Lemmas 3.2, 4.5, 4.6 and 4.9. The second statement follows from Lemma 4.8 which shows that any regular homotopy c_t, $0 \leq t \leq 1$, connecting the two multiple geodesics of different multiplicities must have an instant c_τ which is a curve with an opposite self-tangency and Lemma 4.7 which shows that $\kappa(c_\tau) > 2\pi$. □

We end the paper by demonstrating that Theorem 1.1 (a) does not hold for Riemann surfaces with metrics of non-constant curvature with $K \leq 0$.

Remark 5.1 Consider the upper half-plane with coordinates (x, y), $y > 0$ and metric

$$ds^2 = e^{2f(y)}(dx^2 + dy^2),$$

where $f\colon (0,\infty) \to \mathbb{R}$ is a (weakly) convex function such that its derivative f' satisfies $f'(y) = 0$ for $y \in [1,2] \cup [3,4]$, and such that its second derivative f'' satisfies $f''(y) > 0$ for $y \in (0,1) \cup (4,\infty)$ and such that $f(y) = -\log(y)$ for y in some neighborhood of 0 and of ∞.

Let $A_0 = \mathbb{R}\times[1,2]$, $A_1 = \mathbb{R}\times[3,4]$, and $B = \mathbb{R}\times(0,\infty) - (A_0 \cup A_1)$. Then $K(p) = 0$ for $p \in A_0 \cup A_1$ and $K(p) < 0$ for $p \in B$. Let c_0 and c_1 be convex curves in A_0 and A_1, respectively. Then c_0 and c_1 are regularly homotopic. We claim that for every regular homotopy c_t, $0 \le t \le 1$ connecting c_0 to c_1, there exists an instant c_τ with $\kappa(c_\tau) > 2\pi$. To see this note that as long as the curve c_t stays in A_0 it must remain convex, otherwise $\kappa > 2\pi$. In particular, the curve must remain embedded. Let t_0 be the last moment when the curve lies completely inside the closure of A_0. Since embeddedness is an open condition we see that c_τ is embedded for all $\tau > t_0$ sufficiently close to t_0. Note that $c_{t_0} \cap \partial A_0 \neq \varnothing$. Pick some line l parallel to ∂A_0, and close to a point in $c_{t_0} \cap \partial A_0$ and such that c_{t_0} intersects it transversely in two points. (The existence of such a line follows from Sard's lemma.) Then also c_τ meets l transversely for τ sufficiently close to t_0. Applying the Gauss–Bonnet theorem to the two curves bounded by the bounded segment of l cut out by c_τ and the two remaining pieces of c_τ we find $\kappa(c_\tau) > 2\pi$.

References

[1] **V I Arnol'd**, *Plane curves, their invariants, perestroikas and classifications*, from: "Singularities and bifurcations", Adv. Soviet Math. 21, Amer. Math. Soc., Providence, RI (1994) 33–91 MR1310595

[2] **T Ekholm**, *Regular homotopy and total curvature II: sphere immersions into 3-space*, Algebr. Geom. Topol. 6 (2006) 493–512

[3] **M Gromov**, *Partial differential relations*, Ergebnisse series 9, Springer, Berlin (1986) MR864505

[4] **M W Hirsch**, *Immersions of manifolds*, Trans. Amer. Math. Soc. 93 (1959) 242–276 MR0119214

[5] **A V Inshakov**, *Invariants of types j^+, j^- and* st *of smooth curves on two-dimensional manifolds*, Funktsional. Anal. i Prilozhen. 33 (1999) 35–46, 96 MR1724268

[6] **J Jost**, *Riemannian geometry and geometric analysis*, Universitext, Springer, Berlin (1998) MR1625976

[7] **L A Lyusternik**, **A I Fet**, *Variational problems on closed manifolds*, Doklady Akad. Nauk SSSR (N.S.) 81 (1951) 17–18 MR0044760

[8] **S Smale**, *The classification of immersions of spheres in Euclidean spaces*, Ann. of Math. (2) 69 (1959) 327–344 MR0105117

[9] **V Tchernov**, *Arnold-type invariants of wave fronts on surfaces*, Topology 41 (2002) 1–45 MR1871239

Department of mathematics, USC
Los Angeles CA 90803, USA

tekholm@usc.edu

Received: 8 February 2005 Revised: 22 February 2006