

Geometric filling curves on surfaces

Ara Basmajian*, Hugo Parlier** and Juan Souto***

Abstract. This note is about a type of quantitative density of closed geodesics on closed hyperbolic surfaces. The main results are upper bounds on the length of the shortest closed geodesic that ε -fills the surface.

1. Introduction

Closed geodesics on hyperbolic surfaces provide a concrete link between algebraic, geometric and topological approaches to understanding the geometry of these surfaces and their moduli spaces. One feature of the hyperbolic geometry is that they are dense in both the surface and the unit tangent bundle. This is in strong contrast to the set of simple closed geodesics or those with bounded intersection number which are nowhere dense and in fact Hausdorff dimension 1 [2]. In this article, we investigate a type of quantitative density of closed geodesics.

Given a closed hyperbolic surface X and $\varepsilon > 0$, we're interested in finding the shortest closed geodesic that is ε -dense, by which we mean that all points of X are at distance at most ε from the geodesic. For $g \geq 2$, we denote by \mathcal{M}_g the moduli space of closed hyperbolic surfaces of genus g . Our main result is the following.

Theorem 1.1. *For all $X \in \mathcal{M}_g$ there exists a constant $C_X > 0$ and such that for all $\varepsilon \leq \frac{1}{2}$ there exists a closed geodesic γ_ε that is ε -dense on X and such that*

$$\ell(\gamma_\varepsilon) \leq C_X \frac{1}{\varepsilon} \log \left(\frac{1}{\varepsilon} \right)$$

Another measure of complexity for a closed geodesic is its self-intersection number. Length and self-intersection numbers of a curve are of course related (see for instance [1]). Instead of minimizing length, one can try and minimize self-intersection for ε filling curves. As a corollary of Theorem 1.1 we obtain the following.

*Research supported by a PSC-CUNY Grant and a Simons foundation grant

**Research supported by Swiss National Science Foundation grant number PP00P2.153024

***Research supported by the National Science Foundation Grant No. DMS-1440140

2010 Mathematics Subject Classification: Primary: 30F10. Secondary: 32G15, 53C22.

Key words and phrases: closed geodesics, hyperbolic surfaces

Corollary 1.2. *For all $X \in \mathcal{M}_g$ there exists a constant $C_X > 0$ such that for all $\varepsilon \leq \frac{1}{2}$ there exists a closed geodesic γ_ε that is ε -dense on X and such that*

$$i(\gamma_\varepsilon, \gamma_\varepsilon) \leq C_X \frac{1}{\varepsilon^2} \left(\log \left(\frac{1}{\varepsilon} \right) \right)^2$$

The main ingredient in the proof of Theorem 1.1 is a more technical result (Theorem 2.4) which shows the existence of a closed geodesic of bounded length that contains a given set of geodesic segments on X in its ε -neighborhood. Theorem 1.1 then follows by finding an appropriate set of geodesic segments that fill the surface. This tool can also be used to find a similar quantitative density result in the unit tangent bundle of X (Theorem 3.2).

The growth rate of length in Theorem 1.1 is perhaps not optimal but if not it is off by at most a factor of $\log \left(\frac{1}{\varepsilon} \right)$. Indeed if a closed geodesic is ε -dense then the area of its ε -neighborhood is the area of the surface. In \mathbb{H} , for small ε , the area of a ε -neighborhood of a geodesic segment of length ℓ is roughly $\varepsilon\ell$. And putting these things together tells us that if a geodesic γ ε -fills, it must be of length at least $\frac{\text{area}(X)}{\varepsilon}$. Understanding the $\log \left(\frac{1}{\varepsilon} \right)$ discrepancy between the upper and lower bounds seems to be an interesting problem.

Acknowledgements.

The authors acknowledge support from U.S. National Science Foundation grants DMS 1107452, 1107263, 1107367 RNMS: Geometric structures And Representation varieties (the GEAR Network). This material is based upon work supported by the National Science Foundation under Grant No. DMS-1440140 while the third author was in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2016 semester.

2. Geodesics in \mathbb{H} and on surfaces

We begin with a simple lemma about hyperbolic polygons and angles of geodesics traversing them. By angle between two geodesics we mean the minimal angle, so in particular angles are always less than $\frac{\pi}{2}$.

Lemma 2.1. *Let P be a convex polygon in \mathbb{H} . Then there exists a $\theta_P > 0$ such that any geodesic segment with endpoints on distinct sides of ∂P forms an angle of at least θ_P in one of its endpoints.*

Proof. The statement follows by a compactness argument since a geodesic segment cannot have angle 0 in both endpoints, but there is a direct argument that also shows that this minimal angle corresponds to a specific geometric quantity which we now describe.

Consider the set of all triangles with vertices the vertices of two consecutive sides of P . We'll call these triangles the *ears* of P (see Figure 1). Consider the set of all the (inner) angles of all of the ears (three angles for each ear) and set θ_P to be their minimum value.

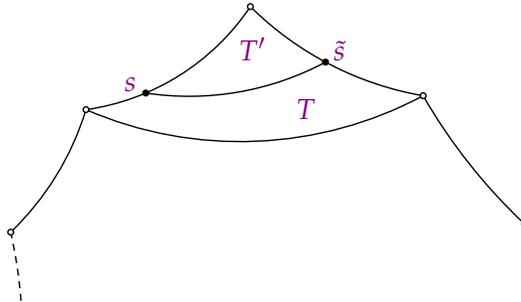


Figure 1: The ear T containing T'

Now consider a geodesic segment c leaving from a side s of ∂P forming an angle θ of at most θ_P . It stays entirely in an ear T (one of the sides of the ear is s) so it intersects a side \tilde{s} of P adjacent to s . The triangle T' formed by c with segments of s and \tilde{s} is contained in T and is thus of lesser area. As it also shares an angle with T , the sum of its two remaining angles cannot be strictly less than $2\theta_P$, otherwise by Gauss-Bonnet the area of T' would be greater than that of T . As $\theta \leq \theta_P$, the remaining angle is at least θ_P .

Note that θ_P is optimal as the inner sides of the ears are admissible geodesic segments. \square

The following is just an observation about geodesics in \mathbb{H} . The proof we give, and many of the following proofs, use hyperbolic trigonometry. We refer the reader to [3] or any other standard hyperbolic geometry textbook for the formulas.

Lemma 2.2. Let $\frac{\pi}{2} \geq \theta_0 > 0$, and set

$$m(\theta_0) := \operatorname{arccosh} \left(\frac{2}{\sin^2(\theta_0)} - 1 \right)$$

If c is a geodesic segment in \mathbb{H} of length at least $m(\theta_0)$ between two (complete) geodesics γ_1, γ_2 such that

$$\angle(c, \gamma_i) \geq \theta_0$$

for $i = 1, 2$, then γ_1 and γ_2 are disjoint.

Proof. Fix θ_0 . To check whether segments of given length ℓ with the angle condition always give rise to disjoint geodesics, we can consider the "worst case scenario" which is when both geodesics form an angle of θ_0 with the segment. The limit case in this worst case

scenario between intersecting and not intersecting is when the two geodesics γ_1 and γ_2 are ultra-parallel and in this case we have a triangle with an ideal vertex and two angles of θ_0 in endpoints of a segment c of length ℓ . In this case ℓ is easy to compute via standard hyperbolic trigonometry: it is

$$m(\theta_0) = \operatorname{arccosh} \left(\frac{2}{\sin^2(\theta_0)} - 1 \right)$$

The distance between γ_1 and γ_2 being an increasing function of ℓ , we can conclude that any c of length greater or equal to $m(\theta_0)$ with the angle condition will lie between disjoint geodesics. \square

Note that the proof also shows that the condition is sharp. The next lemma is more technical and contains several quantifiers, but is again a somewhat elementary statement about geodesics in \mathbb{H} .

Lemma 2.3. *Let $\frac{\pi}{2} \geq \theta_0 > 0$ be a fixed constant. Let c be a geodesic segment in \mathbb{H} and γ the complete geodesic containing c . Fix $\varepsilon > 0$ and let γ_1 and γ_2 be geodesics that intersect γ such that the intersection points p_1, p_2 lie on different sides of c .*

Suppose for $i = 1, 2$ that $\angle(\gamma_i, \gamma) \geq \theta_0$ and

$$d(c, p_i) \geq \log \left(\frac{1}{\varepsilon} \right) + \log \left(\frac{4}{\sin(\theta_0)} \right)$$

Then any geodesic δ intersecting both γ_1 and γ_2 satisfies

$$c \subset B_\varepsilon(\delta)$$

Proof. Given θ_0 and the length of c , we'll consider the "worst case scenario" by which we mean the situation, under the assumptions of the lemma, where a geodesic with endpoints on γ_1 and γ_2 is as far away as possible from c .

In order for it to be the worst case scenario a number of things need to happen. First of all, given any setup for c and γ_1 and γ_2 the geodesic furthest away from c with endpoints on γ_1 and γ_2 is a limit case and actually is a geodesic δ ultra-parallel to both γ_1 and γ_2 . Furthermore, decreasing the distance between c and p_i pushes geodesics away, so we can suppose that

$$d(c, p_i) = \log \left(\frac{1}{\varepsilon} \right) + \log \left(\frac{4}{\sin(\theta_0)} \right) =: r_\varepsilon$$

Finally, if one of the angles is greater than θ_0 then decreasing the angle also pushes c away from the limit case so we can suppose that both angles of intersection are exactly θ_0 . We can thus suppose that we are in the situation of Figure 2.

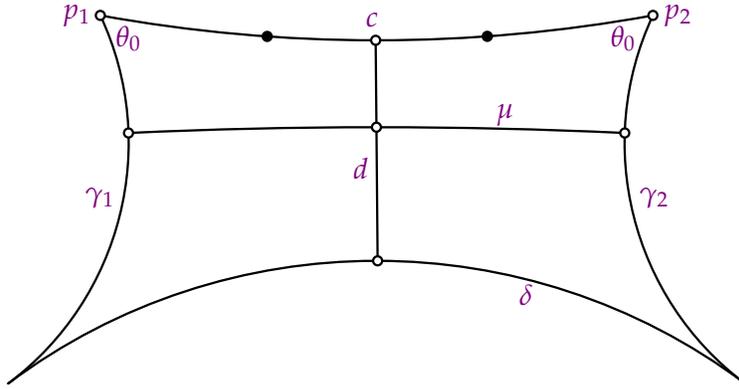


Figure 2: The worst case scenario: the segment $c \subset \gamma$ lies between the two black vertices.

This situation has a lot of symmetry which we'll now use. Let μ be the minimal geodesic path between γ_1 and γ_2 and let d be the the distance between γ and δ . We denote by h the maximal distance between points of c and δ . It is this quantity that we need to bound in function of the other parameters. It lies in a symmetric quadrilateral Q with c as one of its sides

We denote by d'' the distance between γ and μ and by d' the distance between μ and δ . By looking at the quadrilaterals separated by μ , observe that $d' > d''$ and thus that $d' > \frac{d}{2}$. We consider another symmetric quadrilateral Q' , somewhat similar to Q but this time with the height of length d replaced with a height of length d' (see Figure 3).

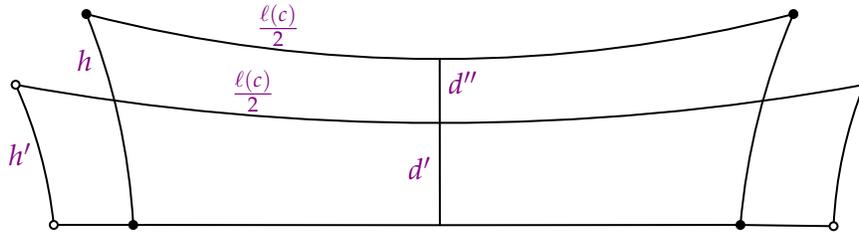


Figure 3: The quadrilaterals Q and Q'

We denote by h' the length of the side corresponding to h and we claim that in fact h' enjoys the property of $h' > \frac{h}{2}$. This will be useful because h' is somewhat easier to compute.

To prove this, using the left right symmetries of Q and Q' , we first restrict ourselves to the left half of each quadrilateral (see Figure 3).

Using hyperbolic trigonometry, we have

$$\sinh(h') = \sinh(d') \cosh\left(\frac{\ell(c)}{2}\right) > \sinh\left(\frac{d}{2}\right)$$

whereas

$$\sinh(h) = \sinh(d) \cosh\left(\frac{\ell(c)}{2}\right)$$

Thus

$$\frac{\sinh(h')}{\sinh(h)} > \frac{\sinh\left(\frac{d}{2}\right)}{\sinh(d)} = \frac{1}{2 \cosh\left(\frac{d}{2}\right)}$$

But if $h' \leq \frac{h}{2}$ we would have

$$\frac{\sinh(h')}{\sinh(h)} \leq \frac{\sinh\left(\frac{h}{2}\right)}{\sinh(h)} = \frac{1}{2 \cosh\left(\frac{d}{2}\right)}$$

a contradiction showing $h' > \frac{h}{2}$.

We now seek to bound h' and to do so, we begin by computing $\ell(\mu)$. Using for instance the upper left quadrilateral of Figure 2 and hyperbolic trigonometry, we have

$$\cosh\left(\frac{\ell(\mu)}{2}\right) = \cosh(r_\varepsilon + \ell(c)/2) \sin(\theta_0)$$

Using the lower left quadrilateral we have

$$\sinh(d') = \frac{1}{\sinh\left(\frac{\ell(\mu)}{2}\right)}$$

We now return to the quadrilateral pictured in Figure 4 to compute h' .

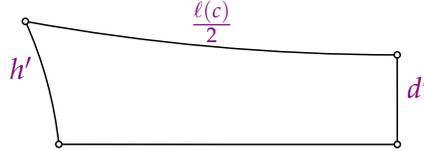


Figure 4: The left half of Q'

We have

$$\sin(\theta_0) \cosh(r_\varepsilon + \ell(c)/2) \sinh(h') = \cosh\left(\frac{\ell(c)}{2}\right)$$

from which we deduce

$$h' = \operatorname{arcsinh}\left(\frac{1}{\sin(\theta_0)} \frac{\cosh\left(\frac{\ell(c)}{2}\right)}{\cosh\left(r_\varepsilon + \frac{\ell(c)}{2}\right)}\right)$$

We want to show that $h' \leq \frac{\varepsilon}{2}$ thus that

$$\frac{1}{\sin(\theta_0)} \frac{\cosh\left(\frac{\ell(c)}{2}\right)}{\cosh\left(r_\varepsilon + \frac{\ell(c)}{2}\right)} \leq \sinh\left(\frac{\varepsilon}{2}\right) \quad (1)$$

The left hand of Inequality 1 can be manipulated to show that

$$\frac{1}{\sin(\theta_0)} \frac{\cosh\left(\frac{\ell(c)}{2}\right)}{\cosh\left(r_\varepsilon + \frac{\ell(c)}{2}\right)} = \frac{1}{\sin(\theta_0)} \frac{1 + e^{-\ell(c)}}{e^{r_\varepsilon} + e^{-r_\varepsilon - \ell(c)}}$$

Note that

$$\frac{1 + e^{-\ell(c)}}{e^{r_\varepsilon} + e^{-r_\varepsilon - \ell(c)}} < \frac{2}{e^{r_\varepsilon}}$$

and thus Inequality 1 certainly holds provided

$$\frac{1}{\sin(\theta_0)} \frac{2}{e^{r_\varepsilon}} \leq \sinh\left(\frac{\varepsilon}{2}\right)$$

and thus will hold if

$$\frac{1}{\sin(\theta_0)} \frac{2}{e^{r_\varepsilon}} \leq \frac{\varepsilon}{2}$$

Expressed differently this last inequality becomes

$$r_\varepsilon \geq \log\left(\frac{1}{\varepsilon}\right) + \log\left(\frac{4}{\sin(\theta_0)}\right)$$

As the above inequality is in fact an equality, this proves the lemma. \square

We now can proceed to prove the main tool for our results.

Theorem 2.4. *For any $X \in \mathcal{M}_g$, there exists a constant K_X such that the following holds. For all $1 > \varepsilon > 0$ and any finite collection $\{c_i\}_{i=1}^N$ of geodesic segments of average length \bar{c} on X , there exists a closed geodesic γ such that*

$$\ell(\gamma) \leq N \left(K_X + \bar{c} + 2 \log\left(\frac{1}{\varepsilon}\right) \right)$$

and for all $i = 1, \dots, N$

$$c_i \subset B_\varepsilon(\gamma)$$

Proof. We take a filling closed geodesic γ_0 on X of minimal length (among filling geodesics). Thus $X \setminus \gamma_0$ consists of a finite collection of polygons $\{P_i\}_{i \in I}$ and we denote by D the maximum of their intrinsic diameters. Furthermore we denote by θ_0 the minimum of $\{\theta_{P_i}\}_{i \in I}$ where the θ_{P_i} s are from Lemma 2.1.

As X is orientable, the geodesic γ_0 has two sides which we think of as being $+$ and $-$ (it does not matter which is which but we fix it). We take μ_+ , resp. μ_- , to be a geodesic arc from γ to itself, orthogonal γ in both end points, and which leaves and returns to the $+$ side, resp. $-$ side. The existence of μ_+ and μ_- might not be obvious but can be shown

as follows. Take a cover $\tilde{X} \rightarrow X$ of finite index so that a lift $\tilde{\gamma}_0$ of γ_0 is simple. In \tilde{X} , it suffices to find orthogonal arcs that return to the same side by completing $\tilde{\gamma}_0$ into a pants decomposition and by taking appropriate orthogonal arcs in the pants containing $\tilde{\gamma}_0$. The image in X of these arcs provides μ_+ and μ_- .)

We note that $\theta_0, D, \ell(\gamma_0), \ell(\mu_+)$ and $\ell(\mu_-)$ are all quantities that depend only on X .

We'll begin by constructing a closed piecewise geodesic segment containing all of the c_i s with endpoints lying on γ_0 . We begin by extending each c_i to a segment c'_i with endpoints on γ_0 as follows. Following Lemmas 2.3 and 2.2, we set

$$r_\varepsilon := \max\left\{\log\left(\frac{1}{\varepsilon}\right) + \log\left(\frac{4}{\sin(\theta_0)}, m(\theta_0)\right)\right\}.$$

We first extend each c_i by r_ε in both directions.

We continue to extend the segments in both directions until both endpoints lie on γ_0 at an angle $\theta \geq \theta_0$ for the first time. The extended segment we denote by c'_i .

We now show how the length of c'_i is bounded. As γ_0 fills, a first intersection point will occur within an extension of length at most D on both ends. If the angle is too small in one of the endpoints, by extending by at most an additional D there is a second intersection point. The arc between two successive intersection points is an arc with endpoints on one of the polygons of $X \setminus \gamma_0$. By Lemma 2.2, one of the two angles must be at least θ_0 .

Note that because the quantities D, θ_0 and $m(\theta_0)$ only depend on X we have that the resulting geodesic arc c'_i satisfies

$$\ell(c'_i) \leq \ell(c_i) + 2 \log\left(\frac{1}{\varepsilon}\right) + C_X$$

where C_X is a constant only depending on X .

We think of the segments c'_i as being cyclically ordered and we construct a closed curve as follows.

We orient each c'_i arbitrarily. To determine what happens between the endpoint of c'_i and the starting point of c'_{i+1} , what will be important is on which sides of γ_0 they start and end on.

If c'_{i+1} begins on the opposite side that c'_i ends on, we join the endpoint of c'_i to the starting point of c'_{i+1} by the shortest subarc of γ_0 which does this and which will be of length at most $\frac{\gamma_0}{2}$.

If c'_{i+1} begins on the same side that c'_i ends on, we'll use either μ_+ or μ_- which we also think of as oriented. Suppose c'_i ends on $+$ (the other case is symmetric). We join the endpoint of

c'_i to the starting point of μ_- by a shortest arc of γ_0 that does this and is of length at most $\frac{\ell(\gamma_0)}{2}$. Now we join the endpoint of μ_- to the initial point of c'_{i+1} similarly, i.e., by a shortest arc of γ_0 of length at most $\frac{\ell(\gamma_0)}{2}$. All in all, the length of these three additional arcs is at most $\ell(\gamma_0) + \max(\ell(\mu_+), \ell(\mu_-))$.

In this way, we have constructed a piecewise geodesic closed curve γ' and we denote by γ the unique geodesic in its homotopy class. Its length is thus upper bounded by

$$\sum_{k=1}^N \ell(c'_k) + N (\ell(\gamma_0) + \max(\ell(\mu_+), \ell(\mu_-)))$$

Putting this bound together with the bound on the length of each c'_i , we set

$$K_X := C_X + \ell(\gamma_0) + \max(\ell(\mu_+), \ell(\mu_-))$$

to obtain the desired bound on the length of γ .

What remains to be seen is that γ is non-trivial and contains all of the segments c_i in its ε -neighborhood. To see this we consider a lift in the universal cover.

Note that the piecewise geodesic γ' consists of an even number of geodesic arcs where arcs of γ_0 are followed by either c'_i s or μ_+ or μ_- (we'll refer to these as the non γ_0 arcs).

We begin by taking a specific lift of γ' (for instance by taking a lift of c'_1 and then constructing the full lift from there). Note that each lift of a c'_i lies between two copies of γ_0 which are disjoint by the length condition imposed on c'_i and by Lemma 2.2. As they are orthogonal in their endpoints, this also holds for lifts of μ_+ and μ_- .

We think of the lift of γ' as being oriented. We were careful about what sides the non γ_0 arcs left from and in particular we ensured that successive non γ_0 arcs left from different sides. In the lift, this means that the piecewise geodesic never backtracks and traverses each lift of γ_0 corresponding to the lift of a γ_0 arc. Thus the successive lifts of γ_0 form a sequence of nested geodesics.

That tells us two things: the lift of γ' is simple and it limits in both directions to distinct endpoints on the boundary of \mathbb{H} . We denote by $\tilde{\gamma}$ the unique geodesic in \mathbb{H} with these endpoints (this is a lift of γ). By standard arguments, this ensures that γ' , and thus γ , is non trivial.

We now look at $\tilde{\gamma}$. Given a lift of c'_i , we look at the two lifts of γ_0 surrounding it which we'll denote γ'_i and γ''_i . By construction c'_i has the proper length on both ends of the lift of c_i and angle condition to ensure that any geodesic intersecting both γ'_i and γ''_i contains c_i in a ε -neighborhood. By the above argument, $\tilde{\gamma}$ crosses both γ'_i and γ''_i and this is what we wanted to show. \square

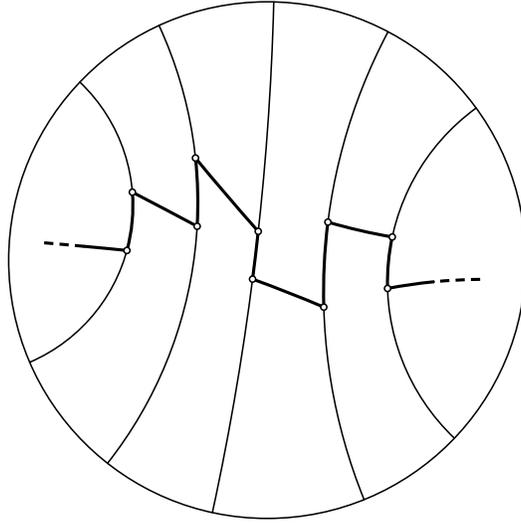


Figure 5: Lifting γ' to \mathbb{H} : the complete geodesics are lifts of γ_0

3. Quantitative density

We now apply Theorem 2.4 to prove results about dense geodesics. Recall that a closed geodesic γ is ε -dense on X if $d_X(x, \gamma) \leq \varepsilon$ for all $x \in X$.

Theorem 1.1. *For all $X \in \mathcal{M}_g$ there exists a constant $C_X > 0$ such that for all $\varepsilon \leq \frac{1}{2}$ there exists a closed geodesic γ_ε that is ε -dense on X and such that*

$$\ell(\gamma_\varepsilon) \leq C_X \frac{1}{\varepsilon} \log \left(\frac{1}{\varepsilon} \right)$$

Proof. The strategy is to find a collection of geodesic segments that $\frac{\varepsilon}{2}$ -fill X and then to apply Theorem 2.4 to find a closed geodesic that contains the segments in a $\frac{\varepsilon}{2}$ -neighborhood.

We begin by setting

$$R_X := \min\{1, \text{inrad}(X)\}$$

where $\text{inrad}(X)$ is the minimum injectivity radius of X . We consider a maximal collection points $\{p_k\}_{k \in I}$ all pairwise distance at least R_X (I is a finite index set). Note that the balls of radius R_X around the p_k cover X . Further note that the balls of radius $\frac{R_X}{2}$ around the p_k are disjoint and thus

$$|I| \leq \frac{\text{area}(X)}{\text{area}(B_{R_X/2})}$$

In particular, there is a bound on $|I|$ which only depends on X . Now consider a ball of radius R_X in \mathbb{H} . As $R_X \leq 1$, the perimeter of this ball is less than

$$2\pi \sinh(1)$$

We want to fill it by geodesic segments such that all points are distance at most $\frac{\varepsilon}{2}$ from one of the segments. To do so we place points $q_k, k \in J$ on the perimeter, all exactly ε -apart except for possibly the last one which might be closer to its two neighbors. Note there are at most

$$\frac{2\pi \sinh(1)}{\varepsilon}$$

Now consider any collection of disjoint "parallel" segments constructed by first taking two neighbors and joining them by a geodesic segment, and then taking the geodesic between their other neighbors and so forth. There are half as many segments as there were points. Taking into account that we might have had an odd number of points, we've constructed at most

$$\frac{\pi \sinh(1)}{\varepsilon} + 1$$

segments. Now any point in the disk lies between two geodesic segments and is at a distance at most $\frac{\varepsilon}{2}$ from one of them (this is because the distance between the endpoints of these segments is at most ε).

Doing this for each of our balls, we have at most

$$A_X \frac{1}{\varepsilon}$$

segments where A_X is a constant depending only on X . Each of them are length at most 2.

We now apply Theorem 2.4 to these segments asking that the closed geodesic γ_ε contain every segment in its $\frac{\varepsilon}{2}$ neighborhood. The estimate follows. \square

Remark 3.1. In the above proof, one might be tempted to proceed more directly by using a covering of X by balls of radius ε . Although this would produce the desired ε -filling geodesic, its length would be on the order of $\frac{1}{\varepsilon^2} \log\left(\frac{1}{\varepsilon}\right)$.

Using the same method as above, one can find an estimate on the length of a shortest closed geodesic that is ε -dense in the unit tangent bundle $UT(X)$ of X . For this one needs a notion of distance. One way of defining this is by associating oriented segments to vectors. Given X and given a constant $l \leq \text{inrad}(X)$, we say that a curve γ is ε -dense in $UT(X)$ if every oriented geodesic segment on X of length l is contained in a ε -neighborhood of γ .

Theorem 3.2. *For all $X \in \mathcal{M}_g$ there exists a constant $C_X > 0$ such that for all $\varepsilon \leq \frac{1}{2}$ there exists a closed geodesic γ_ε that is ε -dense in $UT(X)$ and such that*

$$\ell(\gamma_\varepsilon) \leq C_X \frac{1}{\varepsilon^2} \log\left(\frac{1}{\varepsilon}\right)$$

Proof. The proof is identical to the proof of Theorem 1.1 with the following exceptions.

When choosing the set of balls that cover X , we choose the balls to be of radius $2R_X$ (they may not be embedded but we can consider them in the universal cover and project to X). This is to ensure that we capture the totality of oriented segments of length ℓ .

Now when constructing the geodesic segments, instead of choosing roughly $\frac{1}{\varepsilon}$ segments that lie between the points q_k , we choose all oriented geodesic segments that lie between them. As there are roughly $\frac{1}{\varepsilon}$ points q_k , there are roughly $\frac{1}{\varepsilon^2}$ segments this time, hence the difference in the estimate. \square

We end the paper with a corollary of Theorem 1.1, mentioned in the introduction.

Corollary 1.2. *For all $X \in \mathcal{M}_g$ there exists a constant $C_X > 0$ and an ε_0 such that for all $\varepsilon \leq \varepsilon_0$ there exists a closed geodesic γ_ε that is ε -dense on X and such that*

$$i(\gamma_\varepsilon, \gamma_\varepsilon) \leq C_X \frac{1}{\varepsilon^2} \left(\log \left(\frac{1}{\varepsilon} \right) \right)^2$$

Proof. This follows directly from the length estimates on γ_ε from Theorem 1.1 as any geodesic of length ℓ can have at most roughly ℓ^2 self-intersection points by the following argument. Take the geodesic of length ℓ and break it up into segments of length ℓ_0 where ℓ_0 is less than the minimal injectivity radius of X . Now any two segments of length ℓ_0 can intersect at most once. \square

Addresses:

A.B.: Department of Mathematics, Graduate Center and Hunter College, CUNY

Email: abasmajian@gc.cuny.edu

H.P.: Department of Mathematics, University of Luxembourg

Email: hugo.parlier@uni.lu

J.S.: IRMAR, University of Rennes 1

Email: juan.souto@univ-rennes1.fr

References

- [1] Ara Basmajian. Universal length bounds for non-simple closed geodesics on hyperbolic surfaces. *J. Topol.*, 6(2):513–524, 2013.
- [2] Joan S. Birman and Caroline Series. Geodesics with bounded intersection number on surfaces are sparsely distributed. *Topology*, 24(2):217–225, 1985.
- [3] Peter Buser. *Geometry and spectra of compact Riemann surfaces*, volume 106 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, 1992.