CONFORMAL SURFACE EMBEDDINGS AND EXTREMAL LENGTH

JEREMY KAHN, KEVIN M. PILGRIM, AND DYLAN P. THURSTON

ABSTRACT. Given two Riemann surfaces with boundary and a homotopy class of topological embeddings between them, there is a conformal embedding in the homotopy class if and only if the extremal length of every simple closed multi-curve is decreased under the embedding. Furthermore, the homotopy class has a conformal embedding that misses an open disk if and only if extremal lengths are decreased by a definite ratio. This ratio remains bounded away from one under finite covers.

CONTENTS

1. INTRODUCTION

Let R and S be two Riemann surfaces of finite topological type, possibly with boundary, and let $f\colon R \hookrightarrow S$ be a topological embedding. The goal of this paper is to give conditions for f to be homotopic to a *conformal* embedding, possibly with extra nice properties. We give an answer in terms of ratios of extremal lengths of simple multi-curves.

For us, surfaces S are of finite topological type and not necessarily connected, and embeddings between surfaces are required to respect certain markings; see Definition 2.1. A simple multi-curve on S is an embedded 1-manifold in S; see Definition 2.2. The *extremal length* $\mathrm{EL}_S[C]$ of a simple curve C is a measure of the fattest annulus that can be embedded in S

2020 *Mathematics Subject Classification.* Primary 30F60; Secondary 31A15, 32G15.
Key words and phrases. Riemann surfaces with boundary, conformal embeddings, extremal length.

with core curve isotopic to C. See Section 3.2 for more on extremal length, in particular for multi-curves.

Definition 1.1. For $f\colon R \hookrightarrow S$ a topological embedding of Riemann surfaces, the *stretch factor* of f is the maximal ratio of extremal lengths between the two surfaces:

$$\mathrm{SF}[f] \coloneqq \sup_{C \in \mathcal{C}^+(R)} \frac{\mathrm{EL}_S[f(C)]}{\mathrm{EL}_R[C]},$$

where the supremum runs over all simple closed multi-curves C with $\mathrm{EL}_R[C] \neq 0$.

We will show that $\mathrm{SF}[f]$ is achieved by a ratio of extremal lengths of two measured foliations, not multi-curves. But f does not induce a natural continuous map between measured foliations (Example 4.4), so Definition 1.1 is stated in terms of multi-curves.

Theorem 1. *Let R and S be Riemann surfaces and $f\colon R \hookrightarrow S$ be a topological embedding so that no component of $f(R)$ is contained in a disk or a once-punctured disk. Then f is homotopic to a conformal embedding if and only if $\mathrm{SF}[f] \leq 1$.*

The key part of Theorem 1 is due to Ioffe [Iof75]. In fact, his results show that if $\mathrm{SF}[f] \geq 1$, it is related to the quasi-conformal constant.

Proposition 1.2. *Let $f\colon R \hookrightarrow S$ be a topological embedding of Riemann surfaces. If $\mathrm{SF}[f] \geq 1$, then $\mathrm{SF}[f]$ is equal to the smallest quasi-conformal constant of any quasi-conformal embedding homotopic to f.*

We can also characterize conformal embeddings with some extra "room".

Definition 1.3. Let $f\colon R \hookrightarrow S$ be a conformal embedding between Riemann surfaces. We say that f is a *strict* embedding if its image omits a non-empty open subset of each component of S. An *annular extension* of a Riemann surface S is a surface $\widehat{S}$ obtained by attaching a non-empty conformal annulus to each boundary component, with the boundary of S smoothly embedded in $\widehat{S}$. An *annular conformal embedding* is one that extends to a conformal embedding $\widehat{R} \hookrightarrow S$ for some annular extension $\widehat{R}$ of R.

Remark 1.4. A similar relation for subsets of $\widehat{\mathbb{C}}$ is sometimes written $f(R) \Subset S$.

Theorem 2. *Let R and S be Riemann surfaces, with S connected, and let $f\colon R \hookrightarrow S$ be a topological embedding so that no component of $f(R)$ is contained in a disk or a once-punctured disk. Then the following conditions are equivalent:*

(1) f is homotopic to a strict conformal embedding;
(2) f is homotopic to an annular conformal embedding;
(3) there is a neighborhood N of S in Teichmüller space so that, for all $S' \in N$, f is homotopic to a conformal embedding of R in S'; and
(4) $\mathrm{SF}[f] < 1$.

Remark 1.5. When $\mathrm{SF}[f] = 1$, the embedding guaranteed by Theorem 1 is instead a Teichmüller embedding in the sense of Definition 4.1 (with $K = 1$), as studied by Fortier Bourque [FB18].

In condition (3), $\mathrm{SF}[f]$ is related to the size of the largest ball in Teichmüller space for which we can find conformal images of R.

Definition 1.6. Let $f\colon R \subset S$ be a topological embedding of Riemann surfaces. Let $\mathcal{T}_R(S)$ be the subset of the Teichmüller space $\mathcal{T}(S)$ for which there is a conformal embedding of R in the homotopy class $[f]$. (This does not contain the Riemann surface structure S itself if $\mathrm{SF}[f] > 1$.)

Proposition 1.7. *Let $f\colon R \hookrightarrow S$ be a topological embedding of Riemann surfaces. Then*

$$d(S, \partial\mathcal{T}_R(S)) = \frac{1}{2}\left|\log(\mathrm{SF}[f])\right|.$$

We can also control the behavior of the stretch factor under taking covers. Proposition 1.2 guarantees that when $\mathrm{SF}[f] \geq 1$, the stretch factor is unchanged under taking finite covers (see Proposition 6.3). We can control what happens when $\mathrm{SF}[f] < 1$, as well.

Definition 1.8. For $f\colon R \hookrightarrow S$ a topological embedding of Riemann surfaces and $p\colon \widetilde{S} \to S$ a covering map, the corresponding *cover* of f is the pull-back map $\widetilde{f}$ in the diagram

$$\begin{array}{ccc} \widetilde{R} & \overset{\widetilde{f}}{\dashrightarrow} & \widetilde{S} \\ \vdots\, q & & \downarrow p \\ R & \xrightarrow{f} & S. \end{array}$$

Explicitly, we have

$$\begin{aligned} \widetilde{R} &\coloneqq \left\{ (r, \widetilde{s}) \in R \times \widetilde{S} \mid f(r) = p(\widetilde{s}) \right\} \\ \widetilde{f}(r, \widetilde{s}) &\coloneqq \widetilde{s} \\ q(r, \widetilde{s}) &\coloneqq r. \end{aligned}$$

Then $\widetilde{f}$ is a topological embedding and q is a covering map. We may also say that $\widetilde{f}$ is a cover of f, without specifying p.

Definition 1.9. For $f\colon R \hookrightarrow S$ a topological embedding of Riemann surfaces, the *lifted stretch factor* $\widetilde{\mathrm{SF}}[f]$ is

$$\widetilde{\mathrm{SF}}[f] \coloneqq \sup_{\substack{\widetilde{f} \text{ finite} \\ \text{cover of } f}} \mathrm{SF}[\widetilde{f}].$$

Theorem 3. *Let $f\colon R \hookrightarrow S$ be a topological embedding of Riemann surfaces. If $\mathrm{SF}[f] \geq 1$, then $\widetilde{\mathrm{SF}}[f] = \mathrm{SF}[f]$. If $\mathrm{SF}[f] < 1$, then*

$$\mathrm{SF}[f] \leq \widetilde{\mathrm{SF}}[f] < 1.$$

The hard part of Theorem 3 is showing that $\widetilde{\mathrm{SF}}[f]$ is strictly less than 1 when $\mathrm{SF}[f] < 1$.

By Proposition 1.7, $\widetilde{\mathrm{SF}}[f] < 1$ is equivalent to saying that $\mathcal{T}_{\widetilde{R}}(\widetilde{S})$ contains a ball of uniform size around $\widetilde{S}$ for every finite cover of f.

Theorem 3 is used in later work [Thu16, Thu20] to give a positive characterization of post-critically finite rational maps among topological branched self-covers of the sphere. This provides a counterpoint to W. Thurston's characterization [DH93], which characterizes rational maps in terms of an obstruction.

1.1. **History.** The maximum of the ratio of extremal lengths has appeared before, usually in the context of closed surfaces, where it gives Teichmüller distance, as first proved by Kerckhoff (see Theorem 3.21 below). For surfaces with boundary the behavior is quite different, as the stretch factor can be less than one.

In the special case when the target S is a closed torus, there is very precise information about when R conformally embeds inside of S [Shi87, Shi93]. Shiba proves that in this case $\mathcal{T}_R(S)$ is a disk with respect to the Teichmüller metric. Masumoto gives similar information about the case when the domain R is a once-holed torus [Mas16].

There has been earlier work on portions of Theorem 2. In particular, Earle and Marden [EM78] showed that, with extra topological restrictions on the embedding $R \hookrightarrow S$, if f is homotopic to a strict conformal embedding then it is homotopic to an annular conformal embedding.

It is tempting to look for an analogue of Theorem 1 using hyperbolic length instead of extremal length, given that, by the Schwarz lemma, hyperbolic length is decreased under conformal inclusion. However, the results are false for hyperbolic length in almost all cases [Mas00, FB16].

These results were first announced in a research report by the last author [Thu16].

1.2. **Organization.** Section 2 reviews background material and specifies our definitions for topological surfaces. Section 3 does the same for Riemann surfaces and extremal length, as well as giving elementary properties of the stretch factor. Section 4 proves Theorem 1, largely based on a theorem of Ioffe. Section 5 extends this to prove Theorem 2. Section 6 gives the further extension to prove Theorem 3. In Section 6, we also prove Theorem 4, an estimate on areas of subsurfaces with respect to quadratic differentials; this may be of independent interest. Section 7 gives some directions for future research, and in the process gives another way to get an upper bound on $\widetilde{\mathrm{SF}}[f]$.

1.3. **Acknowledgments.** We thank Matt Bainbridge, Maxime Fortier Bourque, and Frederick Gardiner for many helpful conversations. Aaron Cohen, Russell Lodge, Insung Park, and Maxime Scott gave useful comments on earlier drafts, as did the anonymous referee. JK was supported by NSF grant DMS-1352721. KMP was supported by Simons Foundation Collaboration Grants #245269 and #615022. DPT was supported by NSF grants DMS-1358638, DMS-1507244, and DMS-2110143.

2. Topological Setting

Definition 2.1. By a (smooth) *surface* S we mean a smooth, oriented, compact 2-manifold with boundary, together with a distinguished finite set P of points in the interior of S, the *punctures*. The boundary ∂S of S is a finite union of circles. By a slight abuse of terminology, by the interior S° of S we mean $S \setminus (P \cup \partial S)$. If we want to emphasize that we are talking about the compact version of S, we will write $\overline{S}$.

A surface is *small* if it is the sphere with 0, 1, or 2 punctures or the unit disk with 0 or 1 punctures. These are the surfaces that have no non-trivial curves by the definition below.

By a *topological map* $f\colon R \to S$ between surfaces we mean an orientation-preserving continuous map from R° to S° that extends to a continuous map from $\overline{R}$ to $\overline{S}$. In particular, the image of a puncture is a puncture or a regular point, and embeddings need not be one-to-one on ∂R. Homotopies are taken within the same space of maps.

Definition 2.2. A *multi-curve* C on a surface S is a smooth 1-manifold with boundary $X(C)$ together with an immersion from the interior of $X(C)$ into S° that maps $\partial X(C)$ to ∂S. We do not assume that $X(C)$ is connected; if it is, C is said to be *connected* or a *curve*. We will mostly be concerned with *simple* multi-curves, those for which the immersion is an embedding. An *arc* is a curve for which $X(C)$ is an interval, and a *loop* is a curve for which $X(C)$ is a circle. A multi-curve is *closed* if it has no arc components.

A (multi-)curve is *trivial* if it is contained in a disk or once-punctured disk of S.

Equivalence of multi-curves is the equivalence relation generated by

(1) homotopy within the space of all maps taking $\partial X(C)$ to ∂S (not necessarily immersions),
(2) reparametrization of the 1-manifold $X(C)$ (including orientation reversal), and
(3) dropping trivial components.

The equivalence class of C is denoted $[C]$. The space of simple multi-curves on S up to homotopy is denoted $\mathcal{C}^{\pm}(S)$. If $\partial S \neq \emptyset$, then we distinguish two subsets of $\mathcal{C}^{\pm}(S)$:

- $\mathcal{C}^{+}(S) \subset \mathcal{C}^{\pm}(S)$ is the subset of closed multi-curves and
- $\mathcal{C}^{-}(S) \subset \mathcal{C}^{\pm}(S)$ is the subset with no loops parallel to a boundary component.

A *weighted multi-curve* $C = \sum a_i C_i$ is a multi-curve in which each connected component is given a positive real coefficient a_i. When considering equivalence of weighted multi-curves, we add the further relation that two parallel components may be merged and their weights added. We write $\mathcal{C}_{\mathbb{R}}(S)$ or $\mathcal{C}_{\mathbb{Q}}(S)$ for the space of weighted multi-curves with real or rational weights, respectively.

Definition 2.3. A (positive) *measured foliation* F on a surface S is a singular 1-dimensional foliation on $\overline{S}$, tangent to ∂S, with a non-zero transverse measure. F is allowed to have k-prong singularities, as described, for instance, by Fathi-Laudenbach-Poénaru [FLP79], and summarized below.

- At points of S°, we allow k-prong singularities for $k \geq 3$. (If there are only 2 prongs, it is not a singularity.) This is also called a zero of order $k-2$.
- At punctures, we allow k-prong singularities for $k \geq 1$. This is also called a zero of order $k-1$.
- At points of ∂S, we allow k-prong singularities for $k \geq 3$. This is also called a zero of order $k-2$. If we double the surface, it becomes a $(2k-2)$-prong singularity.

We also admit the empty (zero) measured foliation as a degenerate case. A *singular leaf* of a measured foliation is a leaf that ends at a singularity. A *saddle connection* is a singular leaf that ends at singularities in both directions. If a saddle connection connects two distinct singularities, and at least one of the singularities is in the interior, it is possible to *collapse* it to form a new measured foliation. *Whitehead equivalence* of measured foliations is the equivalence relation generated by homotopy and collapsing saddle connections. We denote the Whitehead equivalence class of a measured foliation by $[F]$, and the set of Whitehead equivalence classes of measured foliations by $\mathcal{MF}^{+}(S)$.

From a multi-curve $C \in \mathcal{C}^{-}(S)$ and a measured foliation F on S, we can form the intersection number

$$i([C],[F]) \coloneqq \inf_{C_1 \in [C]} \int_t |F(C_1'(t))|\, dt.$$

Proposition 2.4. *The map*

$$\begin{aligned} \mathcal{MF}^+(S) &\to \mathbb{R}^{\mathcal{C}^-(S)} \\ [F] &\mapsto \big(i([C],[F])\big)_{[C]\in\mathcal{C}^-(S)} \end{aligned}$$

is an injection, with image a finite-dimensional manifold determined by its projection onto finitely many factors.

Proof sketch. This is standard, although we could not find a single reference that covered the full generality of all surface types we consider.

In the harder case of a closed surface of genus g, from a pair of pants decomposition of S with $3g-3$ curves with some extra marking, there are $9g-9$ curves C_i so that the intersection numbers $i([C_i],[F])$ determine $[F]$, and furthermore realize $\mathcal{MF}(S)$ as a topological ball [FLP79, Exposé 6]. If S is a genus g with k punctures (and no boundary), there is a similar maximal set of disjoint closed curves with $3g-3+k$ curves. With markings, for each pants curve we can, as before, find 3 associated curves, giving $9g-9+3k$ curves C_i so that $i([C_i],[F])$ determine $[F]$.

If S has non-empty boundary, intersections with closed curves do not determine $[F]$, since we allow components parallel to the boundary. We instead look at intersections with arcs. It suffices to take the C_i to be a complete system of non-intersecting arcs in $\mathcal{C}^-(S)$. This gives a decomposition of S into regions that are either hexagons (with alternate boundary components coming from the arcs) or bigons with a puncture. For either type of elementary piece, the intersections $i([C_i],[F])$ for the arcs C_i appearing on the boundary determines $[F]$ within the piece, as seen below.

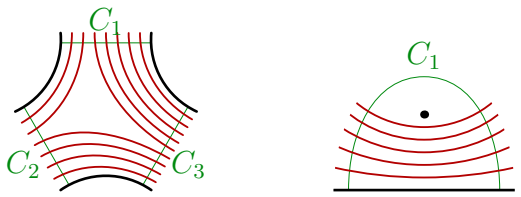

Furthermore, there is no ambiguity about how to glue these foliations for adjacent pieces, and all possibilities for $i([C_i],[F])$ satisfying natural triangle inequalities are possible, so the image is a manifold. (Note we have omitted routine details about doing Whitehead equivalences to make the foliation sit nicely with respect to this decomposition and the like; all of this is done carefully in the case of closed surfaces by Fathi, Laudenbach, and Poénaru [FLP79].) □

Proposition 2.4 can be used to define a topology on $\mathcal{MF}^+(S)$, which we will use.

Proposition 2.5. *The projection map from all measured foliations (not up to equivalence, with its natural function topology) to $\mathcal{MF}^+(S)$ is continuous.*

Proof sketch. For any non-zero measured foliation F_0 and $[C]\in\mathcal{C}^-(S)$, there is a quasi-transverse representative $C_0\in[C]$, which automatically satisfies $i(C_0,F_0)=i([C],[F_0])$. If F_1 is any measured foliation close to F_0, then an analysis of the behavior near singularities shows that there is a representative $C_1\in[C]$ so that C_1 is close to C_0 and C_1 is quasi-transverse with respect to F_1. Then $i([C],[F_1])=i(C_1,F_1)$ and $i(C_1,F_1)$ is close to $i(C_0,F_0)$. □

We can also use Proposition 2.4 to define a map from $\mathcal{C}^+(S)$ to $\mathcal{MF}^+(S)$, sending $[C]\in\mathcal{C}^+(S)$ to the unique measured foliation $[F_C]\in\mathcal{MF}^+(S)$ so that $i([C'],[C])=i([C'],[F_C])$

for all $C' \in \mathcal{C}^-(S)$. This map is an embedding on equivalence classes of weighted simple multi-curves.

Definition 2.6. A *train track* T on a surface S is a graph G embedded in S, so that at each vertex of G (called a *switch*) the incident edges are partitioned into two non-empty subsets that are non-crossing in the cyclic order on the incident vertices. In drawings, the elements of each subset are drawn tangent to each other.

The *complementary regions* of a train track are naturally surfaces with cusps on the boundary. A *taut* train track is a train track with no complementary components that are disks with no cusps or one cusp, or once-punctured disks with no cusps.

Remark 2.7. Many authors (e.g., Penner and Harer [PH92] and Mosher [Mos03]) include our notion of tautness in the definition of a train track, often in a stronger form forbidding bigons (disks with two cusps) and once-punctured monogons as well.

Definition 2.8. The space of positive *transverse measures* or *weights* on a train track T on a surface S is the space $\mathcal{M}(T)$ of assignments of positive numbers ("widths") to edges of the train track so that, at each vertex, the sum of weights on the two sides of the vertex are equal. If $\mathcal{M}(T)$ is non-empty, then T is said to be *recurrent*. We have subspaces $\mathcal{M}_{\mathbb{Q}}(T)$ and $\mathcal{M}_{\mathbb{Z}}(T)$ for transverse measures on T with rational or integral values, respectively. For any train track, there is a natural map $\mathcal{M}_{\mathbb{Z}}(T) \to \mathcal{C}^+(S)$, where we replace an edge of T of weight k by k parallel strands, joining the strands in the natural way at the switches.

Lemma 2.9. *Let T be a recurrent taut train track on S. Then there is a natural continuous map $\mathcal{M}(T) \to \mathcal{MF}^+(S)$ extending $\mathcal{M}_{\mathbb{Z}}(T) \to \mathcal{C}^+(S)$.*

We will denote the map $\mathcal{M}(T) \to \mathcal{MF}^+(S)$ by $w \mapsto T(w)$. If $F = T(w)$ for some w, we say that T *carries* F.

(For convenience in the proof we are assuming the weights on T are strictly positive, but in fact the lemma extends to non-negative weights.)

Proof. Pick a small regular neighborhood $N(T)$ of T, arranged so that $S \setminus N(T)$ has a cusp near each corner where T has a cusp, as illustrated in Figure 1. A weight $w \in \mathcal{M}(T)$ gives a canonical measured foliation $F_N(w)$ on $N(T)$, where an arc cutting across $N(T)$ transverse to a edge e has measure $w(e)$.

Next pick a graph $\Gamma \subset \overline{S} \setminus N(T)$ so that

- Γ contains ∂S,
- Γ has a 1-valent vertex at each cusp of $S \setminus N(T)$ and at each puncture,
- all other vertices of Γ have valence 2 or more, and
- Γ is a spine for $S \setminus N(T)$, i.e., $S \setminus N(T)$ deformation retracts onto Γ.

(The condition that T be taut guarantees that we can find such a Γ). Since Γ is a spine, there is a deformation retraction $\overline{S} \setminus N(T) \to \Gamma$. We can use this to construct a homeomorphism $\phi\colon N(T) \to \overline{S} \setminus \Gamma$ that is the identity on $T \subset N(T)$ and extends to a continuous map $\partial N(T) \to \Gamma$ without backtracking. Then $[\phi(F_N(w))]$ is the desired measured foliation $T(w)$.

As a measured foliation (not up to Whitehead equivalence), $\phi(F_N(w))$ depends continuously on w by construction. The quotient map to the Whitehead equivalence class is continuous. □

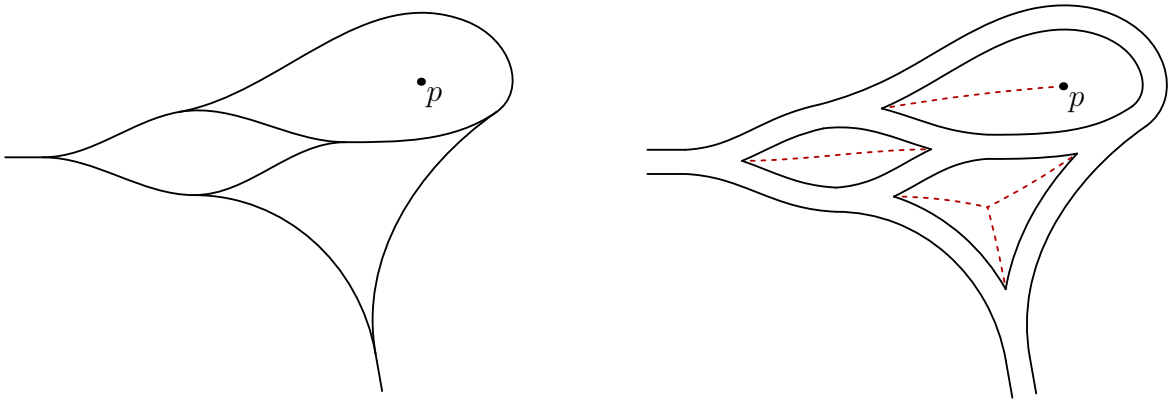

FIGURE 1. Left: A portion of a taut train track T. The point p is a puncture. Right: A neighborhood $N(T)$ of T, together with a spine Γ for $\overline{S} \setminus T$, shown dashed in red.

In Lemma 2.9, if a complementary region of T is a bigon or once-punctured monogon, the corresponding spine is necessarily an interval. Lemma 2.9 is false without the assumption that T is taut; see Example 4.4.

Lemma 2.10. *Every measured foliation F is carried by a taut train track T. Furthermore, T can be chosen so that if F has k zeros on a boundary component, the corresponding complementary component of T has at least k cusps.*

Notice that the number of zeros on a boundary component is not invariant under Whitehead equivalence.

Proof. The techniques here are standard; see, e.g., [Mos03, Proposition 3.6.1], or [PH92, Corollary 1.7.6] for a different approach. Since the definitions we use are slightly different, we sketch the argument.

Pick a set of intervals I_j on S that are transverse to F and cut every leaf of F. These intervals will become the switches of the train track. Let $I = \bigcup_j I_j$.

Divide the leaves of F into *segments* between singularities of F and intersections with I. A *regular* segment is one that intersects I in interior points on both ends. There are only a finite number of non-regular segments (since the number of singularities of F and ends of I is finite), while for any regular segment, nearby segments are isotopic relative to I. There are thus a finite number of classes of parallel regular segments.

Now construct a train track T by taking the union of I and one element of each class of parallel regular segments, and replacing each interval I_j with a single vertex v_j, joined to the same regular segments by connecting arcs. At each switch, the incident edges are divided according to the sides of the corresponding I_j.

Let Γ be the union of the non-regular segments. The components of the complement of T correspond to the connected components of Γ, which is a graph with vertices of valence 1 at cusps of T and possibly at punctures of S, and all other vertices of valence ≥ 2. (That is, Γ is a spine as in the proof of Lemma 2.9.) It follows that T is taut. T carries F with weights equal to the width of the families of parallel segments. If F has k zeros on a boundary component, then T has at least k cusps by construction. □

Proposition 2.11. *$\mathcal{C}^+_{\mathbb{Q}}(S)$ is dense in $\mathcal{MF}^+(S)$.*

Proof. For a given measured foliation F, we will produce a sequence of weighted multi-curves approximating $[F] \in \mathcal{MF}^+(S)$. By Lemma 2.10, $F = T(w)$ for a taut train track T and weight $w \in \mathcal{M}(T)$. Pick a sequence of rational weights $w_n \in \mathcal{M}_{\mathbb{Q}}(T)$ approximating w, and clear denominators to write $w_n = \lambda_n w'_n$ where $w'_n \in \mathcal{M}_{\mathbb{Z}}(T)$. Then $w'_n(T)/\lambda_n$ is a weighted multi-curve approximating F. □

Remark 2.12. On a connected surface S with no boundary, Proposition 2.11 can be strengthened to say that simple curves are projectively dense in measured foliations, as well [Ker80, FLP79]. This strengthening is false for surfaces with boundary. For instance, a pair of pants has only three distinct non-trivial simple curves, but a 3-dimensional space of measured foliations.

3. Conformal Setting

3.1. **Riemann surfaces.**

Definition 3.1. A *Riemann surface* (with boundary) is a smooth surface S, as in Definition 2.1, together with a complex structure on $\overline{S}$, i.e., a fiberwise linear map $J\colon T\overline{S} \to T\overline{S}$ with $J^2 = -\mathrm{id}$.

Convention 3.2. For us, a Riemann surface need not be connected. We only consider surfaces of finite topological type.

Since the complex structure is on $\overline{S}$, not just on S, the complex structure on S° necessarily has a removable singularity near every puncture. On the other hand, a component of ∂S gives an end of the complex structure with a non-removable singularity.

Our definitions of "topological embedding" was chosen in a slightly odd way so that there is a quasi-conformal embedding in every homotopy class. We prove that for reference.

Lemma 3.3. *In the homotopy class of any topological embedding $f\colon R \hookrightarrow S$, there is a quasi-conformal embedding.*

Proof. By standard techniques, the topological embedding $\bar{f}\colon \bar{R} \to \bar{S}$ is homotopic to a smooth orientation-preserving embedding. Specifically, given any topological embedding $\bar{f}$ in the sense of Definition 2.1, shrink the image of $\bar{f}$ slightly so that it is a homeomorphism onto its image. Then use the fact that any homeomorphism between surfaces is isotopic to a diffeomorphism; this was proved for closed surfaces by Epstein [Eps66], and can be extended to open surfaces by looking at the boundary curves [Put16].

Then, since $\bar{R}$ is compact, that smooth map is K-quasi-conformal for some $K < \infty$. □

Definition 3.4. A (holomorphic) *quadratic differential* q on a Riemann surface S is a holomorphic section of the square of the holomorphic cotangent bundle of S°. If z is a local coordinate on S°, we can write $q = \phi(z)\,(dz)^2$ where $\phi(z)$ is holomorphic.

Naturally associated with a quadratic differential we have several objects:

- Local coordinates given by integrating a branch of $\sqrt{q}$ away from the zeros of q. The transition maps are translations or half-turns followed by translations, giving S the structure of a half-translation surface.
- A horizontal measured foliation $\mathcal{F}_h(q) = |\mathrm{Im}\,\sqrt{q}|$. The tangent vectors to the foliation are those vectors $v \in TS$ with $q(v) \geq 0$, and the transverse length of a multi-curve C is

$$\mathcal{F}_h(q)(C) = \int_t \left|\mathrm{Im}\,\sqrt{q(C'(t))}\right| dt,$$

i.e., the total variation of the y coordinate in the half-translation coordinates.

- Similarly, a vertical measured foliation $\mathcal{F}_v(q) = |\mathrm{Re}\,\sqrt{q}|$.
- A locally Euclidean metric $|q|$ on S°, possibly with cone singularities of cone angle $k\pi$ with $k \geq 3$. The length of a multi-curve C with respect to $|q|$ is

$$\ell(C) = \int_t \sqrt{|q(C'(t))|}\, dt.$$

- An area measure A_q on S, the volume measure of $|q|$.

The vector space of finite-area quadratic differentials on S that extend analytically to ∂S (but not necessarily to the punctures) is denoted $\mathcal{Q}(S)$. The finite area constraint implies that at a puncture of S, every $q \in \mathcal{Q}(S)$ has at most a simple pole. That is, if z is a local coordinate on $\overline{S}$ with a puncture at $z = 0$, we can locally write $q = \phi(z)/z\,(dz)^2$ where $\phi(z)$ is holomorphic.

If S has non-empty boundary, then $\mathcal{Q}(S)$ is infinite-dimensional. There is a finite-dimensional subspace $\mathcal{Q}^{\mathbb{R}}(S)$, consisting of those quadratic differentials that are real on vectors tangent to ∂S. There is a convex cone $\mathcal{Q}^+(S) \subset \mathcal{Q}^{\mathbb{R}}(S)$ consisting of those quadratic differentials that are non-negative on ∂S. For non-zero q in $\mathcal{Q}^+(S)$, we have $[\mathcal{F}_h(q)] \in \mathcal{MF}^+(S)$.

3.2. **Extremal length.**

Definition 3.5. For C a multi-curve on a Riemann surface S, pick a Riemannian metric g in the distinguished conformal class. Then the *length* $\ell_g[C]$ is the minimum Riemannian length with respect to g of any rectifiable representative of $[C]$. The *extremal length* of C is

$$\mathrm{EL}_S[C] := \sup_\rho \frac{\ell_{\rho g}[C]^2}{A_{\rho g}(S)}, \tag{3.6}$$

where the supremum runs over all Borel-measurable conformal scale factors $\rho\colon S \to \mathbb{R}_{\geq 0}$ of finite, positive area. (The scaled quantity ρg may give a pseudo-metric rather than a metric, as, e.g., ρ can be 0 on an open subset of S. But we can still define length and area in a natural way.) The definition makes it clear that extremal length does not depend on the metric g within its conformal equivalence class, so we may talk about extremal length on S without reference to g.

When the Riemann surface is clear from context, we suppress it from the notation.

More generally, if $C = \sum_i a_i C_i$ is a weighted multi-curve, then its length is $\ell_g[C] = \sum_i a_i \ell_g[C_i]$, i.e., the corresponding weighted linear combination of lengths of curves, and its extremal length is still defined by Equation (3.6).

We need multi-curves in Definition 3.5, as the main theorems of this paper are false if restricted to curves rather than multi-curves; see Remark 2.12.

We will be interested in simple closed multi-curves C. We must check that extremal length is well-defined on equivalence classes of simple multi-curves. Invariance under homotopy and reparametrization is obvious. Trivial components of a multi-curve C have no effect on $\ell_{\rho g}[C]$ since ρg has finite area, so they also have no effect on extremal length. Finally, let C_0 be a simple multi-curve with parallel components, and let C_1 be the weighted multi-curve with integer weights obtained by merging parallel components and taking the weight to be the number of merged components. Then it is easy to see from the definitions that $\mathrm{EL}[C_0] = \mathrm{EL}[C_1]$.

Furthermore, EL scales quadratically: $\mathrm{EL}[kC] = k^2\,\mathrm{EL}[C]$.

Lemma 3.7. *For any non-trivial multi-curve C on a Riemann surface S, $\mathrm{EL}[C] > 0$. In particular, if S is not small, there is a simple closed curve with non-zero extremal length.*

Proof. Take any finite-area Riemannian metric g on S in the given conformal class. Then, since C has at least one non-trivial component, $\ell_g[C] > 0$, so $\mathrm{EL}[C] > 0$. □

We next give some other interpretations of extremal length for simple multi-curves. First, recall that for a conformal annulus

$$A = \big([0,\ell] \times [0,w]\big)/\big((0,x) \sim (\ell,x)\big),$$

its *modulus* $\mathrm{Mod}(A)$ is w/ℓ. Define the *extremal length* of A to be $\mathrm{EL}(A) \coloneqq 1/\mathrm{Mod}(A) = \ell/w$. Then we can see $\mathrm{EL}[C]$ for a simple multi-curve C as finding the fattest set of conformal annuli around C, in the sense that we minimize total extremal length, as follows.

Proposition 3.8. *Let $C = \bigcup_{i=1}^k C_i$ be a simple closed multi-curve on a Riemann surface S with components C_i. For $i = 1, \dots, k$, let A_i be a (topological) annulus, and let $A = \bigcup_{i=1}^k A_i$. Then*

$$\mathrm{EL}[C] = \inf_{\omega, f} \sum_{i=1}^{k} \mathrm{EL}_\omega(A_i), \tag{3.9}$$

where the infimum runs over all conformal structures ω on A (which amounts to a choice of modulus for each A_i) and over all conformal embeddings $f\colon A \hookrightarrow S$ so that the image of the core curve of A_i is isotopic to C_i.

More generally, if $C = \sum_{i=1}^k a_i C_i$ is a weighted simple closed multi-curve on S, then, with notation as above,

$$\mathrm{EL}[C] = \inf_{\omega, f} \sum_{i=1}^{k} a_i^2 \,\mathrm{EL}_\omega(A_i), \tag{3.10}$$

where the supremum runs over the same set.

We delay the proof of Proposition 3.8 a little.

We can also give a characterization of EL in terms of Jenkins-Strebel differentials.

Definition 3.11. A *Jenkins-Strebel* quadratic differential q on S is one where almost every leaf of $\mathcal{F}_h(q)$ is closed. In this case, the quadratic differential gives a canonical decomposition of S into annuli foliated by the closed leaves.

Theorem 3.12. *Let $C = \bigcup_i a_i C_i$ be a weighted simple closed multi-curve on a Riemann surface S so that no C_i is trivial. Then there is a unique Jenkins-Strebel differential $q_C \in \mathcal{Q}^+(S)$ so that $\mathcal{F}_h(q_C)$ can be decomposed as a disjoint union of annuli A_i with each A_i being a union of leaves of transverse measure a_i and core curve isotopic to C_i. With respect to $|q_C|$, each A_i is isometric to a right Euclidean cylinder.*

For a proof, see, e.g., Strebel [Str84, Theorem 21.1], who attributes the theorem to Hubbard-Masur [HM79] and Renelt [Ren76]. This theorem is one of three different standard theorems on the existence of Jenkins-Strebel differentials.

Proposition 3.13. *For C a weighted simple closed multi-curve on S with no trivial components, let $q = q_C$ be the associated quadratic differential from Theorem 3.12. Then*

$$\mathrm{EL}[C] = A_{|q|}(S).$$

Proposition 3.13 should be standard, but we have been unable to locate it in the literature. We provide a short proof, an easy application of Beurling's criterion.

Proof. We use $|q|$ as the base metric in Equation (3.6) (abusing notation slightly since $|q|$ is not smooth). Let $\ell_i = \ell_{|q|}[C_i]$. Since $|q|$ is a locally CAT(0) metric and local geodesics in locally CAT(0) spaces are globally length-minimizing, ℓ_i is the length in $|q|$ of the core curve of the annulus A_i. (This also follows from Teichmüller's Lemma [Str84, Theorem 14.1].) Then, since $A_q(S) = \sum_i a_i \ell_i$ by the construction of q and $\ell_{|q|}[C] = \sum_i a_i \ell_i$ by definition of $\ell_{|q|}[C]$, we have

$$\frac{\ell_{|q|}[C]^2}{A_q(S)} = A_q(S),$$

so $\mathrm{EL}[C] \geq A_q(S)$.

For the other direction, let ρ be the scaling factor relative to $|q|$ for another metric in the conformal class. For each i and $t \in [0, a_i]$, let $C_i(t)$ be the curve on A_i at distance t from the lower boundary, let $s_i(t) = \int_{C_i(t)} \rho(x)\, dx$, and let $S_i = \min_{t\in[0,a_i]} s_i(t)$. Then, using the Cauchy-Schwarz inequality, we have

$$\begin{aligned}
\ell_{\rho|q|}[C] &\leq \sum\nolimits_i a_i S_i \\
A_{\rho|q|}(S) &= \iint_S \rho^2\, dA_q \geq \frac{1}{A_q(S)} \left(\iint_S \rho\, dA_q \right)^2 \geq \frac{1}{A_q(S)} \left(\sum\nolimits_i a_i S_i \right)^2 \\
\frac{\ell_{\rho|q|}[C]^2}{A_{\rho|q|}(S)} &\leq A_q(S).
\end{aligned}$$

□

Proof of Proposition 3.8. The functional $\sum_{i=1}^k a_i^2 \,\mathrm{EL}(A_i)$ on the space of disjoint embeddings of annuli A_i homotopic to C_i is minimized when the A_i are the annuli from the decomposition of $\mathcal{F}_h(q_C)$ from Theorem 3.12 [Str84, Theorem 20.5]. In this case the value of the functional is $A_{q_C}(S)$, which is equal to $\mathrm{EL}[C]$ by Proposition 3.13. □

More generally, we can work with arbitrary measured foliations, rather than simple closed multi-curves.

Theorem 3.14 (Heights theorem)**.** *Let $[F] \in \mathcal{MF}^+(S)$ be a measured foliation on a Riemann surface S. Then there is a unique quadratic differential $q_F \in \mathcal{Q}^+(S)$ so that $[\mathcal{F}_h(q_F)] = [F]$. Furthermore, q_F depends continuously on F.*

Proofs of Theorem 3.14 have been given by many authors [HM79, Ker80, MS84, Wol96]. Of these, Marden and Strebel [MS84] treat surfaces with boundary. By analogy with Proposition 3.13, we define

$$\mathrm{EL}[F] \coloneqq A_{q_F}(S). \tag{3.15}$$

McMullen has given a definition of $\mathrm{EL}[F]$ closer to Definition 3.5 [McM12, Section 4.4].

3.3. **Stretch factors.** We now turn to a few elementary facts about stretch factors, as already defined in Definition 1.1.

Proposition 3.16. *If $f\colon R \hookrightarrow S$ is a topological embedding of Riemann surfaces where R is not a small surface, then $\mathrm{SF}[f]$ is defined and finite.*

Proof. By Lemma 3.7, there is a simple closed curve C on S with $\mathrm{EL}[C] > 0$, so the set over which we take a supremum is non-empty. The quasi-conformal constant K from Lemma 3.3 then gives an upper bound on the ratio of extremal lengths; this upper bound is usually taken as a definition of K-quasi-conformal. □

Definition 3.17. By analogy with Definition 2.2, we say that a subsurface S' of a surface S is *trivial* if S' is contained in a disk or once-punctured disk inside S.

Proposition 3.18. *For $f\colon R \hookrightarrow S$ a topological embedding of Riemann surfaces where S is not small, $\mathrm{SF}[f] = 0$ if and only if $f(R)$ is trivial as a subsurface of S.*

Proof. If $f(R)$ is trivial in S, it is immediate that $SF[f] = 0$. Otherwise, there is some simple closed curve C on R so that $f(C)$ is nontrivial in S. It follows that C is nontrivial in R, and $\mathrm{SF}[f] \geq \mathrm{EL}[f(C)]/\mathrm{EL}[C] > 0$. □

Proposition 3.19. *If $f\colon S_1 \hookrightarrow S_2$ and $g\colon S_2 \hookrightarrow S_3$ are two topological embeddings of Riemann surfaces, then*

$$\mathrm{SF}[f \circ g] \leq \mathrm{SF}[f] \cdot \mathrm{SF}[g].$$

Proof. Immediate from the definition. □

3.4. **Teichmüller space.** We can assemble the Riemann surface structures on an underlying smooth surface S into the (reduced) *Teichmüller space* $\mathcal{T}(S)$, meaning the space of Riemann surfaces T together with a homeomorphisms $\phi_T\colon S \to T$, considered up to isotopies, taking the boundary to itself but not required to fix it pointwise. The *Teichmüller distance* between two points in $\mathcal{T}(S)$ is defined by

$$d(T, T') := \frac{1}{2} \log K,$$

where K is the minimal stretching of any quasi-conformal homeomorphism f from T to T' so that $(\phi_{T'})^{-1} \circ f \circ \phi_T$ is isotopic to the identity. (Note that this definition uses homeomorphisms, rather than the embeddings used in most of the paper.)

It is a standard result that there is a map f realizing the minimal stretching K, and that its Beltrami differential has the form

$$\mu_f = \frac{K-1}{K+1} \frac{\overline{q}}{|q|} \tag{3.20}$$

for some quadratic differential $q \in \mathcal{Q}^{\mathbb{R}}(T)$. Concretely, we stretch the Euclidean metric $|q|$ along $\mathcal{F}_h(q)$ by a factor of K. (Since q is only real and not positive on ∂T, the foliation $\mathcal{F}_h(q)$ will not in general be in $\mathcal{MF}^+(T)$.) This is usually stated and proved for closed surfaces; the case with boundary follows by considering $T \cup \overline{T}$, the double of T along its boundary, solving the problem in that context, and observing that the quadratic differential of the optimal map is unique, and is thus equivariant with respect to the anti-holomorphic involution that interchanges T and $\overline{T}$; it is therefore real on ∂T.

It follows from Equation (3.20) that

$$\frac{\mathrm{EL}_{T'}(f_* \mathcal{F}_h(q))}{\mathrm{EL}_T(\mathcal{F}_h(q))} = K,$$

and that this is the maximal ratio of extremal lengths. We can approximate $\mathcal{F}_h(q)$ by a weighted multi-curve, possibly with some arc components. We can therefore write the

distance in terms of ratios of extremal lengths. If $f\colon T \to T'$ is a homeomorphism, define a version of the stretch factor by

$$\mathrm{SF}^{\pm}[f] := \sup_{C \in \mathcal{C}^{\pm}_{\mathbb{R}}} \frac{\mathrm{EL}_{T'}[f(C)]}{\mathrm{EL}_{T}[C]}.$$

That is, we allow arc components of the weighted multi-curve; extremal length extends in the natural way to these multi-curves. (Note that if C has arc components, $f(C)$ is only well-defined for homeomorphisms f.) We have $\mathrm{SF}[f] \leq \mathrm{SF}^{\pm}[f]$, since the supremum is over a larger set.

Theorem 3.21. *The Teichmüller distance between $T, T' \in \mathcal{T}(S)$ is*

$$d(T, T') = \frac{1}{2} \log \mathrm{SF}^{\pm}[\mathrm{id}_{T,T'}].$$

Theorem 3.21 was stated and proved by Kerckhoff [Ker80, Theorem 4] for closed surfaces. He furthermore restricted to simple curves (not multi-curves); the technique for the reduction to simple curves cannot be made equivariant with respect to the map interchanging the two components of the mirror of T.

4. Slit maps and Ioffe's theorem

The following terminology is adapted from Ioffe [Iof75] and Fortier Bourque [FB18].

Definition 4.1. On a connected surface S with a non-zero measured foliation F, a *slit* is a finite union of closed segments of leafs of F. (The leaf segments can meet at singularities of F, and so the slit may be a graph.) A *slit complement* in F is the complement of a slit, and a *topological slit map* with respect to F is the inclusion of a slit complement into S. (This is the inclusion of the interiors $R^\circ \hookrightarrow S^\circ$, which extends on a non-injective map $\overline{R} \to \overline{S}$.)

If $f\colon R \hookrightarrow S$ is a slit map with respect to $F \in \mathcal{MF}^+(S)$, then there is a natural pull-back measured foliation $f^*F \in \mathcal{MF}^+(R)$.

If R and S are Riemann surfaces, a *Teichmüller embedding* of dilatation $K \geq 1$ is an embedding $f\colon R \hookrightarrow S$ with quadratic differentials $q_R \in \mathcal{Q}^+(R)$ and $q_S \in \mathcal{Q}^+(S)$ so that f is a topological slit map with respect to $\mathcal{F}_h(q_S)$ and, in the natural coordinates determined by q_R and q_S, the map f has the form $f(x+iy) = Kx+iy$. Note that a Teichmüller embedding is K-quasi-conformal, and that $f^*\mathcal{F}_h(q_S) = \mathcal{F}_h(q_R)$.

Theorem 4.2 (Ioffe [Iof75])**.** *Let R and S be Riemann surfaces, with S connected, and let $f\colon R \hookrightarrow S$ be a topological embedding so that no component of R has trivial image in S. Suppose that f is not homotopic to a conformal embedding. Then there is a quasi-conformal embedding with minimal dilatation in $[f]$. Furthermore, there are unique quadratic differentials $q_R \in \mathcal{Q}^+(R)$ and $q_S \in \mathcal{Q}^+(S)$ so that all quasi-conformal embeddings with minimal dilatation are Teichmüller embeddings with respect to the same quadratic differentials on R and S.*

Remark 4.3. The Teichmüller embedding is not in general unique, but Fortier Bourque proves that two different Teichmüller embeddings differ by translations with respect to the two measured foliations [FB18, Theorem 3.7].

Ioffe's theorem gives a pair of measured foliations on R and S. To relate to Theorem 1, we need to approximate both of these measured foliations by simple closed multi-curves. This is more subtle than it appears at first, since the natural map $f_*\colon \mathcal{C}^+(R) \to \mathcal{C}^+(S)$ does

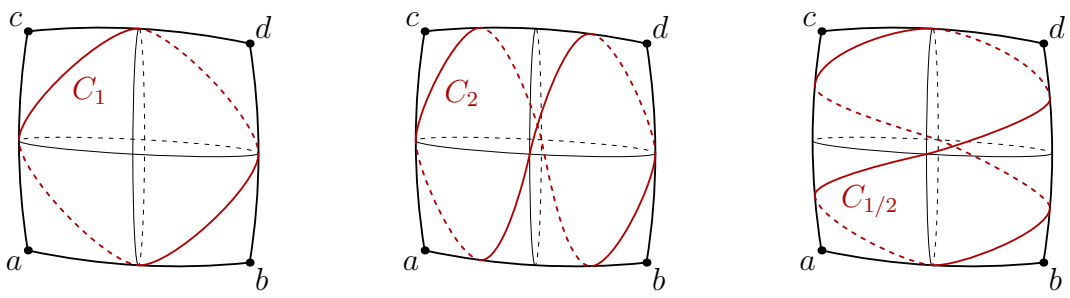

FIGURE 2. Some of the curves C_s in Example 4.4.

not generally extend to a continuous map $\mathcal{MF}^+(R) \to \mathcal{MF}^+(S)$, as the following example shows.

Example 4.4. Let $S = S^2 \setminus \{D_a, D_b, D_c\}$ be the sphere minus three disks and let $R = S^2 \setminus \{D_a, D_b, D_c, D_d\}$ be the subsurface obtained by removing another disk. Pick a set of disjoint arcs $\gamma_{a,b}$, $\gamma_{a,c}$, $\gamma_{b,d}$, and $\gamma_{c,d}$ on R between the respective boundary components. For $s = p/q$ a positive rational number, there is a natural simple curve C_s at slope s with

$$i(\gamma_{a,c}, C_s) = i(\gamma_{b,d}, F_s) = q$$
$$i(\gamma_{a,b}, C_s) = i(\gamma_{c,d}, F_s) = p,$$

as illustrated in Figure 2. Set $F_s \coloneqq (1/q) \cdot [C_s]$ for $s \in \mathbb{Q}_+$, so that

$$i(\gamma_{a,c}, F_s) = i(\gamma_{b,d}, F_s) = 1$$
$$i(\gamma_{a,b}, F_s) = i(\gamma_{c,d}, F_s) = s.$$

Then F_s extends to a continuous family of foliations for $s \in \mathbb{R}_+$.

For $s \in \mathbb{Q}_+$, if we push forward F_s by the inclusion map f, we get a multiple of a simple curve on S. There are only three simple closed curves on S, the curves C_a, C_b, and C_c around the respective boundary component. Which one we get depends only on the parity of p and q, where $s = p/q$ in lowest terms:

$$f_*[F_s] = \frac{1}{q} \cdot \begin{cases} [C_a] & p \text{ odd, } q \text{ odd} \\ [C_b] & p \text{ odd, } q \text{ even} \\ [C_c] & p \text{ even, } q \text{ odd.} \end{cases} \tag{4.5}$$

This map f_* has no continuous extension to $\mathbb{R}_+$.

Example 4.6. We can improve Example 4.4 to avoid dealing with curves around boundary components. Let S' be the surface obtained from the previous surface S by gluing a pair of pants to ∂D_a, ∂D_b, and ∂D_c, and similarly glue a pair of pants to R to get R'. Then S' is a surface of genus two and R' is a surface of genus two minus a disk. Then F_s can be viewed as a continuous family of foliations on R', and Equation (4.5) still holds.

Despite Example 4.4, we can still do simultaneous approximations, using the techniques of Proposition 2.11.

Proposition 4.7. *Let $f\colon R \to S$ be a topological slit map with respect to $F_S \in \mathcal{MF}^+(S)$. Let $F_R = f^*F_S$. Then there is a sequence of simple closed multi-curves C_n on R and weights*

λ_n so that

$$\lim_{n\to\infty} \lambda_n F[C_n] = F_R$$
$$\lim_{n\to\infty} \lambda_n F[f(C_n)] = F_S.$$

Proof. By Lemma 2.10, $[F_R] = T_R(w)$ for some weight w on a taut train track T_R on R. Fix a boundary component B of R, and let β be a curve parallel to B slightly pushed in to R. If $f(\beta)$ bounds a disk in S, the corresponding slit of F_S is a tree which must have at least two endpoints. Each endpoint of the tree contributes a zero to F_R on B, so F_R has at least two zeros on B.

Likewise, if $f(\beta)$ bounds a once-punctured disk in S, the corresponding slit of F_S is a tree with at least two endpoints. At most one of these endpoints may be at the puncture, so F_R has at least one zero on B.

Let $T_S = f(T_R)$. The second part of the statement of Lemma 2.10 guarantees that T_S is taut, and so $F_S = T_S(w)$. (The new disks in T_S that were not disks in T_R have at least two cusps, and the new once-punctured disks have at least one cusp.)

As in the proof of Proposition 2.11, choose a sequence of rational weights $w_n \in \mathcal{M}_{\mathbb{Q}}(T_R)$ approaching w, and choose scalars λ_n so that $w'_n := w_n/\lambda_n$ is integral. Then $T_R(w'_n)$ is a multi-curve $[C_n]$ with $\lambda_n[C_n]$ approaching $[F_R]$. We also have $[f(C_n)] = T_S(w'_n)$, so by Lemma 2.9, $\lambda_n[f(C_n)]$ approaches F_S. □

Proof of Theorem 1. If $S = \bigcup_i S_i$ is not connected, with $R_i = f^{-1}(S_i)$, then the stretch factor is a supremum over all embeddings $R_i \hookrightarrow S_i$, as $\frac{a+b}{c+d} < \max\left(\frac{a}{c}, \frac{b}{d}\right)$. On the other hand, R conformally embeds in S if and only if R_i conformally embeds in S_i in the given homotopy class for all i. So from now on we assume that S is connected.

If $f\colon R \hookrightarrow S$ is homotopic to a conformal embedding, then Proposition 3.8 guarantees that for all multi-curves $[C] \in \mathcal{C}^+(R)$, we have $\mathrm{EL}_S[f(C)] \le \mathrm{EL}_R[C]$, as we have more maps in computing $\mathrm{EL}_S[f(C)]$, so smaller infimum in Equation (3.10). Thus $\mathrm{SF}[f] \le 1$.

Conversely, suppose f is not homotopic to a conformal embedding. Then by Theorem 4.2, f is homotopic to a Teichmüller map g of dilatation K with respect to quadratic differentials $q_R \in \mathcal{Q}^+(R)$ and $q_S \in \mathcal{Q}^+(S)$. Apply Proposition 4.7 to find a sequence of simple closed multi-curves C_n on R and weights λ_n so that $\lambda_n[C_n]$ approximates $\mathcal{F}_h(q_R)$ and $\lambda_n[f_*C_n]$ approximates $\mathcal{F}_h(q_S)$. By Theorem 3.14, the quadratic differentials corresponding to $\lambda_n[C_n]$ approach q_R and the quadratic differentials corresponding to $\lambda_n[f(C_n)]$ approach q_S, and therefore

$$\mathrm{SF}[f] \ge \lim_{n\to\infty} \frac{\mathrm{EL}_S[f(C_n)]}{\mathrm{EL}_R[C_n]} = \frac{\mathrm{EL}_S[\mathcal{F}_h(q_S)]}{\mathrm{EL}_R[\mathcal{F}_h(q_R)]} = \frac{\|q_S\|}{\|q_R\|} = K > 1. \tag{4.8}$$

This completes the proof of Theorem 1. □

When the stretch factor is larger than 1, we find it exactly (Proposition 1.2) with the following standard fact.

Lemma 4.9. *Let $f\colon R \hookrightarrow S$ be a quasi-conformal embedding of Riemann surfaces with quasi-conformal constant $\le K$, and let C be any closed multi-curve on R. Then*

$$\mathrm{EL}_S[f(C)] \le K\,\mathrm{EL}_R[C].$$

Proof of Proposition 1.2. We can again assume that S is connected. If $\mathrm{SF}[f] = 1$, the result is trivial: By Theorem 1, there is a conformal embedding, which has quasi-conformal constant

equal to 1. If $\mathrm{SF}[f] > 1$, then by Theorem 1, the map f is not homotopic to a conformal embedding. Ioffe's Theorem 4.2 constructs a K-quasi-conformal map. We have $\mathrm{SF}[f] \leq K$ by Lemma 4.9, and $\mathrm{SF}[f] \geq K$ by Equation (4.8). □

5. Strict embeddings

We now turn to Theorem 2, on embeddings with stretch factor strictly less than 1. We start with some preliminary lemmas.

Lemma 5.1. *Let $f\colon R \hookrightarrow S$ be a strict conformal embedding. Then there is a constant $K < 1$ so that for any $q \in \mathcal{Q}^+(S)$,*

$$A_q(f(R)) \leq K A_q(S).$$

Proof. For any non-zero quadratic differential q on S, the ratio $A_q(f(R))/A_q(S)$ is less than 1, as the open set missed by the image of f has some non-zero area with respect to q. Then $A_q(f(R))/A_q(S)$ is a continuous function on the projective space $P\mathcal{Q}^+(S)$. Since $P\mathcal{Q}^+(S)$ is compact, the result follows. □

Later, in Theorem 4, we will strengthen Lemma 5.1 considerably.

Lemma 5.2. *Let R be a compact Riemann surface with a quadratic differential $q \in \mathcal{Q}^+(R)$ that is strictly positive on ∂R. Let $\widehat{R}_t$ be the annular extension of R obtained by gluing a Euclidean cylinder of width t onto each boundary component of R with respect to the locally Euclidean metric given by q. Then*

$$\lim_{t\to 0} \mathrm{SF}[\widehat{R}_t \to R] = 1,$$

where $\mathrm{SF}[\widehat{R}_t \to R]$ is the stretch factor of the obvious homotopy class of topological embeddings.

Proof. By Proposition 1.2, it suffices to construct a family of quasi-conformal maps $f_t\colon \widehat{R}_t \to R$ with quasi-conformal constant K_t that approaches 1 as t approaches 0. The assumption that q is positive on ∂R guarantees that near each component C_i of ∂R there is an annulus A_i foliated by leaves of $\mathcal{F}_h(q)$, with circumference r_i and width w_i (with respect to the Euclidean metric induced by q). Let $B_{i,t}$ be the annulus added to this boundary component in $\widehat{R}_t$, and let $\iota_t\colon \widehat{R}_t \to R$ be the affine map of $A_i \cup B_{i,t}$ onto A_i and the identity outside of $A_i \cup B_{i,t}$. Then ι_t has quasi-conformal constant equal to

$$\max_i \frac{w_i + t}{w_i},$$

which goes to 1 as $t \to 0$ as desired. □

Proof of Theorem 2. **(2) ⇒ (1):** An annular conformal embedding is also a strict conformal embedding, so this is clear.

(1) ⇒ (4): Suppose that f is a strict conformal embedding, and let $K < 1$ be the constant from Lemma 5.1. For any multi-curve $[C] \in \mathcal{C}^+(R)$, let $q \in \mathcal{Q}^+(S)$ be the quadratic differential that realizes extremal length for $[f(C)]$, and consider the pull-back metric $\mu = f^*|q|$ on R. Since (R, μ) and $(f(R), |q|)$ are isometric, but there are more curves in the homotopy class $[C]$ on S than those that lie in $f(R)$, we have $\ell_\mu[C] \geq \ell_{|q|}[f(C)]$. Therefore,

$$\mathrm{EL}_R[C] \geq \frac{\ell_\mu[C]^2}{A_\mu(R)} \geq \frac{\ell_{|q|}[f(C)]^2}{K A_{|q|}(S)} = K^{-1}\,\mathrm{EL}_S[f(C)]$$

implying that $\frac{\mathrm{EL}[f(C)]}{\mathrm{EL}[C]} \leq K$. Since C was arbitrary, $\mathrm{SF}[f] \leq K$.

(4) ⇒ (2): Suppose that $\mathrm{SF}[f] < 1$. Pick a quadratic differential $q \in \mathcal{Q}^+(R)$ that is real and strictly positive on ∂R. Let $\widehat{R}_t$ be the family of annular extensions of R with respect to q as in Lemma 5.2, and let $\widehat{f}_t : \widehat{R}_t \to S$ be the composite topological embeddings. Then by Proposition 3.19,

$$\mathrm{SF}[\widehat{f}_t] \leq \mathrm{SF}[\widehat{R}_t \to R] \cdot \mathrm{SF}[f].$$

It follows from Lemma 5.2 that for t sufficiently small, $\mathrm{SF}[\widehat{f}_t] \leq 1$, so by Theorem 1, $\widehat{f}_t$ is homotopic to a conformal embedding.

(4) ⇔ (3): This is a consequence of Proposition 1.7, which we prove next. □

Proof of Proposition 1.7. We first handle the harder case when $\mathrm{SF}[f] \leq 1$. By compactness of balls in Teichmüller space, it suffices to show, on one hand, that if $d(S, S') < -\frac{1}{2}\log \mathrm{SF}[f]$, then there is a conformal embedding of R in S'; and, on the other hand, that there are surfaces S' with $d(S, S')$ arbitrarily close to $-\frac{1}{2}\log \mathrm{SF}[f]$ so that R does not conformally embed in S'.

For the first part, suppose $d(S, S') < -\frac{1}{2}\log \mathrm{SF}[f]$. Let $\mathrm{id}_{S,S'}$ be the identity map from the marking. Then

$$\mathrm{SF}[\mathrm{id}_{S,S'} \circ f] \leq \mathrm{SF}[f] \cdot \mathrm{SF}[\mathrm{id}_{S,S'}] \leq \mathrm{SF}[f] \cdot \mathrm{SF}^{\pm}[\mathrm{id}_{S,S'}] = \mathrm{SF}[f] \cdot e^{2d(S,S')} < 1.$$

as desired.

To get the other direction of the inequality, pick $\varepsilon > 0$, and set $K = e^{\varepsilon}/\mathrm{SF}[f]$ and $\lambda = \frac{K-1}{K+1}$. Find a simple closed multi-curve C on R near the supremum defining SF:

$$\frac{\mathrm{EL}_S[f(C)]}{\mathrm{EL}_R[C]} > e^{-\varepsilon}\,\mathrm{SF}[f].$$

Let $q = q_{f(C)} \in \mathcal{Q}^+(S)$ be the associated Jenkins-Strebel quadratic differential, and set $\mu = \lambda \cdot \overline{q}/|q|$ to be an associated Beltrami differential. Let S' be S stretched by μ, so that

$$d(S, S') = \frac{1}{2}\log \mathrm{SF}^{\pm}[\mathrm{id}_{S,S'}] = \frac{1}{2}\log \mathrm{SF}[\mathrm{id}_{S,S'}] = \frac{1}{2}\frac{\mathrm{EL}_{S'}[f(C)]}{\mathrm{EL}_S[f(C)]} = \frac{\log K}{2} = -\frac{1}{2}\log \mathrm{SF}[f] + \frac{\varepsilon}{2}.$$

We also have

$$\mathrm{SF}[\mathrm{id}_{S,S'} \circ f] \geq \frac{\mathrm{EL}_{S'}[f(C)]}{\mathrm{EL}_R[C]} > \left(e^{\varepsilon}/\mathrm{SF}[f]\right)\left(e^{-\varepsilon}\,\mathrm{SF}[f]\right) = 1$$

so $S' \notin \mathcal{T}_R(S)$. Since ε can be chosen arbitrarily small, we get the desired result.

For the case $\mathrm{SF}[f] > 1$, Ioffe's Theorem 4.2 applies and gives us a quadratic differential q_S on S. Stretching along this quadratic differential gives us a structure S_1 at distance $d = \frac{1}{2}\log \mathrm{SF}[f]$ from S for which there is a conformal embedding. For a structure S' at distance less than d, consider the map $f_{S'}$ from R to the surface with the structure S'; we then have

$$\mathrm{SF}[f] \leq \mathrm{SF}[\mathrm{id}_{S,S'}] \cdot \mathrm{SF}[f_{S'}] \leq e^{2d}\,\mathrm{SF}[f_{S'}].$$

Combined, this implies that $\mathrm{SF}_{R,S'}[f] > 1$. Since this holds for every such S', we have $d(S, \partial\mathcal{T}_R(S)) = d(S, S_1)$ as desired. □

Remark 5.3. It follows from the proof that the stretching to a nearest point on $\partial\mathcal{T}_R(S)$ is horizontal on the boundary.

Remark 5.4. The nearest point to S on $\partial\mathcal{T}_R(S)$ is not always unique, as we can see from the fact that $\widetilde{\mathrm{SF}}[f] \neq \mathrm{SF}[f]$ in Examples 6.6 and 6.7 below. Indeed, let $f\colon R \hookrightarrow S$ be a conformal embedding and let $\widetilde{f}\colon \widetilde{R} \hookrightarrow \widetilde{S}$ be a regular covering with $\mathrm{SF}[f] < \mathrm{SF}[\widetilde{f}] < 1$. Then if there were a unique nearest point $\widetilde{S}'$ to $\widetilde{S}$ on $\partial\mathcal{T}_{\widetilde{R}}(\widetilde{S})$, it would be invariant under the deck transformations, and so would descend to give a point S' on $\partial\mathcal{T}_R(S)$, contradicting $\mathrm{SF}[f] < \mathrm{SF}[\widetilde{f}]$.

6. Behavior under finite covers

We now turn to the behavior of the stretch factor under finite covers. We start with some easy statements.

Lemma 6.1. *Let $\pi\colon \widetilde{S} \to S$ be a covering map of Riemann surfaces of finite degree d. For C a weighted multi-curve on S, define $\pi^{-1}C$ to be the full inverse image of C, with the same weights. Then $\mathrm{EL}_{\widetilde{S}}[\pi^{-1}C] = d\,\mathrm{EL}_S[C]$.*

Proof. By Proposition 3.13, $\mathrm{EL}_S[C] = A_{q_C}(S)$, where q_C is the Jenkins-Strebel quadratic differential corresponding to C. Then $f^*(q_C)$ is a Jenkins-Strebel quadratic differential corresponding to $\pi^{-1}(C)$, and so

$$\mathrm{EL}_{\widetilde{S}}[\pi^{-1}(C)] = A_{f^*(q_C)}(S) = dA_{q_C}(S) = d\,\mathrm{EL}_S[C]. \qquad \square$$

Lemma 6.2. *For $\widetilde{f}$ a finite cover of $f\colon R \hookrightarrow S$, we have $\mathrm{SF}[\widetilde{f}] \geq \mathrm{SF}[f]$.*

Proof. Follows from Lemma 6.1 and the definition of SF, as the supremum involved in computing $\mathrm{SF}[\widetilde{f}]$ is over a larger set. $\square$

Proposition 6.3. *If $f\colon R \hookrightarrow S$ is a topological embedding of Riemann surfaces with $\mathrm{SF}[f] \geq 1$ and $\widetilde{f}$ is a finite cover of f in the sense of Definition 1.8, then $\mathrm{SF}[\widetilde{f}] = \mathrm{SF}[f]$.*

Proof. If $\mathrm{SF}[f] = 1$, the result follows from Lemma 6.2 and Theorem 1.

If $\mathrm{SF}[f] > 1$, by Proposition 1.2 $\mathrm{SF}[f]$ is the minimal quasi-conformal constant of any map homotopic to f, which by Theorem 4.2 is given by a Teichmüller embedding g. Let $\widetilde{g}$ be the corresponding cover of g. Then $\widetilde{g}$ is also a Teichmüller embedding with the same quasi-conformal constant, and so $\mathrm{SF}[\widetilde{f}]$ is the quasi-conformal constant of $\widetilde{g}$. $\square$

Remark 6.4. Proposition 6.3 relies on $\widetilde{f}$ being a cover of finite degree of f. McMullen [McM89, Corollary 1.2] shows that, in the case that R and S are closed surfaces, f is a Teichmüller map, and $\widetilde{f}$ is a non-amenable cover of f, then $\widetilde{f}$ does *not* minimize the quasi-conformal distortion in its bounded homotopy class.

Proposition 6.5. *For $\widetilde{f}$ a finite cover of $f\colon R \hookrightarrow S$, the quantity $\mathrm{SF}[\widetilde{f}]$ is less than one, equal to one, or greater than one exactly when $\mathrm{SF}[f]$ is less than one, equal to one, or greater than one.*

Proof. If $\mathrm{SF}[f] < 1$, by Theorem 2, f is homotopic to a strict conformal embedding. Since a cover of a strict conformal embedding is a strict conformal embedding, we have $\mathrm{SF}[\widetilde{f}] < 1$. The other cases follow from Proposition 6.3. $\square$

Although there is some good behavior, it is not true in general that $\mathrm{SF}[\widetilde{f}] = \mathrm{SF}[f]$.

Example 6.6. Let R and S both be disks with two points removed, with $f\colon R \to S$ a strict conformal embedding and $g\colon S \to R$ a homotopy inverse. The surfaces R and S have, up to equivalence and scale, only one non-trivial simple closed multi-curve (the boundary-parallel curve), so $\mathrm{SF}[f] = 1/\mathrm{SF}[g]$. Also, $SF[f] < 1$, since f was assumed to be a strict conformal embedding. Now take any non-trivial finite cover $\widetilde{R}$ of R and the corresponding cover $\widetilde{S}$ of S. Let the corresponding topological embeddings be $\widetilde{f}\colon \widetilde{R} \to \widetilde{S}$ and $\widetilde{g}\colon \widetilde{S} \to \widetilde{R}$. Since $\mathrm{SF}[g] > 1$, by Proposition 6.3 we have $\mathrm{SF}[\widetilde{g}] = \mathrm{SF}[g]$, with the supremum in the definition of stretch factor realized by a symmetric multi-curve. By Theorem 4.2, the quadratic differentials realizing this stretch factor are *unique*, so for *any* non-symmetric multi-curve C on $\widetilde{S}$ (or equivalently $\widetilde{R}$), we have

$$\mathrm{SF}[g] = \mathrm{SF}[\widetilde{g}] > \frac{\mathrm{EL}_{\widetilde{R}}[C]}{\mathrm{EL}_{\widetilde{S}}[C]}.$$

But then

$$\mathrm{SF}[\widetilde{f}] \geq \frac{\mathrm{EL}_{\widetilde{S}}[C]}{\mathrm{EL}_{\widetilde{R}}[C]} > 1/\mathrm{SF}[g] = \mathrm{SF}[f].$$

Example 6.7. The previous example can be improved to give an examples with arbitrarily large gap between SF and $\widetilde{\mathrm{SF}}$: for any $0 < \varepsilon < \delta < 1$, there is an embedding $f\colon R \hookrightarrow S$ and two-fold cover $\widetilde{f}$ so that $\mathrm{SF}[f] < \varepsilon$ and $\mathrm{SF}[\widetilde{f}] > \delta$. This example is due to Maxime Fortier-Bourque. Let R_t be the disk with two punctures obtained by doubling a $t \times 1$ rectangle along three of its sides, and let S_t be the double cover of R_t branched along one of the two punctures. Then for $s < t$ the embedding $S_s \hookrightarrow S_t$ is a cover of the embedding $R_s \hookrightarrow R_t$.

Let C_1 be the only non-trivial simple closed curve on R_t, the curve parallel to the boundary as shown on the left of Figure 3. Let C_2 be the non-symmetric curve on S_t shown on the right of Figure 3. By construction, $\mathrm{EL}_{R_t}[C_1] = 2/t$. As $t \to \infty$, the surface S_t approaches a sphere with 4 punctures, specifically the double of a square. The curve C_2 is non-trivial on the 4-punctured sphere, and so its extremal length approaches a definite value:

$$\lim_{t\to\infty} \mathrm{EL}_{S_t}[C_2] = 2.$$

Thus, for $t \gg s \gg 0$, we have

$$\mathrm{SF}[R_s \hookrightarrow R_t] = \frac{2/t}{2/s} = \frac{s}{t}$$
$$\mathrm{SF}[S_s \hookrightarrow S_t] \geq \frac{\mathrm{EL}_{S_t}[C_2]}{\mathrm{EL}_{S_s}[C_2]} \to 1,$$

as desired.

With a little more care, one can show that $\mathrm{EL}_{S_t}[C_2] \approx 2(1 + Ke^{\pi t/2})$ for some constant K. This uses the uniformization of S_∞ to the double of a square by the composition of $z \mapsto \sin(\pi i z/2)$ and $z \mapsto \int_{w=0}^{z} dw/\sqrt{w^3 - w}$.

In order to prove Theorem 3, we need some extra control: a strengthening of Lemma 5.1.

Theorem 4. *Let $f\colon R \hookrightarrow S$ be a annular conformal embedding of Riemann surfaces. Then there is a constant $K < 1$ so that for any quadratic differential $q \in \mathcal{Q}(S)$,*

$$A_{f^*q}(R) \leq K A_q(S).$$

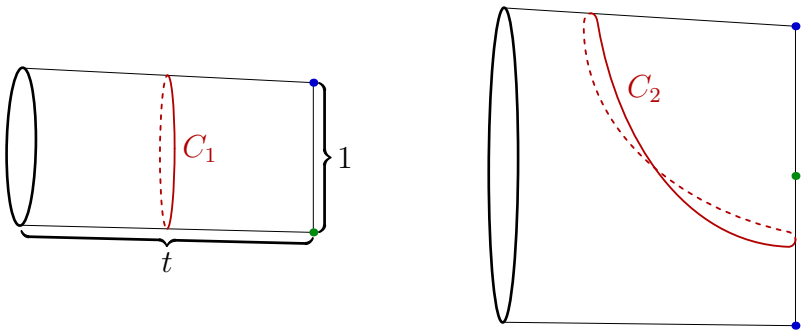

FIGURE 3. The surfaces from Example 6.7. Left: The family of surfaces R_t, with the unique non-trivial simple closed curve C_1. Right: the double cover S_t, with the curve C_2.

Furthermore, the constant K can be chosen uniformly under finite covers, in the sense that for any finite cover $\widetilde{f}\colon \widetilde{R} \to \widetilde{S}$ of f and any quadratic differential $\widetilde{q} \in \mathcal{Q}(\widetilde{S})$,

$$A_{f^*\widetilde{q}}(\widetilde{R}) \leq K A_{\widetilde{q}}(\widetilde{S}).$$

The technique in Lemma 5.1 will not work to prove Theorem 4, as $\mathcal{Q}(S)$ is infinite-dimensional. (That bound is also not uniform under covers.) As in Lemma 5.1, K depends on the actual embedding, not just the homotopy class of the embedding.

When S is a disk, Theorem 4 is not hard. For $a \in \mathbb{C}$ and $r > 0$, we denote by $\mathbb{D}(a, r) = \{z : |z - a| < r\}$ the open disk of radius r about a.

Proposition 6.8. *Let $\Omega \subset \mathbb{D}$ be an open subset of the disk so that $\overline{\Omega} \cap \partial\mathbb{D} = \emptyset$. For any quadratic differential $q \in \mathcal{Q}(\mathbb{D})$,*

$$A_q(\Omega) \leq r^2 A_q(\mathbb{D}),$$

where r is large enough so that $\Omega \subset \mathbb{D}(0, r)$.

Proposition 6.8 is a special case of Proposition 6.9 below, but we give a separate proof because we can give a precise constant.

Proof. Let r_0 be the smallest value so that $\Omega \subset \mathbb{D}(0, r_0) \subset \mathbb{D}$, and let $q \in \mathcal{Q}(\mathbb{D})$ be arbitrary. For $0 \leq r \leq 1$, we will show that $A_q(\mathbb{D}(0, r)) \leq r^2 \cdot A_q(\mathbb{D})$, so that $K = r_0^2$ suffices. Define

$$I(r) = \int_{\theta=0}^{2\pi} |q(re^{i\theta})|\, d\theta$$
$$J(r) = \int_{s=0}^{r} sI(s)\, ds = A_q(\mathbb{D}(0, r)),$$

where we are writing $q = q(z)\,(dz)^2$ with $q(z)$ a holomorphic function. The function $z \mapsto |q(z)|$ is subharmonic, so if $s < r$, we have $I(s) \leq I(r)$. (We would have equality between the corresponding integrals if $|q(z)|$ were harmonic; see, e.g., [Bur79, p. 142]). We therefore have $J(r) = \int_{s=0}^{r} sI(s)\, ds \leq r^2 I(r)/2$, and so

$$\frac{d}{dr}\frac{J(r)}{r^2} = \frac{rJ'(r) - 2J(r)}{r^3} \geq \frac{r^2 I(r) - r^2 I(r)}{r^3} = 0.$$

It follows that $J(r)/r^2 \leq J(1)$, as desired. □

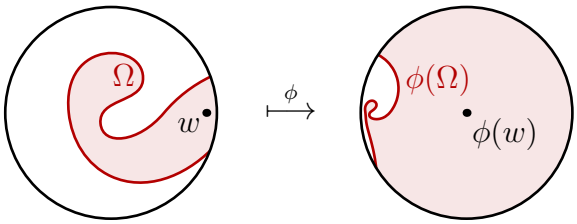

FIGURE 4. Möbius transformations to make the area of a quadratic differential be concentrated near a point w that is close to $\partial\mathbb{D}$.

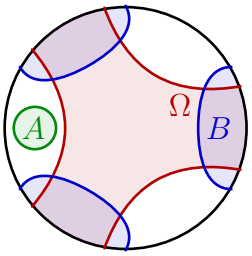

FIGURE 5. The schematic setup of Proposition 6.9.

Proposition 6.8 is false if $\overline{\Omega}$ is allowed to intersect $\partial\mathbb{D}$. Suppose Ω contains a neighborhood of a segment of $\partial\mathbb{D}$, and let w be a point very close to this segment. By a conformal automorphism ϕ of $\mathbb{D}$, we can take w to the center of the disk. Then $\big(d\phi(z)\big)^2$ will have its measure concentrated near $w \in \Omega$, as illustrated in Figure 4.

The following proposition says that this is all that can happen: if the mass of q on Ω gets large, then the mass of q is concentrating near $\partial\mathbb{D}$.

Proposition 6.9. *Let $\Omega \subset \mathbb{D}$ be an open subset of the disk with an open set A in its complement, and let $B \subset \overline{\mathbb{D}}$ be a neighborhood of $\overline{\Omega} \cap \partial\overline{\mathbb{D}}$, as illustrated in Figure 5. Then, for every $\varepsilon > 0$, there is a $\delta > 0$ so that if $q \in \mathcal{Q}(\mathbb{D})$ is such that $q \neq 0$ and*

$$\frac{A_q(\Omega)}{A_q(\mathbb{D})} > 1 - \delta,$$

then

$$\frac{A_q(B)}{A_q(\mathbb{D})} > 1 - \varepsilon.$$

The proposition implies that given a sequence $q_n \in \mathcal{Q}(\mathbb{D})$, if the percentage of the $|q_n|$-area of $\mathbb{D}$ occupied by Ω tends to 1, then the percentage of the $|q_n|$-area occupied by the set B of "thickened ends of Ω" also tends to 1. Figure 4 again provides an example of how this happens.

We give two versions of the proof, one shorter, and the other more explicit and giving (poor) bounds on the constants.

Proof of Proposition 6.9, version 1. If there are no such bounds as in the statement of the proposition, there is an $0 < \varepsilon < 1$ and a sequence of quadratic differentials $q_n \in \mathcal{Q}(\mathbb{D})$ so

that

$$A_{q_n}(\mathbb{D}) = 1 \tag{6.10}$$

$$A_{q_n}(B) < 1 - \varepsilon \tag{6.11}$$

$$A_{q_n}(\Omega) > 1 - 1/n. \tag{6.12}$$

Consider A_{q_n} as a measure on $\overline{\mathbb{D}}$. Since the space of measures of unit area on the closed disk is compact in the weak topology, after passing to a subsequence we may assume that A_{q_n} converges (weakly) to some limiting measure μ (of total mass 1) on $\overline{\mathbb{D}}$. Since holomorphic functions on the disk that are also in $L^1(\mathbb{D})$ form a normal family, after passing to a further subsequence, we may assume that the sequence q_n converges locally uniformly to some holomorphic function q_∞ on $\mathbb{D}$. The restriction of μ to the open disk is then A_{q_∞}. But $A_{q_n}(A) < 1/n$, so $A_{q_\infty}(A) = 0$, so q_∞ is identically 0 on A and therefore on the entire open disk. Hence μ is supported on $\partial\overline{\mathbb{D}}$. Equation (6.12) implies that the support of μ is also contained in $\overline{\Omega}$, and hence in $\overline{\Omega} \cap \partial\overline{\mathbb{D}}$. But this contradicts Equation (6.11). □

Proof of Proposition 6.9, version 2. Apply a Möbius transformation so that A contains 0. We may then assume that $\Omega \subset \mathbb{D} \setminus \overline{\mathbb{D}(0, 2r_0)}$ for some $0 < r_0 < 1/2$. We identify the space $\mathcal{Q}(\mathbb{D})$ of integrable holomorphic quadratic differentials on $\mathbb{D}$ with the Banach space of L^1-integrable holomorphic functions on $\mathbb{D}$, so that $A_q(\mathbb{D}) = \int_{\mathbb{D}} |q| = \|q\|$.

Suppose $q \in \mathcal{Q}(\mathbb{D})$ satisfies $A_q(\mathbb{D}) = 1$. We will quantitatively show that the q-area of a small ball controls the q-area of a big ball. Suppose s is chosen close to 1 with $r_0 < s < 1$. Suppose $|z| \le s$. The Cauchy Integral Formula applied to the concentric circles comprising the disk $\mathbb{D}(z, 1-s)$ shows that

$$|q(z)| \le \frac{1}{\pi(1-s)^2} \int_{\mathbb{D}(z,1-s)} |q| = \frac{1}{\pi(1-s)^2} A_q(\mathbb{D}(z, 1-s)),$$

i.e., $|q|$ is subharmonic. Using the assumption that $A_q(\mathbb{D}) = 1$, this implies

$$|z| \le s \implies |q(z)| \le K(s) := \frac{1}{\pi(1-s)^2}. \tag{6.13}$$

Similar reasoning shows

$$|z| \le r_0 \implies |q(z)| \le \frac{1}{\pi r_0^2} A_q(\mathbb{D}(0, 2r_0)). \tag{6.14}$$

For $0 < t < 1$, let $M_q(t)$ be $\max\{|q(z)| : |z| = t\}$. The Hadamard Three Circles Theorem [Con78, Theorem 6.3.13] implies that $\log M_q$ is a convex function of $\log t$. Thus if r and r_1 are chosen so that $r_0 \le r \le r_1 < s$ then

$$\begin{aligned}
\log M_q(r) &\le \log M_q(r_0) + \frac{\log M_q(s) - \log M_q(r_0)}{\log s - \log r_0}(\log r - \log r_0) \\
&\le \log M_q(r_0) + \frac{\log K(s) - \log M_q(r_0)}{\log s - \log r_0}(\log r_1 - \log r_0) \\
&= \left(1 - \frac{\log r_1 - \log r_0}{\log s - \log r_0}\right) \log M_q(r_0) + \log K(s) \frac{\log r_1 - \log r_0}{\log s - \log r_0} \\
&= K_1 \log M_q(r_0) + K_2
\end{aligned}$$

where K_1 and K_2 are constants, with $K_1 > 0$, depending only on r_0, r_1, and s, and not on q. It follows from (6.14) that there are positive constants c_1 and c_2 depending only on r_0, r_1, and s with

$$A_q(\mathbb{D}(0, r_1)) < c_2 A_q(\mathbb{D}(0, 2r_0))^{c_1}. \tag{6.15}$$

Now suppose that δ is small, $0 < \delta < 1$, and $A_q(\mathbb{D} \setminus \Omega) < \delta$. Note that this implies that $A_q(\mathbb{D}(0, 2r_0)) < \delta$. Given $0 < r_1 < 1$, let E be the annulus $\mathbb{D} \setminus \mathbb{D}(0, r_1)$. From the definition of B, there is some r_1 with $0 < r_1 < 1$ close to 1 for which $E \cap \Omega \subset E \cap B$. Choose s so that $r_0 < r_1 < s < 1$; we are in the setup of the previous paragraph. We have

$$\begin{aligned} 1 - c_2\delta^{c_1} &< A_q(E) && \text{by (6.15)} \\ &= A_q(E \cap (\mathbb{D} \setminus \Omega)) + A_q(E \cap \Omega) \\ &< A_q(\mathbb{D} \setminus \Omega) + A_q(E \cap B) \\ &< \delta + A_q(B) \end{aligned}$$

and so $A_q(B) > 1 - c_2\delta^{c_1} - \delta$, which tends to 1 as δ tends to 0, as required. □

We also need an analogue of Proposition 6.9 for the once-punctured disk. (In fact it is true in more generality.)

Proposition 6.16. *Let $\mathbb{D}^\times$ be the punctured unit disk $\mathbb{D} \setminus \{0\}$, let $\Omega \subset \mathbb{D}^\times$ be an open subset with an open set A in its complement, and let $B \subset \overline{\mathbb{D}}^\times$ be an open neighborhood of $\overline{\Omega} \cap \partial\overline{\mathbb{D}}$. Then, for every $\varepsilon > 0$, there is a $\delta > 0$ so that if $q \in \mathcal{Q}(\mathbb{D}^\times)$ is such that $q \neq 0$ and*

$$\frac{A_q(\Omega)}{A_q(\mathbb{D})} > 1 - \delta,$$

then

$$\frac{A_q(B)}{A_q(\mathbb{D})} > 1 - \varepsilon.$$

Proof. Let $s\colon \mathbb{D} \to \mathbb{D}$ be the squaring map $s(z) = z^2$. We can apply Proposition 6.9 to the tuple $(s^{-1}(\Omega), s^{-1}(A), s^{-1}(B))$. For every quadratic differential $q \in \mathcal{Q}(\mathbb{D}^\times)$ with at most a simple pole at 0, s^*q is a quadratic differential on $\mathbb{D}^\times$ with no pole, and can thus be considered as a quadratic differential on $\mathbb{D}$. Since for any $X \subset \mathbb{D}^\times$,

$$A_{s^*q}(s^{-1}(X)) = 2A_q(X),$$

the area bounds for s^*q on $s^{-1}(\Omega)$ and $s^{-1}(B)$ imply the same bounds for q on Ω and B, as desired. □

Proof of Theorem 4. For simplicity, if S has no boundary or has non-negative Euler characteristic, remove disks from $S \setminus R$ until it has boundary and negative Euler characteristic. Then enlarge R until it is equal to S minus an ε-neighborhood of ∂S, and think about R as a subset of S.

Now choose a maximal set of simple, non-intersecting and non-parallel arcs $\{\gamma_i\}_{i=1}^k$ on S. These will divide S into a collection of half-pants (i.e., hexagons) and once-punctured bigons; arrange the arcs so that they divide R in the same way, as illustrated in Figure 6. Let $\{P_j\}_{j=1}^\ell$ be the connected components of $S \setminus \bigcup \gamma_i$, and let G_i be small disjoint tubular neighborhoods of the γ_i inside S. Let $P_j' = P_j \cap R$ and let $G = \bigcup_i G_i$. As detailed below, we can apply Propositions 6.9 or 6.16 to each triple $\left(P_j, P_j', P_j \cap G\right)$ to show that if the area of a sequence

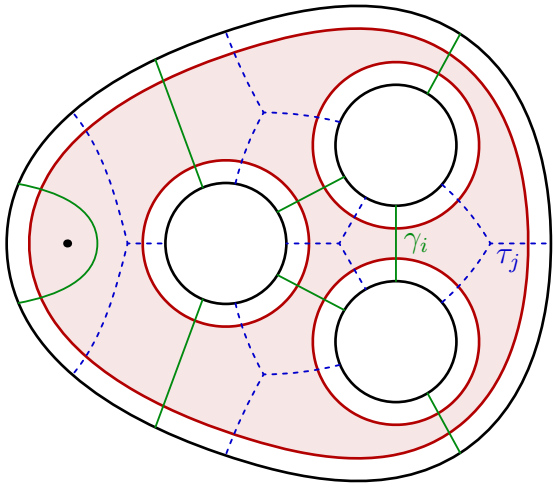

FIGURE 6. Two decompositions of S and R. In this example, S is a sphere with 4 holes and one puncture and R is a smaller copy of S shaded in red. The arcs γ_i (solid, in green) divide the two surfaces into half-pants and a once-punctured bigon. The tripods τ_j (dashed, in blue) divide the two surfaces into rectangles and a once-punctured bigon.

of quadratic differentials q_n on S is concentrating within R, then it is actually concentrating within G.

We also pick another decomposition of R and S into disks. Within each half-pants among the P_j, pick a tripod τ_j with ends on the three components of $P_j \cap \partial S$ and intersecting ∂R in three points, as in Figure 6; ensure that τ_j is disjoint from $\overline{G}$. Let $\{Q_i\}_{i=1}^k$ be the connected components of $S \setminus \bigcup_j \tau_j$. Each Q_i is a rectangle or a once-punctured bigon. Pick a small tubular neighborhood T_j of τ_j, small enough that each T_j and G_i are disjoint. Let $Q_i' = Q_i \cap R$ and $T = \bigcup_j T_j$. Propositions 6.9 and 6.16 will again show that if the area of a sequence of quadratic differentials on S is concentrating within R, then it is concentrating within T; but this is a contradiction, as G and T are disjoint.

We now give the concrete estimates alluded to above. Since all areas are with respect to an arbitrary quadratic differential $q \in \mathcal{Q}(S)$, we will omit it from the notation for brevity. For each j, the triple $(\overline{P_j}, P_j', G \cap P_j)$ is either a triple like $(\mathbb{D}, \Omega, B)$ as in the statement of Proposition 6.9 or a triple like $(\mathbb{D}^\times, \Omega, B)$ as in the statement of Proposition 6.16. We can thus find δ_j according to the propositions so that if $A(P_j') > (1-\delta)A(P_j)$, then $A(G \cap P_j) > (3/4)A(P_j)$. Let $\delta \coloneqq \min_j \delta_j$ and $\delta' \coloneqq \delta/4$.

Claim 6.17. If $A(R) > (1-\delta')A(S)$, then $A(G) > \frac{1}{2} \cdot A(S)$.

Proof. Let $J \subset \{1, \dots, \ell\}$ be the subset of indices j so that $A(P_j') > (1-\delta)A(P_j)$, and let

$$P_J \coloneqq \bigcup_{j \in J} P_j \qquad P_J' \coloneqq P_J \cap R$$

$$P_{\overline{J}} \coloneqq \bigcup_{j \notin J} P_j \qquad P_{\overline{J}}' \coloneqq P_{\overline{J}} \cap R.$$

Then we have

$$(1-\delta')A(S) - A(P'_J) < A(P'_{\overline{J}}) \leq (1-\delta)A(P_{\overline{J}}) < (1-\delta)(A(S) - A(P'_J))$$

which simplifies to

$$A(P'_J) > \frac{\delta - \delta'}{\delta} A(S) = \frac{3}{4} \cdot A(S).$$

On the other hand, by the choice of δ, we have $A(G \cap P'_J) > (3/4)A(P_J)$, so

$$A(G) \geq A(G \cap P'_J) > \frac{3}{4} \cdot A(P_J) \geq \frac{3}{4} \cdot A(P'_J) > \frac{3}{4} \cdot \frac{3}{4} \cdot A(S) > \frac{1}{2} \cdot A(S). \qquad \square$$

An exactly parallel argument shows that there is a $\delta'' > 0$ so that if $A(R) > (1-\delta'')A(S)$, then $A(T) > \frac{1}{2} \cdot A(S)$. Since $G \cap T = \emptyset$, this implies that $A(R) \leq (1 - \min(\delta', \delta'))A(S)$, proving the first statement of the theorem.

Note that the crucial constants δ' and δ'' were defined as a minimum over the triples $(P_j, P'_j, G \cap P_j)$ and $(Q_i, Q'_i, T \cap Q_i)$. On a finite cover $\widetilde{f}\colon \widetilde{R} \hookrightarrow \widetilde{S}$ of f, we can take arcs $\widetilde{\gamma}_i$ and tripods $\widetilde{\tau}_j$ to be lifts of γ_i and τ_j, respectively. Then the triples on the $\widetilde{S}$ are lifts of the triples on S, and the same estimate works in $\widetilde{f}$. $\square$

Proof of Theorem 3. If $\mathrm{SF}[f] \geq 1$, we have already proved the result in Proposition 6.3. If $\mathrm{SF}[f] < 1$, by Theorem 2 we may assume that f is an annular conformal embedding. Let K be the constant from Theorem 4 for the map f. We must show that for any finite cover $\widetilde{f}\colon \widetilde{R} \to \widetilde{S}$ of f and any simple closed multi-curve $\widetilde{C}$ on $\widetilde{R}$,

$$\frac{\mathrm{EL}_{\widetilde{R}}[\widetilde{f}(\widetilde{C})]}{\mathrm{EL}_{\widetilde{S}}[\widetilde{C}]} < K.$$

Let $\widetilde{q}$ be the quadratic differential realizing the extremal length of $[\widetilde{f}(\widetilde{C})]$. Then, as in the proof of Theorem 2,

$$\mathrm{EL}_{\widetilde{R}}[\widetilde{C}] \geq \frac{\ell_{\widetilde{f}^*|\widetilde{q}|}[\widetilde{C}]^2}{A_{\widetilde{f}^*|\widetilde{q}|}(\widetilde{R})} \geq \frac{\ell_{|\widetilde{q}|}[\widetilde{f}(\widetilde{C})]^2}{K A_{|\widetilde{q}|}(\widetilde{S})} = K^{-1}\, \mathrm{EL}_{\widetilde{S}}[\widetilde{f}(\widetilde{C})]. \qquad \square$$

7. Future challenges

There are several obvious questions raised by Theorems 1, 2, and 3. The first is an analogue of Proposition 6.3 when $\mathrm{SF}[f] < 1$.

Problem 7.1. Give an intrinsic characterization of $\widetilde{\mathrm{SF}}[f]$ for general maps $f\colon R \to S$ between Riemann surfaces as an infimum, not just when $\widetilde{\mathrm{SF}}[f] \geq 1$.

To elaborate a little, SF and $\widetilde{\mathrm{SF}}$ are defined as maxima. It would be much easier to find upper bounds (as in the hard direction of Theorem 3) if there were an alternate definition of $\widetilde{\mathrm{SF}}$ as a minimum. For example, there are two characterizations of extremal length: as a maximum over metrics (Definition 3.5) and as a minimum over embeddings of annuli (Proposition 3.8).

When $\mathrm{SF}[f] \geq 1$, Proposition 6.3 serves this role. When $\mathrm{SF}[f] < 1$, there are many different conformal embeddings $R \hookrightarrow S$ in the homotopy class $[f]$. The space of such conformal embeddings is path-connected [FB18]. One could attempt to find a canonical embedding by, for instance, gluing annuli to the boundary components of R [EM78]. But this embedding

seems ill-suited to give tight bounds on $\mathrm{SF}[f]$ or $\widetilde{\mathrm{SF}}[f]$. Ideally one would want a notion of "map with quasi-conformal constant less than one", but that is nonsensical.

Instead, it seems likely we need to consider some sort of "smeared" maps: maps from R to probability distributions on S.

Problem 7.2. Find an energy of smeared maps $g\colon R \to \mathcal{M}(S)$ whose minimum value is $\widetilde{\mathrm{SF}}[f]$.

As an example of what we mean, we give one way to get an explicit upper bound on $\widetilde{\mathrm{SF}}[f]$.

Definition 7.3. A homotopy class of topological embeddings $[f]\colon R \hookrightarrow S$ between Riemann surfaces is *conformally loose* if, for all $y \in \overline{S}$, there is a conformal embedding $g \in [f]$ so that $y \notin \overline{g(R)}$.

Since $\overline{S}$ is compact, if $[f]\colon R \to S$ is conformally loose we can find finitely many conformal embeddings $f_i \in [f], i = 1, \dots, n$ so that

$$\bigcap_{i=1}^{n} \overline{f_i(R)} = \emptyset. \tag{7.4}$$

In this case, we say that $[f]$ is *n-loose*.

Proposition 7.5. *If $[f]\colon R \hookrightarrow S$ is n-loose, then $\widetilde{\mathrm{SF}}[f] \leq 1 - 1/n$.*

Proof. If f is n-loose, then all covers are also n-loose. So it suffices to prove that $\mathrm{SF}[f] \leq 1 - 1/n$.

Let $(f_i)_{i=1}^n$ be the n different embeddings from Equation (7.4). For a simple multi-curve $C \in \mathcal{C}^+(R)$, let $q = q_{f(C)} \in \mathcal{Q}^+(S)$ be the quadratic differential corresponding to $f(C)$ from Theorem 3.12. For at least one i, we will have

$$\frac{A_q(f_i(R))}{A_q(S)} \leq 1 - 1/n$$

by Lemma 7.6 below. Then the argument from case (1) $\Rightarrow$ (4) of the proof of Theorem 2 shows that $\mathrm{EL}_R[C] \leq (1 - 1/n)\,\mathrm{EL}_S[f(C)]$, as desired. □

Lemma 7.6. *If $A_1, \dots, A_n \subset X$ are n subsets of a measure space X so that $\bigcap_{i=1}^n A_i = \emptyset$, then for at least one i we must have $\mu(A_i) \leq (1 - 1/n)\mu(X)$.*

Proof. This follows from the continuous pigeonhole principle. □

In the language of Problem 7.2, if $[f]$ is n-loose, then the averaged map

$$g(x) = \frac{1}{n} \sum_{i=1}^{n} f_i(x)$$

is a smeared map from R to S. Likewise, if $\widetilde{f}\colon \widetilde{R} \to \widetilde{S}$ is n-loose where $q\colon \widetilde{R} \to R$ is a finite cover of degree k, then the averaged map

$$g(x) = \frac{1}{nk} \sum_{q(\widetilde{x})=x} \sum_{i=1}^{n} \widetilde{f}_i(\widetilde{x})$$

is a smeared map from R to S.

Conjecture 7.7. *If $f\colon R \to S$ is a strict conformal embedding of Riemann surfaces where S has no punctures, there is some finite cover $\widetilde{f}$ of f that is conformally loose.*

If $[f]$ maps a puncture x of R to a puncture y of S, a neighborhood of y is in the image of every map in $[f]$, so $[f]$ can never be conformally loose. In this case we could pass to a branched double cover as in the proof of Proposition 6.16.

Remark 7.8. In Problems 7.1 and 7.2, it may be that $\widetilde{\mathrm{SF}}[f]$ is not the most natural quantity to consider; there may be a more natural quantity that bounds $\widetilde{\mathrm{SF}}[f]$ from above and is less than one when $\widetilde{\mathrm{SF}}[f]$ is less than one.

References

[Bur79] Robert Burckel, *An introduction to classical complex analysis*, vol. 1, Birkhäuser, 1979.

[Con78] John B. Conway, *Functions of one complex variable*, second ed., Graduate Texts in Mathematics, vol. 11, Springer-Verlag, 1978.

[DH93] Adrien Douady and John H. Hubbard, *A proof of Thurston's topological characterization of rational functions*, Acta Math. **171** (1993), no. 2, 263–297, doi:10.1007/BF02392534.

[EM78] Clifford J. Earle and Albert Marden, *Conformal embeddings of Riemann surfaces*, J. Analyse Math. **34** (1978), 194–203.

[Eps66] D. B. A. Epstein, *Curves on 2-manifolds and isotopies*, Acta Math. **115** (1966), 83–107.

[FB16] Maxime Fortier Bourque, *The converse of the Schwarz lemma is false*, Ann. Acad. Sci. Fenn. Math. **41** (2016), 235–241, arXiv:1411.5913, doi:10.5186/aasfm.2016.4115.

[FB18] ———, *The holomorphic couch theorem*, Invent. Math. **212** (2018), no. 2, 319–406, arXiv:1503.05473, doi:10.1007/s00222-017-0769-6.

[FLP79] A. Fathi, F. Laudenbach, and V. Poénaru, *Travaux de Thurston sur les surfaces*, Astérisque, no. 66–67, Soc. Math. France, Paris, 1979.

[HM79] John Hubbard and Howard Masur, *Quadratic differentials and foliations*, Acta Math. **142** (1979), no. 3-4, 221–274.

[Iof75] M. S. Ioffe, *Extremal quasiconformal imbeddings of Riemann surfaces*, Sibirsk. Mat. Ž. **16** (1975), no. 3, 520–537, 644.

[Ker80] Steven P. Kerckhoff, *The asymptotic geometry of Teichmüller space*, Topology **19** (1980), no. 1, 23–41.

[Mas00] Makoto Masumoto, *Hyperbolic lengths and conformal embeddings of Riemann surfaces*, Israel J. Math. **116** (2000), 77–92.

[Mas16] ———, *Holomorphic mappings of once-holed tori*, J. Anal. Math. **129** (2016), 69–90, doi:10.1007/s11854-016-0015-y.

[McM89] Curt McMullen, *Amenability, Poincaré series and quasiconformal maps*, Invent. Math. **97** (1989), no. 1, 95–127.

[McM12] Curtis T. McMullen, *Riemann surfaces, dynamics and geometry*, Course notes, Harvard University, 2012.

[Mos03] Lee Mosher, *Train track expansions of measured foliations*, Preprint, 2003, available from `http://andromeda.rutgers.edu/~mosher/`.

[MS84] Albert Marden and Kurt Strebel, *The Heights theorem for quadratic differentials on Riemann surfaces*, Acta Math. **153** (1984), no. 3-4, 153–211.

[PH92] Robert C. Penner and John L. Harer, *Combinatorics of train tracks*, Annals of Mathematics Studies, no. 125, Princeton University Press, Princeton, NJ, 1992.

[Put16] Andrew Putman, Answer to *Who proved that two homotopic embeddings of one surface in another are isotopic?*, MathOverflow, 2016, `http://mathoverflow.net/a/248866` (version: 2016-09-02).

[Ren76] Heinrich Renelt, *Konstruktion gewisser quadratischer differentiale mit hilfe von dirichletintegralen*, Mathematische Nachrichten **73** (1976), 125–142.

[Shi87] Masakazu Shiba, *The moduli of compact continuations of an open Riemann surface of genus one*, Trans. Amer. Math. Soc. **301** (1987), no. 1, 299–311.

[Shi93] ———, *The Euclidean, hyperbolic, and spherical spans of an open Riemann surface of low genus and the related area theorems*, Kodai Math. J. **16** (1993), no. 1, 118–137.

[Str84] Kurt Strebel, *Quadratic differentials*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3), vol. 5, Springer-Verlag, 1984.

[Thu16] Dylan P. Thurston, *From rubber bands to rational maps: A research report*, Res. Math. Sci. **3** (2016), Art. 15, arXiv:1502.02561, doi:10.1186/s40687-015-0039-4.

[Thu20] ———, *A positive characterization of rational maps*, Ann. of Math. (2) **192** (2020), no. 1, 1–46, arXiv:1612.04424, doi:10.4007/annals.2020.192.1.1.

[Wol96] Michael Wolf, *On realizing measured foliations via quadratic differentials of harmonic maps to* $\mathbb{R}$*-trees*, J. Anal. Math. **68** (1996), 107–120, doi:10.1007/BF02790206.

Brown University, 151 Thayer Street, Providence, RI 02912, USA
Email address: jeremy_kahn@brown.edu

Indiana University, 831 E. Third St., Bloomington, Indiana 47405, USA
Email address: pilgrim@indiana.edu

Indiana University, 831 E. Third St., Bloomington, Indiana 47405, USA
Email address: dpthurst@indiana.edu