

NON-CONVEX BALLS IN THE TEICHMÜLLER METRIC

MAXIME FORTIER BOURQUE & KASRA RAFI

Abstract

We prove that the Teichmüller space of surfaces of genus \mathbf{g} with \mathbf{p} punctures contains balls which are not convex in the Teichmüller metric whenever its complex dimension $(3\mathbf{g} - 3 + \mathbf{p})$ is greater than 1.

1. Introduction

Let \bar{S} be a closed oriented surface and $P \subset \bar{S}$ a finite set. The Teichmüller space of $S = \bar{S} \setminus P$ is the set of conformal structures on \bar{S} up to biholomorphisms homotopic to the identity rel P . The Teichmüller metric on this space $\mathcal{T}(S)$ measures how much diffeomorphisms of \bar{S} homotopic to the identity rel P must distort angles with respect to different conformal structures. This metric is complete, uniquely geodesic, Finsler, and agrees with the Kobayashi metric on $\mathcal{T}(S)$. However, its local geometry is quite subtle. Indeed, we prove that:

Theorem 1.1. *There exist non-convex balls in $\mathcal{T}(S)$ whenever its complex dimension is greater than 1.*

Note that for any $X \in \mathcal{T}(S)$, the balls of sufficiently small radius centered at X are convex. This is true in any Finsler manifold [Whi33] [Tra37].

Motivation. If S is a once-punctured torus or a four-times-punctured sphere, then $\mathcal{T}(S)$ is isometric to \mathbb{H}^2 , the hyperbolic plane with constant curvature -4 . This led Kravetz to argue that in general, $\mathcal{T}(S)$ is non-positively curved in the sense of Busemann [Kra59]. However, Linch [Lin71] found a flaw in Kravetz's reasoning and soon after, Masur [Mas75] showed that the result was false: there exist distinct geodesic rays starting from the same point in $\mathcal{T}(S)$ and staying a bounded distance apart whenever $\dim_{\mathbb{C}} \mathcal{T}(S) > 1$. In particular, Teichmüller space is not CAT(0) nor δ -hyperbolic.

In any proper geodesic metric space \mathbb{X} , we have the implications

- \mathbb{X} is non-positively curved in the sense of Busemann
- \Rightarrow the distance to any point is strictly convex along any geodesic not containing that point
- \Rightarrow closed balls are strictly convex
- \Rightarrow the convex hull of any finite set is compact.

The question of whether the third statement held for Teichmüller space was originally motivated by the Nielsen realization problem, which Kravetz thought he had solved with his erroneous result. Masur's paper [Mas75] rendered the problem open again. If balls had been strictly convex, then a positive solution to the Nielsen realization problem would have followed immediately. In light of Theorem 1.1, this approach was doomed to fail. Thankfully, Kerckhoff solved Nielsen's problem many years ago using the convexity of hyperbolic length along earthquake paths [Ker83]. See also [Wol87] for a solution using Weil–Petersson geometry. Whether the fourth statement holds for Teichmüller space is an open question of Masur.

Sketch of proof. Given $X \in \mathcal{T}(S)$ and a simple closed curve $\alpha \subset S$, the extremal length $\text{EL}(\alpha, X)$ is the smallest c such that a Euclidean cylinder of height 1 and circumference c embeds conformally in X in the homotopy class of α . Similarly, the extremal length of a multicurve is the least possible sum of circumferences of disjoint embedded cylinders of height 1 (see Section 2). The first step of the proof of Theorem 1.1 is to reduce it to a statement about extremal length.

Lemma 1.2. *If every ball in $\mathcal{T}(S)$ is convex, then for every multicurve $\gamma \subset S$ and every Teichmüller geodesic $t \mapsto Z_t$ in $\mathcal{T}(S)$ the function $t \mapsto \text{EL}(\gamma, Z_t)$ has no local maximum.*

It was shown in [LR11] that extremal length of a curve is not necessarily convex along Teichmüller geodesics. Indeed, the authors of that paper constructed an example where the function $t \mapsto \text{EL}(\alpha, X_t)$ increases by a definite amount at first and then stays nearly constant on a later interval. The idea of our construction is to take such a pair (α, X_t) with the surface having a puncture, then another copy (β, Y_t) of the same curve and surface but where the time parameter has been reversed and shifted, and to form a connected sum $Z_t = X_t \# Y_t$ via a small slit at the puncture. This is done in such a way that $t \mapsto Z_t$ is still a Teichmüller geodesic.

We then show that $\text{EL}(\alpha + \beta, Z_t)$ converges to $\text{EL}(\alpha, X_t) + \text{EL}(\beta, Y_t)$ as the size of the slit shrinks. If we arrange the time parameter of Y_t so that $\text{EL}(\beta, Y_t)$ is nearly constant when $\text{EL}(\alpha, X_t)$ increases, and decreases when $\text{EL}(\alpha, X_t)$ is nearly constant, then their sum increases on the first

interval and decreases on the second interval. By the convergence of $EL(\alpha + \beta, Z_t)$ to the sum, that quantity also increases during the first interval and decreases later, provided that the slit is small enough. This forces a local maximum in between, and thereby proves the existence of a non-convex ball.

This proof requires the surface S to be the connected sum of two surfaces each of which is sufficiently complicated. It does not work when the complex dimension of $\mathcal{T}(S)$ is less than 4. For those lower complexity cases, our proof is based on rigorous numerical calculations.

Related results. In [LR11], Lenzhen and Rafi proved that balls in $\mathcal{T}(S)$ are quasi-convex. More precisely, they showed that there exists a constant $c = c(S)$ such that for any ball $B \subset \mathcal{T}(S)$, every geodesic segment with endpoints in B stays within distance c of B . In other words, balls cannot fail to be convex arbitrarily badly.

Theorem 1.1 indicates that the Teichmüller metric is positively curved locally, where balls fail to be convex. There are also large-scale manifestations of positive curvature. Namely, there are unbounded regions in Teichmüller space where the Teichmüller metric looks like a sup metric on a product [Min96]. On the other hand, there is a sense in which $\mathcal{T}(S)$ is hyperbolic relative to its thin parts [MM99]. We refer the reader to [Mas09] for a survey on curvature aspects of the Teichmüller metric and to [Raf14] for a coarse description of the Teichmüller metric and its geodesics.

Lastly, Theorem 1.1 should be put in contrast with previous convexity results:

- $\mathcal{T}(S)$ is holomorphically convex [BE64];
- hyperbolic length of a curve is convex along earthquake paths [Ker83] and Weil–Petersson geodesics [Wol87];
- hyperbolic length [Wol06] and extremal length [Miy15] of a curve are log-plurisubharmonic.

Organization. Section 2 starts with some background on Teichmüller theory. We then reformulate of the convexity problem in terms of extremal length in Section 3. Section 4 proves the convergence of extremal length under pinching deformations. Examples of local maxima for extremal length are constructed in Section 5 for surfaces with enough topology. Finally, Section 6 presents the numerical results which settle the lower complexity cases.

ACKNOWLEDGEMENTS. The authors thank Jeremy Kahn for suggesting the proof of Lemma 3.3, Curtis McMullen for useful comments, and Vincent Delecroix and David Dumas for advice on computer-assisted proofs. The first author was partially supported by the Fonds de recherche du Québec – Nature et technologies. The second author was partially supported by NSERC grant # 435885.

2. Preliminaries

Teichmüller space. A point in Teichmüller space $\mathcal{T}(S)$ is a *marked analytically finite complex structure* on S . This means a Riemann surface X together with an orientation-preserving homeomorphism $f : S \rightarrow X$ which extends to a homeomorphism $\bar{f} : \bar{S} \rightarrow \bar{X}$, where \bar{X} is a closed Riemann surface containing X . Two points (X, f) and (Y, g) are identified if there exists a conformal isomorphism $h : X \rightarrow Y$ homotopic to $g \circ f^{-1}$. We will write $X \in \mathcal{T}(S)$, keeping the marking f implicit.

Teichmüller distance. A linear map $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ is K -*quasiconformal* if it preserves signed angles up to a factor $K \geq 1$. Equivalently, a linear map is K -quasiconformal if it has positive determinant and sends circles to ellipses with major axis to minor axis ratio at most K .

A homeomorphism between Riemann surfaces is K -*quasiconformal* if its distributional partial derivatives are locally square-integrable and if its matrix of partial derivatives is K -quasiconformal almost everywhere. The *dilatation* $\text{Dil}(h)$ of a quasiconformal homeomorphism h is the smallest K for which it is K -quasiconformal. All quasiconformal homeomorphisms considered in this paper will be piecewise smooth.

Given X and Y in $\mathcal{T}(S)$ with markings f and g , the *Teichmüller distance* between them is defined as

$$d(X, Y) = \inf_h \frac{1}{2} \log \text{Dil}(h)$$

where the infimum is taken over all quasiconformal homeomorphisms $h : X \rightarrow Y$ homotopic to $g \circ f^{-1}$.

Half-translation structures. A *half-translation* in \mathbb{C} is either a translation or a rotation of angle π about a point, i.e., a map of the form $z \mapsto \pm z + c$. A *half-translation surface* Φ is a collection of polygons in \mathbb{C} with sides identified in pairs via half-translations, with at most finitely many points removed. The Euclidean metric descends to a metric on Φ , which is flat except perhaps at finitely many singularities where the cone angle is a positive integer multiple of π . We require that there be no π -angle cone points, i.e., if such singularities arise, they should be removed. This is to make the surface non-positively curved.

A *half-translation structure on S* is an orientation-preserving homeomorphism $f : S \rightarrow \Phi$ where Φ is a half-translation surface. Two half-translation structures $f : S \rightarrow \Phi$ and $g : S \rightarrow \Psi$ are equivalent if there is an isometry $h : \Phi \rightarrow \Psi$ homotopic to $g \circ f^{-1}$ which preserves the horizontal direction.

There is a natural projection π from the space $\mathcal{QD}(S)$ of half-translation structures on S to $\mathcal{T}(S)$ since half-translation structures are in particular complex structures. A *half-translation structure on a Riemann surface $X \in \mathcal{T}(S)$* is one that projects to X under π . The set $\pi^{-1}(X)$ of half-translation structures on X is in bijection with the set of

non-zero integrable holomorphic quadratic differentials on X . Given a quadratic differential q on X , one obtains half-translation charts by integrating the 1-form \sqrt{q} . Conversely, given a half-translation structure, the differential dz^2 in \mathbb{C} descends to a holomorphic quadratic differential on the underlying Riemann surface. See [Str84] for the definition and basic properties of quadratic differentials. We will switch back and forth between the two terminologies as convenient.

Teichmüller flow. The group $\mathrm{GL}^+(2, \mathbb{R})$ of orientation-preserving linear automorphisms of \mathbb{R}^2 acts on $\mathcal{QD}(S)$ since it conjugates the group of half-translations to itself. For every $t \in \mathbb{R}$, the linear map

$$\mathcal{G}_t = \begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$$

is e^{2t} -quasiconformal. The action of the diagonal subgroup $\{\mathcal{G}_t \mid t \in \mathbb{R}\}$ on $\mathcal{QD}(S)$ is called the *Teichmüller flow*. A *Teichmüller line* is the projection to $\mathcal{T}(S)$ of the \mathcal{G}_t -orbit of a half-translation structure Φ , parametrized by $t \mapsto \pi(\mathcal{G}_t\Phi)$.

Teichmüller proved that every Teichmüller line is a distance-minimizing geodesic for the Teichmüller distance. He also proved that through any two distinct points in $\mathcal{T}(S)$ passes a unique Teichmüller line.

Extremal length. A *conformal metric* on a Riemann surface X is a Borel measurable function $\rho : TX \rightarrow \mathbb{R}_{\geq 0}$ such that $\rho(\lambda v) = |\lambda|\rho(v)$ for every $\lambda \in \mathbb{C}$ and every tangent vector $v \in TX$. In other words, it is a choice of scale at each point.

Let Γ be a set of 1-manifolds in a Riemann surface X . The length of the set Γ with respect to a conformal metric ρ is

$$\ell_\rho(\Gamma) = \ell(\Gamma, \rho) = \inf_{\gamma \in \Gamma} \int_\gamma \rho$$

and the area of ρ is $\int_X \rho^2$. The *extremal length* of Γ in X is defined as

$$(2.1) \quad \mathrm{EL}(\Gamma, X) = \sup_\rho \frac{\ell_\rho(\Gamma)^2}{\mathrm{area}(\rho)}$$

where the supremum is over all conformal metrics ρ of finite positive area.

Typically, one takes Γ to be the free homotopy class of a simple closed curve α in X . We will abuse notation and write length or extremal length of a curve to mean the length or extremal length of its homotopy class. The basic example is when X is an upright Euclidean cylinder of circumference c and height h , and α is the curve wrapping once around X . In this case, the optimal metric ρ is the Euclidean one and we get that $\mathrm{EL}(\alpha, X) = c/h$. We will write $\mathrm{EL}(X)$ instead of $\mathrm{EL}(\alpha, X)$ since the core curve α is unique up to homotopy.

Pulling-back metrics shows that extremal length does not increase under conformal embeddings. Thus if X is any Riemann surface and $C \subset X$ is an embedded cylinder, then $\text{EL}(C) \geq \text{EL}(\alpha, X)$ where α is the core curve in C . If X is analytically finite and α is *essential*, meaning that it is not homotopic to a point or a puncture in X , then the equality $\text{EL}(C) = \text{EL}(\alpha, X)$ is achieved for a unique embedded annulus $C \subset X$ homotopic to α . Furthermore, there exists a unique half-translation structure $\Phi \in \pi^{-1}(X)$ in which C is an upright Euclidean cylinder of height 1 and the equality $\text{EL}(\alpha, X) = \ell_\rho(\alpha)^2 / \text{area}(\rho)$ holds if and only if ρ is equal almost everywhere to a positive constant multiple of the Euclidean metric on Φ . These results are due to Jenkins [Jen57].

Let $\mathcal{C}_0(S)$ be the 0-skeleton of the curve complex of S , i.e., the set of homotopy classes of essential simple closed curves in S . Kerckhoff's formula [Ker80] states that for any two points $X, Y \in \mathcal{T}(S)$ we have

$$(2.2) \quad d(X, Y) = \sup_{\alpha \in \mathcal{C}_0(S)} \frac{1}{2} \log \frac{\text{EL}(\alpha, Y)}{\text{EL}(\alpha, X)}.$$

That the Teichmüller distance is at least as large as the right-hand side follows from the fact that extremal length does not increase by more than a factor K under K -quasiconformal homeomorphisms. This property (for all sets of curves Γ) is often taken as the definition of quasiconformal maps.

Multicurves. A (*weighted*) *multicurve* in S is a formal positive linear combination of essential simple closed curves on S that are pairwise disjoint and pairwise homotopically distinct. The set of homotopy classes of multicurves in S will be denoted $\mathbb{R}_+ \times \mathcal{C}(S)$; it is the cartesian product of \mathbb{R}_+ with the curve complex.

The length of a multicurve $\alpha = \sum_{j \in J} w_j \alpha_j$ with respect to a conformal metric ρ is the weighted sum of the lengths of its components:

$$\ell_\rho(\alpha) = \sum_{j \in J} w_j \ell_\rho(\alpha_j).$$

The definition 2.1 of extremal length as a supremum of length squared divided by area then extends verbatim to multicurves.

Given $X \in \mathcal{T}(S)$ and a multicurve $\alpha = \sum_{j \in J} w_j \alpha_j$, we also have

$$(2.3) \quad \text{EL}(\alpha, X) = \inf \sum_{j \in J} w_j^2 \text{EL}(C_j)$$

where the infimum is taken over all collections of cylinders C_j embedded conformally and disjointly in X , with C_j homotopic to α_j [KPT15]. Again, the infimum is achieved by a unique collection of cylinders C_j and there is a half-translation structure on X in which each C_j is foliated by horizontal trajectories and has height w_j [Ren76]. Such a half-translation structure obtained by gluing cylinders along their boundaries is known as a *Jenkins-Strebel differential*.

There is a topology on weighted multicurves defined using intersection numbers. With respect to this topology, weighted simple closed curves are dense, and for every $X \in \mathcal{T}(S)$ the map $\alpha \mapsto \text{EL}(\alpha, X)$ is continuous. In fact, Kerckhoff showed that this map extends continuously to all measured foliations [Ker80].

3. Horoballs

The goal of this section is to rephrase the problem of the convexity of balls in terms of extremal length. To this end, we look at sublevel sets of extremal length functions, which we call *horoballs*.

Definition 3.1. Given $\alpha \in \mathcal{C}_0(S)$ and $c > 0$, we define the associated *open horoball* as

$$H(\alpha, c) = \{ X \in \mathcal{T}(S) : \text{EL}(\alpha, X) < c \}$$

and the associated *closed horoball* as

$$\overline{H}(\alpha, c) = \{ X \in \mathcal{T}(S) : \text{EL}(\alpha, X) \leq c \}.$$

REMARK. One can define horoballs for any measured foliation. We emphasize that we only consider horoballs associated with simple closed curves here.

Note that the closure of an open horoball is the corresponding closed horoball, and the interior of a closed horoball is the corresponding open horoball. This follows from the fact that the extremal length of a curve α is continuous in the second variable and does not have local minima in $\mathcal{T}(S)$. Indeed, every point $X \in \mathcal{T}(S)$ lies on a geodesic along which the extremal length of α increases exponentially, given by the Jenkins-Strebel differential on X with a single cylinder homotopic to α . In fact, the boundary of any horoball is a hypersurface in $\mathcal{T}(S)$ homeomorphic to Euclidean space [GM91].

It follows directly from Kerckhoff's formula that closed balls are intersections of closed horoballs.

Lemma 3.2. *Every closed ball in $\mathcal{T}(S)$ is a countable intersection of closed horoballs.*

Proof. Let $X \in \mathcal{T}(S)$, let $r \geq 0$ and let $\overline{B}(X, r)$ be the closed ball of radius r centered at X . By equation 2.2, $d(X, Y) \leq r$ if and only if

$$\text{EL}(\alpha, Y) \leq e^{2r} \text{EL}(\alpha, X)$$

for every $\alpha \in \mathcal{C}_0(S)$, which shows that

$$\overline{B}(X, r) = \bigcap_{\alpha \in \mathcal{C}_0(S)} \overline{H}(\alpha, e^{2r} \text{EL}(\alpha, X)).$$

q.e.d.

In hyperbolic space, horoballs are limits of larger and larger balls with centers escaping to infinity. More precisely, any open horoball is the union of all the open balls that share a given normal vector. The same description holds in Teichmüller space.

Lemma 3.3. *Every open horoball in $\mathcal{T}(S)$ is a nested union of open balls.*

Proof. Let $\alpha \in \mathcal{C}_0(S)$ and let $c > 0$. Pick an arbitrary point $X \in \partial H(\alpha, c)$ and consider the geodesic X_t defined by the half-translation structure Φ on X in which almost all *vertical* trajectories are homotopic to α , so that $\text{EL}(\alpha, X_t) = e^{-2t}c$. We will show that

$$H(\alpha, c) = \bigcup_{t>0} B(X_t, t).$$

If $d(X_t, Y) < t$ then

$$\text{EL}(\alpha, Y) < e^{2t} \text{EL}(\alpha, X_t) = c$$

by Kerckhoff's formula, which shows that $B(X_t, t) \subset H(\alpha, c)$ for every $t > 0$. The triangle inequality implies that $B(X_s, s) \subset B(X_t, t)$ whenever $0 < s < t$, i.e., the union is nested.

Let $Y \in H(\alpha, c)$ and let $b = \text{EL}(\alpha, Y)$. We need to show that $Y \in B(X_t, t)$ when t is large enough, which amounts to constructing a K -quasiconformal homeomorphism between Y and X_t with $K < e^{2t}$. The construction is essentially the same as the one in [Mas75] showing that certain geodesic rays in Teichmüller space stay a bounded distance apart.

Let Y_s be the geodesic through Y corresponding to the half-translation structure in which almost all vertical trajectories are homotopic to α , but with the time parameter shifted so that $\text{EL}(\alpha, Y_s) = e^{-2s}c$. Then fix an $s < 0$ such that

$$\frac{c}{b} - e^{2s} > 1.$$

Let α_Y (respectively α_X) be the closed vertical trajectory in the half-translation structure on Y_s (respectively X_s) that splits the α -cylinder in two equal parts. By [Mas75, Lemma 2], there exists a quasiconformal homeomorphism $f : Y_s \rightarrow X_s$ that respects the markings and sends α_Y isometrically onto α_X . Let L be the quasiconformal dilatation of f .

We rescale the flat metric on Y and Y_s so that the circumference of the α -cylinder is 1. After rescaling, the Teichmüller map $Y_s \rightarrow Y$ becomes a horizontal stretch by some factor bigger than 1. Now Y_s is a cylinder with boundary identifications, and stretching a cylinder lengthwise is the same as cutting it in the middle and inserting another piece of cylinder to make it longer. In other words, Y can be obtained by cutting Y_s open along the core curve α_Y and gluing back a cylinder

of modulus¹

$$\frac{1}{\text{EL}(\alpha, Y)} - \frac{1}{\text{EL}(\alpha, Y_s)} = \frac{1}{b} - \frac{e^{2s}}{c}$$

without twisting. Similarly X_t can be obtained from X_s by inserting a cylinder of modulus

$$\frac{1}{\text{EL}(\alpha, X_t)} - \frac{1}{\text{EL}(\alpha, X_s)} = \frac{e^{2t}}{c} - \frac{e^{2s}}{c}$$

in the middle.

From this cut-and-paste decomposition of Y and X_t , we can define a marking-preserving homeomorphism $g_t : Y \rightarrow X_t$ by using f on the complement of the middle cylinder and using the horizontal stretch map of magnitude

$$\frac{e^{2t} - e^{2s}}{\frac{c}{b} - e^{2s}} < e^{2t} - e^{2s} < e^{2t}$$

on the middle cylinder. Then g_t is K_t -quasiconformal where

$$K_t = \max \left\{ \frac{e^{2t} - e^{2s}}{\frac{c}{b} - e^{2s}}, L \right\}.$$

If t is large enough, then $L < e^{2t}$ and hence $d(X_t, Y) \leq \frac{1}{2} \log K_t < t$.

q.e.d.

The previous lemmata imply that the convexity of balls is equivalent to the convexity of horoballs.

Theorem 3.4. *The following are equivalent in $\mathcal{T}(S)$:*

- every closed horoball is convex;
- every closed ball is convex;
- every open ball is convex;
- every open horoball is convex.

Proof. If every closed horoball is convex, then every closed ball is convex by Lemma 3.2, since an arbitrary intersection of convex sets is convex.

If every closed ball is convex, then every open ball is convex. Indeed, if an open ball is not convex, then there is a smaller closed ball with the same center which is non-convex.

If every open ball is convex, then every open horoball is convex by Lemma 3.3, since nested unions of convex sets are convex.

Suppose that a closed horoball is non-convex. Then the open horoballs of slightly higher level for the same simple closed curve are also non-convex. Thus if every open horoball is convex, then every closed horoball is convex. q.e.d.

¹The modulus of a Euclidean cylinder is the reciprocal of its extremal length, i.e., the distance between its boundary components once it has been rescaled to have circumference 1.

Therefore, to show the existence of a non-convex ball, we need to find a non-convex horoball. More explicitly, we need to find a simple closed curve $\alpha \in \mathcal{C}_0(S)$ and three points $X, Y, Z \in \mathcal{T}(S)$ appearing in that order along a geodesic such that the extremal length of α in Y is strictly larger than in both X and Z . We can weaken this criterion by allowing α to be a multicurve and replacing the 3-point condition by a 4-point condition. This will be useful later.

Lemma 3.5. *Suppose that there exists four points X, Y, Z, W appearing in that order along a geodesic in $\mathcal{T}(S)$ and a multicurve $\alpha \in \mathbb{R}_+ \times \mathcal{C}(S)$ such that*

$$\text{EL}(\alpha, X) < \text{EL}(\alpha, Y) \quad \text{and} \quad \text{EL}(\alpha, Z) > \text{EL}(\alpha, W).$$

Then there exists a non-convex ball in $\mathcal{T}(S)$.

Proof. Since extremal length depends continuously on the first variable and since weighted simple closed curves are dense in $\mathbb{R}_+ \times \mathcal{C}(S)$, there exists a weighted simple closed curve $w\beta$ such that

$$\text{EL}(w\beta, X) < \text{EL}(w\beta, Y) \quad \text{and} \quad \text{EL}(w\beta, Z) > \text{EL}(w\beta, W).$$

As extremal length is homogeneous of degree 2 in the first variable, we also have

$$\text{EL}(\beta, X) < \text{EL}(\beta, Y) \quad \text{and} \quad \text{EL}(\beta, Z) > \text{EL}(\beta, W).$$

Let $c = \max(\text{EL}(\beta, X), \text{EL}(\beta, W))$. Then the geodesic segment from X to W has its endpoints in $\overline{H}(\beta, c)$ and passes through Y and Z . At least one of Y or Z lies outside $\overline{H}(\beta, c)$, so that the closed horoball $\overline{H}(\beta, c)$ is non-convex. By Theorem 3.4, this implies the existence of a non-convex ball. q.e.d.

If extremal length of a multicurve has a strict local maximum along a geodesic, then we can clearly find 4 points satisfying the hypotheses of the above lemma, so that non-convex balls exist. It turns out that every local maximum is a strict local maximum (see below), which implies Lemma 1.2. In practice, the hypotheses of Lemma 3.5 are easier to check than the existence of a local maximum, but Lemma 1.2 was simpler to state for the introduction.

Lemma 3.6. *Any local maximum of extremal length along a geodesic is a strict local maximum.*

Proof. Suppose there is a multicurve $\alpha \in \mathbb{R}_+ \times \mathcal{C}(S)$ and a geodesic $t \mapsto X_t$ in $\mathcal{T}(S)$ such that the function $f(t) = \text{EL}(\alpha, X_t)$ has a local maximum at T and let $M = f(T)$. If the local maximum is not strict, then $f^{-1}(M)$ accumulates at T . Since f is real-analytic [Miy15], it must be constant by the identity principle. But this is impossible. Indeed, take ρ_t to be the Euclidean metric defining the geodesic X_t . Then the area of ρ_t is constant, whereas the length $\ell(\alpha, \rho_t)$ is unbounded in

at least one direction. Indeed, $\ell(\alpha, \rho_t)$ is bounded below by the intersection number between α and the vertical foliation on X_t as well as by the intersection number with the horizontal foliation. These intersection numbers depend exponentially on t and at least one of them is non-zero, so that it diverges as $t \rightarrow \infty$ or as $t \rightarrow -\infty$. Since $f(t) \geq \ell(\alpha, \rho_t)^2 / \text{area}(\rho_t)$, this is a contradiction. q.e.d.

Rather than exhibiting 4 collinear points verifying the inequalities of Lemma 3.5, we will construct a sequence of collinear points X_n, Y_n, Z_n and W_n that degenerate in a controlled way as $n \rightarrow \infty$ and such that the desired inequalities hold in the limit. We thus need to show that extremal length behaves well under mild degeneration.

4. Convergence of extremal length under pinching

Let $R = \sqcup_{j \in J} R_j$ be a subsurface of S obtained by cutting S along a multicurve and possibly forgetting some of the pieces. Each connected component R_j of R is homeomorphic to a punctured surface R'_j . The Teichmüller space $\mathcal{T}(R)$ is defined as the Cartesian product $\prod_{j \in J} \mathcal{T}(R'_j)$.

Definition 4.1. Let $X_n \in \mathcal{T}(S)$ and $Y \in \mathcal{T}(R)$. We say that X_n converges conformally to Y as $n \rightarrow \infty$ if there exist nested surfaces $Y_n \subset Y$ exhausting Y and K_n -quasiconformal embeddings $f_n : Y_n \rightarrow X_n$ homotopic to the inclusion map $R \subset S$ such that $K_n \rightarrow 1$ as $n \rightarrow \infty$.

We emphasize that in this definition, the ends of Y are all required to be punctures. Informally speaking, X_n converges conformally to Y if there is a multicurve in X_n that gets pinched and in the process, some pieces survive to form Y . We don't care about the other pieces, meaning that they don't need to stabilize as $n \rightarrow \infty$. Thus conformal convergence in the above sense is more general than convergence in the augmented Teichmüller space, where every piece is required to stabilize [Abi77].

There are different equivalent ways to formulate conformal convergence. We can say that X_n converges conformally to Y if

- for every simple closed curve α in R (including peripheral ones), the hyperbolic length of α in X_n converges to the hyperbolic length of α in Y ;
- for every $j \in J$, the covering space of X_n associated with the subsurface R_j , equipped with its hyperbolic metric, converges in the Gromov-Hausdorff topology to the corresponding component Y_j of Y with respect to some choices of basepoints;
- for every $j \in J$, ρ_n^j converges up to conjugacy to ρ^j , where $\rho_n : \pi_1(S) \rightarrow \text{PSL}(2, \mathbb{R})$ is the representation defining X_n , ρ_n^j is its restriction to $\pi_1(R_j)$ coming from the inclusion $R_j \subset S$ and $\rho_j : \pi_1(R_j) \rightarrow \text{PSL}(2, \mathbb{R})$ is the representation defining Y_j .

For our purposes, the definition in terms of nearly conformal embeddings is the most convenient. The statement we want to prove is that conformal convergence implies convergence of extremal length for multicurves supported on the limiting surface. Extremal length on a disconnected surface is defined in the usual way, as the supremum of weighted length squared divided by area over all conformal metrics. A standard argument shows this equals the sum of the extremal lengths on connected components (cf. [Ahl10, p. 55]).

Lemma 4.2. *Let $Y = \sqcup_{j \in J} Y_j$ be a disjoint union of Riemann surfaces, let α be a multicurve on Y and let $\alpha^j = \alpha \cap Y_j$. Then*

$$\text{EL}(\alpha, Y) = \sum_{j \in J} \text{EL}(\alpha^j, Y_j).$$

Proof. First observe that the extremal length of α on Y is the same as its extremal length on the union Z of the components which it intersects. Indeed, given a metric ρ on Y , the ratio $\ell(\alpha, \rho)^2 / \text{area}(\rho)$ does not decrease if we modify ρ to be zero outside Z . This shows that $\text{EL}(\alpha, Y) \leq \text{EL}(\alpha, Z)$ and the reverse inequality follows by extending any conformal metric on Z to be zero on $Y \setminus Z$.

If α^j is empty, then clearly $\text{EL}(\alpha^j, Y_j) = 0$. In proving the above formula, we may therefore assume that α^j is non-empty for each $j \in J$.

For each $j \in J$, let ρ_j be any metric on Y_j such that $\ell(\alpha^j, \rho_j)$ and $\text{area}(\rho_j)$ are finite and positive. By rescaling ρ_j , we may assume that $\ell(\alpha^j, \rho_j) = \text{area}(\rho_j)$. Let ρ be the metric on Y which is equal to ρ_j on Y_j . Then

$$\ell(\alpha, \rho) = \sum_{j \in J} \ell(\alpha^j, Y_j) = \sum_{j \in J} \text{area}(\rho_j) = \text{area}(\rho)$$

which implies that

$$\text{EL}(\alpha, Y) \geq \frac{\ell(\alpha, \rho)^2}{\text{area}(\rho)} = \sum_{j \in J} \frac{\ell(\alpha, \rho_j)^2}{\text{area}(\rho_j)}.$$

We can replace the right-hand side by its supremum over all non-degenerate metrics ρ_j to get

$$\text{EL}(\alpha, Y) \geq \sum_{j \in J} \text{EL}(\alpha^j, Y_j).$$

Conversely, let σ be any conformal metric on Y and let σ_j be its restriction to Y_j . Then for each $j \in J$ we have

$$\text{EL}(\alpha^j, Y_j) \geq \frac{\ell(\alpha^j, \sigma_j)^2}{\text{area}(\sigma_j)}.$$

Summing over all j yields

$$\sum_{j \in J} \text{EL}(\alpha^j, Y_j) \geq \sum_{j \in J} \frac{\ell(\alpha^j, \sigma_j)^2}{\text{area}(\sigma_j)} \geq \frac{\left(\sum_{j \in J} \ell(\alpha^j, \sigma_j) \right)^2}{\sum_{j \in J} \text{area}(\sigma_j)}$$

where the second inequality follows from the Cauchy-Schwarz inequality. Finally, observe that

$$\sum_{j \in J} \ell(\alpha^j, \sigma_j) = \ell(\alpha, \sigma) \quad \text{and} \quad \sum_{j \in J} \text{area}(\sigma_j) = \text{area}(\sigma)$$

so that

$$\sum_{j \in J} \text{EL}(\alpha^j, Y_j) \geq \frac{\ell(\alpha, \sigma)^2}{\text{area}(\sigma)}.$$

Since the inequality holds for any conformal metric σ , it holds for the supremum as well and we have

$$\sum_{j \in J} \text{EL}(\alpha^j, Y_j) \geq \text{EL}(\alpha, Y).$$

q.e.d.

This lemma implies that equation 2.3, which says that extremal length is the infimum of weighted sums of extremal lengths of embedded cylinders, still holds for disconnected surfaces. We use this in the proof of convergence of extremal length under degeneration.

Theorem 4.3. *Let α be a multicurve in R and suppose that X_n converges conformally to Y as $n \rightarrow \infty$. Then $\text{EL}(\alpha, X_n) \rightarrow \text{EL}(\alpha, Y)$ as $n \rightarrow \infty$.*

Proof. Let $K > 1$. We will show that if n is large enough, then

$$\frac{1}{K^2} \text{EL}(\alpha, Y) \leq \text{EL}(\alpha, X_n) \leq K^2 \text{EL}(\alpha, Y),$$

starting with the second inequality.

Write α as a weighted sum of simple closed curves $\sum_{i \in I} w_i \alpha_i$ and let $C = \sqcup_{i \in I} C_i$ be the collection of cylinders in Y such that $\text{EL}(\alpha, Y) = \sum_{i \in I} w_i^2 \text{EL}(C_i)$. For each $i \in I$, let $A_i \subset C_i$ be a compactly contained essential cylinder (for example a straight subcylinder) such that $\text{EL}(A_i) \leq K \text{EL}(C_i)$.

Let Y_n be a nested exhaustion of Y and let $f_n : Y_n \rightarrow X_n$ be quasi-conformal embeddings as in Definition 4.1. If n is large enough, then Y_n contains $\cup_{i \in I} A_i$ and f_n is K -quasiconformal. Then by equation 2.3

we have

$$\begin{aligned} \text{EL}(\alpha, X_n) &\leq \sum_{i \in I} w_i^2 \text{EL}(f_n(A_i)) \\ &\leq \sum_{i \in I} w_i^2 K \text{EL}(A_i) \\ &\leq K^2 \sum_{i \in I} w_i^2 \text{EL}(C_i) = K^2 \text{EL}(\alpha, Y). \end{aligned}$$

For the reverse inequality, let ρ be the conformal metric realizing $\text{EL}(\alpha, Y)$. Our goal is to construct a good enough conformal metric on X_n from ρ .

We may assume that α intersects every component of Y (otherwise ignore the superfluous components). By the previous lemma and Renelt's theorem [Ren76], $\rho = \sqrt{|q|}$ for a holomorphic quadratic differential q on Y with at most simple poles at the punctures. Let \bar{Y} be the completion of Y in the metric ρ and let $Q = \bar{Y} \setminus Y$ be the set of punctures of Y . Since Q is finite, there exists a $\delta_0 > 0$ such that the δ_0 -balls around the points of Q are embedded and pairwise disjoint. Here we are using the fact that ρ has isolated singularities so that the distance it induces on \bar{Y} via path integrals is compatible with the underlying topology.

Let $\mu = 2\delta_0$. Then any homotopically non-trivial arc from Q to itself has ρ -length at least μ . For $\delta > 0$, let $\mathcal{N}_\delta(Q)$ be the open δ -neighborhood of Q in the metric ρ and let ρ_δ be the metric on Y which agrees with ρ outside of $\mathcal{N}_\delta(Q)$ and is identically zero on $\mathcal{N}_\delta(Q) \cap Y$.

Claim 4.4. *For every $\beta \in \mathcal{C}_0(R)$ the length $\ell(\beta, \rho_\delta)$ converges to $\ell(\beta, \rho)$ as $\delta \rightarrow 0$.*

Proof of claim. It is clear that $\ell(\beta, \rho_\delta) \leq \ell(\beta, \rho)$ since $\rho_\delta \leq \rho$. We will show that

$$\ell(\beta, \rho) \leq \frac{\mu + \delta}{\mu - 2\delta} \ell(\beta, \rho_\delta)$$

whenever $\delta < \delta_0$.

Let γ be a piecewise smooth curve homotopic to β on Y . Our goal is to find a curve $\tilde{\gamma}$ homotopic to γ such that

$$|\tilde{\gamma}| \leq \frac{\mu + \delta}{\mu - 2\delta} \left| \gamma \setminus \overline{\mathcal{N}_\delta(Q)} \right|$$

where the length $|\cdot|$ is measured with respect to ρ .

We modify γ in two steps, each time making it shorter outside $\overline{\mathcal{N}_\delta(Q)}$. We estimate its total length at the end. To avoid unnecessary notation, we denote the modified curve by γ again instead of $\tilde{\gamma}$.

We may assume that each component of $\gamma \setminus \overline{\mathcal{N}_\delta(Q)}$ is homotopically non-trivial rel endpoints after collapsing each component of $\mathcal{N}_\delta(Q)$. Otherwise, we can homotope those trivial subarcs to the boundary of

$\overline{\mathcal{N}_\delta(Q)}$, which shortens γ . We can also assume that $\gamma \setminus \overline{\mathcal{N}_\delta(Q)}$ has only finitely many components, because each such component has length at least $\mu - 2\delta$. If $\gamma \setminus \overline{\mathcal{N}_\delta(Q)}$ has infinite length, then there is nothing to show. Thus $\gamma \cap \overline{\mathcal{N}_\delta(Q)}$ has finitely many components as well. Each such component can be homotoped, keeping the endpoints fixed, to a path in $\overline{\mathcal{N}_\delta(Q)}$ of ρ -length at most 3δ . This does not change the length of the portion of γ lying outside $\overline{\mathcal{N}_\delta(Q)}$.

If γ is disjoint from $\overline{\mathcal{N}_\delta(Q)}$, then $|\gamma| = \left| \gamma \setminus \overline{\mathcal{N}_\delta(Q)} \right|$ and we are done.

Otherwise γ breaks up into $\gamma \setminus \overline{\mathcal{N}_\delta(Q)}$ and $\gamma \cap \overline{\mathcal{N}_\delta(Q)}$ and these two sets have the same number of components. Let σ be a component of $\gamma \setminus \overline{\mathcal{N}_\delta(Q)}$ and τ a component of $\gamma \cap \overline{\mathcal{N}_\delta(Q)}$. Then $|\sigma| \geq \mu - 2\delta$ and $|\tau| \leq 3\delta$, so that

$$|\sigma| + |\tau| \leq \frac{\mu + \delta}{\mu - 2\delta} |\sigma|.$$

Adding these inequalities over distinct σ - τ pairs whose union is γ yields

$$|\gamma| \leq \frac{\mu + \delta}{\mu - 2\delta} \left| \gamma \setminus \overline{\mathcal{N}_\delta(Q)} \right|,$$

which shows that

$$\ell(\beta, \rho) \leq \frac{\mu + \delta}{\mu - 2\delta} \left| \gamma \setminus \overline{\mathcal{N}_\delta(Q)} \right|.$$

Since γ was arbitrary, the right-hand side can be replaced with the infimum $\frac{\mu + \delta}{\mu - 2\delta} \ell(\beta, \rho_\delta)$. q.e.d.

The analogous result for multicurves follows by linearity. Thus there exists a $\delta > 0$ such that

$$\ell(\alpha, \rho_\delta)^2 \geq \frac{1}{K} \ell(\alpha, \rho)^2.$$

We fix such a δ for the rest of the proof.

Let n be large enough so that Y_n contains $Y \setminus \mathcal{N}_\delta(Q)$ and so that $f_n : Y_n \rightarrow X_n$ is K -quasiconformal. We may assume that f_n is smooth except at finitely many points.

We define a conformal metric ρ_n on X_n by

$$\rho_n(v) = \max_{\theta \in [0, 2\pi]} \rho(df_n^{-1}(e^{i\theta}v))$$

if v is a tangent vector based at a point in $f_n(Y \setminus \mathcal{N}_\delta(Q))$ and $\rho_n(v) = 0$ otherwise.

Claim 4.5. *For every $\beta \in \mathcal{C}_0(R)$, we have $\ell(\beta, \rho_n) \geq \ell(\beta, \rho_\delta)$.*

Proof of claim. Let γ be a curve homotopic to β in X_n . If γ is contained in the image $f_n(Y \setminus \mathcal{N}_\delta(Q))$, then

$$\begin{aligned}
\int \rho_n(\gamma'(t)) |dt| &\geq \int \rho((df_n^{-1}(\gamma'(t))) |dt| \\
&= \int \rho((f_n^{-1} \circ \gamma)'(t)) |dt| \\
&\geq \ell(\beta, \rho) \\
&\geq \ell(\beta, \rho_\delta).
\end{aligned}$$

Otherwise, we can homotope γ to a curve $\tilde{\gamma}$ which is not longer, yet is contained in $f_n(Y \setminus \mathcal{N}_\delta(Q))$. q.e.d.

Once again, the analogous result for multicurves on R follows immediately.

Claim 4.6. *The areas satisfy $\text{area}(\rho_n) \leq K \text{area}(\rho)$.*

Proof of claim. Since f_n is K -quasiconformal, we have

$$\max_{\theta \in [0, 2\pi]} \rho(df_n^{-1}(e^{i\theta}v)) \leq K \min_{\theta \in [0, 2\pi]} \rho(df_n^{-1}(e^{i\theta}v)).$$

This shows that $\rho_n^2 \leq K(f_n)_*\rho^2$ on $f_n(Y \setminus \mathcal{N}_\delta(Q))$, which means that

$$\begin{aligned}
\text{area}(\rho_n) &= \int_{X_n} \rho_n^2 \\
&\leq K \int_{f_n(Y \setminus \mathcal{N}_\delta(Q))} (f_n)_*\rho^2 \\
&= K \int_{Y \setminus \mathcal{N}_\delta(Q)} \rho^2 \\
&\leq K \text{area}(\rho).
\end{aligned}$$

q.e.d.

Combining the above inequalities yields

$$\text{EL}(\alpha, X_n) \geq \frac{\ell(\alpha, \rho_n)^2}{\text{area}(\rho_n)} \geq \frac{1}{K} \frac{\ell(\alpha, \rho_\delta)^2}{\text{area}(\rho)} \geq \frac{1}{K^2} \frac{\ell(\alpha, \rho)^2}{\text{area}(\rho)} = \frac{1}{K^2} \text{EL}(\alpha, Y)$$

for all large enough n , which is what we wanted to show. q.e.d.

We now possess all the necessary tools to construct Teichmüller geodesics along which extremal length of a multicurve increases first and then decreases later.

5. The examples

5.1. The sphere with seven punctures. Given non-negative lengths l_1, \dots, l_n and heights h_1, \dots, h_n , consider the staircase-shaped polygon $P(l_1, h_1, \dots, l_n, h_n) \subset \mathbb{R}^2$ with j -th step of length l_j and height h_j . This is illustrated in Figure 1 for $n = 4$. We allow either l_1 or h_n to be infinite, in which case P is a horizontal or vertical half-infinite strip

ending in a staircase. If all lengths and heights are finite, then we put the bottom-left corner of P at the origin.

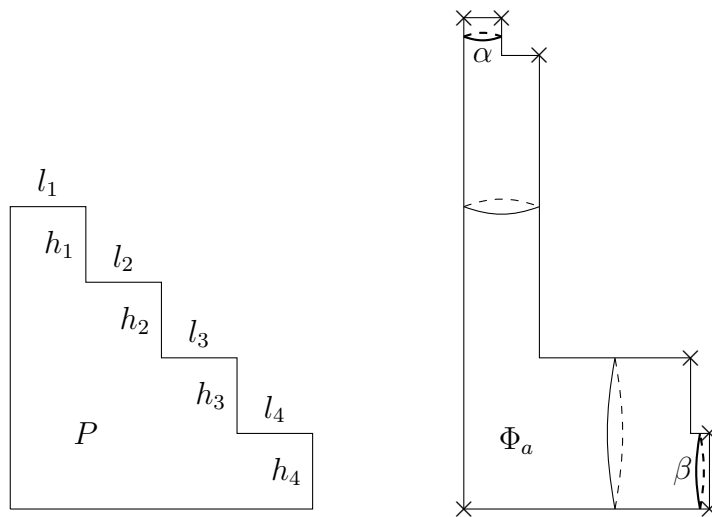


Figure 1. The staircase-shaped polygon P and the half-translation surface Φ_a .

For $0 < a < 1$, let

$$P_a = P(1, 1, 1, 1/a^3, 1/a^2, 1/a, a, 1/a)$$

and let $\bar{\Phi}_a$ be the double of P_a across its boundary. More precisely, take P_a and its image P_a^* under a horizontal reflection σ and glue them along their boundary using σ . It helps to think of P_a as the front of $\bar{\Phi}_a$ and P_a^* as the back. The sphere $\bar{\Phi}_a$ has 7 cone points of angle π (coming from the 7 interior right angles in P_a), which we remove in order to get the half-translation surface Φ_a .

Let α be the simple closed curve that separates the two highest punctures on Φ_a from the rest, i.e., the double of the middle horizontal line in the highest step of the staircase P_a . Similarly, let β be the double of the middle vertical line in the step furthest to the right in P_a (see Figure 1).

As we apply Teichmüller flow to the half-translation surface Φ_a , the extremal length of α increases rapidly at first and then stays nearly constant for a long time. The extremal length of β does the opposite: it remains nearly constant for a long time then decreases rapidly. Also, since α and β are separated by cylinders of very large modulus, the extremal length of $\alpha + \beta$ is roughly equal to the sum of the individual extremal lengths. The net effect is that $\text{EL}(\alpha + \beta, \mathcal{G}_t \Phi_a)$ increases at first and decreases some time later.

Let us be more precise. Consider the points

$$\begin{aligned} X_a &= \begin{pmatrix} 1/e & 0 \\ 0 & e \end{pmatrix} \Phi_a & Y_a &= \Phi_a \\ Z_a &= \begin{pmatrix} 1/a & 0 \\ 0 & a \end{pmatrix} \Phi_a & W_a &= \begin{pmatrix} e/a & 0 \\ 0 & a/e \end{pmatrix} \Phi_a. \end{aligned}$$

appearing at times -1 , 0 , $\log(1/a)$ and $\log(1/a) + 1$ along the Teichmüller line $t \mapsto \mathcal{G}_t \Phi_a$.

Observe that $Z_a = \tau Y_a$ and $W_a = \tau X_a$, where τ is the reflection about the line $y = x$ in the plane. Indeed, the definition of P_a was arranged so that

$$\begin{pmatrix} 1/a & 0 \\ 0 & a \end{pmatrix} P_a = P(1/a, a, 1/a, 1/a^2, 1/a^3, 1, 1, 1) = \tau P_a.$$

We claim that from the point of view of the multicurve $\alpha + \beta$, the surfaces X_a, Y_a, Z_a and W_a all have conformal limits as $a \rightarrow 0$. Let Ψ^α be the double of $P(1, 1, 1, \infty)$ minus the π -angle singularities, let Ψ^β be the double of $P(\infty, 1, 0, 1)$ minus the three distinguished vertices, and let $\Psi = \Psi^\alpha \sqcup \Psi^\beta$ (see Figure 2). We consider Ψ as a half-translation surface with infinite area. The conformal limits of X_a, Y_a, Z_a and W_a as $a \rightarrow 0$ are

$$X_0 = \mathcal{G}_{-1} \Psi, \quad Y_0 = \Psi, \quad Z_0 = \tau \Psi \quad \text{and} \quad W_0 = \tau \mathcal{G}_{-1} \Psi$$

respectively.




Figure 2. The conformal limit Ψ at time $t = 0$. Observe that Ψ^β is conformally invariant under Teichmüller flow.

Lemma 5.1. *For each $\Lambda \in \{X, Y, Z, W\}$, the surface Λ_a converges conformally to Λ_0 as $a \rightarrow 0$.*

Proof. We prove that $Y_a = \Phi_a$ converges conformally to $Y_0 = \Psi$ as $a \rightarrow 0$. The proof of conformal convergence at other times is similar.

It is clear that the top half of P_a converges to $P(1, 1, 1, \infty)$ as $a \rightarrow 0$. Indeed, for any $L > 0$ the top portion of $P(1, 1, 1, \infty)$ of height L embeds isometrically into P_a provided that $1/a^2 > L$. Moreover, the isometric embedding maps vertices to vertices. By doubling this embedding, we obtain a conformal embedding of a large portion of Ψ^α into Φ_a .

If we apply an homothety of factor a to P_a , we get the polygon

$$Q_a = P(a, a, a, 1/a^2, 1/a, 1, a^2, 1).$$

We will show that the right half of Q_a converges conformally to the unbounded polygon $P(\infty, 1, 0, 1)$. For this, we will use a theorem of Radó which says that if a sequence of parametrized Jordan curves $\gamma_n : S^1 \rightarrow \widehat{\mathbb{C}}$ converges uniformly to a Jordan curve $\gamma_\infty : S^1 \rightarrow \widehat{\mathbb{C}}$, then the corresponding (appropriately normalized) Riemann maps converge uniformly on the closed unit disk [Pom92, p.26] [Gol69, p.59].

For $a \geq 0$, let $R_a = P(\infty, 1, a^2, 1)$. For concreteness, we take the finite vertices of R_a to be located at 0 , i , $-a^2 + i$ and $-a^2 + 2i$ in \mathbb{C} . Let $\gamma_a : \mathbb{R} \cup \{\infty\} \rightarrow \widehat{\mathbb{C}}$ be the Jordan curve $\partial R_a \cup \{\infty\}$ parametrized counter-clockwise by arclength with $\gamma_a(0) = 0$ and $\gamma_a(\infty) = \infty$. Then γ_a converges uniformly to γ_0 as $a \rightarrow 0$.

Let $f_a : R_0 \rightarrow R_a$ be the unique conformal homeomorphism such that $f_a(0) = 0$, $f_a(2i) = -a^2 + 2i$ and $f_a(\infty) = \infty$. Then f_a converges uniformly to the identity by Radó's theorem. In particular, $f_a^{-1}(i) = iv_a$ converges to i as $a \rightarrow 0$. Define $g_a : R_0 \rightarrow R_0$ by

$$g_a(x + iy) = \begin{cases} x + iv_a y & \text{if } y \in [0, 1] \\ x + i((2 - v_a)(y - 1) + v_a) & \text{if } y \in (1, 2]. \end{cases}$$

This map is piecewise linear, fixes 0 , $2i$ and ∞ and sends i to $f_a(i)$. The quasiconformal dilatation of g_a is equal to $\max\left\{v_a, \frac{1}{v_a}, 2 - v_a, \frac{1}{2 - v_a}\right\}$, which tends to 1 as $a \rightarrow 0$. Thus when a is small, $f_a \circ g_a$ is a nearly conformal homeomorphism from R_0 to R_a taking 0 , i , $2i$ and ∞ to 0 , i , $-a^2 + 2i$ and ∞ respectively.

Let $L < 0$ and let $C = \{z \in R_0 \mid \operatorname{Re} z > L\}$. The images $f_a \circ g_a(C)$ stay bounded away from ∞ since $f_a \circ g_a$ converges uniformly to the identity. Let's say that $\operatorname{Re} z \geq u$ for all $z \in f_a \circ g_a(C)$ and all $a \geq 0$. Then $f_a \circ g_a(C)$ embeds isometrically in Q_a in the obvious way provided that $1/a > |u|$. By doubling all these objects and maps, we obtain a quasiconformal embedding of the double of C into Φ_a with dilatation arbitrarily close to 1 when a is small.

q.e.d.

It only remains to prove estimates for the extremal length on these limiting surfaces.

Lemma 5.2. *EL(α, X_0) and EL(β, W_0) are bounded above by $2/e^2$.*

Proof. Recall that the component of X_0 containing α is $\mathcal{G}_{-1}\Psi^\alpha$. Take the top 1×1 square in $P(1, 1, 1, \infty)$ without its horizontal sides. Its double is an open Euclidean cylinder of circumference 2 and height 1 homotopic to α in Ψ^α . This cylinder gets stretched to one of circumference $2/e$ and height e , hence extremal length $2/e^2$, under the map \mathcal{G}_{-1} . The inequality $\text{EL}(\alpha, X_0) \leq 2/e^2$ thus follows from the monotonicity of extremal length under conformal embeddings.

Now, the reflection τ maps X_0 anti-conformally onto W_0 and sends α to β so that $\text{EL}(\alpha, X_0) = \text{EL}(\beta, W_0)$. q.e.d.

Lemma 5.3. $\text{EL}(\alpha, Y_0)$ and $\text{EL}(\beta, Z_0)$ are bounded below by $2/3$.

Proof. Take ρ to be the Euclidean metric on the top part T of height 2 in the component of Y_0 containing α (this is a union of 6 unit squares, 3 in the front, 3 in the back) extended to be identically zero elsewhere. Then ρ has area 6. Moreover, any curve γ homotopic to α on Y_0 has length at least 2 in the metric ρ . If γ is not contained in T , then some point p on γ is at height less than -2 . But some point q on γ has to be at height at least -1 since it is homotopic to α (it has to cross the seam between the front and back of Y_0 joining the two punctures on the top right). In this case, the length of γ is at least twice the height difference between q and the bottom of T (because there is a subarc from p to q then from q to p), i.e., at least 2. A similar argument also applies if γ is contained in T (it then has to cross the left seam in addition to the other one). We conclude that the extremal length of α on Y_0 is at least $2^2/6 = 2/3$.

The extremal length of β on Z_0 is the same as the extremal length of α on Y_0 by symmetry. q.e.d.

Lemma 5.4. We have $\text{EL}(\beta, X_0) = \text{EL}(\beta, Y_0)$ and $\text{EL}(\alpha, Z_0) = \text{EL}(\alpha, W_0)$.

Proof. The equality $\text{EL}(\beta, X_0) = \text{EL}(\beta, Y_0)$ is due to the fact that the component $\mathcal{G}_{-1}\psi^\beta$ of X_0 containing β is conformally equivalent to the corresponding component Ψ^β of Y_0 . Indeed, recall that Ψ^β is the double of $P(\infty, 1, 0, 1)$. The image of the latter by \mathcal{G}_{-1} is $P(\infty, e, 0, e)$ which is homothetic to the first polygon by a factor e . This homothety doubles to a conformal isomorphism between Ψ^β and $\mathcal{G}_{-1}\Psi^\beta$ preserving the marked points and the curve β .

Similarly, $\text{EL}(\alpha, Z_0) = \text{EL}(\alpha, W_0)$ since the connected component of Z_0 containing α is conformally invariant under Teichmüller flow. q.e.d.

These are all the ingredients we need to prove the desired behavior for the extremal length of $\alpha + \beta$ along the geodesic $\mathcal{G}_t\Phi_a$.

Theorem 5.5. *If a is small enough, then*

$$\text{EL}(\alpha + \beta, X_a) < \text{EL}(\alpha + \beta, Y_a) \quad \text{and} \quad \text{EL}(\alpha + \beta, Z_a) > \text{EL}(\alpha + \beta, W_a).$$

Proof. For each $\Lambda \in \{X, Y, Z, W\}$ we have that

$$\text{EL}(\alpha + \beta, \Lambda_a) \rightarrow \text{EL}(\alpha + \beta, \Lambda_0)$$

as $a \rightarrow 0$ by Lemma 5.1 and Theorem 4.3. Since each Λ_0 is disconnected, we also have

$$\text{EL}(\alpha + \beta, \Lambda_0) = \text{EL}(\alpha, \Lambda_0) + \text{EL}(\beta, \Lambda_0).$$

by Lemma 4.2.

By the previous three lemmata we have

$$\begin{aligned} \text{EL}(\alpha + \beta, X_0) &= \text{EL}(\alpha, X_0) + \text{EL}(\beta, X_0) \\ &\leq \frac{2}{e^2} + \text{EL}(\beta, X_0) \\ &< \frac{2}{3} + \text{EL}(\beta, Y_0) \\ &\leq \text{EL}(\alpha, Y_0) + \text{EL}(\beta, Y_0) \\ &= \text{EL}(\alpha + \beta, Y_0) \end{aligned}$$

and

$$\begin{aligned} \text{EL}(\alpha + \beta, Z_0) &= \text{EL}(\alpha, Z_0) + \text{EL}(\beta, Z_0) \\ &\geq \frac{2}{3} + \text{EL}(\beta, Z_0) \\ &> \frac{2}{e^2} + \text{EL}(\beta, W_0) \\ &\geq \text{EL}(\alpha, W_0) + \text{EL}(\beta, W_0) \\ &= \text{EL}(\alpha + \beta, W_0). \end{aligned}$$

The analogous inequalities must hold for small enough $a > 0$ by convergence.

q.e.d.

By Lemma 3.5, this implies the existence of non-convex balls in $\mathcal{T}(S_{0,7})$, where $S_{\mathbf{g},\mathbf{p}}$ is the closed surface of genus \mathbf{g} with \mathbf{p} points removed.

5.2. Increasing the genus. We modify the above construction to get a surface of genus 1 with 4 punctures. As before, we start with the polygon $P_a = P(1, 1, 1, 1/a^3, 1/a^2, 1/a, a, 1/a)$ for $0 < a < 1$ and take a copy P_a^* of P_a with reverse orientation. We think of P_a as the front of the surface to be constructed and P_a^* as the back. We glue each the side of P_a to the corresponding side of P_a^* except for the highest two horizontal sides. Call these sides A and B and let A^* and B^* be the corresponding sides of P_a^* . Then we glue A to B^* and B to A^* to obtain $\bar{\Phi}_a$. In other words, we glue the circle $A \cup A^*$ to $B \cup B^*$ in an orientation-reversing manner but with a half-twist. This creates a handle and a singularity of angle 4π . Then we remove the 4 singularities

of angle π from $\bar{\Phi}_a$ to obtain the half-translation surface Φ_a . The curves α and β are as before. We can also cut the top squares in P_a and P_a^* along their diagonal, rotate and glue $A \cup A^*$ to $B \cup B^*$ to obtain another useful representation of Φ_a . See Figure 3.




Figure 3. Left: adding a handle to Φ_a . Right: another representation of Φ_a obtained by cut-and-paste.

As in the previous subsection, we let $X_a = \mathcal{G}_{-1}\Phi_a$, $Y_a = \Phi_a$, $Z_a = \mathcal{G}_{\log(1/a)}\Phi_a$ and $W_a = \mathcal{G}_{\log(1/a)+1}\Phi_a$. The claim is that all of these have conformal limits as $a \rightarrow 0$.

Let Υ be two copies of the polygon $P(1, 1, 1, \infty)$ glued in the same pattern as described above, i.e., as in Figure 3. Also let two copies of the polygon $P(1, 0, 1, \infty)$ with corresponding vertical sides glued together, the segment $[0, 1]$ on the front glued to $[1, 2]$ on the back, and vice versa. Denote the resulting surface Ω .

Let $Y_0 = \Upsilon \sqcup \Psi^\beta$, $X_0 = \mathcal{G}_{-1}Y_0$, $Z_0 = \Omega \sqcup \tau\Psi^\alpha$ and $W_0 = \mathcal{G}_1Z_0$, where Ψ^α , Ψ^β and τ are as in the previous subsection.

Lemma 5.6. *For each $\Lambda \in \{X, Y, Z, W\}$, the surface Λ_a converges conformally to Λ_0 as $a \rightarrow 0$.*

Proof. For each Λ , the convergence from the point of view of the bottom right subsurface containing β holds for the same reasons as before.

From the point of view of α , it is clear that Y_a converges conformally to Υ as $a \rightarrow 0$ since any compact subset of Υ eventually embeds isometrically into Y_a . Similarly, the top left portion of X_a converges conformally to $\mathcal{G}_{-1}\Upsilon$ as $a \rightarrow 0$.

The only part left to prove is that Z_a and W_a converge to Ω from the point of view of α . We prove this for Z_a , the other case being similar.

For $a \geq 0$ and $L > 0$, let

$$T_a^L = \{(x, y, \varepsilon) \in \mathbb{R}^2 \times \{+, -\} : |x| \leq 1, -L < y \leq a^2|x|\} / \sim$$

where $(x, a^2|x|, +) \sim (-x, a^2|x|, -)$ for every $x \in [-1, 1]$. This is a torus with one hole obtained by gluing two M -shapes together. Note that $T_0^\infty = \Omega$.

If we rescale Z_a by a factor a so that its left vertical chimney has circumference 4, we see that T_a^L embeds conformally into Z_a provided

that $1/a \geq L$. This uses the alternative gluing pattern for Φ_a with diagonal lines.

Consider the piecewise linear homeomorphism $f_a : T_0^\infty \rightarrow T_a^\infty$ defined by

$$f_a(x, y, \varepsilon) = (x, a^2|x| + y, \varepsilon).$$

On each piece of T_0^∞ where x and ε have constant sign, the map f_a is a vertical shear. Its dilatation tends to 1 as $a \rightarrow 0$.

Let $L \in (0, \infty)$. If $1/a \geq L$, then the restriction of f_a to T_0^L followed by the conformal embedding of T_a^L into Z_a provides a quasiconformal embedding with dilatation arbitrarily close to 1. Since the subsurfaces T_0^L exhaust Ω , we are done.

q.e.d.

We leave it to the reader to check that the extremal length estimates of Lemma 5.2, Lemma 5.3 and Lemma 5.4 hold for this example as well. In the same way as before, we deduce that

$$\text{EL}(\alpha + \beta, X_a) < \text{EL}(\alpha + \beta, Y_a) \quad \text{and} \quad \text{EL}(\alpha + \beta, Z_a) > \text{EL}(\alpha + \beta, W_a)$$

provided that a is small enough. Hence there exist non-convex balls in $\mathcal{T}(S_{1,4})$.

In the same fashion, we can further replace the 3 punctures on the bottom right of Φ_a by a handle, which shows that $\mathcal{T}(S_{2,1})$ contains non-convex balls.

We can also produce examples in any higher topological complexity as follows. Suppose that $3\mathbf{g} - 3 + \mathbf{p} > 4$ and let $\mathbf{h} = \min(2, \mathbf{g})$ and $\mathbf{q} = 7 - 3\mathbf{h}$. Let Φ_a be the half-translation surface constructed above of genus \mathbf{h} with \mathbf{q} punctures. In the bottom left corner of Φ_a , we may remove $\mathbf{p} - \mathbf{q}$ points, cut $\mathbf{g} - \mathbf{h}$ horizontal slits, and glue each one back to itself in an $ABA^{-1}B^{-1}$ pattern to form a handle. The resulting half-translation surface $\tilde{\Phi}_a$ has genus \mathbf{g} and \mathbf{p} punctures. Moreover, the conformal limits of $\mathcal{G}_{-1}\tilde{\Phi}_a$, $\tilde{\Phi}_a$, $\mathcal{G}_{\log(1/a)}\tilde{\Phi}_a$ and $\mathcal{G}_{\log(1/a)+1}\tilde{\Phi}_a$ for the top left and bottom right subsurfaces are all unchanged. Indeed, the images of the nearly conformal embeddings used to prove conformal convergence were all disjoint from the bottom left corner. The same proof carries over and we obtain:

Theorem 5.7. *There exist non-convex balls in $\mathcal{T}(S_{\mathbf{g},\mathbf{p}})$ whenever $3\mathbf{g} - 3 + \mathbf{p} \geq 4$.*

This leaves out 5 cases with $\dim_{\mathbb{C}} \mathcal{T}(S_{\mathbf{g},\mathbf{p}}) = 3\mathbf{g} - 3 + \mathbf{p} > 1$: $S_{0,5}$, $S_{0,6}$, $S_{1,2}$, $S_{1,3}$ and $S_{2,0}$. Note that the above strategy of proof cannot be applied to $S_{0,5}$. Indeed, we would need to split $S_{0,5}$ into two subsurfaces each containing an essential simple closed curve. But no matter how we cut $S_{0,5}$, one component is a sphere with at most 3 holes, hence has no essential simple closed curve. Although the limiting argument does

not carry over, the idea of playing a horizontal curve against a vertical curve is still fruitful.

6. Lower complexity cases

6.1. The Schwarz-Christoffel formula. Consider the polygon $L_a = P(1, a, a, 1)$ where $a > 0$ and P is the staircase-shaped polygon from section 5.1. This is an L -shape obtained by removing the top right a by a square from a $(1 + a)$ by $(1 + a)$ square. We mark each of the 5 internal right angles in L_a . Let α be the arc crossing the vertical leg in L_a , let β be the arc crossing the horizontal leg, and let $\gamma = \alpha + \beta$.




Figure 4. The polygon L_a .

Since all the sides in L_a are either horizontal or vertical, its double Φ_a —topologically a sphere with 5 marked points—is a half-translation surface. We want to study the behavior of the extremal length of the double $\hat{\gamma}$ of γ in Φ_a under the Teichmüller flow.

Lemma 6.1. *There exists a conformal homeomorphism $h : \mathcal{G}_t L_a \rightarrow M$ where $M = R_\alpha \cup R_\beta$ is a stack of two rectangles of height 1 that line up on their right side such that*

- 1) *the inverse images of corners of M with interior angle $\pi/2$ are marked points in $\mathcal{G}_t L_a$;*
- 2) *$h(\alpha)$ joins the left side of R_α to the right side of M ;*
- 3) *$h(\beta)$ joins the left side of R_β to the right side of M .*

Moreover, we have

$$\text{EL}(\hat{\gamma}, \mathcal{G}_t \Phi_a) = 2 \text{EL}(\gamma, \mathcal{G}_t L_a) = 2 \text{area}(M).$$

Proof. As mentioned in Section 2, the extremal length $\text{EL}(\hat{\gamma}, \mathcal{G}_t \Phi_a)$ is realized by a unique Jenkins-Strebel half-translation structure $\mathcal{G}_t \Phi_a \rightarrow \Psi$ partitioned into two horizontal cylinders C_α and C_β of height 1 each, homotopic to the doubles $\hat{\alpha}$ and $\hat{\beta}$ of the arcs α and β . Then

$$\text{EL}(\hat{\gamma}, \mathcal{G}_t \Phi_a) = \text{EL}(C_\alpha) + \text{EL}(C_\beta) = \text{area}(C_\alpha) + \text{area}(C_\beta) = \text{area}(\Psi).$$

Let $J : \mathcal{G}_t\Phi_a \rightarrow \mathcal{G}_t\Phi_a$ be the anti-conformal involution exchanging \mathcal{G}_tL_a with its mirror image. Then $J(C_\alpha)$ and $J(C_\beta)$ are disjoint cylinders homotopic to $\widehat{\alpha}$ and $\widehat{\beta}$ respectively having the same extremal length as C_α and C_β . By uniqueness of the extremal cylinders, the latter are invariant under J . It follows that Ψ is also symmetric with respect to J . Indeed, $\overline{J^*\Psi}$ is a half-translation structure on $\mathcal{G}_t\Phi_a$ partitioned into two horizontal cylinders of height 1 homotopic to $\widehat{\alpha}$ and $\widehat{\beta}$, and is thus equal to Ψ by uniqueness.

Any anti-conformal involution of a Euclidean cylinder $S^1 \times I$ which reverses the orientation of its core curve comes from a reflection of S^1 about a diameter. Thus $R_\alpha = C_\alpha \cap \mathcal{G}_tL_a$ and $R_\beta = C_\beta \cap \mathcal{G}_tL_a$ are Euclidean rectangles of height 1 in the half-translation structure Ψ . Let $h : \mathcal{G}_tL_a \rightarrow M$ be the restriction of the conformal isomorphism $\mathcal{G}_t\Phi_a \rightarrow \Psi$. Then $M = R_\alpha \cup R_\beta$ with R_α and R_β glued isometrically along some part of their horizontal boundary.

The Gauss-Bonnet theorem tells us that Ψ has an angle defect of 4π . Since Ψ has at most 5 cone points of angle π , it has at most one cone point of angle 3π . Such a cone point has to lie on the circle of symmetry of Ψ , otherwise there would be two. Thus M has no singularities in its interior, which means that it is really a polygon. The preimages of the right angles in M by h have to be marked points in \mathcal{G}_tL_a , for after doubling M the right angles give rise to π -angle singularities of Ψ . Since there are only 5 marked points in \mathcal{G}_tL_a , the rectangles have to line up on one side. If we rotate M so that R_α is on top, then they line up on the right side, where there is no marked point separating α from β .

Let ρ be the Euclidean metric on M . Then

$$\begin{aligned} \text{EL}(\gamma, \mathcal{G}_tL_a) &= \text{EL}(h(\gamma), M) \\ &\geq \frac{\ell(h(\gamma), \rho)^2}{\text{area}(\rho)} = \frac{(\ell(R_\alpha) + \ell(R_\beta))^2}{\text{area}(M)} = \text{area}(M). \end{aligned}$$

On the other hand, if the ratio $\ell(\gamma, \sigma)^2 / \text{area}(\sigma)$ was strictly bigger than $\text{area}(M)$ for some conformal metric σ on \mathcal{G}_tL_a , then by doubling we would get

$$\text{EL}(\widehat{\gamma}, \mathcal{G}_t\Phi_a) \geq \frac{\ell(\widehat{\gamma}, \widehat{\sigma})^2}{\text{area}(\widehat{\sigma})} = \frac{(2\ell(\gamma, \sigma))^2}{2\text{area}(\sigma)} > 2\text{area}(M) = \text{area}(\Psi),$$

a contradiction. Alternatively, one can prove that ρ is extremal using the standard length-area argument [KPT15]. q.e.d.

If we show that for some $a > 0$ the function $t \mapsto \text{EL}(\gamma, \mathcal{G}_tL_a)$ increases and later decreases, then the same holds for the function $t \mapsto \text{EL}(\widehat{\gamma}, \mathcal{G}_t\Phi_a)$ and this implies the existence of non-convex balls in $\mathcal{T}(S_{0,5})$.

Another relevant observation is that the reflection τ in the diagonal line $y = x$ maps \mathcal{G}_tL_a to $\mathcal{G}_{-t}L_a$ anti-conformally and sends the homotopy class of γ to itself so that the function $t \mapsto \text{EL}(\gamma, \mathcal{G}_tL_a)$ is even.




Figure 5. The conformal homeomorphism in Lemma 6.1

Therefore, all we have to show is that there exists positive a and t such that $\text{EL}(\gamma, \mathcal{G}_0 L_a) > \text{EL}(\gamma, \mathcal{G}_t L_a)$.

Let $f : \overline{\mathbb{H}^2} \rightarrow \mathcal{G}_t L_a$ be a conformal homeomorphism. Then f extends by Schwarz reflection to a conformal homeomorphism $f : \widehat{\mathbb{C}} \rightarrow \mathcal{G}_t \overline{\Phi}_a$. The pull-back $q = f^* dz^2$ is a meromorphic quadratic differential on $\widehat{\mathbb{C}}$ with a simple pole at the preimage of each marked point and a simple zero at the preimage of the inward corner in $\mathcal{G}_t L_a$. Moreover, q is symmetric with respect to complex conjugation. We thus have

$$q = \frac{A(z-b)}{\prod_{j=0}^4 (z-z_j)} dz^2 = (f'(z))^2 dz^2$$

for some A , b and z_j in \mathbb{R} . It follows that

$$f(z) = \sqrt{A} \int_0^z \sqrt{\frac{(\zeta-b)}{\prod_{j=0}^4 (\zeta-z_j)}} d\zeta + f(0).$$

This is a special case of the Schwarz-Christoffel formula for conformal maps onto polygons [DT02, p.10]. For the formula to make sense, one has to pick a consistent choice of square root, which we can do on $\overline{\mathbb{H}^2}$.

Let $g = h \circ f : \overline{\mathbb{H}^2} \rightarrow M$ where $h : \mathcal{G}_t L_a \rightarrow M$ is as in Lemma 6.1. By the same reasoning as above, g has the form

$$g(z) = C \int_0^z \sqrt{\frac{(\zeta-p)}{\prod_{j=0}^4 (\zeta-z_j)}} d\zeta + D$$

for some constants p , C and D . The area of M can then be recovered from its side lengths, which are integrals of the above form.

REMARK. One can use the Schwarz-Christoffel formula to prove the first part of Lemma 6.1. Indeed, for any choice of $p \in \mathbb{R}$, the map

$$G_p(z) = \int_0^z \sqrt{\frac{(\zeta-p)}{\prod_{j=0}^4 (\zeta-z_j)}} d\zeta$$

is a conformal homeomorphism from $\overline{\mathbb{H}^2}$ to a polygon with angle $\pi/2$ at each vertex $G_p(z_j)$ and angle $3\pi/2$ at $G_p(p)$. Suppose that

$$(6.1) \quad |G_{z_1}(z_0) - G_{z_1}(z_1)| \leq |G_{z_1}(z_1) - G_{z_1}(z_2)|.$$

Then by the intermediate value theorem, there exists a point p between z_1 and z_2 such that

$$|G_p(z_0) - G_p(z_1)| = |G_p(p) - G_p(z_2)|.$$

Indeed, $|G_p(z_0) - G_p(z_1)|$ is bounded away from zero for $p \in [z_1, z_2]$ whereas $|G_p(p) - G_p(z_2)|$ tends to zero as $p \rightarrow z_2$. If the reverse of inequality (6.1) holds, then there is a p between z_0 and z_1 such that

$$|G_p(z_0) - G_p(p)| = |G_p(z_1) - G_p(z_2)|.$$

In either case, after rescaling we get that $G_p(\overline{\mathbb{H}^2})$ is a stack of two rectangles of height 1.

The problem of calculating $\text{EL}(\gamma, \mathcal{G}_t L_a)$ has now been reduced to finding the correct parameters z_0, \dots, z_4, b and p (all of which depend on a and t). The Schwarz-Christoffel Toolbox [**Dri**] for MATLAB is designed to solve this parameter problem numerically. We used this to compute $\text{EL}(\gamma, \mathcal{G}_t L_a)$ for $a = 1/4$ at $10^3 + 1$ equally spaced values of t in the interval $[-0.275, 0.275]$ and obtained Figure 6.




Figure 6. Graph of $t \mapsto \text{EL}(\gamma, \mathcal{G}_t L_a)$ for $a = 1/4$.

The figure clearly shows a decrease from time $t = 0$ to $t \approx 0.159$. However, the Schwarz-Christoffel Toolbox does not come with any certified error estimates. Moreover, the apparent decrease of extremal length is rather small: it drops from about 3.87 to 3.856, which represents less than 1% decrease.

In order to turn this into a rigorous proof, we do the following. We take the approximate parameters provided by the SC Toolbox, then compute the corresponding Schwarz-Christoffel integrals numerically but with certified precision. Since the initial parameters are inexact, the images of the Schwarz-Christoffel maps are not the polygons we

expect, but we can estimate how far away they are from the correct polygons and deduce bounds for extremal length.

One way to get rigorous bounds on a numerical result is to use interval arithmetic. Roughly speaking, interval arithmetic means that instead of rounding to the nearest representable number, the computer keeps track of correct lower and upper bounds for every operation, yielding a true interval in which the result of a calculation lies.

There exist packages that do numerical integration using interval arithmetic. However, we did not find any that can handle improper integrals. We thus wrote a program in Sage [TSD16] to compute lower and upper bounds on the integrals needed using interval arithmetic. The Sage worksheet and its output are available at <http://github.com/maxforbou/non-convex-balls>.

6.2. Rigorous bounds. Let $k = 5.27110734472$, let

$$f(z) = \int_0^z \frac{d\zeta}{\sqrt{\zeta(\zeta^2 - 1)(\zeta^2 - k^2)}}$$

and let $X = f(\overline{\mathbb{H}^2})$ with marked points at $0 = f(0)$, $f(\pm 1)$ and $f(\pm k)$. The polygon X is an L -shape with angle $\pi/2$ at the marked points and angle $3\pi/2$ at $f(\infty)$. Furthermore, X is symmetric about the diagonal line $y = x$ since the function under the square root is odd. Thus, X is a rescaled copy of L_a , where

$$a = \frac{|f(0) - f(1)|}{|f(1) - f(k)|} - 1.$$

We want to get rigorous bounds on both the shape of X and the extremal length of γ in X . The first thing we need to compute is the integral

$$|f(0) - f(1)| = \left| \int_0^1 \frac{dx}{\sqrt{x(x^2 - 1)(x^2 - k^2)}} \right| = \int_0^1 \frac{dx}{\sqrt{|x(x^2 - 1)(x^2 - k^2)|}}.$$

The main observation is that the integrand

$$F(x) = \frac{1}{\sqrt{|x(x^2 - 1)(x^2 - k^2)|}}$$

is logarithmically convex (hence convex) on $(0, 1)$.

Lemma 6.2. *Suppose that $z_0 < z_1 < z_2 < z_3 < z_4$. Then the function*

$$F(x) = \prod_{j=0}^4 |x - z_j|^{-1/2}$$

is log-convex between any two consecutive z_j 's.

Proof. We compute

$$(\log F)'(x) = -\frac{1}{2} \left(\sum_{j=0}^4 \frac{1}{x - z_j} \right)$$

and

$$(\log F)''(x) = \frac{1}{2} \left(\sum_{j=0}^4 \frac{1}{(x - z_j)^2} \right) > 0.$$

q.e.d.

Therefore, for any compact subinterval $I \subset (0, 1)$ and any partition $\{x_{-n}, \dots, x_n\}$ of I we have

$$\begin{aligned} \sum_{j=-n}^{n-1} (x_{j+1} - x_j) F\left(\frac{x_j + x_{j+1}}{2}\right) &\leq \int_I F(x) dx \\ &\leq \sum_{j=-n}^{n-1} (x_{j+1} - x_j) \left(\frac{F(x_j) + F(x_{j+1})}{2}\right) \end{aligned}$$

by the trapezoid rule. We choose the partition $\{x_{-n}, \dots, x_n\}$ using the tanh-sinh quadrature [TM74] which is well-adapted for this type of singular integral. On a bounded interval (a, b) the quadrature points are defined as

$$x_j = \frac{(a+b)}{2} + \frac{(b-a)}{2} \tanh\left(\frac{\pi}{2} \sinh(j\Delta)\right)$$

where $\Delta > 0$ is a step size to be determined together with n . In this case we took $\Delta = 2^{-13}$ and $n = 2^{15}$.

Let $\delta = x_{-n} = 1 - x_n$ where the x_j 's are sample points for the interval $(0, 1)$. An elementary calculation shows that

$$0 \leq \int_0^\delta \frac{dx}{\sqrt{x(1-x^2)(k^2-x^2)}} \leq \frac{2\sqrt{\delta}}{\sqrt{(1-\delta^2)(k^2-\delta^2)}}.$$

and

$$0 \leq \int_{1-\delta}^1 \frac{dx}{\sqrt{x(1-x^2)(k^2-x^2)}} \leq \frac{2\sqrt{\delta}}{\sqrt{(1-\delta)(2-\delta)(k^2-1)}}.$$

Adding the lower bounds for each of the three subintervals $[0, \delta]$, $[\delta, 1 - \delta]$ and $[1 - \delta, 1]$ yields a certified lower bound on $|f(0) - f(1)|$, and similarly for upper bounds. We use the same method to estimate $|f(1) - f(k)|$.

In order to compute the extremal length $\text{EL}(\gamma, X)$, we consider the conformal homeomorphism

$$g(z) = -i \int_0^z \frac{d\zeta}{\sqrt{(\zeta^2 - 1)(\zeta^2 - k^2)}}$$

between $\overline{\mathbb{H}^2}$ and a rectangle X' with marked points at $0 = g(0)$, $g(\pm 1)$ and $g(\pm k)$. Then $g \circ f^{-1} : X \rightarrow X'$ is a conformal homeomorphism preserving the marked points so that $\text{EL}(\gamma, X) = \text{EL}(\gamma, X')$. Since the above integrand is even, $g(0)$ subdivides X' into two congruent rectangles. After rescaling X' to have height 2, the extremal length is given by area. This means that

$$\text{EL}(\gamma, X) = \text{EL}(\gamma, X') = 2 \frac{|g(1) - g(k)|}{|g(0) - g(1)|}.$$

We get rigorous bounds on $|g(0) - g(1)|$ and $|g(1) - g(k)|$ with the same method as for f . The results are compiled in Table 1.

	lower bound	upper bound	approximation
$ f(0) - f(1) $	0.500482492919	0.500482504323	0.500482496721
$ f(1) - f(k) $	0.400385993317	0.400386005494	0.400385997377
$ g(0) - g(1) $	0.300738235179	0.300738239980	0.30073823678
$ g(1) - g(k) $	0.581911579444	0.581911593793	0.581911584228
$\text{EL}(\gamma, X)$	3.86988751070	3.86988766789	3.86988758368

Table 1. Certified bounds on the side lengths of X and X' . The last column shows the corresponding value calculated with Sage's `nintegral` routine.

We now estimate extremal length after stretching X . Let the parameters z_0, z_1, z_2, p, z_3 and z_4 be equal to -3.33297982345 , -0.26873921366 , 0 , 0.17317940636 , 1 and 2.94288195633 respectively. Then let

$$\phi(z) = \int_0^z \prod_{j=0}^4 (\zeta - z_j)^{-1/2} d\zeta,$$

$$\psi(z) = -i \int_0^z (\zeta - p)^{1/2} \prod_{j=0}^4 (\zeta - z_j)^{-1/2} d\zeta,$$

$Y = \phi(\overline{\mathbb{H}^2})$ and $Y' = \psi(\overline{\mathbb{H}^2})$. The polygon Y is meant to be close to a rescaled version of $\mathcal{G}_t X$ for $t \approx 0.159$ whereas Y' is a stack of two rectangles of nearly the same height, which we use to estimate $\text{EL}(\gamma, Y)$.

Since the integrand $\prod_{j=0}^4 |x - z_j|^{-1/2}$ is convex, we may use the trapezoid rule to compute the side lengths of Y . There are also elementary estimates near the poles like before.

For Y' the integrand is convex on each interval of continuity not adjacent to p . Indeed, if

$$G(x) = |x - p|^{1/2} \prod_{j=0}^4 |x - z_j|^{-1/2}$$

then

$$2(\log G)''(x) = \sum_{j=0}^4 \frac{1}{(x - z_j)^2} - \frac{1}{(x - p)^2}$$

which is positive when $x < z_2$ and when $x > z_3$. We can thus apply the trapezoid rule with tanh-sinh quadrature to bound the side lengths of Y' not adjacent to $\psi(p)$. The length

$$|\psi(z_4) - \psi(z_0)| = |\psi(z_4) - \psi(\infty)| + |\psi(\infty) - \psi(z_0)|$$

is a little bit different since we need to compute integrals over two half-infinite intervals. We use another doubly exponential quadrature on these intervals given by

$$x_j = \exp\left(\frac{\pi}{2} \sinh(j\Delta)\right)$$

for the interval $(0, \infty)$. To estimate the area lost by truncating away from infinity, note that for $x > 2z_4 - p$ we have $|x - p| < 2|x - z_4|$ as well as $|x - z_j| \geq |x - z_4|$ for each j . It follows that $G(x) \leq \sqrt{2}|x - z_4|^{-2}$ and hence

$$\int_a^\infty G(x) dx \leq \sqrt{2}|a - z_4|^{-1}$$

provided that $a \geq 6$. Similarly, we have

$$\int_{-\infty}^b G(x) dx \leq \sqrt{2}|b - z_0|^{-1}$$

provided that $b \leq -7$.

The polygon Y' is not exactly a stack of two rectangles of the same height, but we can still use it to estimate $\text{EL}(\gamma, Y) = \text{EL}(\gamma, Y')$. Using the Euclidean metric on Y' yields the lower bound

$$\text{EL}(\gamma, Y') \geq \frac{\ell(\gamma)^2}{\text{area}(Y')} = \frac{(|\psi(z_0) - \psi(z_1)| + |\psi(z_3) - \psi(z_4)|)^2}{\text{area}(Y')}.$$

Moreover, the sum of the extremal lengths of the horizontal rectangles R_α and R_β in Y' is an upper bound for the extremal length:

$$\text{EL}(\gamma, Y') \leq \frac{|\psi(z_0) - \psi(z_1)|}{|\psi(z_1) - \psi(z_2)|} + \frac{|\psi(z_3) - \psi(z_4)|}{|\psi(p) - \psi(z_3)|}.$$

The last caveat is that Y does not lie exactly along the Teichmüller geodesic through X . Let

$$a = \frac{|f(0) - f(1)|}{|f(1) - f(k)|} - 1$$

and

$$K = \frac{|\phi(z_0) - \phi(z_1)|}{|\phi(z_3) - \phi(z_4)|}$$

and consider the polygon $Z = P(K, a, Ka, 1)$. Then up to rescaling $Z = \mathcal{G}_t X$ for $t = \frac{1}{2} \log K$. Divide each of Y and Z into three rectangles

with sides parallel to the coordinate axes and let $h : Y \rightarrow Z$ be the homeomorphism which is affine on each subrectangle. Then h preserves the marked points and

$$\frac{1}{C} \text{EL}(\gamma, Y) \leq \text{EL}(\gamma, Z) \leq C \text{EL}(\gamma, Y)$$

where $C \geq \exp(2d(Y, Z))$ is the dilatation of h . Note that C can be expressed in terms of the aspect ratios of the three subrectangles in Y and Z . The resulting bounds are shown in Table 2.

	lower bound	upper bound	approximation
$ \phi(z_0) - \phi(z_1) $	1.036823405576	1.036823443983	1.03682341838
$ \phi(z_1) - \phi(z_2) $	0.943128409696	0.943128430640	0.943128416679
$ \phi(z_2) - \phi(z_3) $	1.296029251902	1.296029284584	1.2960292628
$ \phi(z_3) - \phi(z_4) $	0.754502722746	0.754502742641	0.754502729379
$ \psi(z_0) - \psi(z_1) $	1.068955145751	1.068955175385	1.06895515563
$ \psi(z_1) - \psi(z_2) $	0.512964353079	0.512964364188	0.512964356783
$ \psi(z_3) - \psi(z_4) $	0.908877581965	0.908877603159	0.908877589032
$ \psi(z_4) - \psi(z_0) $	1.025928700631	1.025928738891	1.02592871356
$\text{EL}(\gamma, Y)$	3.855692084405	3.855692498209	3.85569234685
$\exp(2d(Y, Z))$	–	1.000000357759	–
$\text{EL}(\gamma, Z)$	3.855690704998	3.855693877617	3.85569234685

Table 2. Certified integrals after stretching

We thus have

$$\text{EL}(\gamma, \mathcal{G}_t X) = \text{EL}(\gamma, Z) < 3.8557 < 3.8698 < \text{EL}(\gamma, X),$$

from which we conclude that $\mathcal{T}(S_{0,5})$ contains non-convex balls.

6.3. Remaining cases. Adding an artificial marked point on the boundary of X between $f(1)$ and $f(k)$ (the right-most side of X) does not change the extremal length of γ at any time. After doubling, this shows the existence of a non-convex ball in $\mathcal{T}(S_{0,6})$.

Recall that there are isometries $\mathcal{T}(S_{0,5}) \cong \mathcal{T}(S_{1,2})$ and $\mathcal{T}(S_{0,6}) \cong \mathcal{T}(S_{2,0})$ arising from the hyperelliptic involutions on $S_{1,2}$ and $S_{2,0}$. This shows that there exist non-convex balls in those two cases as well.

To treat the torus with 3 punctures, we can cut horizontal slits of length $s > 0$ at two punctures in the double of X then glue the two slits together to form a handle. As $s \rightarrow 0$, the extremal length of the double $\widehat{\gamma}$ of γ on the 3 times punctured torus converges to its value on the double of X . The same is true after applying the Teichmüller flow \mathcal{G}_t for any t . It follows that if s is small enough, then the resulting geodesic in $\mathcal{T}(S_{1,3})$ exhibits an increase of extremal length followed by a decrease. This completes the proof of Theorem 1.1.

References

- [Abi77] W. Abikoff, *Degenerating families of Riemann surfaces*, Ann. of Math. (2) **105** (1977), no. 1, 29–44, MR0442293, Zbl 0347.32010.
- [Ahl10] L. V. Ahlfors, *Conformal invariants*, Topics in geometric function theory, Reprint of the 1973 original, With a foreword by Peter Duren, F. W. Gehring and Brad Osgood, AMS Chelsea Publishing, Providence, RI, 2010, MR2730573, Zbl 1211.30002.
- [BE64] L. Bers and L. Ehrenpreis, *Holomorphic convexity of Teichmüller spaces*, Bull. Amer. Math. Soc. **70** (1964), no. 6, 761–764, MR0168800, Zbl 0136.07004.
- [Dri] T. A. Driscoll, *The Schwarz-Christoffel Toolbox*, version 2.3, <http://www.math.udel.edu/~driscoll/SC/>.
- [DT02] T. A. Driscoll and L. N. Trefethen, *Schwarz-Christoffel mapping*, Cambridge Monographs on Applied and Computational Mathematics, vol. 8, Cambridge University Press, Cambridge, 2002, MR1908657, Zbl 1003.30005.
- [GM91] F. Gardiner and H. Masur, *Extremal length geometry of Teichmüller space*, Complex Variables, Theory and Application: An International Journal **16** (1991), 209–237, MR1099913, Zbl 0702.32019.
- [Gol69] G. M. Goluzin, *Geometric theory of functions of a complex variable*, Translations of Mathematical Monographs, Vol. 26, American Mathematical Society, Providence, R.I., 1969, MR0247039, Zbl 0183.07502.
- [Jen57] J. A. Jenkins, *On the existence of certain general extremal metrics*, Ann. of Math. (2) **66** (1957), no. 3, 440–453, MR0090648, Zbl 0082.06301.
- [Ker80] S. P. Kerckhoff, *The asymptotic geometry of Teichmüller space*, Topology **19** (1980), 23–41, MR0559474, Zbl 0439.30012.
- [Ker83] ———, *The Nielsen realization problem*, Ann. of Math. (2) **117** (1983), no. 2, 235–265, MR0690845, Zbl 0528.57008.
- [KPT15] J. Kahn, K. M. Pilgrim, and D. P. Thurston, *Conformal surface embeddings and extremal length*, preprint, [arxiv:1507.05294](https://arxiv.org/abs/1507.05294), 2015.
- [Kra59] S. Kravetz, *On the geometry of Teichmüller spaces and the structure of their modular groups*, Ann. Acad. Sci. Fenn. Ser. A I No. **278** (1959), 35, MR0148906, Zbl 0168.04601.
- [Lin71] M. R. Linch, *On metrics in Teichmueller space*, Thesis (Ph.D.)—Columbia University, ProQuest LLC, Ann Arbor, MI, 1971, MR2620985.
- [LR11] A. Lenzhen and K. Rafi, *Length of a curve is quasi-convex along a Teichmüller geodesic*, J. Differential Geom. **88** (2011), no. 2, 267–295, MR2838267, Zbl 1243.30089.
- [Mas75] H. Masur, *On a class of geodesics in Teichmüller space*, Ann. of Math. (2) **102** (1975), no. 2, 205–221, MR0385173, Zbl 0322.32010.
- [Mas09] ———, *Geometry of Teichmüller space with the Teichmüller metric*, Surveys in differential geometry. Vol. XIV. Geometry of Riemann surfaces and their moduli spaces, Surv. Differ. Geom., vol. 14, Int. Press, Somerville, MA, 2009, pp. 295–313, MR2655331, Zbl 1210.14031.
- [Min96] Y. N. Minsky, *Extremal length estimates and product regions in Teichmüller space*, Duke Math. J. **83** (1996), no. 2, 249–286, MR1390649, Zbl 0861.32015.
- [Miy15] H. Miyachi, *Extremal length functions are log-plurisubharmonic*, preprint, [arxiv:1505.06785](https://arxiv.org/abs/1505.06785), 2015.

- [MM99] H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, *Invent. Math.* **138** (1999), no. 1, 103–149, MR1714338, Zbl 0941.32012.
- [Pom92] Ch. Pommerenke, *Boundary behaviour of conformal maps*, *Grundlehren der Mathematischen Wissenschaften*, vol. 299, Springer-Verlag, Berlin, 1992, MR1217706, Zbl 0762.30001.
- [Raf14] K. Rafi, *Hyperbolicity in Teichmüller space*, *Geom. Topol.* **18** (2014), no. 5, 3025–3053, MR3285228, Zbl 1314.30082.
- [Ren76] H. Renelt, *Konstruktion gewisser quadratischer differentiale mit hilfe von dirichletintegralen*, *Math. Nachr.* **73** (1976), no. 1, 125–1142, MR0499132, Zbl 0374.30017.
- [Str84] K. Strebel, *Quadratic differentials*, *Ergebnisse der Mathematik und ihrer Grenzgebiete (3)*, vol. 5, Springer-Verlag, Berlin, 1984, MR0743423, Zbl 0547.30001.
- [TM74] H. Takahasi and M. Mori, *Double exponential formulas for numerical integration*, *Publ. Res. Inst. Math. Sci.* **9** (1973/74), 721–741, MR0347061, Zbl 0293.65011.
- [Tra37] R. E. Traber, *A fundamental lemma on normal coordinates and its applications*, *Q. J. Math.* **os-8** (1937), no. 1, 142–147, Zbl 0017.03603.
- [TSD16] The Sage Developers, *Sagemath, the Sage Mathematics Software System (Version 7.1)*, 2016, <http://www.sagemath.org>.
- [Whi33] J. H. C. Whitehead, *Convex regions in the geometry of paths—addendum*, *Q. J. Math.* **4** (1933), 226–227, Zbl 0007.36801.
- [Wol87] S. A. Wolpert, *Geodesic length functions and the Nielsen problem*, *J. Differential Geom.* **25** (1987), no. 2, 275–296, MR0880186, Zbl 0616.53039.
- [Wol06] ———, *Convexity of geodesic-length functions: a reprise*, *Spaces of Kleinian groups* (Y. Minsky, M. Sakuma, and C. Series, eds.), vol. 329, Cambridge University Press, 2006, 233–245, MR2258752, Zbl 1126.32010.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TORONTO, 40 ST. GEORGE STREET,
TORONTO, ON, CANADA M5S 2E4

E-mail address: mbourque@math.toronto.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TORONTO, 40 ST. GEORGE STREET,
TORONTO, ON, CANADA M5S 2E4

E-mail address: rafi@math.toronto.edu