

THE HOLOMORPHIC COUCH THEOREM

MAXIME FORTIER BOURQUE

ABSTRACT. We prove that if two conformal embeddings between Riemann surfaces with finite topology are homotopic, then they are isotopic through conformal embeddings. Furthermore, we show that the space of all conformal embeddings in a given homotopy class is homotopy equivalent to a point, a circle, a torus, or the unit tangent bundle of the codomain, depending on the induced homomorphism on fundamental groups. Quadratic differentials play a central role in the proof.

CONTENTS

1. Introduction	1
2. Preliminaries	7
3. Ioffe's theorem	15
4. Slit mappings are almost rigid	25
5. The modulus of extension	32
6. The blob and its boundary	38
7. Extremal length of partial measured foliations	42
8. The blob is semi-smooth	46
9. The blob is a disk	54
10. The deformation retraction	60
11. The remaining cases	65
References	69

1. INTRODUCTION

Loosely speaking, the 1-*parametric h-principle* is said to hold for a class of maps between manifolds if the only obstructions to connecting two maps in the class through maps in the same class are topological [EM02, p.60]. For example, the 1-parametric *h-principle* holds for immersions of S^2 in \mathbb{R}^3 , so that the standard sphere can be turned inside out via immersions. This is known as Smale's paradox. Of course, there are situations where the 1-parametric *h-principle* fails due to geometric obstructions. A famous example is Gromov's symplectic camel theorem, which says that one cannot move a closed 4-dimensional ball through a hole in a wall in \mathbb{R}^4 via symplectic embeddings if the ball is bigger than the hole.

For us, a *Riemann surface* is a connected 1-dimensional complex manifold. A *finite Riemann surface* is a Riemann surface with finitely generated fundamental

Research partially supported by a Postgraduate Scholarship from the Natural Sciences and Engineering Research Council of Canada and grant DMS 0905812 from the National Science Foundation.

group. Equivalently, a finite Riemann surface is a closed Riemann surface with finitely many points and closed disks removed. This should not be confused with the notion of a Riemann surface of *finite type*, which is a closed Riemann surface with finitely many points removed. A *conformal embedding* between Riemann surfaces is an injective holomorphic map.

In this paper, we prove that the 1-parametric h -principle holds for conformal embeddings between finite Riemann surfaces.

Theorem 1.1 (The holomorphic couch theorem). *If two conformal embeddings between finite Riemann surfaces are homotopic, then they are isotopic through conformal embeddings.*

We think of the codomain as a house and the domain as a couch that we want to move around in the house without changing its holomorphic structure. Hence the name “holomorphic couch”.

Given finite Riemann surfaces X and Y , and a topological embedding $h : X \rightarrow Y$, we define $\text{CEmb}(X, Y, h)$ to be the set of all conformal embeddings homotopic to h . We equip this set with the compact-open topology, which is the same as the topology of uniform convergence on compact sets with respect to any metric inducing the topology given on Y . Theorem 1.1 is equivalent¹ to the statement that $\text{CEmb}(X, Y, h)$ is path-connected whenever it is non-empty.

In fact, we prove a stronger result. Namely, we determine the homotopy type of the space $\text{CEmb}(X, Y, h)$. The answer depends on the image of h at the level of fundamental groups. We say that h is *trivial*, *cyclic*, or *generic* if the image of the induced homomorphism $\pi_1(h) : \pi_1(X, x) \rightarrow \pi_1(Y, h(x))$ is trivial, infinite cyclic, or non-abelian, respectively. If Y is a torus, then the image of $\pi_1(h)$ can be isomorphic to \mathbb{Z}^2 , but we do not need to distinguish this case.

Theorem 1.2 (Strong holomorphic couch). *Let $h : X \rightarrow Y$ be a topological embedding between finite Riemann surfaces such that $\text{CEmb}(X, Y, h)$ is non-empty. Then $\text{CEmb}(X, Y, h)$ is homotopy equivalent to either the unit tangent bundle of Y , a circle, or a point, depending on whether h is trivial, cyclic, or generic. This is unless Y is a torus and h is non-trivial, in which case $\text{CEmb}(X, Y, h)$ is homotopy equivalent to a torus.*

If h is generic, then $\text{CEmb}(X, Y, h)$ is contractible. This is the most interesting case; the other cases are either analogous or corollaries (see Section 11). The rest of the introduction outlines the proof of Theorem 1.2 in the case that h is generic, although we state some intermediate results in greater generality.

1.1. Ioffe’s theorem. A *Teichmüller embedding* between finite Riemann surfaces X and Y is an injective continuous map $f : X \rightarrow Y$ for which there exists a constant $K \geq 1$ and non-zero integrable holomorphic quadratic differentials on X and Y that extend to be real and non-negative along the ideal boundary, such that f is locally of the form $x + iy \mapsto Kx + iy$ in natural coordinates and $Y \setminus f(X)$ is a finite union of points and horizontal arcs.

Our main tool is a theorem of Ioffe which says that extremal quasiconformal embeddings and Teichmüller embeddings are one and the same [Iof75].

¹Since X is locally compact Hausdorff, a path $[0, 1] \rightarrow \text{Map}(X, Y)$ is the same as a homotopy $X \times [0, 1] \rightarrow Y$ [Mun00, p.287].

Theorem 1.3 (Ioffe). *Let $f : X \rightarrow Y$ be a quasiconformal embedding between finite Riemann surfaces which is not conformal. Then f has minimal dilatation (i.e. is closest to being conformal) in its homotopy class if and only if it is a Teichmüller embedding.*

We reprove this theorem in Section 3. In the special case where X and Y are closed, this reduces to Teichmüller’s celebrated theorem, since an embedding between closed surfaces is a homeomorphism. What is different from Teichmüller’s theorem, however, is that Teichmüller embeddings are not necessarily unique in their homotopy class—even after ruling out the obvious counterexamples. This is an important issue which we discuss next.

1.2. Slit mappings. A *slit mapping* is a conformal Teichmüller embedding, i.e., with stretch factor $K = 1$. In this case, the quadratic differential on X is the pull-back of the quadratic differential on Y by the slit mapping.

We show that if $\text{CEmb}(X, Y, h)$ contains a slit mapping, then every element of $\text{CEmb}(X, Y, h)$ is a slit mapping and $\text{CEmb}(X, Y, h)$ is naturally homeomorphic to a point, a compact interval, a circle, or a torus.

Theorem 1.4 (Slit mappings are almost rigid). *Let $h : X \rightarrow Y$ be any embedding between finite Riemann surfaces. Suppose that $\text{CEmb}(X, Y, h)$ contains a slit mapping with respect to a quadratic differential ψ on Y . Then every $f \in \text{CEmb}(X, Y, h)$ is a slit mapping with respect to ψ . Moreover, for every $x \in X$, the evaluation map $\text{CEmb}(X, Y, h) \rightarrow Y$ sending f to $f(x)$ is a homeomorphism onto its image. This image is equal to Y if Y is a torus, is a compact horizontal arc or a point if h is generic, and is a horizontal circle if h is cyclic and Y is not a torus.*

A better way to say this is: except in the torus case, any conformal embedding homotopic to a slit mapping differs from the latter by a horizontal translation which can be performed gradually (Theorem 4.1). An analogous statement holds for Teichmüller embeddings of dilatation $K > 1$.

This theorem corrects the wrong statement from [Iof75] that $\text{CEmb}(X, Y, h)$ is a single point if h is generic and $\text{CEmb}(X, Y, h)$ contains a slit mapping. Indeed, there are easy examples showing that $\text{CEmb}(X, Y, h)$ can be a non-degenerate interval (see Subsection 3.3).

We prove Theorem 1.4 in Section 4. Observe that Theorem 1.4 implies Theorem 1.2 whenever $\text{CEmb}(X, Y, h)$ contains a slit mapping. If h is generic and $\text{CEmb}(X, Y, h)$ is non-empty but does not contain a slit mapping, the idea is to enlarge X until it barely fits in Y , then appeal to Theorem 1.4 for the enlarged surface.

1.3. Modulus of extension. Given a finite Riemann surface X with non-empty ideal boundary, we define a 1-parameter family of enlargements of X as follows. We first choose an analytic parametrization $S^1 \rightarrow C$ of each ideal boundary component C of X . Then, for every $r \in (0, \infty]$, we let X_r be the bordered surface $X \cup \partial X$ with a copy of the cylinder $S^1 \times [0, r]$ glued to each ideal boundary component along $S^1 \times \{0\}$ via the fixed parametrization $S^1 \rightarrow C$. We also let $X_0 = X$.

We say that a cyclic embedding is *parabolic* if its image on π_1 is generated by a loop around a puncture. Denote by **(H)** the hypothesis that:

- $h : X \rightarrow Y$ is a non-trivial and non-parabolic embedding between finite Riemann surfaces;

- X has non-empty ideal boundary²;
- $\{X_r\}_{r \in [0, \infty]}$ is a 1-parameter family of enlargements of X ;
- $\text{CEmb}(X, Y, h)$ is non-empty.

Under hypothesis **(H)**, we define the *modulus of extension* $\mathfrak{m}(f)$ of any $f \in \text{CEmb}(X, Y, h)$ as the supremum of the set of $r \in [0, \infty]$ such that f extends to a conformal embedding of X_r into Y . Montel's theorem in complex analysis implies that:

- for every $f \in \text{CEmb}(X, Y, h)$ the supremum $\mathfrak{m}(f)$ is achieved by a unique conformal embedding $\hat{f} : X_{\mathfrak{m}(f)} \rightarrow Y$ extending f ;
- $\text{CEmb}(X, Y, h)$ is compact;
- \mathfrak{m} is upper semi-continuous.

In particular, \mathfrak{m} achieves its maximum value over $\text{CEmb}(X, Y, h)$. Using Ioffe's theorem, it is not too hard to show that if \mathfrak{m} attains its maximum at f , then \hat{f} is a slit mapping. We prove in Section 5 that the same holds if f is only assumed to be a local maximum of \mathfrak{m} .

Theorem 1.5 (Characterization of local maxima). *Assume hypothesis **(H)**. Let $f \in \text{CEmb}(X, Y, h)$ be a local maximum of \mathfrak{m} such that $\mathfrak{m}(f) < \infty$ and let \hat{f} be the conformal extension of f to $X_{\mathfrak{m}(f)}$. Then \hat{f} is a slit mapping. Conversely, if $g : X_r \rightarrow Y$ is a slit mapping such that $g|_X$ is homotopic to h , then $g|_X$ is a global maximum of \mathfrak{m} . The set M of all local maxima of \mathfrak{m} is homeomorphic to a point, a compact interval, a circle, or a torus, and \mathfrak{m} is constant on M .*

The initial motivation for studying \mathfrak{m} was to think of it as a Morse function for the space $\text{CEmb}(X, Y, h)$. In an ideal world, flowing along the gradient of \mathfrak{m} would yield a deformation retraction of $\text{CEmb}(X, Y, h)$ into M . This does not make sense, however, since \mathfrak{m} is not even continuous unless it is constant equal to zero. In any case, the connectedness of $\text{CEmb}(X, Y, h)$ is a direct consequence of Theorem 1.5 (see Section 5).

Theorem 1.6. *Under hypothesis **(H)**, the space $\text{CEmb}(X, Y, h)$ is connected.*

To improve upon this, we show that there are no local obstructions to contractibility.

1.4. Where can one point go? Throughout this subsection, we assume that:

- $h : X \rightarrow Y$ is a generic embedding between finite Riemann surfaces;
- X has non-empty ideal boundary;
- $\text{CEmb}(X, Y, h)$ is non-empty and does not contain any slit mapping;

which we call hypothesis **(H')**. Note that **(H')** \Rightarrow **(H)**.

Given a point $x \in X$, we are interested in set of points in Y where x can be mapped by the elements of $\text{CEmb}(X, Y, h)$. It is convenient to also keep track of how x gets mapped to a given $y \in Y$ in the following sense. If $f \in \text{CEmb}(X, Y, h)$, then by definition there exists a homotopy

$$H : X \times [0, 1] \rightarrow Y$$

from h to f . Since h is generic, the homotopy class rel endpoints of the path $t \mapsto H(x, t)$ from $h(x)$ to $f(x)$ does not depend on the particular choice of H (see Lemma 6.1). Denote the homotopy class of that path by $\text{lift}_x(f)$. If the point

²If X has finite type, then Theorem 1.2 is easy (see Subsection 2.2).

$x \in X$ is kept fixed, $\text{lift}_x(f)$ represents an element of the universal cover of Y based at $h(x)$. Since Y has non-abelian fundamental group, its universal cover is conformally equivalent to the unit disk \mathbb{D} .

The map $\text{lift}_x : \text{CEmb}(X, Y, h) \rightarrow \mathbb{D}$ is continuous, and we call its image $\text{Blob}(x, X, Y, h)$. The blob is simpler than the image of the evaluation map in much the same way as Teichmüller space is simpler than moduli space. Indeed, the blob is as simple as can be.

Theorem 1.7 (The blob is a disk). *Under hypothesis (\mathbf{H}') , $\text{Blob}(x, X, Y, h)$ is homeomorphic to a closed disk for any $x \in X$.*

The proof has four steps:

- the blob is compact and connected (because $\text{CEmb}(X, Y, h)$ is);
- the blob is semi-smooth;
- every semi-smooth subset of \mathbb{R}^2 is a 2-manifold with boundary;
- there are no holes in the blob.

We refer the reader to Sections 6 to 9 for the definition of a semi-smooth set and the breakdown of these steps. Suffice it to say that Ioffe's theorem implies that points on the boundary of $\text{Blob}(x, X, Y, h)$ come from conformal embeddings f whose restriction $f^* : X \setminus \{x\} \rightarrow Y \setminus \{f(x)\}$ is a slit mapping. We then use a variational formula for extremal length to deduce information about the shape of $\text{Blob}(x, X, Y, h)$ near such boundary points.

1.5. Moving one point at a time. We now explain how to deduce that $\text{CEmb}(X, Y, h)$ is contractible from the previous results, still assuming hypothesis (\mathbf{H}') . Pick a countable dense set $\{x_1, x_2, \dots\}$ in X and let $F \in \text{CEmb}(X, Y, h)$ be any conformal embedding. We define a deformation retraction of the space $\text{CEmb}(X, Y, h)$ into $\{F\}$ by moving one point at a time. Given a map f in $\text{CEmb}(X, Y, h)$, we join $\text{lift}_{x_1}(f)$ to $\text{lift}_{x_1}(F)$ by a path γ_1 in $\text{Blob}(x_1, X, Y, h)$. Such a path exists since $\text{Blob}(x_1, X, Y, h)$ is homeomorphic to a closed disk.

For every $t \in [0, 1]$, we then look at where x_2 can go under maps g in $\text{CEmb}(X, Y, h)$ which satisfy $\text{lift}_{x_1}(g) = \gamma_1(t)$. This defines a new kind of blob, call it $\text{Blob}_t(x_2)$. We show that $\text{Blob}_t(x_2)$ moves continuously with t , which allows us to construct a second path γ_2 from $\text{lift}_{x_2}(f)$ to $\text{lift}_{x_2}(F)$ with the property that for every $t \in [0, 1]$, the point $\gamma_2(t)$ belongs to $\text{Blob}_t(x_2)$.

Proceeding by induction, we obtain a sequence of paths $\gamma_1, \gamma_2, \dots$ such that for every $n \in \mathbb{N}$ and every $t \in [0, 1]$, there exists at least one map f_t^n in $\text{CEmb}(X, Y, h)$ such that $\text{lift}_{x_j}(f_t^n) = \gamma_j(t)$ for every $j \in \{1, \dots, n\}$. If we fix t and pass to a subsequence, we get some limit $f_t \in \text{CEmb}(X, Y, h)$ for which $\text{lift}_{x_j}(f_t) = \gamma_j(t)$ for every $j \in \mathbb{N}$. Since any two limits agree on the dense set $\{x_1, x_2, \dots\}$, we actually have convergence $f_t^n \rightarrow f_t$ as $n \rightarrow \infty$ (without passing to a subsequence). By a similar argument, f_t depends continuously on t . We thus found a path from f to F in $\text{CEmb}(X, Y, h)$.

We construct the paths $\gamma_1, \gamma_2, \dots$ carefully enough so that they depend continuously on the initial map f , hence the path $t \mapsto f_t$ also depends continuously on f . The end result is a deformation retraction of $\text{CEmb}(X, Y, h)$ into $\{F\}$. See Section 10 for details.

1.6. Notes and references. The holomorphic couch problem arose in the context of renormalization in complex dynamics. Although Theorem 1.1 does not have any direct application to dynamics, some of the tools used here do (see [Thu16]).

The space $\text{CEmb}(\mathbb{D}, \mathbb{C})$, or rather its subspace \mathcal{S} of conformal embeddings $f : \mathbb{D} \rightarrow \mathbb{C}$ satisfying $f(0) = 0$ and $f'(0) = 1$, was the subject of much interest until the solution of the Bieberbach conjecture by de Branges in 1984. It is easy to see that \mathcal{S} is contractible in the compact-open topology. On the other hand, \mathcal{S} has isolated points when equipped with the topology of uniform convergence of the Schwarzian derivative [Thu86]. The literature on the class \mathcal{S} is quite vast. In comparison, not much has been written about conformal embeddings between general Riemann surfaces. Exceptions include [EM78], [Iof75], [Iof78], [Jen58], [SS54] and [Shi04].

The holomorphic couch problem for embeddings of a multiply-punctured disk in a multiply-punctured sphere was first considered in [Roy54b]. However, the solution presented there relies in part on a rigidity claim [Roy54a] which is known to be false in general [Jen59].

Theorems 1.4 and 1.5 generalize the fact that in the homotopy class of an essential simple closed curve α on a finite Riemann surface Y , there exists a unique embedded annulus $X \subset Y$ of largest modulus, and X is the horizontal cylinder of a quadratic differential ψ on Y (known as the Jenkins–Strebel differential corresponding to α).

In [EM78], Earle and Marden consider a functional similar to our modulus of extension \mathfrak{m} where they keep the annuli disjoint from X . Their construction is more natural than ours as it does not require any choice of parametrization of ∂X . We decided to glue the annuli to X in order to deal with a connected surface. The main theorem in [EM78] is analogous to Theorem 1.5, but its uniqueness statement is false for the same reason that Ioffe’s uniqueness statement is.

There is a plethora of other extremal problems on Riemann surfaces whose solutions involve quadratic differentials (see e.g. [Kru05] and the references therein). These are all examples of “Teichmüller’s principle” [Jen58, p.48].

For the class \mathcal{S} of normalized univalent functions from \mathbb{D} to \mathbb{C} , a suitable version of the blob is actually a round disk. More precisely, for every $z \in \mathbb{D}$ and every $f \in \mathcal{S}$, the quantity $w = \log(f(z)/z)$ satisfies

$$\left| w - \log \frac{1}{1 - |z|^2} \right| \leq \log \frac{1 + |z|}{1 - |z|}$$

and every value w satisfying the inequality is achieved for some $f \in \mathcal{S}$. This was proved by Grunsky in 1932 (see [Dur83, p.323]). The blob for K -quasiconformal homeomorphisms of the disk with prescribed boundary values was studied in [Str98] (see also [EL02] for a generalization to arbitrary hyperbolic surfaces). Our approach for proving that the blob is homeomorphic to a closed disk seems similar to Strebel’s, but the context is different.

The idea of moving one point at a time to get an isotopy is reminiscent of the finite “holomorphic axiom of choice” used by Slodkowski to extend holomorphic motions [Slo91]. Our isotopies are holomorphic in the space variable and continuous in the time variable, whereas holomorphic motions are the other way around.

Acknowledgements. The results of this paper are adapted from the author’s Ph.D. thesis at the Graduate Center of the City University of New York. I thank my advisor Jeremy Kahn for suggesting this problem and for his continued interest and support. I also thank Dylan Thurston, Kevin Pilgrim, Frederick Gardiner

and Patrick Hooper for useful conversations, François Labourie for pointing out the connection with the h -principle, and the anonymous referees for their helpful comments and suggestions.

2. PRELIMINARIES

2.1. Ideal boundary and punctures. A Riemann surface is *hyperbolic* if its universal covering space is conformally isomorphic to the unit disk \mathbb{D} . The only non-hyperbolic Riemann surfaces are the Riemann sphere $\widehat{\mathbb{C}}$, the complex plane \mathbb{C} , the once-punctured plane $\mathbb{C} \setminus \{0\}$ and complex tori. A hyperbolic surface X can be regarded as the quotient of its universal covering space \mathbb{D} by its group of deck transformations Γ . The *limit set* Λ_Γ is the set of accumulation points in $\partial\mathbb{D}$ of the Γ -orbit of any point $z \in \mathbb{D}$, and the *set of discontinuity* is $\Omega_\Gamma = \partial\mathbb{D} \setminus \Lambda_\Gamma$. The *ideal boundary* of X is $\partial X = \Omega_\Gamma/\Gamma$. The union $X \cup \partial X = (\mathbb{D} \cup \Omega_\Gamma)/\Gamma$ is naturally a bordered Riemann surface, since Γ acts properly discontinuously and analytically on $\mathbb{D} \cup \Omega_\Gamma$. If X is a finite hyperbolic surface, then ∂X has finitely many connected components, each homeomorphic to a circle.

A *puncture* in a Riemann surface X is an end corresponding to a proper (preimages of compact sets are compact) conformal embedding $\mathbb{D} \setminus \{0\} \rightarrow X$. For example, \mathbb{C} has one puncture at infinity and $\mathbb{C} \setminus \{0\}$ has two punctures. For hyperbolic surfaces, punctures are the same as cusps, or ends with parabolic monodromy. Every puncture can be *filled*, meaning that one can add the missing point and extend the complex structure there. The set of punctures of X is denoted by \dot{X} .

Given a finite Riemann surface X , we write $\widehat{X} = X \cup \partial X \cup \dot{X}$ for the compact bordered Riemann surface obtained after adding the ideal boundary and filling the punctures (by definition, a non-hyperbolic Riemann surface has empty ideal boundary). Suppose that ∂X is non-empty. Then if we take two copies of \widehat{X} —the second with reversed orientation—and glue them along ∂X with the identity, we get a closed Riemann surface called the *double* of \widehat{X} . Because of this construction, many theorems about Riemann surfaces of finite type are also true for finite Riemann surfaces.

2.2. Automorphisms. Given a Riemann surface X , let $\text{Aut}_0(X)$ be its group of conformal automorphisms homotopic to the identity. It is well-known that:

- $\text{Aut}_0(\widehat{\mathbb{C}})$ acts simply transitively on ordered triples of distinct points in $\widehat{\mathbb{C}}$;
- $\text{Aut}_0(\mathbb{C})$ acts simply transitively on ordered pairs of distinct points in \mathbb{C} ;
- $\text{Aut}_0(\mathbb{C} \setminus \{0\})$ acts simply transitively on $\mathbb{C} \setminus \{0\}$;
- if X is a torus, then $\text{Aut}_0(X)$ acts simply transitively on X ;
- $\text{Aut}_0(\mathbb{D})$ acts simply transitively on the unit tangent bundle of \mathbb{D} (with respect to the hyperbolic metric);
- if $r \in [0, 1)$, then $\text{Aut}_0(\mathbb{D} \setminus r\bar{\mathbb{D}})$ is isomorphic to the circle S^1 .

In all other cases, $\text{Aut}_0(X)$ only contains the identity.

Lemma 2.1. *Let X be a hyperbolic Riemann surface not isomorphic to \mathbb{D} nor $\mathbb{D} \setminus \{0\}$, and let $h : X \rightarrow X$ be a holomorphic map homotopic to the identity. Then h is equal to the identity unless X is an annulus and h is a rotation.*

In this lemma, X is not assumed to be finite and h is not assumed to be either injective or surjective.

Proof. By the Schwarz lemma, h is 1-Lipschitz with respect to the hyperbolic metric. Therefore, if α is a closed geodesic in X , then $h(\alpha)$ is at most as long as α . But geodesics minimize length in their homotopy class, so that $h(\alpha) = \alpha$. In particular, h is an isometry along α .

If X is an annulus, then it contains a unique simple closed geodesic α . We can post-compose h by a rotation r of X so that $r \circ h$ is equal to the identity on α and hence on all of X by the identity principle.

If X is not an annulus, then it contains a closed geodesic α which self-intersects exactly once. Then h fixes this self-intersection point, thus all of α pointwise, and hence all of X pointwise by the identity principle. \square

It follows that $\text{Aut}_0(X)$ is path-connected for any Riemann surface X . More precisely, $\text{Aut}_0(X)$ is

- homotopy equivalent to the unit tangent bundle of X if $\pi_1(X) = \{0\}$;
- homotopy equivalent to S^1 if $\pi_1(X) \cong \mathbb{Z}$;
- homeomorphic to $S^1 \times S^1$ if $\pi_1(X) \cong \mathbb{Z}^2$;
- a point if $\pi_1(X)$ is non-abelian.

Theorems 1.1 and 1.2 can be viewed as generalizations of this. When the domain has finite type, Theorem 1.2 actually follows from the above.

Lemma 2.2. *Suppose that $h : X \rightarrow Y$ is a topological embedding between finite Riemann surfaces, where ∂X is empty and Y is not the sphere with at most 2 punctures nor a torus. Then $\text{CEmb}(X, Y, h)$ contains at most one element.*

Proof. Since conformal embeddings send punctures to punctures or regular points, every f in $\text{CEmb}(X, Y, h)$ extends to a conformal embedding \hat{f} from $X \cup \hat{X}$ to $Y \cup \hat{Y}$. As $\hat{X} = X \cup \hat{X}$ is compact and \hat{f} is open, the map \hat{f} is surjective. In particular, $Y \cup \hat{Y} = \hat{Y}$ is compact. If $f, g \in \text{CEmb}(X, Y, h)$, then the inverses of the extensions $\hat{f}, \hat{g} : \hat{X} \rightarrow \hat{Y}$ are homotopic rel \hat{Y} . The composition $\hat{g} \circ \hat{f}^{-1}$ is thus a conformal automorphism of \hat{Y} homotopic to the identity rel \hat{Y} . If \hat{Y} has genus 0, then \hat{Y} must contain at least 3 points by hypothesis, so that $\hat{f} = \hat{g}$. If \hat{Y} has genus 1, then \hat{Y} contains at least 1 point, so that $\hat{f} = \hat{g}$. If \hat{Y} has genus at least 2, then $\hat{f} = \hat{g}$ by Lemma 2.1. \square

By the same argument, if $\partial X = \emptyset$ and Y is the sphere with at most 2 punctures or a torus, then $\text{CEmb}(X, Y, h)$ is either empty or homeomorphic to $\text{Aut}_0(Y)$. Theorem 1.2 for domains of finite type follows easily.

Proof of Theorem 1.2 in the case where $\partial X = \emptyset$. The image of any element in $\text{CEmb}(X, Y, h)$ has finite complement in Y so that $\pi_1(h)$ surjects onto $\pi_1(Y)$. If $\pi_1(Y)$ is non-abelian, then h is generic and $\text{CEmb}(X, Y, h)$ is a singleton by Lemma 2.2. Otherwise $\text{CEmb}(X, Y, h)$ is homeomorphic to $\text{Aut}_0(Y)$, whose homotopy type was described above. \square

2.3. Montel's theorem. The simplest version of Montel's theorem says that the set of all holomorphic maps from \mathbb{D} to \mathbb{D} is compact. This implies a similar result for holomorphic maps between arbitrary hyperbolic surfaces by lifting to the universal covers.

A sequence of maps $f_n : X \rightarrow Y$ between Riemann surfaces *diverges locally uniformly* if for every compact sets $K \subset X$ and $L \subset Y$, the sets $f_n(K)$ and L are disjoint for all large enough n . A set \mathcal{F} of maps between two Riemann surfaces X

and Y is *normal* if every sequence in \mathcal{F} admits either a locally uniformly convergent subsequence or a locally uniformly divergent subsequence.

Theorem 2.3 (Montel’s theorem). *If X and Y are hyperbolic surfaces, then every set of holomorphic maps from X to Y is normal.*

See [Mil06, p.34]. Note that the limit of a convergent sequence of holomorphic maps is holomorphic. If every map in the sequence is injective, then the limit is either injective or constant. If every map in the sequence is locally injective, then the limit is either locally injective or constant.

2.4. Quasiconformal maps. Let $K \geq 1$. A K -*quasiconformal map*³ between Riemann surfaces is a map f such that in charts, its first partial derivatives in the distributional sense are locally in L^2 and the formal matrix df of partial derivatives satisfies the inequality $\|df\|^2 \leq K \det(df)$ almost everywhere. For almost every point z , the real linear map $d_z f$ sends circles in the tangent plane at z to ellipses of eccentricity $\|d_z f\|^2 / \det(d_z f)$ in the tangent plane at $f(z)$, and this ratio is called the *pointwise dilatation* of f at z . The *dilatation* of f , denoted $\text{Dil}(f)$, is the smallest $K \geq 1$ for which f is K -quasiconformal. This is the same as the essential supremum of the pointwise dilatation of f .

A *Beltrami form* on a Riemann surface X is a map $\mu : TX \setminus \vec{0} \rightarrow \mathbb{C}$ such that $\mu(\lambda \mathbf{v}) = (\bar{\lambda}/\lambda)\mu(\mathbf{v})$ for every $\mathbf{v} \in TX \setminus \vec{0}$ and every $\lambda \in \mathbb{C} \setminus \{0\}$ where $\vec{0}$ denotes the zero section. In charts, the *Wirtinger derivatives* of a quasiconformal map f are

$$\partial f = \frac{1}{2} \left(\frac{\partial f}{\partial x} - i \frac{\partial f}{\partial y} \right) \quad \text{and} \quad \bar{\partial} f = \frac{1}{2} \left(\frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} \right).$$

The ratio $\bar{\partial} f / \partial f$ is naturally a Beltrami form, and is sometimes called the *Beltrami coefficient* of f . The Beltrami coefficient of f encodes the field of ellipses in TX which df sends to circles.

The measurable Riemann mapping theorem of Morrey and Ahlfors–Bers says that every measurable ellipse field with bounded eccentricity is the Beltrami coefficient of a quasiconformal homeomorphism.

Theorem 2.4 (The measurable Riemann mapping theorem). *Let X be a Riemann surface and let μ be a measurable Beltrami form on X such that $\|\mu\|_\infty < 1$. Then there exists a Riemann surface Y and a quasiconformal homeomorphism $f : X \rightarrow Y$ such that $\bar{\partial} f / \partial f = \mu$ almost everywhere. The surface and the quasiconformal homeomorphism are unique up to post-composition with a conformal homeomorphism.*

An important consequence is the following factorization principle. Suppose that $f : X \rightarrow Y$ is a quasiconformal map. Then $f = g \circ h$ where $h : X \rightarrow Z$ is a quasiconformal homeomorphism and $g : Z \rightarrow Y$ is holomorphic. Indeed, we can take h to be the solution of the Beltrami equation with $\mu = \bar{\partial} f / \partial f$ and let $g = f \circ h^{-1}$.

Another fundamental property of K -quasiconformal homeomorphisms is compactness under appropriate normalization [Ahl06, p.32].

³This notion is usually called “quasiregular map”, and the expression “quasiconformal map” is usually reserved for homeomorphisms. It seems more convenient to modify the noun “map” instead, to indicate further attributes. For example we will use “quasiconformal immersion”, “quasiconformal embedding”, or “quasiconformal homeomorphism” for a quasiconformal map which is a local embedding, an embedding, or a homeomorphism respectively.

Theorem 2.5. *For every $K \geq 1$, the space of all K -quasiconformal homeomorphisms from \mathbb{D} to \mathbb{D} fixing 0 is compact.*

Lastly, we will use the fact that quasiconformal homeomorphisms send punctures to punctures. As a consequence, any quasiconformal embedding between finite Riemann surfaces extends to a quasiconformal embedding between the surfaces with their punctures filled.

The reader may consult [Ahl06] for more background on quasiconformal homeomorphisms.

2.5. Quadratic differentials. A *quadratic differential* on a Riemann surface X is a map $\varphi : TX \rightarrow \widehat{\mathbb{C}}$ such that $\varphi(\lambda \mathbf{v}) = \lambda^2 \varphi(\mathbf{v})$ for every $\mathbf{v} \in TX$ and every $\lambda \in \mathbb{C}$. A quadratic differential on X is *holomorphic* (resp. *meromorphic*) if for every open set $U \subset X$, and every holomorphic vector field $\mathbf{v} : U \rightarrow TU$, the function $\varphi \circ \mathbf{v} : U \rightarrow \mathbb{C}$ is holomorphic (resp. meromorphic). All quadratic differentials in this paper will be holomorphic or meromorphic. The *pull-back* $f^*\varphi$ of a quadratic differential φ by a holomorphic map f is defined in the usual way by the formula $f^*\varphi(\mathbf{v}) = \varphi(df(\mathbf{v}))$.

A vector $\mathbf{v} \in TX$ is *horizontal* (resp. *vertical*) for φ if $\varphi(\mathbf{v}) > 0$ (resp. $\varphi(\mathbf{v}) < 0$). Let $I \subset \mathbb{R}$ be an interval. A piecewise smooth arc $\gamma : I \rightarrow X$ is *horizontal* (resp. *vertical*) if $\gamma'(t)$ is horizontal (resp. vertical) wherever it is defined. Such a trajectory is called *regular* if it does not contain any zero or pole of φ . The absolute value $|\varphi|$ is an area form, and its integral $\|\varphi\| = \int_X |\varphi|$ is the *norm* of φ . For a finite Riemann surface X , we denote by $\mathcal{Q}(X)$ the set of all integrable holomorphic quadratic differentials φ on X which extend analytically to the ideal boundary of X , and such that $\varphi(\mathbf{v}) \in \mathbb{R}$ for every vector \mathbf{v} tangent to ∂X . Every $\varphi \in \mathcal{Q}(X)$ extends to a meromorphic quadratic differential on \widehat{X} with at most simple poles on \widehat{X} . The set $\mathcal{Q}^+(X)$ is similarly defined, but with the additional requirements that $\varphi \geq 0$ along ∂X and that φ is not identically zero. The set $\mathcal{Q}(X)$ is a real vector space inside of which $\mathcal{Q}^+(X)$ forms a convex cone.

For every simply connected open set $U \subset \widehat{X}$ where a quadratic differential φ does not have any zero or pole, there exists a locally injective holomorphic map $z : U \rightarrow \mathbb{C}$ such that $\varphi = dz^2$. The map z is unique up to translation and sign and is called a *natural coordinate* when it is injective. If $\varphi \in \mathcal{Q}(X)$, then the atlas of natural coordinates for φ is a *half-translation structure* on \widehat{X} minus the singularities of φ . The transition maps for this atlas have the form $z \mapsto \pm z + c$. Such a structure induces a flat geometry with cone points on \widehat{X} . We return to this geometry at the end of the present section. The standard reference for this material is [Str84].

2.6. Teichmüller's theorem. A *Teichmüller homeomorphism* (usually called Teichmüller map) between finite Riemann surfaces X and Y is a homeomorphism $f : X \rightarrow Y$ such that there exists a constant $K > 1$ and non-zero $\varphi \in \mathcal{Q}(X)$ and $\psi \in \mathcal{Q}(Y)$ such that f is locally of the form $x + iy \mapsto Kx + iy$ in natural coordinates for φ and ψ , up to sign and translation. Such a homeomorphism is K -quasiconformal with constant pointwise dilatation.

The following theorems of Teichmüller are of central importance.

Theorem 2.6 (Teichmüller’s existence theorem). *Let h be a quasiconformal homeomorphism between finite Riemann surfaces. If there is no conformal homeomorphism homotopic to h , then there is a Teichmüller homeomorphism homotopic to h .*

Theorem 2.7 (Teichmüller’s uniqueness theorem). *Let $f : X \rightarrow Y$ be a Teichmüller homeomorphism of dilatation K between finite Riemann surfaces. If g is a K -quasiconformal homeomorphism homotopic to f , then $g \circ f^{-1}$ is a conformal automorphism of Y homotopic to the identity. If Y is not an annulus nor a torus, then $g = f$.*

Teichmüller’s theorem is usually stated and proved for closed Riemann surfaces, but the general case follows from the closed case by doubling across the ideal boundary and by taking a branched cover of degree 2 or 4 ramified at the punctures [Ahl53].

2.7. Teichmüller spaces. Let S be a finite Riemann surface. The *Teichmüller space* $\mathcal{T}(S)$ is defined as the set of pairs (X, f) where X is a finite Riemann surface and $f : S \rightarrow X$ is a quasiconformal homeomorphism, modulo the equivalence relation $(X, f) \sim (Y, g)$ if and only if $g \circ f^{-1}$ is homotopic to a conformal homeomorphism. The equivalence class of (X, f) is denoted $[X, f]$, or just X when the *marking* f is implicit. The *Teichmüller distance* between two points of $\mathcal{T}(S)$ is defined as

$$d([X, f], [Y, g]) = \frac{1}{2} \inf \log \text{Dil}(h)$$

where the infimum is taken over all quasiconformal homeomorphisms h homotopic to $g \circ f^{-1}$. By Teichmüller’s theorem, the infimum is realized by a (usually unique) quasiconformal homeomorphism h which is either conformal or a Teichmüller homeomorphism.

The space $\mathcal{T}(S)$ is a contractible real-analytic manifold of finite dimension. Let $\mathcal{M}(X)$ denote the space of essentially bounded Beltrami forms on $X \in \mathcal{T}(S)$. By the measurable Riemann mapping theorem, the tangent space to $\mathcal{T}(S)$ at X can be identified with the quotient of $\mathcal{M}(X)$ by its subspace $\mathcal{M}_0(X)$ of infinitesimally trivial deformations. There is a natural pairing between $\mathcal{M}(X)$ and $\mathcal{Q}(X)$ given by

$$\langle \mu, \varphi \rangle = \text{Re} \int_X \mu \varphi,$$

and it turns out that $\mathcal{M}_0(X) = \mathcal{Q}(X)^\perp$ with respect to this pairing. The tangent and cotangent spaces to $\mathcal{T}(S)$ at X are thus isomorphic to $\mathcal{M}(X)/\mathcal{Q}(X)^\perp$ and $\mathcal{Q}(X)$ respectively. See [Ear64] and [Ear67] for more details.

Remark. In the literature, $\mathcal{T}(S)$ is called the *reduced* Teichmüller space of S . The (unreduced) Teichmüller space of S is defined similarly, except that pairs (X, f) and (Y, g) are declared equivalent if $g \circ f^{-1}$ is homotopic to a conformal homeomorphism rel ∂X . This unreduced Teichmüller space is always a contractible complex manifold, but has infinite dimension whenever the ideal boundary of S is non-empty. We will only consider reduced Teichmüller spaces in this paper, and will therefore omit the adjective “reduced”.

2.8. Homotopies. If two maps are homotopic, then they induce the same homomorphism between fundamental groups, up to conjugation. The converse also holds under appropriate conditions [Ahl06, p.60] [Ber58, §6].

Lemma 2.8. *Let X be a space which has a universal cover, let Y be a metric space whose universal cover is a uniquely geodesic space in which geodesics depend continuously on endpoints, and let $f_0, f_1 : X \rightarrow Y$ be continuous maps. Suppose that for some $x \in X$ the induced homomorphisms $\pi_1(f_j) : \pi_1(X, x) \rightarrow \pi_1(Y, f_j(x))$ agree up to conjugation by a path between $f_0(x)$ and $f_1(x)$. Then f_0 and f_1 are homotopic.*

Proof. Let \tilde{X} and \tilde{Y} be the universal covers of X and Y , and let α be a path connecting $f_0(x)$ to $f_1(x)$ which conjugates the homomorphisms $\pi_1(f_0)$ and $\pi_1(f_1)$. Given a lift $\tilde{f}_0 : \tilde{X} \rightarrow \tilde{Y}$, the path α allows us to lift f_1 in such a way that \tilde{f}_0 and \tilde{f}_1 are equivariant with respect to the same homomorphism of deck groups. The homotopy from \tilde{f}_0 to \tilde{f}_1 sending $(x, t) \in \tilde{X} \times [0, 1]$ to the point at proportion t along the geodesic from $\tilde{f}_0(x)$ to $\tilde{f}_1(x)$ in \tilde{Y} is continuous and equivariant, so it descends to a homotopy from f_0 to f_1 . \square

This is also true if X is a CW-complex and Y is a $K(\pi, 1)$ [Hat01, p.90]. If X and Y are finite Riemann surfaces and Y is not the sphere, then either hypotheses are satisfied. The most useful consequence for us is that homotopy classes of maps between finite Riemann surfaces are closed.

Corollary 2.9. *Let X and Y be finite Riemann surfaces and let $f_n, f : X \rightarrow Y$ be continuous maps such that $f_n \rightarrow f$ as $n \rightarrow \infty$. Then f_n is homotopic to f for all large enough n .*

Proof. First assume that Y is the Riemann sphere. By the Hopf theorem [Mil97, p.50], two maps $X \rightarrow Y$ are homotopic if and only if they have the same degree. Since topological degree depends continuously on the map, f_n eventually has the same degree as f and is therefore homotopic to it.

Suppose that Y is not the Riemann sphere. Then the universal cover of Y supports a metric which is proper (whose closed balls are compact) and is uniquely geodesic. In such a metric, geodesics depend continuously on endpoints. Thus the hypotheses of Lemma 2.8 are satisfied.

Let β_1, \dots, β_k be loops based at $x \in X$ which generate $\pi_1(X, x)$ and let V be a simply connected neighborhood of $f(x)$ in Y . Let n be large enough so that $f_n(x) \in V$ and so that there is a homotopy between the restrictions $f_n|_{\beta_j}$ and $f|_{\beta_j}$ which keeps the image of the basepoint x inside V for every $j \in \{1, \dots, k\}$. Let α be any path from $f(x)$ to $f_n(x)$ in V . Then $\alpha * f_n(\beta_j) * \bar{\alpha}$ is homotopic to $f(\beta_j)$ as loops based at $f(x)$, for every j . In other words, α conjugates the induced homomorphisms $\pi_1(f)$ and $\pi_1(f_n)$. By Lemma 2.8, f_n is homotopic to f . \square

In Teichmüller theory, one often goes back and forth between punctures and marked points as convenient. This passage is justified by the fact that quasiconformal homeomorphisms send punctures to punctures. Moreover, homotopies defined in the complement of punctures can be modified as to extend to the punctures.

Definition 2.10. Let $h : X \rightarrow Y$ be a quasiconformal embedding between finite Riemann surfaces. Recall that h extends to a quasiconformal embedding \bar{h} from

$X \cup \dot{X}$ to $Y \cup \dot{Y}$. We say that a puncture $p \in \dot{X}$ is *essential* if $\bar{h}(p) \in \dot{Y}$ and is *inessential* if $\bar{h}(p) \in Y$.

Lemma 2.11. *Let X and Y be finite Riemann surfaces, let $f_0, f_1 : X \rightarrow Y$ be quasiconformal embeddings, let \bar{f}_0 and \bar{f}_1 be their extensions to $X \cup \dot{X}$, and let $H : X \times [0, 1] \rightarrow Y$ be a homotopy from f_0 to f_1 . Let $E \subset \dot{X}$ be the set of punctures which are essential with respect to f_0 . Then there exists a homotopy $\bar{H} : X \cup \dot{X} \times [0, 1] \rightarrow Y \cup \dot{Y}$ from \bar{f}_0 to \bar{f}_1 which is constant on $E \times [0, 1]$, maps $X \times [0, 1]$ into Y , and whose restriction to $X \times [0, 1]$ is homotopic to H .*

Proof. For each $p \in \dot{X}$, let $D_p \subset X \cup \dot{X}$ be an embedded closed disk such that $D_p \cap \dot{X} = \{p\}$. Further assume that the disks D_p are all disjoint. The idea is to define $\bar{H} = H$ on $X \setminus \bigcup_{p \in \dot{X}} D_p \times [0, 1]$ and then extend this to a continuous map $\bar{H} : D_p \times [0, 1] \rightarrow Y \cup \dot{Y}$ sending $D_p \setminus \{p\}$ into Y for each $p \in \dot{X}$. If the puncture $p \in \dot{X}$ is inessential, then the annulus $H(\partial D_p \times [0, 1])$ is contractible in Y so that H extends to a continuous map $\bar{H} : D_p \times [0, 1] \rightarrow Y$. If $p \in \dot{X}$ is essential, then the annulus $H(\partial D_p \times [0, 1])$ is contractible in $Y \cup \{f_0(p)\}$ so that H extends to a continuous map $\bar{H} : D_p \times [0, 1] \rightarrow Y \cup \{f_0(p)\}$. We can further choose the extension to satisfy $\bar{H}^{-1}(f_0(p)) = \{p\} \times [0, 1]$. \square

2.9. Geometry of quadratic differentials. Let X be a finite Riemann surface and let $\varphi \in \mathcal{Q}(X) \setminus \{0\}$. Then φ induces a notion of direction as well as a Riemannian metric with singularities $|\varphi|$ on X . This metric is Euclidean except at the zeros of φ , where it has cone points. At a zero of order k , the metric looks like a cone with angle $(k + 2)\pi$. Because the total angle at cone points is at least 2π , the induced metric is locally CAT(0). However, the metric $|\varphi|$ is not complete whenever X is not closed. Nevertheless, given an arc $\gamma : [0, 1] \rightarrow X$ there exists a unique “geodesic representative” $\gamma^\dagger : [0, 1] \rightarrow \widehat{X}$ which is a limit of arcs homotopic to γ in X and has minimal $|\varphi|$ -length among such arcs. To see this, it is convenient to pass to the universal cover first.

Let $\pi : \widetilde{X} \rightarrow X$ be the universal covering map and let $\widetilde{\varphi} = \pi^*\varphi$. The Riemannian metric with singularities $|\widetilde{\varphi}|$ induces a distance d on \widetilde{X} in the usual way: for any $x, y \in \widetilde{X}$ the distance $d(x, y)$ is defined as the infimum of $\int_\gamma \sqrt{|\widetilde{\varphi}|}$ over all piecewise smooth paths γ between x and y . The resulting metric space (\widetilde{X}, d) is CAT(0). It follows that its metric completion (\widetilde{X}', d') is also CAT(0) [BH09, Corollary 3.11]. When X is hyperbolic, \widetilde{X}' is the universal cover \widetilde{X} together with the set of discontinuity and the parabolic limit points. Observe that \widehat{X} is the metric completion of X in the metric $|\varphi|$ and that π extends to a continuous map $\pi' : \widetilde{X}' \rightarrow \widehat{X}$.

Since (\widetilde{X}', d') is a complete CAT(0) space, there is a unique geodesic segment between any two points of \widetilde{X}' . Thus given an arc $\gamma : [0, 1] \rightarrow X$, we can lift it to an arc $\widetilde{\gamma} : [0, 1] \rightarrow \widetilde{X}$, find the geodesic $\widetilde{\gamma}^\dagger : [0, 1] \rightarrow \widetilde{X}'$ between the endpoints of $\widetilde{\gamma}$, then define the “geodesic representative”⁴ of γ to be $\gamma^\dagger = \pi' \circ \widetilde{\gamma}^\dagger$.

By hypothesis, φ extends analytically to the ideal boundary of X (if any). Hence $\widetilde{\varphi}$ extends analytically to the set of discontinuity in \widetilde{X}' . Say that a *singularity of $\widetilde{\varphi}$* is a point in \widetilde{X}' where $\widetilde{\varphi}$ has a zero or is not defined (those are the parabolic

⁴The quotation marks are because γ^\dagger is not necessarily geodesic in \widehat{X} . For example, it may pass through a pole of φ .

limit points). Since (\tilde{X}', d') is locally Euclidean, geodesics are straight lines in natural coordinates away from the singularities of $\tilde{\varphi}$. If a geodesic passes through a singularity, then there should not be any shortcut on either side, which translates into an angle condition at the singularity. More precisely, an arc $\gamma : [0, 1] \rightarrow \tilde{X}'$ is geodesic if and only if

- γ is smooth except at singularities of $\tilde{\varphi}$;
- the argument of $\tilde{\varphi}(\gamma'(t))$ is locally constant where γ is smooth;
- at a singularity of $\tilde{\varphi}$, the cone angle on either side⁵ of γ is at least π .

For example, an arc which is horizontal and does not backtrack is geodesic. Actually, any horizontal arc which does not backtrack minimizes horizontal travel (defined as the integral of $|\operatorname{Re} \sqrt{\tilde{\varphi}}|$) between its endpoints, since it is quasi-transverse to the vertical foliation [HM79, Proposition 2.5].

A *geodesic polygon* in a CAT(0) space is a closed curve which is piecewise geodesic. Self-intersecting polygons are allowed. Given a geodesic polygon P in a CAT(0) space, and three consecutive vertices a , b and c along P , let $\angle abc$ be the Alexandrov angle between the geodesics $[a, b]$ and $[b, c]$ at b [BH09, p.184]. If P is a geodesic polygon in \tilde{X}' , then $\angle abc$ is the minimum between π and the two cone angles on either side of $[a, b] \cup [b, c]$ at b , as measured in the metric induced by $\tilde{\varphi}$.

The next Proposition is a generalization of Strebel's "divergence principle" [Str84, p.77]. It says that if two geodesic rays γ_0 and γ_1 in \tilde{X}' are such that the angles they form with the geodesic $[\gamma_0(0), \gamma_1(0)]$ sum to at least π , then the distance between $\gamma_0(t)$ and $\gamma_1(t)$ is non-decreasing as a function of $t \geq 0$. In the original statement, each angle is assumed to be at least $\pi/2$. The result actually holds in any CAT(0) space.

Lemma 2.12. *Let Q be a geodesic quadrilateral with vertices a, b, c, d in a CAT(0) space (M, μ) . Suppose that $\mu(a, d) = \mu(b, c)$ and $\angle dab + \angle abc \geq \pi$. Then $\mu(c, d) \geq \mu(a, b)$ with equality only if Q is isometric to a (possibly degenerate) Euclidean parallelogram.*

Proof. Consider the geodesic triangles Δabc and Δacd . Let $\Delta a'b'c'$ and $\Delta a'c'd'$ be corresponding comparison triangles in \mathbb{R}^2 . We may assume that these two triangles are on opposite sides of the segment $[a', c']$ so that the quadrilateral Q' they form is simple. We need to show that $|c' - d'| \geq |a' - b'|$. By the law of cosines, it suffices to prove that $\angle d'a'c' \geq \angle b'c'a'$. Indeed, the triangles $\Delta a'b'c'$ and $\Delta a'c'd'$ share the side $[a', c']$ and the sides $[b', c']$ and $[d', a']$ are congruent. There are two cases to consider depending on whether Q' is convex or not.

Suppose first that $\angle d'a'c' + \angle c'a'b' \leq \pi$. Note that angles in $\Delta a'b'c'$ and $\Delta a'c'd'$ are at least as large as corresponding angles in Δabc and Δacd . Moreover, angles are subadditive in the sense that $\angle dab \leq \angle dac + \angle cab$ for instance. Since the sum of $\angle d'a'c'$ and $\angle c'a'b'$ is at most π , we have

$$\angle d'a'b' = \angle d'a'c' + \angle c'a'b' \geq \angle dac + \angle cab \geq \angle dab$$

so that

$$\angle d'a'b' + \angle a'b'c' \geq \angle dab + \angle abc \geq \pi.$$

⁵At a boundary point of \tilde{X}' , the cone angle only makes sense on one side of γ , but we can define the cone angle on the "other side" to be $+\infty$ by convention.

We deduce that

$$\begin{aligned} 0 &\leq \angle d'a'b' + \angle a'b'c' - \pi \\ &= (\angle d'a'c' + \angle c'a'b') + \angle a'b'c' - (\angle c'a'b' + \angle a'b'c' + \angle b'c'a') \\ &= \angle d'a'c' - \angle b'c'a' \end{aligned}$$

as required.

If $\angle d'a'c' = \angle b'c'a'$, then Q' is a parallelogram (which is possibly contained in a line) and all the above inequalities are equations. Thus the angles in Δabc and Δacd are the same as in their comparison triangles, which implies that Δabc and Δacd are isometric to their comparison triangles. By considering the other diagonal of Q , we get that any two adjacent sides of Q lie in a flat triangle. The equality $\angle dab = \angle dac + \angle cab$ means that the triangles Δabc and Δacd line up, in the sense that the union of their convex hulls is convex. Hence Q spans a Euclidean parallelogram.

Now suppose that $\angle d'a'c' + \angle c'a'b' > \pi$. Then

$$\begin{aligned} 0 &\leq \pi + \angle a'b'c' - \pi \\ &< (\angle d'a'c' + \angle c'a'b') + \angle a'b'c' - (\angle c'a'b' + \angle a'b'c' + \angle b'c'a') \\ &= \angle d'a'c' - \angle b'c'a'. \end{aligned}$$

Equality cannot hold in this case. \square

One can deduce from this the well-known fact that if two simple closed geodesics in X (with respect to $|\varphi|$) are freely homotopic, then they span a Euclidean cylinder.

3. IOFFE'S THEOREM

The goal of this section is to reprove Theorem 1.3 from Ioffe, that is, to characterize quasiconformal embeddings that have minimal dilatation in their homotopy class. The main motivation for doing so is that Ioffe's original statement contains a mistake (we pointed this out in the introduction and will return to this issue in Subsection 3.3). Thus we wanted to work out the proof in detail to make sure that the remaining part was correct.

The results are stated for quasiconformal embeddings from one finite Riemann surface to another, but they extend to quasiconformal embeddings from a finite union of finite Riemann surfaces to a finite Riemann surface, which is the context considered in Ioffe's paper [Iof75].

3.1. Compactness. Recall that an embedding h between finite Riemann surfaces is *trivial* if the image of $\pi_1(h)$ is trivial, is *cyclic* if the image of $\pi_1(h)$ is infinite cyclic, and is *generic* if the image of $\pi_1(h)$ is non-abelian. An embedding h is called *parabolic* if the image of $\pi_1(h)$ is infinite cyclic and generated by a loop around a puncture.

The following compactness lemma (cf. [Iof75, Theorem 1.1]) guarantees the existence of extremal quasiconformal embeddings, that is, with minimal dilatation in their homotopy class.

Lemma 3.1. *Let $K \geq 1$ and let $h : X \rightarrow Y$ be a K -quasiconformal embedding between finite Riemann surfaces. The space of all K -quasiconformal embeddings homotopic to h is compact if and only if h is neither trivial nor parabolic.*

Proof. Suppose that Y is equal to $\widehat{\mathbb{C}}$, \mathbb{C} or $\mathbb{C} \setminus \{0\}$. Then h is either trivial or parabolic. Moreover, the space of K -quasiconformal embeddings homotopic to h is closed under post-composition with elements of $\text{Aut}_0(Y)$, which is non-compact. It follows that the space of K -quasiconformal embeddings homotopic to h is also non-compact.

Suppose that Y is a torus. The group $\text{Aut}_0(Y)$ is homeomorphic to Y and is in particular compact. Let W be the set of all K -quasiconformal embeddings homotopic to h and let N be the space of K -quasiconformal embeddings homotopic to h rel x_0 , where x_0 is any point in X . Then W is homeomorphic to $\text{Aut}_0(Y) \times N$, hence is compact if and only if N is. Now N is the same as the space of K -quasiconformal embeddings homotopic to the restriction $h^* : X \setminus \{x_0\} \rightarrow Y \setminus \{h(x_0)\}$. Moreover h non-trivial if and only if h^* is non-parabolic. The case where Y is a torus thus reduces to the case where Y is a once-punctured torus, hence hyperbolic.

For the rest of the proof, we may assume that Y is hyperbolic. This implies that X is hyperbolic as well, since there is no non-constant quasiconformal map from a non-hyperbolic surface to a hyperbolic surface.

If h is trivial then its image is contained in a disk in Y . In other words, h can be written as $h = F \circ g$ where $F : \mathbb{D} \rightarrow Y$ is a conformal embedding and $g : X \rightarrow \mathbb{D}$ is a K -quasiconformal embedding. Consider the sequence $h_n = F_n \circ g$ where $F_n(z) = F(z/n)$. Each h_n is a K -quasiconformal embedding homotopic to h , but the sequence converges to a constant map. Similarly, if h is parabolic, then we can form a sequence of homotopic K -quasiconformal embeddings which diverges to the corresponding puncture.

If h is neither trivial nor parabolic then it is either generic, or cyclic but not parabolic. We treat the generic case first and return to the cyclic case at the end. Suppose that h is generic. Pick any basepoint $b_0 \in X$ and let $\pi_X : (\mathbb{D}, 0) \rightarrow (X, b_0)$ and $\pi_Y : (\mathbb{D}, 0) \rightarrow (Y, h(b_0))$ be universal covering maps with respective deck groups $\text{Deck}(X)$ and $\text{Deck}(Y)$. Let \tilde{h} be the unique lift of h such that $\tilde{h}(0) = 0$. For any $\alpha \in \text{Deck}(X)$, let $\Theta(\alpha)$ be the unique element in $\text{Deck}(Y)$ sending 0 to $\tilde{h}(\alpha \cdot 0)$. Then Θ is a homomorphism.

If $f : X \rightarrow Y$ is a K -quasiconformal embedding homotopic to h , then it lifts to a K -quasiconformal immersion $\tilde{f} : \mathbb{D} \rightarrow \mathbb{D}$ satisfying

$$\tilde{f}(\alpha \cdot z) = \Theta(\alpha) \cdot \tilde{f}(z)$$

for every $z \in \mathbb{D}$ and every $\alpha \in \text{Deck}(X)$. The lift \tilde{f} is uniquely determined by this equation, since there is only one homotopy class of homotopy from h to f according to Lemma 6.1. We do not really need uniqueness of the lift here (any lift would do), which is why the lemma is postponed to a later section. Observe that f is injective if and only if

$$\pi_Y(\tilde{f}(z)) = \pi_Y(\tilde{f}(w)) \Rightarrow \pi_X(z) = \pi_X(w).$$

Let f be a K -quasiconformal embedding homotopic to h and \tilde{f} be its lift as above. Write $\tilde{f} = F \circ g$ where $g : \mathbb{D} \rightarrow \mathbb{D}$ is a K -quasiconformal homeomorphism fixing 0 and $F : \mathbb{D} \rightarrow \mathbb{D}$ is a holomorphic immersion. As mentioned in Section 2, the space of K -quasiconformal homeomorphisms of \mathbb{D} fixing the origin and the space of holomorphic maps from \mathbb{D} to $\overline{\mathbb{D}}$ are both compact. Thus, given a sequence of K -quasiconformal embeddings f_n homotopic to h and corresponding lifts $\tilde{f}_n = F_n \circ g_n$

factored as above, we can pass to a subsequence such that $F_n \rightarrow \tilde{F}$ and $g_n \rightarrow g$ and hence $\tilde{f}_n \rightarrow \tilde{f} := F \circ g$.

We claim that the limit \tilde{f} is not constant. Indeed, if there is a constant $w_0 \in \overline{\mathbb{D}}$ such that $\tilde{f}(z) = w_0$ for every $z \in \mathbb{D}$, then

$$w_0 = \tilde{f}(\alpha \cdot z) = \lim_{n \rightarrow \infty} \tilde{f}_n(\alpha \cdot z) = \lim_{n \rightarrow \infty} \Theta(\alpha) \cdot \tilde{f}_n(z) = \Theta(\alpha) \cdot \tilde{f}(z) = \Theta(\alpha) \cdot w_0$$

for every $\alpha \in \text{Deck}(X)$. However, a Fuchsian group which fixes a point in $\overline{\mathbb{D}}$ is cyclic, contradicting the assumption that h is generic. In particular, the holomorphic map g is not constant, hence has image in \mathbb{D} . By Hurwitz's theorem in complex analysis, g is locally injective. Therefore $\tilde{f} = F \circ g$ is a K -quasiconformal immersion. Moreover, the equality $\tilde{f}(\alpha \cdot z) = \Theta(\alpha) \cdot \tilde{f}(z)$ for every $z \in \mathbb{D}$ and every $\alpha \in \text{Deck}(X)$ implies that \tilde{f} descends to a K -quasiconformal immersion $f : X \rightarrow Y$.

It remains to show that f is injective. If $f(\pi_X(z)) = f(\pi_X(w))$, then $\pi_Y(\tilde{f}(z)) = \pi_Y(\tilde{f}(w))$. Since $\tilde{f}_n \rightarrow \tilde{f}$ and since these maps are open, we can find a sequence z_n converging to z and a sequence w_n converging to w such that $\tilde{f}_n(z_n) = \tilde{f}(z)$ and $\tilde{f}_n(w_n) = \tilde{f}(w)$ for all n large enough. Then $\pi_Y(\tilde{f}_n(z_n)) = \pi_Y(\tilde{f}_n(w_n))$, which implies that $\pi_X(z_n) = \pi_X(w_n)$ since \tilde{f}_n is injective. Taking the limit as $n \rightarrow \infty$, we obtain $\pi_X(z) = \pi_X(w)$.

Lastly, f is homotopic to h because it is a limit of maps which are (see Corollary 2.9). We have shown that if Y is hyperbolic and h is generic, then the space of K -quasiconformal embeddings homotopic to h is sequentially compact.

Suppose now that h is cyclic but not parabolic. The same construction as above still applies, but the image of Θ is a cyclic group generated by a hyperbolic element $\beta \in \text{Deck}(Y)$. Moreover, given an embedding f homotopic to h , its lift \tilde{f} is only defined up to post-composition with powers of β . Let $D \subset \mathbb{D}$ be a fundamental domain for β whose closure is disjoint from the fixed points of β . Let $\{f_n\}_{n \geq 1}$ be a sequence of K -quasiconformal embeddings homotopic to h . By applying an appropriate power of β , we can choose a lift \tilde{f}_n of f_n such that $\tilde{f}_n(0) \in D$. As above, we can extract a subsequence of $\{\tilde{f}_n\}_{n \geq 1}$ converging to some map \tilde{f} . By our normalization, the limit $\tilde{f}(0)$ belongs to the closure \overline{D} , so it is not one of the fixed points of β . It follows that \tilde{f} is not constant. The rest of the argument applies verbatim to conclude that \tilde{f} descends to a K -quasiconformal embedding f homotopic to h . \square

The following corollary is immediate.

Corollary 3.2. *Let $h : X \rightarrow Y$ be a non-trivial and non-parabolic quasiconformal embedding between finite Riemann surfaces. Among the quasiconformal embeddings homotopic to h , there is one with minimal dilatation.*

Proof. Dilatation is lower semi-continuous, hence achieves its minimum on any compact set. \square

Remark. If $h : X \rightarrow Y$ is a quasiconformal embedding between finite Riemann surfaces which is trivial or parabolic, then it is homotopic to a conformal embedding. This is because $h(X)$ is contained in a subset $Z \subset Y$ isomorphic to the disk with at most one puncture or the sphere with at most two punctures, whose Teichmüller space is a point. Using Theorem 2.4, we can apply a quasiconformal deformation on

Z to make h conformal. The hypothesis that h be non-trivial and non-peripheral is thus superfluous in Corollary 3.2.

3.2. Teichmüller embeddings. Let us recall the definition of a Teichmüller embedding.

Definition 3.3. A *Teichmüller embedding* of dilatation $K \geq 1$ between finite Riemann surfaces X and Y is an injective continuous map $f : X \rightarrow Y$ for which there exist quadratic differentials $\varphi \in \mathcal{Q}^+(X)$ and $\psi \in \mathcal{Q}^+(Y)$ such that f has the form $x + iy \mapsto Kx + iy$ in natural coordinates and such that $Y \setminus f(X)$ is a finite union of points and horizontal arcs for ψ . We say that φ and ψ are *initial* and *terminal* quadratic differentials for f . A *slit mapping* is a conformal Teichmüller embedding, i.e. one with $K = 1$.

Remark. The horizontal arcs in the definition are allowed to overlap. In general, the complement of the image of a Teichmüller embedding is an analytic graph (see Figure 4 for example).

Remark. Despite the appellation, a Teichmüller homeomorphism between surfaces with non-empty ideal boundary is not necessarily a Teichmüller embedding, as its initial and terminal quadratic differentials are allowed to be negative along the ideal boundary.

Remark. If Y is a finite Riemann surface, $\psi \in \mathcal{Q}^+(Y)$, and X is the complement of a finite union of points and of horizontal arcs in Y , then $\iota^*\psi$ belongs to $\mathcal{Q}^+(X)$, where $\iota : X \rightarrow Y$ is the inclusion map. This means that $\iota^*\psi$ extends analytically to the ideal boundary of X and that the latter is horizontal. Near the endpoint of a horizontal slit, one needs to take a square root in order to unfold the slit to an ideal boundary component. If we pull-back the quadratic differential $z^k dz^2$ in \mathbb{C} by the squaring map $s(z) = z^2$ from \mathbb{H} to $\mathbb{C} \setminus [0, \infty)$, we get the quadratic differential $4z^{2k+2} dz^2$. In other words, unfolding a slit ending at a singularity of order $k \geq -1$ yields a singularity of order $2k+2 \geq 0$ on the boundary. In particular, if the slit ends at a simple pole then the unfolded quadratic differential is regular at the corresponding point. It is perhaps more natural to count the number of prongs: an n -prong singularity transforms into half of a $2n$ -prong singularity (see Figure 1). The ideal boundary remains horizontal in the process of unfolding so that $\iota^*\psi \in \mathcal{Q}^+(X)$.

Every slit mapping $f : X \rightarrow Y$ arises in this way, in the sense that f gives a conformal homeomorphism from X to $f(X)$ and $f(X) \subset Y$ is obtained by removing finitely many points and horizontal arcs for some $\psi \in \mathcal{Q}^+(Y)$. The quadratic differential $\varphi = f^*\psi$ is redundant data.

Remark. Every Teichmüller embedding can be factored as a Teichmüller homeomorphism $G : X \rightarrow Z$ (which happens to also be a Teichmüller embedding) followed by a slit mapping $F : Z \rightarrow Y$, where the terminal quadratic differential of G agrees with the initial quadratic differential of F .

Remark. If there is a Teichmüller embedding $X \rightarrow Y$, then $\mathcal{Q}^+(Y)$ is non-empty. This imposes some restrictions on Y , namely, it cannot be a sphere with at most 3 punctures or a disk with at most 1 puncture. By the same token, a Teichmüller embedding is always non-trivial and non-parabolic.




FIGURE 1. Unfolding a slit ending at a simple zero (a 3-prong singularity) of a quadratic differential yields half of a zero of order 4 (a 6-prong singularity).

Remark. If f is a Teichmüller embedding of dilatation K with initial and terminal quadratic differentials φ and ψ , then $\bar{\partial}f/\partial f = k\bar{\varphi}/|\varphi|$ on X and $\bar{\partial}(f^{-1})/\partial(f^{-1}) = -k\bar{\psi}/|\psi|$ on $f(X)$, where $k = \frac{K-1}{K+1}$. This is one way of describing f without using coordinates.

We present Ioffe's theorem in two parts. The first part says that every extremal quasiconformal embedding is either conformal or a Teichmüller embedding.

Theorem 3.4 (Extremal embeddings are Teichmüller). *Let $f : X \rightarrow Y$ be a quasiconformal embedding between finite Riemann surfaces with minimal dilatation in its homotopy class. If f is not conformal, then it is a Teichmüller embedding.*

Proof. Suppose that f is not conformal and let

$$\mu := \begin{cases} \bar{\partial}(f^{-1})/\partial(f^{-1}) & \text{on } f(X) \\ 0 & \text{on } Y \setminus f(X). \end{cases}$$

Let $F : Y \rightarrow Y_\mu$ be the solution to the Beltrami equation $\bar{\partial}F/\partial F = \mu$ provided by Theorem 2.4. By construction we have $\bar{\partial}(F \circ f)/\partial(F \circ f) = 0$ so that $F \circ f$ is a conformal embedding. By Teichmüller's theorem (Theorems 2.6 and 2.7), there exists a quasiconformal homeomorphism $G : Y_\mu \rightarrow Y$ homotopic to F^{-1} with minimal dilatation. Moreover, G is either conformal or a Teichmüller homeomorphism. The composition $G \circ F \circ f : X \rightarrow Y$ is a quasiconformal embedding homotopic to f , so that

$$\text{Dil}(f) \leq \text{Dil}(G \circ F \circ f) \leq \text{Dil}(G) \leq \text{Dil}(F^{-1}) = \text{Dil}(f).$$

Thus all the terms in this chain are equal. The equality $\text{Dil}(G) = \text{Dil}(F^{-1})$ implies that F has minimal dilatation in its homotopy class. Since F is not conformal, it is a Teichmüller homeomorphism. This means that there is a non-zero $\psi \in \mathcal{Q}(Y)$ and a constant $k \in (0, 1)$ such that $\mu = -k\bar{\psi}/|\psi|$ almost everywhere. In particular, $Y \setminus f(X)$ has measure zero and f has constant pointwise dilatation.

Thought of as a homeomorphism from X to $f(X)$, the map f has minimal dilatation in its homotopy class and is thus a Teichmüller homeomorphism. Let φ and ω be its initial and terminal quadratic differentials. Since $F \circ f$ is conformal, the directions of maximal stretching for F and f^{-1} must be perpendicular, which means that $\psi = c\omega$ on $f(X)$ for some $c > 0$, which we may assume is equal to 1 by rescaling.

We have to show that $\varphi \in \mathcal{Q}^+(X)$. If not, then $\varphi < 0$ along some segment $I \subset \partial X$. We will explicitly construct a quasiconformal embedding \tilde{f} from X to

$f(X)$ with pointwise dilatation smaller than f near I . We may work in a natural coordinate chart for φ in which I is equal to the vertical segment $[-i, i]$ in the plane and X is to the right of I . There is also a natural chart for ω in which f takes the form $x + iy \mapsto Kx + iy$. Let Δ be the isoceles triangle with base $[-i, i]$ and apex $\delta > 0$. Consider the map $L : \Delta \rightarrow \Delta$ which is affine on the upper and lower




FIGURE 2. Reducing the pointwise dilatation of an embedding near a vertical boundary arc.

halves of Δ , fixes all three vertices of Δ , and sends the midpoint of I to $(1 - 1/K)\delta$. Extend L to be the identity on $X \setminus \Delta$ and let $\tilde{f} = f \circ L$. The linear part of \tilde{f} on the lower half of Δ is equal to

$$\begin{pmatrix} 1 & (K-1)\delta \\ 0 & 1 \end{pmatrix}$$

and the dilatation of this matrix tends to 1 as $\delta \rightarrow 0$. A similar statement holds in the upper half of Δ . Therefore, if δ is small enough, then the embedding \tilde{f} has strictly smaller pointwise dilatation than f on Δ . Moreover, the global dilatation of \tilde{f} is the same as f , so that \tilde{f} also has minimal dilatation in its homotopy class. By the first paragraph of the proof, the pointwise dilatation of \tilde{f} must be constant. This is a contradiction, and hence $\varphi \in \mathcal{Q}^+(X)$.

It remains to show that $f(X)$ is the complement of a graph which is horizontal with respect to ψ . Recall that \hat{X} is the compactification of X obtained by adding its ideal boundary and filling its punctures. The metrics induced by $|\varphi|$ and $|\psi|$ extend to complete metrics on \hat{X} and \hat{Y} . Since $f : X \rightarrow Y$ is K -Lipschitz with respect to these metrics, it extends to a K -Lipschitz map $\hat{f} : \hat{X} \rightarrow \hat{Y}$. Moreover, \hat{f} is surjective since $Y \setminus f(X)$ has measure zero and hence empty interior. Let I be the closure of a connected component of $\partial X \setminus \{\text{zeros of } \varphi\}$. There is a sequence $\{I_n\}$ of arcs in X which are horizontal for φ and converge uniformly to I . Since the image arcs $f(I_n)$ are all horizontal for ψ , they can only accumulate onto horizontal arcs, and thus $\hat{f}(I)$ is horizontal. Therefore, the complement $\hat{Y} \setminus f(X) = \hat{f}(\partial X \cup \hat{X})$ is a union of finitely many points and horizontal arcs for ψ . In particular, the ideal boundary ∂Y is horizontal for ψ so that $\psi \in \mathcal{Q}^+(Y)$. \square

In the last paragraph of the proof we actually showed the following useful criterion.

Lemma 3.5. *Let X and Y be finite Riemann surfaces, let $\varphi \in \mathcal{Q}^+(X)$ and $\psi \in \mathcal{Q}(Y) \setminus \{0\}$, and let $f : X \rightarrow Y$ be an embedding which is locally of the form $x + iy \mapsto Kx + iy$ in natural coordinates. If $f(X)$ is dense in Y , then $\psi \in \mathcal{Q}^+(Y)$ and f is a Teichmüller embedding with respect to φ and ψ .*

The second part of Ioffe's theorem says that every Teichmüller embedding is extremal. The proof is very similar to the proofs of Teichmüller's uniqueness theorem given in [Ber58] and [FM11, Chapter 11], only with additional technicalities due to the lack of compactness.

Theorem 3.6 (Teichmüller embeddings are extremal). *Let $f : X \rightarrow Y$ be a Teichmüller embedding of dilatation K with initial and terminal quadratic differential φ and ψ , and let $g : X \rightarrow Y$ be a K -quasiconformal embedding homotopic to f . Then g is a Teichmüller embedding of dilatation K with initial and terminal quadratic differentials φ and ψ , and $g \circ f^{-1} : f(X) \rightarrow g(X)$ is locally a translation in natural coordinates for ψ .*

Proof. We may assume that all the punctures of X are essential with respect to f . Otherwise, the set I of inessential punctures can be filled and f extended to a Teichmüller embedding $\tilde{f} : X \cup I \rightarrow Y$. Similarly, g extends to a K -quasiconformal embedding $\tilde{g} : X \cup I \rightarrow Y$ homotopic to \tilde{f} . If \tilde{g} is a Teichmüller embedding with respect to ψ , then so is its restriction g .

By rescaling, we may assume that $\|\varphi\| = 1$. This implies that $\|\psi\| = K$ since f multiplies area by a factor K and $f(X)$ has full measure in Y . Recall that \widehat{X} and \widehat{Y} are the metric completions of X and Y with respect to the distance induced by $|\varphi|$ and $|\psi|$. Since f is K -Lipschitz with respect to these metrics, it extends to a K -Lipschitz map $\widehat{f} : \widehat{X} \rightarrow \widehat{Y}$. Since f is assumed to have only essential punctures, \widehat{f} maps \dot{X} into \dot{Y} .

Let $M > 1$. For every $n \in \mathbb{N}$, let $G_n : \widehat{X} \rightarrow \widehat{X}$ be a smooth M -quasiconformal and M -Lipschitz embedding such that

- $G_n(\dot{X}) = \dot{X}$;
- G_n is homotopic to the identity $\text{rel } \dot{X}$;
- $G_n(\widehat{X})$ is contained in $X \cup \dot{X}$;
- $G_n \rightarrow \text{id}$ uniformly in the C^1 norm as $n \rightarrow \infty$.

Here is one way to construct such maps. Given a smooth vector field on \widehat{X} pointing inwards on ∂X and vanishing on \dot{X} , we can let G_n be the corresponding flow at small enough time t_n . Then let $g_n = g \circ G_n : \widehat{X} \rightarrow \widehat{Y}$, which is a KM -quasiconformal embedding. The purpose of this construction is to circumvent the following difficulty: when $X \cup \dot{X}$ is non-compact g has no reason to extend continuously to ∂X while g_n does.

Fix n for a little while. By construction, $g_n|_X$ is homotopic to g and hence to f . By Lemma 2.11, there is a homotopy $H : X \times [0, 1] \rightarrow Y$ from f to $g_n|_X$ that extends to be constant at the punctures (which are assumed to be essential). By flowing ∂X inward at the beginning and end of the homotopy as above, we may further assume that H extends continuously to $\partial X \times [0, 1]$. Let \widehat{H} be the extension of H to $\widehat{X} \times [0, 1]$.

For every $x \in X$, let $\ell(x) = \inf \int_\alpha \sqrt{|\psi|}$ where the infimum is taken over all piecewise smooth paths $\alpha : [0, 1] \rightarrow Y$ that are homotopic to $t \mapsto H(x, t)$ rel endpoints. In general, the infimum need not be realized since the restriction of $|\psi|$ to Y is not complete. However, $\ell(x)$ is equal to the length of the “geodesic representative” γ_x of $t \mapsto H(x, t)$ in the completion \widehat{Y} as explained in Subsection 2.9.

Since H is continuous, the maps $x \mapsto \gamma_x$ and $x \mapsto \ell(x)$ are continuous. Moreover, they extend continuously to \widehat{X} . Indeed, for $x \in \widehat{X} \setminus X$ we can define γ_x as the limit of γ_{x_n} as $n \rightarrow \infty$, where $\{x_n\}_{n \geq 1}$ is a sequence in X converging to x . This limit exists and does not depend on the sequence $\{x_n\}_{n \geq 1}$ since the path $t \mapsto \widehat{H}(x, t)$ is well-defined even for $x \in \widehat{X} \setminus X$. Moreover, its length $\ell(x)$ is the limit of the lengths $\ell(x_n)$. As \widehat{X} is compact, there exists a constant B such that $\ell(x) < B$ for every $x \in \widehat{X}$. This constant B depends on n , but we will make it disappear before changing n .

Let η be a horizontal arc of length $L > 0$ in X . Since f is a Teichmüller embedding of dilatation K , it sends η to a horizontal arc of length KL in Y . Let x_0 and x_1 be the endpoints of η . We can obtain a path homotopic to $f(\eta)$ in Y by taking the concatenation of a piecewise smooth path α_0 of length at most B homotopic to $t \mapsto H(x_0, t)$, the image $g_n(\eta)$, and a piecewise smooth path α_1 of length at most B homotopic to $t \mapsto H(x_1, 1 - t)$. Since horizontal arcs minimize horizontal travel among all homotopic paths, we have

$$\begin{aligned} KL &= \int_{f(\eta)} |\operatorname{Re} \sqrt{\psi}| \leq \int_{\alpha_0} |\operatorname{Re} \sqrt{\psi}| + \int_{g_n(\eta)} |\operatorname{Re} \sqrt{\psi}| + \int_{\alpha_1} |\operatorname{Re} \sqrt{\psi}| \\ &\leq 2B + \int_{g_n(\eta)} |\operatorname{Re} \sqrt{\psi}|. \end{aligned}$$

Let dg_n denote the matrix of partial derivatives of g_n with respect to natural coordinates⁶ and $(dg_n)_{1,1}$ its first entry. If $z = x + iy$ and ζ are natural coordinates for φ and ψ , then $(dg_n)_{1,1} = \operatorname{Re}(\partial(\zeta \circ g_n \circ z^{-1})/\partial x)$. If g_n is absolutely continuous on η , then we have

$$\int_{\eta} |(dg_n)_{1,1}| \cdot \sqrt{|\varphi|} = \int_{g_n(\eta)} |\operatorname{Re} \sqrt{\psi}| \geq KL - 2B.$$

Remove from X all trajectories that go through a puncture of X or a zero of φ and denote the resulting full measure subset by U . For every $x \in U$, there is a unique (possibly closed) bi-infinite horizontal trajectory through x . For every $L > 0$ and every $x \in U$, let $\eta(x, L)$ be the horizontal arc of length L centered at x . Since g_n is quasiconformal, it is absolutely continuous on almost every horizontal trajectory. Upon applying Fubini's theorem, we find

$$\begin{aligned} \int_U |(dg_n)_{1,1}| \cdot |\varphi| &= \int_{x \in U} \left(\frac{1}{L} \int_{\eta(x, L)} |(dg_n)_{1,1}| \cdot \sqrt{|\varphi|} \right) \cdot |\varphi| \\ &\geq \left(K - \frac{2B}{L} \right) \int_U |\varphi|. \end{aligned}$$

Letting $L \rightarrow \infty$, we obtain $\int_U |(dg_n)_{1,1}| \cdot |\varphi| \geq K \int_U |\varphi|$ and hence

$$\int_X |(dg_n)_{1,1}| \cdot |\varphi| \geq K \int_X |\varphi| = K,$$

since U has full measure in X .

Now that we got rid of the constant $B = B(n)$, we can vary the index n . We claim that $\int_X |(dg)_{1,1}| \cdot |\varphi| = \lim_{n \rightarrow \infty} \int_X |(dg_n)_{1,1}| \cdot |\varphi|$ and hence that $\int_X |(dg)_{1,1}| \cdot |\varphi| \geq K$. This is a consequence of the Vitali convergence theorem [RF10, p.94]. In order

⁶The matrix is only defined up to sign, but no matter.

to apply the theorem, we need to check that the functions $|(dg_n)_{1,1}|$ are uniformly integrable. First observe that

$$\int_X \det(dg) \cdot |\varphi| = \int_{g(X)} |\psi| \leq K,$$

so that $\det(dg)$ is integrable. It follows that for every $\varepsilon > 0$, there exists a $\delta > 0$ such that if $A \subset X$ is measurable and $\int_A |\varphi| < \delta$, then $\int_A \det(dg) \cdot |\varphi| < \varepsilon$. Now if $\int_A |\varphi| < \delta/M^2$, then $\int_{G_n(A)} |\varphi| < \delta$ since G_n is M -Lipschitz. By the Cauchy-Schwarz inequality we have

$$\begin{aligned} \left(\int_A |(dg_n)_{1,1}| \cdot |\varphi| \right)^2 &\leq \int_A |(dg_n)_{1,1}|^2 \cdot |\varphi| \leq \int_A \|dg_n\|^2 \cdot |\varphi| \\ &\leq KM \int_A \det(dg_n) \cdot |\varphi| \\ &= KM \int_A \det(d_{G_n(z)}g) \det(d_z G_n) \cdot |\varphi| \\ &= KM \int_{G_n(A)} \det(dg) \cdot |\varphi| < KM\varepsilon, \end{aligned}$$

which shows uniform integrability and proves the claim.

Applying Cauchy-Schwarz to the inequality $K \leq \int_X |(dg)_{1,1}| \cdot |\varphi|$ yields

$$\begin{aligned} K^2 &\stackrel{(a)}{\leq} \left(\int_X |(dg)_{1,1}| \cdot |\varphi| \right)^2 \stackrel{(b)}{\leq} \int_X |(dg)_{1,1}|^2 \cdot |\varphi| \\ &\stackrel{(c)}{\leq} \int_X \|dg\|^2 \cdot |\varphi| \leq K \int_X \det(dg) \cdot |\varphi| \\ &= K \int_{g(X)} |\psi| \stackrel{(d)}{\leq} K \int_Y |\psi| = K^2. \end{aligned}$$

Since the two ends of this chain of inequalities agree, each intermediate inequality is in fact an equation. Equality in (b) implies that $|(dg)_{1,1}|$ is equal to a constant almost everywhere on X , and that constant is equal to K by (a). The inequality (c) is based on

$$|(dg)_{1,1}|^2 \leq |(dg)_{1,1}|^2 + |(dg)_{2,1}|^2 = |dg\left(\frac{1}{0}\right)|^2 \leq \sup_{\|\mathbf{v}\|=1} |dg(\mathbf{v})|^2 = \|dg\|^2.$$

Equality implies that $dg\left(\frac{1}{0}\right) = \pm\left(\frac{K}{0}\right)$. Moreover, since dg stretches maximally in the horizontal direction which is preserved, dg must be diagonal, i.e. $dg = \pm\left(\frac{K}{0} \ 0\right)$ with $0 < * \leq K$. Then the equality $K^2 = \|dg\|^2 = K \det(dg)$ determines that $dg = \pm\left(\frac{K}{0} \ 0\right)$ almost everywhere on X .

Since $df = \left(\frac{K}{0} \ 0\right)$ up to sign as well, we have $d(g \circ f^{-1}) = \pm\left(\frac{1}{0} \ 0\right)$ almost everywhere on $f(X)$. The Beltrami coefficient of $g \circ f^{-1}$ is thus equal to 0 almost everywhere on $f(X)$, so that $g \circ f^{-1}$ is conformal and in particular smooth. Since f is smooth except at the zeros of φ , the same holds for g . Therefore the equality $dg = \pm\left(\frac{K}{0} \ 0\right)$ holds everywhere except at the zeros of φ , and g takes the form $x + iy \mapsto \pm(Kx + iy) + c$ in natural coordinates. Equality in (d) means that $g(X)$ has full measure in Y . By Lemma 3.5, g is a Teichmüller embedding with respect to φ and ψ . Finally, the equality $d(g \circ f^{-1}) = \pm\left(\frac{1}{0} \ 0\right)$ holds everywhere, so that $g \circ f^{-1}$ is a local translation. \square

Theorems 3.4 and 3.6 together imply Theorem 1.3 from the introduction.

Remark. Theorem 3.6 does not say that the quadratic differentials φ and ψ are unique up to scale because that is not the case when $K = 1$. A slit mapping may be so with respect to a large-dimensional family of quadratic differentials.

For example, suppose that Y admits an anti-conformal involution σ and let X be the complement in Y of finitely many arcs contained in the fixed locus of σ . Any element of $\mathcal{Q}^+(Y/\sigma)$ can be doubled to a quadratic differential in $\mathcal{Q}^+(Y)$ which is non-negative along the fixed locus of σ . The inclusion map $X \hookrightarrow Y$ is a slit mapping with respect to any quadratic differential obtained in this way.

However, if $K > 1$ then φ and ψ are unique up to a positive scalar, since the Beltrami coefficients $\bar{\partial}f/\partial f$ and $\bar{\partial}(f^{-1})/\partial(f^{-1})$ on X and $f(X)$ encode the directions of maximal stretching, which are the horizontal and vertical directions of φ and ψ respectively.

3.3. Lack of uniqueness. If f and g are homotopic Teichmüller embeddings between finite Riemann surfaces X and Y , then the inclusion map $f(X) \hookrightarrow Y$ and the composition $g \circ f^{-1} : f(X) \rightarrow Y$ are homotopic slit mappings with respect to the same quadratic differential by Theorem 3.6. If X and Y are closed, then $g \circ f^{-1}$ is a conformal automorphism of Y homotopic to the identity. If Y has genus at least 2, then $g = f$ by Lemma 2.1.

However, the map $g \circ f^{-1}$ does not extend to all of Y in general and indeed, Teichmüller embeddings are not necessarily unique in their homotopy class. There are two obvious ways for uniqueness to fail:

- if Y is a torus, then we can post-compose f with any automorphism of Y isotopic to the identity;
- if $f(X)$ is contained in an annulus $A \subset Y$, then we can post-compose f with rotations of A .

In [Iof75], Ioffe claims these are the only exceptions, but this is wrong⁷. The next simplest example is as follows. Let Y be a round annulus in the plane with a concentric circular arc removed, and let X be the same annulus but with a slightly longer arc removed. Then we can obviously rotate X inside of Y by some amount. This gives a 1-parameter family of slit mappings between triply connected domains. If the annulus is centered at the origin, then the quadratic differential in play is $-dz^2/z^2$.

One might think that every counterexample comes from a torus or an annulus with slits, but this is not the case. Here is a general method for constructing examples of slit mappings which are not unique in their homotopy class, with essentially any codomain Y . Let ω be a holomorphic 1-form on Y such that $\omega^2 \in \mathcal{Q}^+(Y)$. Then we can find finitely many horizontal arcs for ω^2 such that their complement $X \subset Y$ is not rigid as follows. For every point $y \in \hat{Y}$ which is either a zero of ω or a puncture of Y , and every trajectory γ of ω ending at y in forward time, remove a neighborhood of y in γ from Y to obtain X . For all small enough $t > 0$, the time- t flow for the vector field $1/\omega$ is well-defined on X , and is a slit mapping homotopic to the inclusion map yet different from it. The slit annulus example is a special case of this.

⁷The source of the mistake is [Iof75, Lemma 3.2]. Similarly, [Iof78], [EM78], and [GG01] contain minor errors as they build up on the false claim.




FIGURE 3. A 1-parameter family of slit mappings between pairs of pants.

In the next section, we show that this kind of phenomenon is the worse that can happen: we can always get from any slit mapping to any homotopic one by flowing horizontally (except in the torus case), although the quadratic differential does not have to be globally the square of an abelian differential in general.

4. SLIT MAPPINGS ARE ALMOST RIGID

We prove that slit mappings are almost rigid in their isotopy class, in the sense that we can get from any one to any other by flowing horizontally (except in the torus case).

Theorem 4.1. *Let $f_0, f_1 : X \rightarrow Y$ be distinct homotopic slit mappings with terminal quadratic differential ψ . There is an isotopy f_t from f_0 to f_1 through slit mappings with terminal quadratic differential ψ such that for every $x \in X$, the path $t \mapsto f_t(x)$ is a regular geodesic for ψ whose length and slope are independent of x . If Y is not a torus, then the paths $t \mapsto f_t(x)$ are horizontal.*

Remark. The statement is still true if we replace each occurrence of “slit mappings” with “Teichmüller embeddings of dilatation K ”, since this is really a statement about $f_1 \circ f_0^{-1}$, which is a slit mapping either way (Theorem 3.6).

Proof. The setup is the same as in Theorem 3.6. We can assume that f_0 and f_1 have only essential punctures, since inessential punctures can be filled and play no special role whatsoever. As explained earlier, the maps f_0 and f_1 extend to continuous maps \widehat{f}_0 and \widehat{f}_1 from \widehat{X} to \widehat{Y} . By Lemma 2.11, there is a homotopy $H : X \times [0, 1] \rightarrow Y$ from f_0 to f_1 which extends to be constant at the punctures. By pushing ∂X inward at the beginning and end of the homotopy, we may further assume that H extends continuously to $\partial X \times [0, 1]$. To summarize, we have a homotopy $\widehat{H} : \widehat{X} \times [0, 1] \rightarrow \widehat{Y}$ from \widehat{f}_0 to \widehat{f}_1 which sends $\widehat{X} \times [0, 1]$ into \widehat{Y} and whose restriction H to $X \times [0, 1]$ has image in Y .

Given $x \in X$, let $\ell(x)$ be the infimum $|\psi|$ -length of piecewise smooth arcs homotopic rel endpoints to the path $t \mapsto H(x, t)$ in Y . This is equal to the length of the “geodesic representative” γ_x of $t \mapsto H(x, t)$, which has image in \widehat{Y} a priori. Note that γ_x has endpoints $f_0(x)$ and $f_1(x)$ by construction. As explained in the proof of Theorem 3.6, the maps $x \mapsto \gamma_x$ and $x \mapsto \ell(x)$ are continuous and extend continuously to \widehat{X} .

Recall that the initial quadratic differential $\varphi = f_0^* \psi = f_1^* \psi$ is the same for both slit mappings and that $f_1 \circ f_0^{-1}$ is a local translation with respect to ψ by Theorem 3.6. In what follows, a *singularity* of ψ is either a zero of ψ or a puncture of Y . The rationale behind this is that even if ψ is regular at some puncture, its pull-back to the universal cover has infinite angle singularities at the corresponding boundary points.

Claim 4.2. *Suppose that ℓ has a local maximum at $x \in \widehat{X}$. Then $x \notin \dot{X}$ and ℓ is constant in a neighborhood of x . If $x \in X$, then γ_x does not contain any singularity of ψ , and φ does not have a zero at x . Moreover, for every y near x there is an isometric immersion from a Euclidean parallelogram P to Y which sends two opposite sides of P to γ_x and γ_y .*

Proof of Claim. Suppose that $x \in \dot{X}$. Then $\ell(x) = 0$ since H is constant on \dot{X} . Since x is a local maximum, $\ell(y) = 0$ for all y in a neighborhood U of x . This means that $\widehat{f}_0 = \widehat{f}_1$ on U and hence on all of \widehat{X} by the identity principle. This contradicts the hypothesis that $f_0 \neq f_1$.

We need to set up some notation. Let $\pi_X : \widetilde{X} \rightarrow X$ and $\pi_Y : \widetilde{Y} \rightarrow Y$ be the universal covers and let \widetilde{X}' and \widetilde{Y}' be their metric completions with respect to $\widetilde{\varphi} = \pi_X^* \varphi$ and $\widetilde{\psi} = \pi_Y^* \psi$. Recall that π_X and π_Y extend to continuous maps $\pi'_X : \widetilde{X}' \rightarrow \widehat{X}$ and $\pi'_Y : \widetilde{Y}' \rightarrow \widehat{Y}$. The homotopy $\widehat{H} : \widehat{X} \times [0, 1] \rightarrow \widehat{Y}$ lifts to a continuous map $\widetilde{H} : \widetilde{X}' \times [0, 1] \rightarrow \widetilde{Y}'$ under these. For any $(u, j) \in \widetilde{X}' \times \{0, 1\}$, let us write $\widetilde{H}_j(u) := \widetilde{H}(u, j)$. Note that \widetilde{H}_j is a lift of \widehat{f}_j . In particular, the restriction of \widetilde{H}_j to \widetilde{X} is a local translation with respect to $\widetilde{\varphi}$ and $\widetilde{\psi}$. For any $u \in \widetilde{X}'$, let $\widetilde{\gamma}_u$ be the geodesic between $\widetilde{H}_0(u)$ and $\widetilde{H}_1(u)$ in \widetilde{Y}' and let $\widetilde{\ell}(u)$ be its length. By definition, $\gamma_{\pi'_X(u)} = \pi'_X(\widetilde{\gamma}_u)$ and $\ell(\pi'_X(u)) = \widetilde{\ell}(u)$. Let $u \in \widetilde{X}'$ be any lift of x under π'_X . Then $\widetilde{\ell}$ has a local maximum at u . Let B be a ball centered at u in \widetilde{X}' such that $\widetilde{\ell}(v) \leq \widetilde{\ell}(u)$ for every $v \in B$. Let $\eta : [-1, 1] \rightarrow B$ be a geodesic such that $\eta(0) = u$. Let η^- and η^+ be its restrictions to $[-1, 0]$ and $[0, 1]$ respectively, and let $v^\pm = \eta(\pm 1)$. Given two geodesic rays α and β sharing an endpoint in \widetilde{X}' or \widetilde{Y}' , we will denote their Alexandrov angle by $\angle(\alpha, \beta)$ instead of the three point notation from Subsection 2.9 (this angle takes values in $[0, \pi]$).

After possibly shrinking B and η , the images $\widetilde{H}_0(\eta^+)$ and $\widetilde{H}_1(\eta^+)$ are geodesic segments of the same length in \widetilde{Y}' (and similarly for η^-). This is because \widetilde{H}_0 and \widetilde{H}_1 are local isometries in \widetilde{X} . Their extensions to $(\pi'_X)^{-1}(\partial X)$ are also local isometries except perhaps at zeros of φ where they can fold the boundary in two. In any case, \widetilde{H}_0 and \widetilde{H}_1 are isometries along any sufficiently short geodesic rays at u .

Here is the main observation. Let $\sigma \in \{+, -\}$. By Lemma 2.12, if

$$\angle(\widetilde{H}_0(\eta^\sigma), \widetilde{\gamma}_u) + \angle(\widetilde{\gamma}_u, \widetilde{H}_1(\eta^\sigma)) \geq \pi$$

then $\widetilde{\ell}(v^\sigma) \geq \widetilde{\ell}(u)$. But $\widetilde{\ell}(v^\sigma) \leq \widetilde{\ell}(u)$ by hypothesis. Since equality holds, the geodesic quadrilateral Q^σ with sides $\widetilde{\gamma}_u$, $\widetilde{\gamma}_{v^\sigma}$, $\widetilde{H}_0(\eta^\sigma)$ and $\widetilde{H}_1(\eta^\sigma)$ is isometric to a Euclidean parallelogram and $\angle(\widetilde{H}_0(\eta^\sigma), \widetilde{\gamma}_u) + \angle(\widetilde{\gamma}_u, \widetilde{H}_1(\eta^\sigma)) = \pi$.

Case 1: Suppose that $x \in X$.

Then after possibly shrinking the ball B centered at $u \in \widetilde{X}$, the restrictions of \widetilde{H}_0 and \widetilde{H}_1 to B are isometries. This means that $\widetilde{H}_0(\eta)$ and $\widetilde{H}_1(\eta)$ are geodesics and hence that

$$\pi = \angle(\widetilde{H}_j(\eta^+), \widetilde{H}_j(\eta^-)) \leq \angle(\widetilde{H}_j(\eta^+), \widetilde{\gamma}_u) + \angle(\widetilde{\gamma}_u, \widetilde{H}_j(\eta^-))$$

for $j = 0$ and $j = 1$ by subadditivity of angles. Thus if

$$\angle(\widetilde{H}_0(\eta^+), \widetilde{\gamma}_u) + \angle(\widetilde{\gamma}_u, \widetilde{H}_1(\eta^+)) < \pi$$

then

$$\angle(\widetilde{H}_0(\eta^-), \widetilde{\gamma}_u) + \angle(\widetilde{\gamma}_u, \widetilde{H}_1(\eta^-)) > \pi.$$

Lemma 2.12 implies that $\tilde{\ell}(v^-) > \tilde{\ell}(u)$, which is a contradiction. We conclude that $\angle(\tilde{H}_0(\eta^+), \tilde{\gamma}_u) + \angle(\tilde{\gamma}_u, \tilde{H}_1(\eta^+)) \geq \pi$. By the main observation, Q^+ is isometric to a Euclidean parallelogram and $\tilde{\ell}(v^+) = \tilde{\ell}(u)$. Since the geodesic η through u was arbitrary, $\tilde{\ell}$ is constant in a neighborhood of u .

Choose η in such a way that each of $H_0(\eta^+)$, $H_0(\eta^-)$, $H_1(\eta^+)$ and $H_1(\eta^-)$ forms a positive angle with $\tilde{\gamma}_u$. Then the corresponding parallelograms Q^+ and Q^- are non-degenerate and lie on opposite sides of $\tilde{\gamma}_u$. In particular, there is no excess angle on either side of $\tilde{\gamma}_u$ in the metric $|\tilde{\psi}|$. In other words, the interior of $\tilde{\gamma}_u$ does not contain any singularity of $\tilde{\psi}$. Suppose however that $\tilde{\psi}$ has a zero at an endpoint of $\tilde{\gamma}_u$. Then $\tilde{\varphi}$ has a zero at u as both \tilde{H}_0 and \tilde{H}_1 are local translations in a neighborhood of u . Since u is in the interior, the cone angle at u in the metric $|\tilde{\varphi}|$ is at least 3π . Thus there is a whole sector of points $v^+ \in B \setminus \{u\}$ such that $\angle(\tilde{H}_0(\eta^+), \tilde{\gamma}_u) = \pi$ where η^+ is the geodesic $[u, v^+]$. For each such v^+ , we then have $\angle(\tilde{H}_1(\eta^+), \tilde{\gamma}_u) = 0$, which means that $\tilde{H}_1(\eta^+)$ and $\tilde{\gamma}_u$ share a segment. In other words, \tilde{H}_1 collapses a whole sector of B into $\tilde{\gamma}_u$. But this is impossible since \tilde{H}_1 is locally injective. This shows that $\tilde{\gamma}_u$ is completely free of singularities. In particular, $\tilde{\varphi} = (\tilde{H}_0)^*\tilde{\psi}$ does not have a zero at u .

Since $\tilde{\gamma}_u$ is regular and has endpoints in \tilde{Y} , it is contained in \tilde{Y} . Similarly, for any short geodesic ray η^+ at u , the resulting parallelogram Q^+ is contained in \tilde{Y} . The covering map $\pi_Y : \tilde{Y} \rightarrow Y$ is an isometric immersion which sends two parallel sides of Q^+ to γ_x and γ_y where $y = \pi_Y(v^+)$.

Case 2: Suppose that $x \in \partial X$ and that φ has a zero at x .

Then the total cone angle at u in the metric $|\tilde{\varphi}|$ is at least 2π . For any geodesic η through u in B , the main observation still holds. Now for every geodesic ray η^+ from u , there exists a ray η^- such that $\tilde{H}_0(\eta)$ and $\tilde{H}_1(\eta)$ are geodesic, where $\eta = \eta^+ \cup \eta^-$. This is because there is enough angle at u to make sure that the cone angle on either side of $\tilde{H}_j(\eta)$ at $\tilde{H}_j(u)$ is at least π .

The argument from Case 1 applies for any such geodesic η : if the angles at the endpoints of $\tilde{\gamma}_u$ in Q^+ sum to less than π , then $\tilde{\ell}(v^-) > \tilde{\ell}(u)$. Hence the angles sum to at least π and Q^+ is a Euclidean parallelogram by the main observation. This shows that $\tilde{\ell}(v^+) = \tilde{\ell}(u)$. This holds for every $v^+ \in B \setminus \{u\}$, so that $\tilde{\ell}$ is constant there.

Case 3: Suppose that $x \in \partial X$ and that φ does not have a zero at x .

Then the cone angle at u in the metric $|\tilde{\varphi}|$ is π and there is only one geodesic η through u , namely the horizontal one running along the boundary. In this case, neither \tilde{H}_0 nor \tilde{H}_1 folds the boundary of \tilde{X}' in two at u . Indeed, if \tilde{H}_j did fold in two at u , then $\tilde{\psi}$ would have a simple pole at $\tilde{H}_j(u)$. But all the simple poles of ψ in \hat{Y} have been unwrapped to infinite angle singularities in \tilde{Y}' .

Therefore $\tilde{H}_0(\eta)$ and $\tilde{H}_1(\eta)$ are geodesic. Thus the same argument as above applies to show that the corresponding quadrilaterals Q^\pm are parallelograms. If one of Q^+ or Q^- is degenerate, then so is the other one.

Suppose first that Q^\pm are degenerate so that the beginning and end of $\tilde{\gamma}_u$ are contained in $\tilde{H}_0(\eta)$ and $\tilde{H}_1(\eta)$ respectively. Moreover if the beginning is contained in $\tilde{H}_0(\eta^+)$ then the end is contained in $\tilde{H}_1(\eta^-)$, and vice versa. It follows that for any $w \in B \setminus \{u\}$ the geodesic $\tau = [u, w]$ satisfies

$$(1) \quad \angle(\tilde{H}_0(\tau), \tilde{\gamma}_u) + \angle(\tilde{\gamma}_u, \tilde{H}_1(\tau)) = \pi.$$

By the main observation, $\tilde{\gamma}_w$ and $\tilde{\gamma}_u$ are the opposite sides of a Euclidean parallelogram and $\tilde{\ell}(w) = \tilde{\ell}(u)$.

Now suppose that Q^+ and Q^- are non-degenerate. Then the interior of $\tilde{\gamma}_u$ is free of singularities, as Q^+ and Q^- lie on opposite sides of it. This implies that equation (1) holds for any geodesic ray τ at u in B , as parallel transport is well-defined along $\tilde{\gamma}_u$. Therefore $\tilde{\ell}$ is constant on B . \square

It follows from the above claim that ℓ is constant. Indeed, the subset of \hat{X} where ℓ attains its maximum is open by the claim and is closed by continuity, hence equal to \hat{X} . In particular, every point of \hat{X} is a local maximum of ℓ , so the additional conclusions of the claim hold everywhere. Specifically,

- (C1) \dot{X} is empty⁸;
- (C2) φ does not have any zero in X ;
- (C3) for every $x \in X$ and y near x , the geodesics γ_x and γ_y form opposite sides of an isometrically immersed parallelogram in Y .

Let $x \in X$. Parametrize $\gamma_x : [0, 1] \rightarrow Y$ proportionally to arc length in such a way that $\gamma_x(0) = f_0(x)$ and $\gamma_x(1) = f_1(x)$. For every $(x, t) \in X \times [0, 1]$, define $f_t(x) = \gamma_x(t)$. Then f_t has image in Y and depends continuously on t . By (C3), $f_t \circ f_0^{-1}$ is a local translation with respect to ψ for every $t \in [0, 1]$. Thus f_t is a local translation with respect to φ and ψ . It remains to prove that f_t is injective and that its image is the complement of finitely many horizontal arcs and points.

Suppose that Y is a torus. Then $\psi = \omega^2$ for an abelian differential ω . For any $x \in X$, the tangent vector $v(x, t) = \omega(\gamma'_x(t)) \in \mathbb{C}$ does not depend on $t \in [0, 1]$ since γ_x is a smooth geodesic. As nearby geodesics γ_x and γ_y are parallel and of the same length, v is locally constant as a function of $x \in X$, hence constant on X . This means that f_t differs from f_0 by the translation $z \mapsto z + tv$ with respect to ω , and is therefore a slit mapping.

Suppose that Y is an annulus. Then ψ is again the square of an abelian differential ω . By the same argument as in the previous paragraph, $v = \omega(\gamma'_x(t))$ is independent of $(x, t) \in X \times [0, 1]$. If the geodesics γ_x are not horizontal, then they point away from Y on one of the two boundary components, which is absurd. Hence $v \in \mathbb{R}$ and f_t differs from f_0 by the horizontal translation $z \mapsto z + tv$ with respect to ω (i.e., by a rotation of Y).

Our proof that f_t is a slit mapping in the general case is rather indirect. By Lemma 3.5, it suffices to show that f_t is injective and that $f_t(X)$ is dense in Y . To prove this, we first show that the geodesics γ_x are horizontal.

Claim 4.3. *If Y is not a torus, then γ_x is horizontal for every $x \in \hat{X}$.*

Proof. By the same argument as in the torus and annulus case, the unoriented slope $\psi(\gamma'_x(t))$ is constant on $X \times [0, 1]$. Since the map $x \mapsto \gamma_x$ is continuous on \hat{X} and since the limit of a sequence of geodesics of slope s has slope s wherever it is smooth, the slope function is constant on all \hat{X} . Thus if we find a single horizontal geodesic, then all geodesics are horizontal.

If Y is an annulus, then we have already shown that γ_x is horizontal for every $x \in \hat{X}$. We can leverage the same idea to prove that if Y contains a horizontal cylinder, then γ_x is horizontal for some (hence all) $x \in \hat{X}$.

⁸Recall that punctures are assumed to be essential in the proof. This proves that if two slit mappings are homotopic and have an essential puncture, then they are equal.

Let $A \subset Y$ be a maximal horizontal cylinder for ψ . Let α be a closed horizontal trajectory in $A \cap f_0(X)$. Then $f_t \circ f_0^{-1}(\alpha)$ is a closed horizontal trajectory homotopic to α —hence is contained in A —for every $t \in [0, 1]$. Since all but finitely many horizontal trajectories in A are contained in $f_0(X)$, we have that $f_t \circ f_0^{-1}(A \cap f_0(X)) \subset A$ for every $t \in [0, 1]$. In particular, the geodesic γ_x is contained in A for every $x \in f_0^{-1}(A)$.

Observe that $f_0(X) \cap A$ has either one or two connected components. If $f_0(X) \cap A$ is connected, then the same argument as in the annulus case applies: if the geodesics γ_x are not horizontal, then they point away from A on one of its two boundary components, contradicting the fact that they are contained in A . Suppose that $f_0(X) \cap A$ has two components. Since X is connected, $f_0(X)$ intersects each boundary component of A . If γ_x points strictly inward at such a boundary point $f_0(x)$, then for small $t > 0$ the map $f_t \circ f_0^{-1}$ drags points from the complement of A into A , which is impossible. Said differently, if $t > 0$ and $\alpha \subset A \cap f_t(X)$ is a closed horizontal trajectory sufficiently close to $f_0(x)$, then $f_0 \circ f_t^{-1}(\alpha)$ is a closed horizontal trajectory outside of A yet homotopic to α , contradicting the maximality of A . In either case, the geodesics γ_x have to be horizontal.

If Y has a horizontal cylinder then we are done, so suppose that it does not have any. Let $t \in [0, 1]$. Since $f_t : X \rightarrow Y$ is 1-Lipschitz with respect to $|\varphi|$ and $|\psi|$, it extends to a 1-Lipschitz map $\hat{f}_t : \hat{X} \rightarrow \hat{Y}$ between the completions. Let C be a connected component of ∂X , let N be neighborhood of C in \hat{X} such that $N \cap X$ is an annulus, and let $\alpha \subset N \cap X$ be homotopic to C in N . If $f_0(\alpha)$ is not homotopic to a point or a puncture in Y , then $\hat{f}_t(C)$ becomes the “geodesic representative” of $f_0(\alpha)$ after erasing its backtracks (see Figure 4). Since Y contains no horizontal cylinders, the “geodesic representative” of $f_0(\alpha)$ is unique. This means that $\hat{f}_t(C)$ is contained in the same horizontal leaf for every $t \in [0, 1]$, hence that γ_x is horizontal for every $x \in C$.




FIGURE 4. Pulling a horizontal curve tight shrinks it to a point, a puncture, or a geodesic.

We may thus assume that $f_0(\alpha)$ is homotopic to a point or a puncture in Y for every C, N and α as above. This means that $\hat{f}_t(C)$ is completely folded onto itself, as it reduces to a point after erasing its backtracks. Thus $\hat{f}_t(N)$ forms an open neighborhood of $\hat{f}_t(C)$ in \hat{Y} . This holds for every component C of ∂X and every $t \in [0, 1]$. Moreover, ∂Y is empty for otherwise there is some component C of ∂X such that $\hat{f}_0(C)$ intersects ∂Y , forcing $f_0(\alpha)$ to be essential.

Since Y is not a torus nor an annulus, ψ has at least one singularity, say at $y \in \hat{Y}$. As $\hat{f}_0 : \hat{X} \rightarrow \hat{Y}$ is surjective, there is some $x \in \hat{X}$ such that $\hat{f}_0(x) = y$.

We have that $x \in \partial X$ since \dot{X} is empty (C1) and φ does not have any zeros in X (C2). Let C be the component of ∂X containing x . If γ_x is not horizontal, then $\widehat{f}_t(C)$ moves off of its horizontal leaf for small enough $t > 0$. On the other hand, y is contained in the open set $\widehat{f}_t(N)$ for all sufficiently small $t > 0$. But then $y \in f_t(N \setminus C) \subset f_t(X) \subset Y$ so that ψ has a zero at y and hence φ has zeros on $f_t^{-1}(y)$, contradicting (C2). Therefore γ_x is horizontal. \square

Since f_t is a local translation with respect to φ and ψ , the area of its image is at most the area of X , with equality only if f_t is injective. Recall that the set $U \subset X$ of points through which passes a bi-infinite regular horizontal trajectory of φ has full measure in X . For every bi-infinite regular horizontal trajectory $\eta \subset U$, its image $f_0(\eta)$ is a bi-infinite regular horizontal trajectory of ψ . For any $x \in \eta$, the horizontal geodesic γ_x is thus contained in $f_0(\eta)$. Hence $f_t \circ f_0^{-1}$ acts as a translation along $f_0(\eta)$ for every $t \in [0, 1]$, so that $f_t(\eta) = f_0(\eta)$. It follows that $f_t(U) = f_0(U)$. But $f_0(U)$ has full measure in Y since U has full measure in X and $f_0(X)$ has full measure in Y . Therefore $f_t(X)$ has full measure in Y , whose area is equal to that of X . Hence f_t is injective.

It follows that f_t is a slit mapping by Lemma 3.5. Thus $(x, t) \mapsto f_t(x)$ is an isotopy from f_0 to f_1 through slit mappings. If Y is not a torus, then the path $t \mapsto f_t(x) = \gamma_x(t)$ is horizontal for every $x \in X$ by the previous claim. \square

Remark. As observed in footnote 8, if $f_0, f_1 : X \rightarrow Y$ are homotopic slit mappings with an essential puncture, then $f_0 = f_1$. Similarly, if their initial quadratic differential has a zero in X then $f_0 = f_1$. As another example, if their terminal quadratic differential has a horizontal saddle connection whose interior is contained in $f_j(X)$ for $j = 0$ or $j = 1$ then $f_0 = f_1$. This is because the saddle connection has nowhere to go under the horizontal isotopy. The same statements apply to Teichmüller embeddings.

Remark. If Y is not the torus and $f_0, f_1 : X \rightarrow Y$ are distinct homotopic slit mappings, then their terminal quadratic differential ψ is unique up to a positive scalar (cf. Remark 3.2). Indeed, for every $x \in X$ and every $\psi \in \mathcal{Q}^+(Y)$ with respect to which f_0 and f_1 are slit mappings, $f_0(x)$ and $f_1(x)$ are on the same horizontal leaf of ψ . Thus if f_t is *any* isotopy between f_0 and f_1 through slit mappings, then $t \mapsto f_t(x)$ is horizontal for every $x \in X$ and every $\psi \in \mathcal{Q}^+(Y)$ with respect to which f_0 and f_1 are slit mappings. But if two quadratic differentials have the same horizontal trajectories, then they are positive multiples of each other. The contrapositive is worth mentioning explicitly: if the terminal quadratic differential of a slit mapping f is not unique up to scale, then f is the unique conformal embedding in its homotopy class. For example, if $\dim \mathcal{Q}^+(Y) > 1$ and $X \subset Y$ has finite complement, then the inclusion map $X \hookrightarrow Y$ is the only conformal embedding in its homotopy class (cf. Lemma 2.2).

Let $h : X \rightarrow Y$ be a topological embedding between finite Riemann surfaces. Recall that the space $\text{CEmb}(X, Y, h)$ is the set of all conformal embeddings homotopic to h equipped with the compact-open topology. We now deduce Theorem 1.4 from the introduction, which describes the homeomorphism type of $\text{CEmb}(X, Y, h)$ whenever the latter contains a slit mapping.

Proof of Theorem 1.4. If $\text{CEmb}(X, Y, h)$ contains a slit mapping with terminal differential ψ , then every $f \in \text{CEmb}(X, Y, h)$ is a slit mapping with respect to ψ by Theorem 3.6. Let $x \in X$. The evaluation map $\text{CEmb}(X, Y, h) \rightarrow Y$ sending f to $f(x)$ is continuous by definition of the compact-open topology. We now show that it is injective.

Suppose that $f_0(x) = f_1(x)$ for distinct $f_0, f_1 \in \text{CEmb}(X, Y, h)$ and some $x \in X$. Our goal is to show that $f_0 = f_1$. If not, let f_t be the isotopy from f_0 to f_1 provided by Theorem 4.1. Then $\gamma : S^1 \rightarrow Y$ defined by $\gamma(t) = f_t(x)$ is a closed regular geodesic for ψ .

Suppose that Y is a torus. For each $t \in [0, 1]$, there is a unique translation $g_t \in \text{Aut}_0(Y)$ such that $g_t(\gamma(t)) = \gamma(0)$. Then $t \mapsto g_t \circ f_t$ is an isotopy from f_0 to f_1 which is constant at x . Thus the restrictions $f_0^*, f_1^* : X \setminus \{x\} \rightarrow Y \setminus \{f_0(x)\}$ are homotopic slit mappings with an essential puncture. It follows that $f_0 = f_1$ by Remark 4. Alternatively, the proof of Theorem 4.1 show that f_0 and f_1 differ by an element of $\text{Aut}_0(Y)$, which has to be the identity since it fixes the point $f_0(x) = f_1(x)$.

Suppose that Y is not a torus. Then γ is horizontal. Let $A \subset Y$ be a cylinder foliated by closed horizontal trajectories homotopic to γ which is symmetric about γ . Since $f_0(x)$ is a fixed point of the local translation $f_1 \circ f_0^{-1}$, we have that $f_1 \circ f_0^{-1}$ is either equal to the identity near $f_0(x)$ or a half-turn around $f_0(x)$. If $f_1 \circ f_0^{-1}$ is equal to the identity near $f_0(x)$, then $f_0 = f_1$ by the identity principle. Otherwise, let σ be the conformal involution of A which fixes $f_0(x)$ and permutes its two boundary components. By the identity principle $f_1 \circ f_0^{-1} = \sigma$ on the connected component U of $A \cap f_0(X)$ containing $f_0(x)$. Let $\alpha \subset U$ be a closed horizontal trajectory and let $\beta = f_0^{-1}(\alpha)$. Then $f_0(\beta) = \alpha$ and $f_1(\beta) = \sigma(\alpha)$ are not homotopic in Y , for they have reverse orientations. This is a contradiction, which proves that the first situation prevails.

Thus the evaluation map is injective. Since $\text{CEmb}(X, Y, h)$ contains a slit mapping, h is non-trivial and non-parabolic by Remark 3.2. It follows that $\text{CEmb}(X, Y, h)$ is compact by Lemma 3.1. Any injective continuous map from a compact space to a Hausdorff space is a homeomorphism onto its image. Let $V(x)$ be the image of evaluation map at $x \in X$. If Y is a torus, then $V(x) = Y$ since $\text{Aut}_0(Y)$ acts on $\text{CEmb}(X, Y, h)$ by composition on the left and acts transitively on Y .

Suppose that Y is not a torus. Then $V(x)$ is path-connected and is contained in a regular horizontal trajectory by Theorem 4.1. As $V(x)$ is compact, it is either a closed trajectory, a geodesic segment, or a point. If h is cyclic, then the image of any $f \in \text{CEmb}(X, Y, h)$ is contained in some annulus $A \subset Y$. Since f can be post-composed with rotations of A , the set $V(x)$ contains a circle, hence is equal to a circle.

Conversely, if $V(x)$ is a horizontal circle then its free homotopy class in Y does not depend on $x \in X$ since X is path-connected and $V(x)$ depends continuously on x . Therefore $\bigcup_{x \in X} V(x)$ is contained in a horizontal cylinder $A \subset Y$, which means that $f(X) \subset A$ for every $f \in \text{CEmb}(X, Y, h)$ so that h is cyclic. This shows that if h is generic (so that Y is not a torus), then $V(x)$ is either a segment or a point. \square

As pointed out in the introduction, if $\text{CEmb}(X, Y, h)$ contains a slit mapping then the conclusions of Theorem 1.2 follow immediately from Theorem 1.4. In that case h is necessarily non-trivial. If Y is a torus then $\text{CEmb}(X, Y, h)$ is homeomorphic to a torus, hence homotopy equivalent to one. If Y is not a torus and h is cyclic, then

$\text{CEmb}(X, Y, h)$ is homeomorphic to a circle. If h is generic, then $\text{CEmb}(X, Y, h)$ is homeomorphic to a point or a compact interval, hence contractible.

5. THE MODULUS OF EXTENSION

In this section, we characterize local maxima of the modulus of extension (Theorem 1.5) and use that to prove that the space of all conformal embeddings in a given homotopy class is connected under some conditions (Theorem 1.6).

Let $h : X \rightarrow Y$ be a topological embedding between finite Riemann surfaces where X has non-empty ideal boundary. We will impose further conditions on h in a moment, but for now the only hypothesis is that $\partial X \neq \emptyset$.

For each connected component C of ∂X , choose an analytic parametrization $\zeta_C : S^1 \rightarrow C$. For every $r \in (0, \infty]$ and every component C of ∂X , glue a copy of the cylinder $S^1 \times [0, r)$ to $X \cup \partial X$ along $S^1 \times \{0\}$ using the map ζ_C (see Figure 5). We denote the resulting surface by X_r and also let $X_0 = X$. If $\rho \leq r$, then the inclusion $[0, \rho) \subset [0, r)$ induces a conformal embedding $X_\rho \subset X_r$. We call the directed family of Riemann surfaces $\{X_r\}$ a *1-parameter family of enlargements of X* . Note that for every $r \in [0, \infty]$, there is a homeomorphism $H_r : X_r \rightarrow X$ which when followed by the inclusion $X \subset X_r$ is homotopic to the identity. We will abuse notation and write $\text{CEmb}(X_r, Y, h)$ instead of $\text{CEmb}(X_r, Y, h \circ H_r)$.




FIGURE 5. The surface X_r is obtained by gluing a cylinder of modulus r to each ideal boundary component of X .

Given a conformal embedding $f \in \text{CEmb}(X, Y, h)$, we define the *modulus of extension of f* as

$$\mathbf{m}(f) = \sup\{r \in [0, \infty] \mid f \text{ extends to a conformal embedding } X_r \hookrightarrow Y\}.$$

This number depends on the 1-parameter family of enlargements $\{X_r\}_{r \in [0, \infty]}$ which is fixed once and for all. It is easy to see that the supremum is realized and that the conformal extension of f to $X_{\mathbf{m}(f)}$ is unique.

Lemma 5.1. *For every $f \in \text{CEmb}(X, Y, h)$, there is a unique conformal embedding $\widehat{f} : X_{\mathbf{m}(f)} \rightarrow Y$ extending f .*

Proof. Let $r = \mathbf{m}(f)$, let $r_n \nearrow r$, and let $f_n : X_{r_n} \rightarrow Y$ be a conformal embedding extending f . Every $x \in X_r$ is contained in X_{r_n} when n is large enough since $X_{r_n} \nearrow X_r$. Define $\widehat{f}(x) = f_n(x)$. This does not depend on n since $f_j = f_{j+1}$ on X_{r_j} by the identity principle. The function \widehat{f} is holomorphic and injective on X_r because each f_n is holomorphic and injective. The uniqueness of \widehat{f} follows from the identity principle. \square

Similarly, \mathbf{m} is upper semi-continuous on $\text{CEmb}(X, Y, h)$. Again, this holds without any extra assumptions on h .

Lemma 5.2. *The modulus of extension \mathbf{m} is upper semi-continuous.*

Proof. Suppose that $f_n \rightarrow f$ in $\text{CEmb}(X, Y, h)$. Pass to a subsequence so that $r_n = \mathbf{m}(f_n)$ converges to some $r \in [0, \infty]$. We have to show that $\mathbf{m}(f) \geq r$. If $r = 0$, then there is nothing to prove so we assume that $r > 0$.

If Y is not hyperbolic, then its group $\text{Aut}_0(Y)$ of conformal automorphisms homotopic to the identity acts simply transitively on k -tuples of distinct points for some $k \in \{1, 2, 3\}$. Pick some k -tuple $P \subset X$ of distinct points. Then there exists a sequence of automorphisms $T_n \in \text{Aut}_0(Y)$ converging to the identity such that $T_n \circ f_n(P) = f(P)$ for every n . Observe that $\mathbf{m}(f_n) = \mathbf{m}(T_n \circ f_n)$, so we might as well assume from the start that $f_n(P) = f(P)$ for every n . Under this assumption, we can consider f_n and f as maps from $X \setminus P$ to $Y \setminus f(P)$. Now $Y \setminus f(P)$ is hyperbolic, so the problem is reduced to that case.

Assume that Y is hyperbolic. Let $\rho \in (0, r)$ and let \widehat{f}_n be the conformal extension of f_n to X_{r_n} . If n is large enough, then $\rho \leq r_n$, and we let g_n be the restriction of \widehat{f}_n to X_ρ . Since g_n cannot diverge locally uniformly (its restriction to X converges to f), we may pass to a subsequence such that $g_n \rightarrow g$ for some $g \in \text{CEmb}(X_\rho, Y, h)$ by Montel's theorem. The restriction of g to X is necessarily equal to f . In other words, g is a conformal embedding extending f , so that $\mathbf{m}(f) \geq \rho$. Since $\rho \in (0, r)$ was arbitrary we have $\mathbf{m}(f) \geq r$. \square

Since \mathbf{m} is upper semi-continuous, it attains its maximum whenever the space $\text{CEmb}(X, Y, h)$ is compact. This happens precisely when h is non-trivial and non-parabolic by Lemma 3.1. Let us recall hypothesis **(H)** from the introduction:

- $h : X \rightarrow Y$ is a non-trivial and non-parabolic embedding between finite Riemann surfaces;
- X has non-empty ideal boundary;
- $\{X_r\}_{r \in [0, \infty]}$ is a 1-parameter family of enlargements of X ;
- \mathbf{m} is the associated modulus of extension;
- $\text{CEmb}(X, Y, h)$ is non-empty.

Note that this hypothesis prevents Y from being the sphere with at most 2 punctures or the disk with at most 1 puncture. Under hypothesis **(H)**, \mathbf{m} achieves its maximum at some $f \in \text{CEmb}(X, Y, h)$. We will see that any such maximum is the restriction of a slit mapping. We first need to show that any limit of Teichmüller embeddings is itself a Teichmüller embedding.

Definition 5.3. Let $X_n \in \mathcal{T}(X)$ and $Y_n \in \mathcal{T}(Y)$ be such that $X_n \rightarrow X$ and $Y_n \rightarrow Y$ as $n \rightarrow \infty$, and let $\sigma_n : X_n \rightarrow X$ and $\tau_n : Y_n \rightarrow Y$ be the Teichmüller homeomorphisms that respect the markings. Let $h_n : X_n \rightarrow Y_n$ and $h : X \rightarrow Y$ be any maps. We say that $h_n \rightarrow h$ as $n \rightarrow \infty$ if $\tau_n \circ h_n \circ \sigma_n^{-1} \rightarrow h$ locally uniformly on X as $n \rightarrow \infty$.

Lemma 5.4. *Let $h : X \rightarrow Y$ be a non-trivial and non-parabolic embedding between finite Riemann surfaces. Suppose that $X_n \rightarrow X$ in $\mathcal{T}(X)$ and $Y_n \rightarrow Y$ in $\mathcal{T}(Y)$, and let $f_n : X_n \rightarrow Y_n$ be a sequence of Teichmüller embeddings homotopic to h . Then there exists a Teichmüller embedding $f : X \rightarrow Y$ homotopic to h such that $f_n \rightarrow f$ after passing to a subsequence.*

Note that passing to a subsequence might be necessary due to the non-uniqueness of Teichmüller embeddings.

Proof. By Lemma 3.1, we may pass to a subsequence such that $f_n \rightarrow f$ for some quasiconformal embedding $f : X \rightarrow Y$ homotopic to h . By Theorem 3.6, each f_n has minimal dilatation in its homotopy class. It follows that f has minimal dilatation in its homotopy class. If f is not conformal, then it is a Teichmüller embedding by Theorem 3.4, and we are done.

Suppose that f is conformal. We need to show that f is a slit mapping. Let ψ_n be the terminal quadratic differential of f_n , and let $g_n : Y_n \rightarrow Y'_n$ be a Teichmüller homeomorphism of dilatation e^2 and initial quadratic differential ψ_n . Since $d(Y_n, Y'_n) = 1$ for every n , $Y_n \rightarrow Y$ as $n \rightarrow \infty$, and closed balls in $\mathcal{T}(Y)$ are compact, we may pass to a subsequence such that $Y'_n \rightarrow Y'$ as $n \rightarrow \infty$, for some $Y' \in \mathcal{T}(Y)$ with $d(Y, Y') = 1$. Let $g : Y \rightarrow Y'$ be a Teichmüller homeomorphism that respects the markings. By a standard argument similar to the one above, $g_n \rightarrow g$ after passing to yet another subsequence (this is only necessary if Y is an annulus or a torus).

By construction, $g_n \circ f_n : X_n \rightarrow Y'_n$ is a Teichmüller embedding. Moreover, $g_n \circ f_n \rightarrow g \circ f$. As in the first paragraph of this proof, $g \circ f$ has minimal dilatation in its homotopy class. This dilatation is equal to $e^2 > 1$, so that $g \circ f$ is a Teichmüller embedding. Since f is conformal, the terminal quadratic differential of $g \circ f$ is equal to the terminal quadratic differential of g . It follows that $f = g^{-1} \circ (g \circ f)$ is a slit mapping with respect to the initial quadratic differential of g . \square

We obtain the following characterization of the global maxima of \mathfrak{m} .

Proposition 5.5. *Assume hypothesis (H). Let f be a global maximum of \mathfrak{m} with $\mathfrak{m}(f) < \infty$ and let \widehat{f} be the conformal extension of f to $X_{\mathfrak{m}(f)}$. Then \widehat{f} is a slit mapping.*

Proof. Let $R = \mathfrak{m}(f)$. For every $r > R$, there is no conformal embedding $g : X_r \rightarrow Y$ whose restriction to X is homotopic to h , for otherwise we would have $\mathfrak{m}(g|_X) \geq r > R = \mathfrak{m}(f)$. By Corollary 3.2 and Theorem 3.4, there exists a Teichmüller embedding $g_r : X_r \rightarrow Y$ whose restriction to X is homotopic to h . It is easy to see that X_r moves continuously in $\mathcal{T}(X)$ as a function of $r \in [0, \infty)$. By Lemma 5.4, we can extract a limiting Teichmüller embedding $g : X_R \rightarrow Y$ from some subsequence of g_r as $r \rightarrow R$. Since \widehat{f} is conformal and homotopic to g , Theorem 3.6 implies that \widehat{f} is itself a Teichmüller embedding. A conformal Teichmüller embedding is a slit mapping. \square

Remark. Observe that every end of the surface X_∞ is a puncture since the cylinder $S^1 \times [0, \infty)$ is conformally isomorphic to $\mathbb{D} \setminus \{0\}$. Thus if $\mathfrak{m}(f) = \infty$, then the extension \widehat{f} extends further to a conformal homeomorphism between \widehat{X}_∞ and \widehat{Y} . In particular, $Y \setminus \widehat{f}(X_\infty)$ is finite, so that \widehat{f} is a slit mapping with respect to any $\psi \in \mathcal{Q}^+(Y)$. This is unless Y is the triply-punctured sphere in which case $\mathcal{Q}^+(Y)$ is empty. Thus the hypothesis $\mathfrak{m}(f) < \infty$ in the above theorem is superfluous unless Y is the triply-punctured sphere. In the latter case \widehat{f} is nevertheless unique in its homotopy class by Lemma 2.2.

We can in fact strengthen Proposition 5.5 by replacing the word “global” with “local”. The proof requires another lemma similar to Lemma 5.4.

Lemma 5.6. *Let $h : X \rightarrow Y$ be a non-trivial and non-parabolic embedding between finite Riemann surfaces. Suppose that $Y_n \nearrow Y$, where the inclusion $\iota_n : Y_n \hookrightarrow Y$ is homotopic to a homeomorphism, and let $f_n : X \rightarrow Y_n$ be a sequence of Teichmüller embeddings such that $\iota_n \circ f_n$ is homotopic to h . Then there exists a Teichmüller embedding $f : X \rightarrow Y$ homotopic to h such that $f_n \rightarrow f$ after passing to a subsequence.*

The difference with Lemma 5.4 is that Y_n is allowed to live in a different Teichmüller space than Y . For example, a surface with punctures can be exhausted by surfaces with holes. A priori, the terminal quadratic differentials could converge to a quadratic differential with double poles at the punctures of Y and then the limiting map would not be a Teichmüller embedding. What prevents this from happening is that the domain X is fixed.

Proof. Let K_n be the dilatation of f_n . The sequence K_n is non-increasing and thus converges to a limit $K \geq 1$. After passing to a subsequence, f_n converges to a K -quasiconformal embedding $f : X \rightarrow Y$.

We claim that f has minimal dilatation in its homotopy class. If not, let $g : X \rightarrow Y$ be a quasiconformal embedding homotopic to f such that $\text{Dil}(g) < \text{Dil}(f)$. If $r > 0$ is small enough, then there is a quasiconformal homeomorphism $\sigma : X_r \rightarrow X$ of dilatation strictly smaller than $\text{Dil}(f)/\text{Dil}(g)$. Thus the quasiconformal embedding F consisting of the inclusion $X \hookrightarrow X_r$ followed by $g \circ \sigma$ has dilatation strictly less than $\text{Dil}(f)$. Then $F(X) \subset Y_n$ whenever n is large enough. Indeed, the ends of $F(X)$ which are not punctures are contained in a compact subset of Y since they are surrounded by the union of collars $g \circ \sigma(X_r \setminus X)$. Thus the only way to go to infinity in the closure $\overline{F(X)}$ relative to Y is via punctures of $F(X)$ that are also punctures of Y . For every such puncture p and every $n \in \mathbb{N}$, a neighborhood of p in Y is contained in Y_n . This is because f_n is quasiconformal and as such it cannot map punctures of X to ends of Y_n with finite modulus. Since $Y_n \nearrow Y$, the inclusion $F(X) \subset Y_n$ holds for n large enough. But the inequality $\text{Dil}(F) < \text{Dil}(f) \leq \text{Dil}(f_n)$ contradicts the hypothesis that f_n is a Teichmüller embedding and hence has minimal dilatation in its homotopy class.

If f is not conformal, then it is a Teichmüller embedding by Theorem 3.4, and we are done. Thus suppose that f is conformal but is not a slit mapping. By Proposition 5.5, there exists an $r > 0$ and a conformal embedding $g : X_r \rightarrow Y$ whose restriction to X is homotopic to f . By the above argument, we have $g(X_{r/2}) \subset Y_n$ if n is large enough, and thus $g(X) \subset Y_n$ with complement having non-empty interior. On the other hand, the restriction $g|_X : X \rightarrow Y_n$ is conformal and homotopic to the Teichmüller embedding $f_n : X \rightarrow Y_n$. By Theorem 3.6, $g|_X$ is a slit mapping so that $Y_n \setminus g(X)$ has empty interior, contradiction. \square

We come to the main result of this section, which is that every local maximum of \mathfrak{m} is the restriction of a slit mapping.

Theorem 5.7. *Assume hypothesis (H). Let f be a local maximum of \mathfrak{m} with $\mathfrak{m}(f) < \infty$, and let \widehat{f} be the conformal extension of f to $X_{\mathfrak{m}(f)}$. Then \widehat{f} is a slit mapping.*

Proof. Let $R = \mathfrak{m}(f)$. We first show that the complement $Y \setminus \widehat{f}(X_R)$ is horizontal for some meromorphic quadratic differential on Y , and is in particular an analytic graph. Let $\{x_1, x_2, \dots\}$ be a dense subset of X .

Claim 5.8. *There exists an $n \in \mathbb{N}$ such that if $g \in \text{CEmb}(X, Y, h)$ satisfies $g(x_j) = f(x_j)$ for every $j \in \{1, \dots, n\}$, then $\mathbf{m}(g) \leq \mathbf{m}(f)$.*

Proof of Claim. Suppose on the contrary that for every $n \in \mathbb{N}$ there exists an element g_n of $\text{CEmb}(X, Y, h)$ satisfying $g_n(x_j) = f(x_j)$ for every $j \in \{1, \dots, n\}$ such that $\mathbf{m}(g_n) > \mathbf{m}(f)$. As $\text{CEmb}(X, Y, h)$ is compact, every subsequence of $\{g_n\}_{n=1}^\infty$ has a subsequence converging to some $g \in \text{CEmb}(X, Y, h)$. Any limit g agrees with f on the dense set $\{x_1, x_2, \dots\}$, and hence is equal to f . Thus $g_n \rightarrow f$ with $\mathbf{m}(g_n) > \mathbf{m}(f)$. This contradicts the hypothesis that f is a local maximum of \mathbf{m} . \square

Let n be as in the claim, and let $P = \{x_1, x_2, \dots, x_n\}$. Then for every $r > R$, there is no conformal embedding $g : X_r \rightarrow Y$ homotopic to $f \text{ rel } P$. By Corollary 3.2 and Theorem 3.4, there exists a Teichmüller embedding $g_r : X_r \setminus P \rightarrow Y \setminus f(P)$ homotopic to $f \text{ rel } P$.

Let g be any limit of any subsequence of g_r as $r \searrow R$. Then $g : X_R \setminus P \rightarrow Y \setminus f(P)$ is a Teichmüller embedding by Lemma 5.4. Since \hat{f} is conformal and homotopic to $g \text{ rel } P$, Theorem 3.6 implies that \hat{f} is itself a slit mapping, considered as a map from $X_R \setminus P$ to $Y \setminus f(P)$. Therefore the complement $Y \setminus \hat{f}(X_R)$ is a finite union of horizontal arcs for some meromorphic quadratic differential on Y , possibly with simple poles on the set $f(P)$.

Let $\Gamma = Y \setminus \hat{f}(X_R)$, let $\{y_1, y_2, \dots\}$ be a dense subset of the graph Γ minus its vertices, and fix a Riemannian metric on Y .

Claim 5.9. *There exists a $k \in \mathbb{N}$ such that for every $r > R$ and every $\varepsilon > 0$, there is no conformal embedding $g : X_r \rightarrow Y$ homotopic to h whose image is disjoint from the balls $\overline{B}(y_1, \varepsilon), \dots, \overline{B}(y_k, \varepsilon)$.*

Proof of Claim. Suppose that for every $k \in \mathbb{N}$ there exist an $r_k > R$, an $\varepsilon_k > 0$, and a conformal embedding $g_k : X_{r_k} \rightarrow Y$ whose restriction to X is homotopic to h such that $g_k(X_{r_k})$ is disjoint from the balls $\overline{B}(y_1, \varepsilon_k), \dots, \overline{B}(y_k, \varepsilon_k)$. We may assume that $r_k \rightarrow R$ and $\varepsilon_k \rightarrow 0$. Let g be any limit of any subsequence of the sequence $\{g_k\}$. Then $g(X_R)$ is disjoint from the set $\{y_1, y_2, \dots\}$ and hence from its closure Γ , so that $\hat{f}^{-1} \circ g : X_R \rightarrow X_R$ is a conformal embedding homotopic to the identity. If X_R is not an annulus, then Lemma 2.1 implies that $g = \hat{f}$ and hence $g_k \rightarrow \hat{f}$. If X_R is an annulus, then we may pre-compose each g_k by a rotation so that we still get $g_k \rightarrow \hat{f}$. Since $\mathbf{m}(g_k|_X) \geq r_k > R = \mathbf{m}(f)$, this contradicts the hypothesis that f is a local maximum of \mathbf{m} . \square

Let k be as in the last claim, and let $Q = \{y_1, \dots, y_k\}$. For each $\varepsilon > 0$, let $Y_\varepsilon = Y \setminus (\overline{B}(y_1, \varepsilon) \cup \dots \cup \overline{B}(y_k, \varepsilon))$. Let $\iota_\varepsilon : Y_\varepsilon \rightarrow Y \setminus Q$ be a homeomorphism homotopic to the inclusion map, and let $h_\varepsilon = \iota_\varepsilon^{-1} \circ \hat{f}$. The embedding $h_\varepsilon : X_R \rightarrow Y_\varepsilon$ followed by the inclusion $Y_\varepsilon \hookrightarrow Y$ is homotopic to h . By the claim, for every $r > R$, there is no conformal embedding $g : X_r \rightarrow Y_\varepsilon$ homotopic to h_ε . Therefore, there is a Teichmüller embedding $g'_\varepsilon : X_r \rightarrow Y_\varepsilon$ homotopic to h_ε . Letting $r \searrow R$, we can extract a limiting Teichmüller embedding $g_\varepsilon : X_R \rightarrow Y_\varepsilon$ by Lemma 5.4.

Since $Y_\varepsilon \nearrow (Y \setminus Q)$ as $\varepsilon \searrow 0$, we can apply Lemma 5.6 and obtain a Teichmüller embedding $g : X_R \rightarrow Y \setminus Q$ as a limit of a subsequence of $\{g_\varepsilon\}$. Since $\hat{f} : X_R \rightarrow Y \setminus Q$ is homotopic to g , it is a slit mapping with respect to some $\psi \in \mathcal{Q}^+(Y \setminus Q)$. Thus ψ is meromorphic on Y with at most simple poles on Q . Moreover, the graph $\Gamma = Y \setminus \hat{f}(X_R)$ is horizontal for ψ . Since every point of Q is contained in the

interior of an edge of Γ , the quadratic differential ψ cannot have simple poles on Q . Indeed, there is only one horizontal ray emanating from any simple pole. Therefore, ψ is holomorphic on Y and the map $\hat{f}: X_R \rightarrow Y$ is an honest slit mapping. \square

Once again, the hypothesis $\mathbf{m}(f) < \infty$ is not necessary unless Y is the triply-punctured sphere.

As a consequence of Theorem 5.7, every conformal embedding which is not a slit mapping can be approximated by a sequence of conformal embeddings each of which extends by some amount.

Corollary 5.10. *Assume hypothesis (H) and let $g \in \text{CEmb}(X, Y, h)$. If g is not a slit mapping, then there is a sequence $\{g_n\}$ converging to g in $\text{CEmb}(X, Y, h)$ such that g_n extends to a conformal embedding of X_{r_n} into Y for some $r_n > 0$.*

Proof. If $\mathbf{m}(g) > 0$, then we can take $g_n = g$. If $\mathbf{m}(g) = 0$ but g is not a local maximum of \mathbf{m} , then there exists a sequence $g_n \rightarrow g$ with $\mathbf{m}(g_n) > 0$. If $\mathbf{m}(g) = 0$ and g is a local maximum of \mathbf{m} , then g is a slit mapping by the previous theorem. \square

A strong converse to Theorem 5.7 holds due to Theorem 3.6.

Lemma 5.11. *Assume hypothesis (H). Suppose that $g: X_r \rightarrow Y$ is a slit mapping such that $g|_X$ is homotopic to h . Then $g|_X$ is a global maximum of \mathbf{m} .*

Proof. First observe that $\mathbf{m}(g|_X) \geq r$. Suppose that $\mathbf{m}(f) \geq \mathbf{m}(g|_X)$ for some element f of $\text{CEmb}(X, Y, h)$ and let \hat{f} be the maximal extension of f . Then $\hat{f}|_{X_r}$ is homotopic to g . By Theorem 3.6, $\hat{f}|_{X_r}$ is a slit mapping. In particular, the complement of $\hat{f}(X_r)$ has empty interior in Y . Therefore $X_{\mathbf{m}(f)} \setminus X_r$ is empty so that $\mathbf{m}(f) \leq r \leq \mathbf{m}(g|_X)$. \square

Furthermore, the almost rigidity of slit mappings implies that the set of local maxima of \mathbf{m} is path-connected.

Lemma 5.12. *Assume hypothesis (H). The set M of all local maxima of \mathbf{m} is homeomorphic to either a point, a compact interval, a circle, or a torus, and \mathbf{m} is constant on M .*

Proof. Suppose first that there is some $f \in M$ such that $\mathbf{m}(f) < \infty$. Then by Theorem 5.7, the maximal extension \hat{f} is a slit mapping. By Lemma 5.11, f is a global maximum of \mathbf{m} . In particular, $\mathbf{m}(g) < \infty$ for every $g \in M$ and thus every $g \in M$ is a global maximum of \mathbf{m} . In particular, \mathbf{m} is constant on M , say equal to R . The map $M \rightarrow \text{CEmb}(X_R, Y, h)$ defined by $f \mapsto \hat{f}$ is a homeomorphism with inverse $g \mapsto g|_X$. By Theorem 1.4, the evaluation map $\text{CEmb}(X_R, Y, h) \rightarrow Y$ sending f to $f(x)$ is a homeomorphism onto its image for every $x \in X_R$, and its image is either a point, a compact interval, a circle, or a torus.

Otherwise, \mathbf{m} is constant equal to ∞ on M . In this case M is homeomorphic to $\text{CEmb}(X_\infty, Y, h)$, which is the same as $\text{Aut}_0(Y)$. This is either a point or a torus (see Subsection 2.2). \square

Theorem 5.7, Lemma 5.11 and Lemma 5.12 together imply Theorem 1.5 from the introduction. The fact that \mathbf{m} has a connected plateau of local maxima easily implies that the space $\text{CEmb}(X, Y, h)$ is connected.

Proof of Theorem 1.6. Suppose that $\text{CEmb}(X, Y, h) = E_0 \cup E_1$ where E_0 and E_1 are disjoint non-empty closed sets. Then each of E_0 and E_1 is both compact and open. Since \mathfrak{m} is upper semi-continuous, the restriction $\mathfrak{m}|_{E_j}$ attains its maximum at some $f_j \in E_j$. Then f_j is a local maximum of \mathfrak{m} since E_j is open. By Lemma 5.12, f_0 and f_1 are both contained in a connected subset M of $\text{CEmb}(X, Y, h)$. But $M = (M \cap E_0) \cup (M \cap E_1)$ is disconnected. Contradiction. \square

6. THE BLOB AND ITS BOUNDARY

We say that a continuous map $h : X \rightarrow Y$ between any Riemann surfaces is *generic* if the induced homomorphism $\pi_1(h) : \pi_1(X, x) \rightarrow \pi_1(Y, h(x))$ has non-abelian image. This implies that X and Y are hyperbolic. Given such a generic map h , let $\text{Map}(X, Y, h)$ be the space of all continuous maps $f : X \rightarrow Y$ homotopic to h . The following lemma shows that for every $x \in X$ and every $f \in \text{Map}(X, Y, h)$, there is a well-defined way to lift the image point $f(x)$ to the universal cover of Y .

Lemma 6.1. *Let $h : X \rightarrow Y$ be a generic map between Riemann surfaces. Let $H : X \times [0, 1] \rightarrow Y$ be a homotopy from h to some $f \in \text{Map}(X, Y, h)$. Then for every $x \in X$, the homotopy class rel endpoints of the path $t \mapsto H(x, t)$ does not depend on the choice of H .*

Proof. We use two standard facts about hyperbolic surfaces:

- every abelian subgroup of $\pi_1(Y, h(x))$ is cyclic;
- every non-trivial element in $\pi_1(Y, h(x))$ is the positive power of a unique primitive element.

Let G be any other homotopy from h to f . By composing H with G run backwards, we get a homotopy from h to itself, hence a map $F : X \times S^1 \rightarrow Y$. Suppose that the loop $\gamma(t) = F(x, t)$ is not trivial in $\pi_1(Y, h(x))$. Then it is equal to β^k for some primitive element β and some $k > 0$.

Let α be any loop in X based at x . Then the map $S^1 \times S^1 \rightarrow Y$ given by $(s, t) \mapsto F(\alpha(s), t)$ induces a homomorphism of \mathbb{Z}^2 into $\pi_1(Y, h(x))$. The image of this homomorphism is cyclic, and contains both $[h \circ \alpha]$ and $[\gamma]$. From the existence and uniqueness of primitive roots in $\pi_1(Y, h(x))$, it follows that $[h \circ \alpha] = \beta^j$ for some $j \in \mathbb{Z}$. Since α was arbitrary, we deduce that the image of the homomorphism $\pi_1(h) : \pi_1(X, x) \rightarrow \pi_1(Y, h(x))$ is contained in the cyclic group $\langle \beta \rangle$. This contradicts the hypothesis that h is generic. We conclude that the loop $\gamma(t) = F(x, t)$ is null-homotopic. Equivalently, the paths $t \mapsto H(x, t)$ and $t \mapsto G(x, t)$ are homotopic rel endpoints. \square

Definition 6.2. Let $h : X \rightarrow Y$ be a generic map between Riemann surfaces, let $f \in \text{Map}(X, Y, h)$ and let $x \in X$. We define $\text{lift}_x(f)$ to be the homotopy class rel endpoints of the path $t \mapsto H(x, t)$ in Y where H is any homotopy from h to f . This is well-defined by Lemma 6.1. By definition, $\text{lift}_x(f)$ is an element of the universal cover \tilde{Y} of Y based at $h(x)$. This universal cover is isomorphic to \mathbb{D} , but for our purposes it will be better to think of it as the set of homotopy classes of paths in Y starting at $h(x)$.

Remark. If $h : X \rightarrow Y$ is a cyclic map between Riemann surfaces, i.e., such that the image of $\pi_1(h)$ is an infinite cyclic subgroup $C \leq \pi_1(Y, h(x))$, and Y is not a

torus then a similar construction defines a lift from $\text{Map}(X, Y, h)$ to the annulus cover $A \rightarrow Y$ associated to C .

Lemma 6.3. *Let $h : X \rightarrow Y$ be a generic map between Riemann surfaces, where X is finite. Then $\text{lift}_x : \text{Map}(X, Y, h) \rightarrow \tilde{Y}$ is continuous.*

Proof. Since $\text{Map}(X, Y, h)$ is metrizable, it suffices to prove sequential continuity. Let $f_n, f \in \text{Map}(X, Y, h)$ be such that $f_n \rightarrow f$ as $n \rightarrow \infty$. Let $K \subset X$ be a compact deformation retract of X containing x . Let $\varepsilon > 0$ be smaller than the minimum of the injectivity radius of Y over $f(K)$ with respect to the hyperbolic metric and let n be large enough so that $|f_n - f| < \varepsilon$ on K . For every $(\xi, t) \in K \times [0, 1]$, let $F_n(\xi, t)$ be the point at proportion t along the unique shortest length geodesic between $f(\xi)$ and $f_n(\xi)$ in Y . This gives a continuous homotopy from $f|_K$ to $f_n|_K$. By composing the deformation retraction $X \rightarrow K$ with F_n , we get a homotopy G_n from f to f_n moving points of K by distance at most ε .

Given any homotopy H from h to f , the concatenation $H * G_n$ (that is, H followed by G_n) is a homotopy from h to f_n . Thus $\text{lift}_x(f_n)$ can be represented as $\alpha * \beta_n$ where α is any representative of $\text{lift}_x(f)$ and $\beta_n(t) = G_n(x, t)$. Since the geodesic β_n is contained in the ball of radius ε centered at $\alpha(1)$ in Y , and ε can be taken arbitrarily small, we have $\text{lift}_x(f_n) \rightarrow \text{lift}_x(f)$ as $n \rightarrow \infty$. \square

We now get back to conformal embeddings and look at where a given point can go under all conformal embeddings in a given homotopy class.

Definition 6.4. Let $h : X \rightarrow Y$ be a generic embedding between finite Riemann surfaces and let $x \in X$. We define $\text{Blob}(x, X, Y, h)$ to be the image of $\text{CEmb}(X, Y, h)$ by the map lift_x from Definition 6.2. This is a subset of the universal cover \tilde{Y} of Y based at $h(x)$. Given $c \in \tilde{Y}$, we will write $\pi_Y(c)$ for its projection to Y , that is, the endpoint $\gamma(1)$ of any representative $\gamma \in c$.

Our previous results imply that $\text{Blob}(x, X, Y, h)$ is at most 1-dimensional whenever $\text{CEmb}(X, Y, h)$ contains a slit mapping.

Proposition 6.5. *Let $h : X \rightarrow Y$ be a generic embedding between finite Riemann surfaces and let $x \in X$. If $\text{CEmb}(X, Y, h)$ contains a slit mapping, then $\text{Blob}(x, X, Y, h)$ is homeomorphic to a point or a compact interval.*

Proof. By Theorem 1.4, the evaluation map $\text{ev}_x : \text{CEmb}(X, Y, h) \rightarrow Y$ sending f to $f(x)$ is a homeomorphism onto its image $V(x)$. Moreover, $V(x)$ is either a point or a compact interval since Y is not a torus and h is not cyclic. The restriction of the universal covering map $\pi_Y : \tilde{Y} \rightarrow Y$ to $\text{Blob}(x, X, Y, h)$ is a homeomorphism onto its image $V(x)$ with inverse $\text{lift}_x \circ \text{ev}_x^{-1}$. \square

We will see that $\text{Blob}(x, X, Y, h)$ is not much more complicated when $\text{CEmb}(X, Y, h)$ does not contain any slit mapping. Let us recall hypothesis **(H')** from the introduction:

- $h : X \rightarrow Y$ is a generic embedding between finite Riemann surfaces;
- $\partial X \neq \emptyset$;
- $\text{CEmb}(X, Y, h)$ is non-empty and does not contain any slit mapping.

Assume hypothesis **(H')**. Since $\text{CEmb}(X, Y, h)$ is compact (Lemma 3.1) and connected (Theorem 1.6), and since lift_x is continuous (Lemma 6.3), the set $\text{Blob}(x, X, Y, h)$

is compact and connected for any $x \in X$. Our goal is to show that the blob is homeomorphic to a closed disk (Theorem 1.7). The strategy of the proof is to analyze the boundary of the blob. We will show that every point in $\partial \text{Blob}(x, X, Y, h)$ is attained by a special kind of map in $\text{CEmb}(X, Y, h)$ which we call a slit mapping rel x .

Definition 6.6. Let $h : X \rightarrow Y$ be a topological embedding between finite Riemann surfaces and let $x \in X$. We say that a map $f \in \text{Map}(X, Y, h)$ is a *Teichmüller embedding rel x* if the restriction $f^* : X \setminus \{x\} \rightarrow Y \setminus \{f(x)\}$ is a Teichmüller embedding. A *slit mapping rel x* is a Teichmüller embedding rel x which is conformal.

The distinction to make here is that the initial and terminal quadratic differentials of f are allowed to have simple poles at x and $f(x)$ respectively. For example, there are no slit mappings into the triply-punctured sphere, but plenty of slit mappings relative to a point.

In order to characterize the boundary points of the blob, we first need a standard construction for pushing a point around on a Riemann surface. To paraphrase [FM11, p.97]: imagine placing your finger on a surface and pushing it along a smooth path, dragging the rest of the surface along as you go. The diffeomorphism obtained at the end is called a *point-pushing diffeomorphism*.

Lemma 6.7. *Let Y be any Riemann surface and let $\gamma : [0, 1] \rightarrow Y$ be a smooth immersion. There exists an isotopy $H : Y \times [0, 1] \rightarrow Y$ through quasiconformal diffeomorphisms such that $H(y, 0) = y$ for every $y \in Y$ and $H(\gamma(0), t) = \gamma(t)$ for every $t \in [0, 1]$.*

Proof. First assume that γ embedded. Extend the vector field $\gamma_*(\partial/\partial t)$ to a smooth vector field V supported in a tubular neighborhood of $\gamma([0, 1])$, and define H to be the vector flow along V . For each $t \in [0, 1]$, the map $y \mapsto H(y, t)$ is a diffeomorphism which is the identity outside a compact set, hence is quasiconformal.

If γ is not embedded, break it up into finitely many embedded subarcs γ_j , then construct an isotopy H_j on each corresponding subinterval $[a_{j-1}, a_j]$ using the above method. The isotopy $H(y, t)$ is defined as $H_1(y, t)$ for $t \in [a_0, a_1]$, then $H_2(H_1(y, a_1), t)$ for $t \in [a_1, a_2]$ and so on, picking up where we left off at each step. \square

Given a smooth immersion $\gamma : [0, 1] \rightarrow Y$ into a Riemann surface and an isotopy H as in Lemma 6.7, we say that $y \mapsto \mathcal{P}_\gamma(y) := H(y, 1)$ is a *point-pushing diffeomorphism* along γ . By definition, \mathcal{P}_γ is isotopic to the identity and satisfies $\mathcal{P}_\gamma(\gamma(0)) = \gamma(1)$. Of course, \mathcal{P}_γ depends on the specific choice of H but its isotopy class rel $\gamma(0)$ only depends on the homotopy class of γ by the Birman exact sequence [FM11, Theorem 4.6].

Here are two elementary observations. If $h : X \rightarrow Y$ is a generic map between Riemann surfaces and $\gamma : [0, 1] \rightarrow Y$ is an immersed curve such that $\gamma(0) = h(x)$, then $\text{lift}_x(\mathcal{P}_\gamma \circ h) = [\gamma]$ by definition. Secondly, if two maps f and g in $\text{Map}(X, Y, h)$ are homotopic rel x , then obviously $\text{lift}_x(f) = \text{lift}_x(g)$. The converse is also true, as can be seen using point-pushing on Y . Thus lift_x detects when two homotopic maps are homotopic rel x .

We say that a map $f \in \text{Map}(X, Y, h)$ *realizes* a point $c \in \tilde{Y}$ if $\text{lift}_x(f) = c$. We can now prove that points outside the blob are realized by Teichmüller embeddings rel x , provided that there is some quasiconformal embedding in the homotopy class of h .

Lemma 6.8. *Let $h : X \rightarrow Y$ be a generic quasiconformal embedding between finite Riemann surfaces, let $x \in X$, and let $c \in \tilde{Y} \setminus \text{Blob}(x, X, Y, h)$. Then there exists a unique Teichmüller embedding f rel x homotopic to h such that $\text{lift}_x(f) = c$.*

Proof. Let $\gamma : [0, 1] \rightarrow Y$ be a smooth immersed curve in c , for example the hyperbolic geodesic, and let \mathcal{P}_γ be a quasiconformal point-pushing diffeomorphism along γ . Then $\mathcal{P}_\gamma \circ h$ is a quasiconformal embedding homotopic to h such that $\text{lift}_x(\mathcal{P}_\gamma \circ h) = [\gamma] = c$. Let $f^* : X \setminus \{x\} \rightarrow Y \setminus \{\gamma(1)\}$ be a quasiconformal embedding homotopic to the restriction of $\mathcal{P}_\gamma \circ h$ with minimal dilatation (Corollary 3.2) and let $f : X \rightarrow Y$ be its extension. Since f is homotopic to $\mathcal{P}_\gamma \circ h$ rel x , we have $\text{lift}_x(f) = \text{lift}_x(\mathcal{P}_\gamma \circ h) = c$.

Suppose that f^* is conformal. Then $f \in \text{CEmb}(X, Y, h)$ so that $c = \text{lift}_x(f)$ belongs to $\text{Blob}(x, X, Y, h)$, contrary to our assumption. By Theorem 3.4, f^* is a Teichmüller embedding, so that f is a Teichmüller embedding rel x .

If g is another Teichmüller embedding rel x homotopic to h such that $\text{lift}_x(g) = c$, then g is homotopic to f rel x . Thus the restrictions g^* and f^* to $X \setminus \{x\} \rightarrow Y \setminus \{\gamma(1)\}$ are homotopic Teichmüller embeddings with an essential puncture, and we conclude that $g = f$ by Remark 4. \square

We deduce that points on the boundary of the blob are realized by slit mappings rel x .

Proposition 6.9. *Assume (\mathbf{H}') , let $x \in X$ and let $c \in \partial \text{Blob}(x, X, Y, h)$. Then there is a unique $f \in \text{CEmb}(X, Y, h)$ such that $\text{lift}_x(f) = c$. Moreover, f is a slit mapping rel x . If φ and ψ are initial and terminal quadratic differentials for f rel x , then φ has a simple pole at x and ψ has a simple pole at $f(x)$.*

Proof. Since $\text{Blob}(x, X, Y, h)$ is closed, c belongs to $\text{Blob}(x, X, Y, h)$. Hence there exists some $f \in \text{CEmb}(X, Y, h)$ such that $\text{lift}_x(f) = c$.

Let $\{c_n\}_{n=1}^\infty$ be a sequence in $\tilde{Y} \setminus \text{Blob}(x, X, Y, h)$ such that $c_n \rightarrow c$ as $n \rightarrow \infty$. Then $Y \setminus \{\pi_Y(c_n)\}$ converges to $Y \setminus \{\pi_Y(c)\}$ as $n \rightarrow \infty$, where we can take markings to be the identity outside of a small neighborhood of $\pi_Y(c)$. By the previous lemma, there exists a Teichmüller embedding $f_n : X \rightarrow Y$ rel x homotopic to h such that $\text{lift}_x(f_n) = c_n$. By Lemma 5.4, we can extract a subsequence of the restrictions $f_n^* : X \setminus \{x\} \rightarrow Y \setminus \{\pi_Y(c_n)\}$ converging to a Teichmüller embedding $g : X \setminus \{x\} \rightarrow Y \setminus \{\pi_Y(c)\}$ which is homotopic to the restriction f^* of f . As f^* is conformal, it is a slit mapping by Theorem 3.6, so that f is a slit mapping rel x . For any $g \in \text{CEmb}(X, Y, h)$ realizing c , we have that g is homotopic to f rel x . Hence $g = f$ by Theorem 3.6 and Remark 4.

Suppose that ψ does not have a pole at y . Then f is an honest slit mapping from X to Y . But we assumed that $\text{CEmb}(X, Y, h)$ does not contain any slit mapping. Therefore ψ has a simple pole at $f(x)$ and $\varphi = f^*\psi$ has a simple pole at x . \square

Assume hypothesis (\mathbf{H}') and let $x \in X$. Suppose that $c \in \partial \text{Blob}(x, X, Y, h)$ and that $f \in \text{CEmb}(X, Y, h)$ realizes c . We say that any quadratic differential $\psi \in \mathcal{Q}^+(Y \setminus \{f(x)\})$ with respect to which f is a slit mapping rel x realizes c as well. Even though f is unique by Proposition 6.9, ψ need not be unique up to scale (see Remark 3.2). Nevertheless, the set of quadratic differentials realizing a given point $c \in \partial \text{Blob}(x, X, Y, h)$ is convex.

Lemma 6.10. *Assume hypothesis (\mathbf{H}') . Suppose that ψ_0 and ψ_1 in $\mathcal{Q}^+(Y \setminus \{y\})$ realize $c \in \partial \text{Blob}(x, X, Y, h)$ where $y = \pi_Y(c)$. Then for every $\alpha, \beta > 0$ the quadratic differential $\alpha\psi_0 + \beta\psi_1$ belongs to $\mathcal{Q}^+(Y \setminus \{y\})$, realizes c , and has a simple pole at y .*

Proof. Let f be the slit mapping realizing c . Then $\psi := \alpha\psi_0 + \beta\psi_1$ is ≥ 0 along any arc in $\widehat{Y} \setminus f(X)$ so that f is a slit mapping rel x with respect to ψ . If ψ does not have a pole at y , then f is a slit mapping from X to Y , contradicting hypothesis (\mathbf{H}') . \square

We will see that any $\psi \in \mathcal{Q}^+(Y \setminus \{\pi_Y(c)\})$ which realizes c tells us something about the shape of $\text{Blob}(x, X, Y, h)$ near c . More precisely, ψ can be used to find a region $U \subset \widetilde{Y}$ of forbidden values for the maps in $\text{CEmb}(X, Y, h)$, that is, disjoint from $\text{Blob}(x, X, Y, h)$. The idea is that if we push $\pi_Y(c)$ in some directions, then a certain quantity goes up, while it cannot increase under conformal embedding. This quantity is extremal length.

7. EXTREMAL LENGTH OF PARTIAL MEASURED FOLIATIONS

There are several equivalent ways to define the extremal length of a measured foliation. We follow the approach developed in [GL10a] and [GL10b].

A map will be called *almost-smooth* if it is continuous and continuously differentiable except perhaps at finitely many points.

Definition 7.1. A *partial measured foliation* $F = \{(U_j, v_j)\}_{j \in J}$ on a Riemann surface X is a collection of open sets $U_j \subset X$ together with almost-smooth functions $v_j : U_j \rightarrow \mathbb{R}$ satisfying

$$v_j = \pm v_k + c_{jk}$$

on $U_j \cap U_k$, where c_{jk} is locally constant. The set $U = \bigcup_j U_j$ is called the *support* of F . Since $|dv_j| = |dv_k|$ on $U_j \cap U_k$, these patch up to a well-defined object $|dv|$ on U called the *transverse measure* of F . The *leaves* of F are the maximal connected subsets of U on which each v_j is locally constant. We will write F or $|dv|$ interchangeably.

Remark. This is much weaker than usual notion of measured foliation [FLP12]. For one thing, the support does not have to be the complement of a finite set. Moreover, the leaves are not required to be 1-dimensional submanifolds; they can be thick. In practice, we will only deal with partial measured foliations which are measured foliations on a subsurface.

Remark. We could relax the regularity condition on the functions v_j and only assume that they belong to the Sobolev space $W^{1,2}(U_j)$. This would be more natural from the point of view of quasiconformal maps. For the sake of simplicity we will stick to the almost-smooth condition.

For example, if φ is a quadratic differential and U_j is a simply connected domain on which φ does not have any singularities, then the function

$$v_j(z) = \text{Im} \int_{z_0}^z \sqrt{\varphi}$$

is well-defined on U_j up to an additive constant and sign. The resulting partial measured foliation $\mathcal{F}^h \varphi = |dv| = |\text{Im} \sqrt{\varphi}|$ is the *horizontal foliation* of φ . Its leaves are the regular horizontal trajectories of φ .

The *Dirichlet energy* of a partial measured foliation $F = |dv|$ with support U is

$$\text{Dir}(F) := \int_U |\nabla v|^2 dA$$

where the gradient ∇v (only defined up to sign) is computed with respect to any conformal Riemannian metric on X with corresponding area form dA . Alternatively, we can write

$$\text{Dir}(F) = \int_U \left(\frac{\partial v}{\partial x} \right)^2 + \left(\frac{\partial v}{\partial y} \right)^2 dx dy$$

where $z = x + iy$ is any conformal coordinate.

For example, if φ is a holomorphic quadratic differential on X and $\mathcal{F}^h \varphi$ is its horizontal foliation, then

$$\text{Dir}(\mathcal{F}^h \varphi) = \int_X |\varphi| = \|\varphi\|$$

as can be seen by computing the Dirichlet energy in natural coordinates where $\varphi = dz^2$ and $v(x + iy) = \pm y + c$.

A *cross-cut* on a Riemann surface X is a proper arc $\alpha : (0, 1) \rightarrow X$. Two cross-cuts are *homotopic* if there is a homotopy through cross-cuts between them. The *height* of a homotopy class c of closed curves or cross-cuts in X with respect to a partial measured foliation $F = |dv|$ is the quantity

$$\text{height}(c, F) := \inf_{\alpha \in c} \int_{\alpha} |dv|,$$

where the infimum is taken over all piecewise smooth curves $\alpha \in c$, and where $|dv|$ is extended to be zero outside its support.

Given partial measured foliations F and G on a Riemann surface X , we say that G *dominates* F if

$$\text{height}(c, G) \geq \text{height}(c, F)$$

for every homotopy class c of closed curves or cross-cuts in X . Two partial measured foliations are *measure equivalent* if they dominate each other, i.e., if they have the same heights.

Definition 7.2. The *extremal length* of a partial measured foliation F on a Riemann surface X is defined as

$$\text{EL}(F, X) := \inf \{ \text{Dir}(G) : G \text{ dominates } F \}.$$

This is not the standard definition of extremal length. However, if F is a measured foliation (not partial) on a finite Riemann surface X , then $\text{EL}(F, X)$ is the norm of the unique holomorphic quadratic differential $\Phi(F, X)$ on X with the same heights as F (this is often taken as a definition). The existence and uniqueness of $\Phi(F, X)$ is due to Hubbard and Masur [HM79], but what interests us here is its minimizing property.

Theorem 7.3 (The minimal norm property). *Let X be a finite Riemann surface, let $\varphi \in \mathcal{Q}(X) \setminus \{0\}$, and let $\mathcal{F}^h \varphi$ be the horizontal foliation of φ . Then*

$$\text{EL}(\mathcal{F}^h \varphi, X) = \text{Dir}(\mathcal{F}^h \varphi) = \|\varphi\|.$$

That is,

$$\text{Dir}(\mathcal{F}^h \varphi) \leq \text{Dir}(G)$$

for every partial measured foliation G on X which dominates $\mathcal{F}^h\varphi$.

This is proved in [MS84, Theorem 3.2] (see also [Str84, Chapter VII]). There the result is stated for partial measured foliations G whose support has a discrete complement, but this hypothesis is not used anywhere in the proof. The idea of the proof is to look at the vertical foliation $\mathcal{F}^v\varphi$ which splits up into rectangles, cylinders and minimal components. If λ is a vertical leaf which is a cross-cut or a closed curve, then $\int_\lambda d\mathcal{F}^h\varphi = \text{height}(\lambda, \mathcal{F}^h\varphi) \leq \int_\lambda dG$. There is a similar inequality for minimal components, but one needs to pick a horizontal transversal and do a surgery on vertical leaves to obtain closed curves. The result is obtained by integrating these 1-dimensional inequalities against the leaf space and applying the Cauchy-Schwarz inequality. This is a sophisticated version of the so-called ‘‘length-area argument’’.

7.1. Pushing forward. If $f : X \rightarrow Y$ is an almost-smooth embedding between Riemann surfaces and $F = \{(U_j, v_j)\}_{j \in J}$ is a partial measured foliation on X , then the *push-forward* $f_*F = \{(f(U_j), v_j \circ f^{-1})\}_{j \in J}$ is a partial measured foliation on Y . This justifies our preference for partial measured foliations over measured foliations: we will push them forward by conformal embeddings.

Remark. The notion of heights induces a topology on the set of (measure equivalence classes of) partial measured foliations on a finite Riemann surface. The reader should be warned that the push-forward operator is not continuous with respect to that topology [KPT15, Example 4.4].

A standard calculation shows that Dirichlet energy increases by a factor at most K under almost-smooth K -quasiconformal embeddings [Ahl06, p.14]. As a consequence, extremal length increases by a factor at most K under such maps. In particular, extremal length does not increase under conformal embeddings.

Lemma 7.4. *Let $f : X \rightarrow Y$ be an almost-smooth K -quasiconformal embedding between finite Riemann surfaces and let $\varphi \in \mathcal{Q}(X) \setminus \{0\}$. Then*

$$\text{EL}(f_*\mathcal{F}^h\varphi, Y) \leq K \text{EL}(\mathcal{F}^h\varphi, X).$$

Proof. Let $\zeta = \sigma + i\tau$ be a conformal coordinate on $f(X)$ and let $z = x + iy$ be a conformal coordinate at $f^{-1}(\zeta)$. Since f is K -quasiconformal, we have

$$dxdy = \det(df^{-1})d\sigma d\tau \geq K^{-1}\|df^{-1}\|^2 d\sigma d\tau.$$

Let $v(z) = \text{Im} \int_{z_0}^z \sqrt{\varphi}$ so that $|dv| = |\text{Im} \sqrt{\varphi}| = \mathcal{F}^h\varphi$. We compute

$$\begin{aligned} \text{Dir}(f_*|dv|) &= \int_{f(X)} |\nabla(v \circ f^{-1})(\zeta)|^2 d\sigma d\tau \\ &= \int_{f(X)} |(df^{-1})(\nabla v)(f^{-1}(\zeta))|^2 d\sigma d\tau \\ &\leq \int_{f(X)} \|df^{-1}\|^2 |(\nabla v)(f^{-1}(\zeta))|^2 d\sigma d\tau \\ &\leq K \int_{f(X)} \det(df^{-1}) |(\nabla v)(f^{-1}(\zeta))|^2 d\sigma d\tau \\ &= K \int_X |(\nabla v)(z)|^2 dxdy = K \text{Dir}(|dv|). \end{aligned}$$

It follows that

$$\text{EL}(f_*|dv|, Y) \leq \text{Dir}(f_*|dv|) \leq K \text{Dir}(|dv|) = K \text{EL}(|dv|, X),$$

where the last equality holds by Theorem 7.3. \square

Remark. It would be desirable to have this for any partial measured foliation F on X , but one runs into a difficulty: it is not clear that f_*G dominates f_*F if G dominates F . See Lemma 7.6 for a similar statement which sounds obvious but is not entirely straightforward.

The previous inequality is sharp as the case of Teichmüller embeddings illustrates.

Lemma 7.5. *Let $f : X \rightarrow Y$ be a Teichmüller embedding of dilatation K with initial and terminal quadratic differentials φ and ψ . Then $f_*\mathcal{F}^h\varphi$ is measure equivalent to $\mathcal{F}^h\psi$ on Y and we have*

$$\text{EL}(\mathcal{F}^h\psi, Y) = \text{EL}(f_*\mathcal{F}^h\varphi, Y) = K \text{EL}(\mathcal{F}^h\varphi, X).$$

Proof. We have $\mathcal{F}^h\psi = f_*\mathcal{F}^h\varphi$ on $f(X)$ since $f(x + iy) = Kx + iy$ in natural coordinates. Moreover, since $Y \setminus f(X)$ is a finite union of horizontal arcs and points, the integral of $\mathcal{F}^h\psi$ is zero along any piece of curve contained in this set. For any piecewise smooth curve α in Y , we thus have

$$\begin{aligned} \int_{\alpha} d\mathcal{F}^h\psi &= \int_{\alpha} \chi_{f(X)} d\mathcal{F}^h\psi + \int_{\alpha} \chi_{Y \setminus f(X)} d\mathcal{F}^h\psi \\ &= \int_{\alpha} df_*\mathcal{F}^h\varphi + \int_{\alpha} \chi_{Y \setminus f(X)} d\mathcal{F}^h\psi \\ &= \int_{\alpha} df_*\mathcal{F}^h\varphi \end{aligned}$$

where χ_A is the characteristic function of the set A . Therefore $\mathcal{F}^h\psi$ and $f_*\mathcal{F}^h\varphi$ are measure equivalent, which implies that they have the same extremal length. By Lemma 7.4, the inequality

$$\text{EL}(f_*\mathcal{F}^h\varphi, Y) \leq K \text{EL}(\mathcal{F}^h\varphi, X)$$

holds and by Theorem 7.3 we have $\text{EL}(\mathcal{F}^h\psi, Y) = \|\psi\|$ and $\text{EL}(\mathcal{F}^h\varphi, X) = \|\varphi\|$. We also have $\|\psi\| = K\|\varphi\|$ since f stretches horizontally by a factor K and $f(X)$ has full measure in Y . Putting everything together, we get

$$\|\psi\| = \text{EL}(\mathcal{F}^h\psi, Y) = \text{EL}(f_*\mathcal{F}^h\varphi, Y) \leq K \text{EL}(\mathcal{F}^h\varphi, X) = K\|\varphi\| = \|\psi\|$$

and hence equality holds. \square

We will need a sufficient condition for when the push-forwards of a measured foliation by two homotopic embeddings are measure equivalent. We say that an embedding $f : X \rightarrow Y$ between finite Riemann surfaces is *tame* if it is almost-smooth and extends to a continuous map $\hat{f} : \hat{X} \rightarrow \hat{Y}$ which is piecewise smooth along ∂X . For example, every Teichmüller embedding is tame.

Lemma 7.6. *Let $f, g : X \rightarrow Y$ be homotopic tame embeddings between finite Riemann surfaces and let $\varphi \in \mathcal{Q}^+(X)$. Then $f_*\mathcal{F}^h\varphi$ and $g_*\mathcal{F}^h\varphi$ are measure equivalent on Y .*

It is important that we take φ in $\mathcal{Q}^+(X)$ and not just in $\mathcal{Q}(X) \setminus \{0\}$. We want the transverse measure of each boundary component to be zero.

Proof. Let $F = \mathcal{F}^h\varphi$. By symmetry, it suffices to show that f_*F dominates g_*F . Let α be a piecewise smooth closed curve or cross-cut in Y . Given $\varepsilon > 0$, we have to find a closed curve or cross-cut β homotopic to α such that

$$(2) \quad \int_{\beta} g_*F \leq \int_{\alpha} f_*F + \varepsilon.$$

If α is contained in $f(X)$, then we can just take $\beta = g \circ f^{-1}(\alpha)$. The difficulty is when α intersects $Y \setminus f(X)$.

If $\int_{\alpha} f_*F = \infty$ there is nothing to show, so assume the integral is finite. Then there is only a finite number n of components of $f^{-1}(\alpha)$ which are essential in X in the sense that they cannot be homotoped into an arbitrarily small neighborhood of $\partial X \cup \dot{X}$. Let $\delta = \varepsilon/(8n)$ and let X^δ be \dot{X} minus a δ -neighborhood of $\partial X \cup \dot{X}$ in the metric $|\varphi|$. The map $g \circ f^{-1} : f(X^\delta) \rightarrow g(X^\delta)$ can be extended to an almost-smooth homeomorphism $H : Y \rightarrow Y$ homotopic to the identity (see [Mas75, Lemma 2]).

We will take β to be a modified version of $H(\alpha)$. First, each inessential component of $H(\alpha) \cap g(X)$ can be homotoped within $g(X)$ to an arc contained in $g(X \setminus X^\delta)$ having arbitrarily small transverse measure with respect to g_*F . Even if there are infinitely many such inessential components, we can make sure that their total transverse measure is at most $\varepsilon/2$. Let γ be the curve obtained after having done this.

Next, we homotope each essential component c of $\gamma \cap g(X \setminus X^\delta)$, within $g(X \setminus X^\delta)$, to an arc having transverse measure at most 2δ , while keeping endpoints fixed. This is possible since the height of each component of $X \setminus X^\delta$ is δ . Moreover, no matter how many times $g^{-1}(c)$ winds around the annular component of $X \setminus X^\delta$ in which it is contained, we can push this winding part toward $\partial X \cup \dot{X}$. In doing so, the transverse measure of the winding part tends to zero and is thus eventually less than δ , for a total of at most 2δ . There are $2n$ such components—two for each essential component of $\gamma \cap g(X)$ —for a total transverse measure of $(2n)(2\delta) = \varepsilon/2$.

Let β be the curve obtained after having done these two modifications. Then inequality (2) holds since $\beta \setminus g(X)$ has measure zero, $\beta \cap g(X \setminus X^\delta)$ contributes at most $\varepsilon/2 + \varepsilon/2 = \varepsilon$, and $\beta \cap g(X^\delta)$ is the image by $g \circ f^{-1}$ of a subset of α so its transverse measure with respect to g_*F is at most the transverse measure of α with respect to f_*F . \square

8. THE BLOB IS SEMI-SMOOTH

In this section, we show that each quadratic differential realizing a point on the boundary of $\text{Blob}(x, X, Y, h)$ determines a vector normal to the blob and vice versa. We use this to prove that the blob satisfies a certain regularity condition near its boundary which we call semi-smoothness. We first need the following formula for the derivative of extremal length [Gar84, Theorem 8].

Theorem 8.1 (Gardiner's formula). *Let S be a finite Riemann surface and let F be a measured foliation on S . Then the function⁹ $Z \mapsto \text{EL}(F, Z)$ is differentiable*

⁹This is a slight abuse of notation. Recall that a point in $\mathcal{T}(S)$ is an equivalence class of marking $f : S \rightarrow Z$, which we may assume is almost-smooth. By the expression $\text{EL}(F, Z)$ we really mean $\text{EL}(f_*F, Z)$.

on $\mathcal{T}(S)$. Its derivative at $Z \in \mathcal{T}(S)$ in the direction $\mu \in T_Z\mathcal{T}(S)$ is

$$(3) \quad 2\langle \mu, \varphi \rangle$$

where $\varphi \in \mathcal{Q}(Z)$ is the unique quadratic differential such that $\mathcal{F}^h\varphi$ is measure equivalent to F .

We apply this formula along a curve $Z_t = Y \setminus \{\gamma(t)\}$ for some analytic path $\gamma : I \rightarrow Y$, where Y is a finite hyperbolic surface. It is implicit here that the change of marking $Z_0 \rightarrow Z_t$ is given by point-pushing along γ . In this case, the pairing $\langle \mu, \varphi \rangle$ is proportional to the real part of the residue of φ in the direction of $\gamma'(0)$. See [McM13] for a similar but more general calculation.

Lemma 8.2. *Let Y be a finite hyperbolic surface, let $\gamma : (-1, 1) \rightarrow Y$ be an analytic arc and let $Z_t = Y \setminus \{\gamma(t)\}$. Then the derivative μ of Z_t at $t = 0$ satisfies*

$$(4) \quad \langle \mu, q \rangle = -\pi \operatorname{Re}[\operatorname{Res}_{\gamma(0)}(q \cdot \gamma'(0))]$$

for every $q \in \mathcal{Q}(Z_0)$.

Proof. We may assume that γ is embedded by restricting to a subinterval. Since γ is analytic, the vector field $\gamma_*(\partial/\partial t)$ can be extended to a holomorphic vector field \bar{v} on an embedded round disk $D \subset Y$ centered at $y = \gamma(0)$. Let E be another round disk of smaller radius centered at y and let ϕ be a smooth bump function which is equal to 1 on E and 0 outside D . Then $\phi\bar{v}$ is a smooth vector field defined on all of Y . Let Φ_t be the time- t flow for $\phi\bar{v}$ and let $\nu(t) = \bar{\partial}\Phi_t/\partial\Phi_t$. Then for small enough t , we have $\Phi_t(\gamma(0)) = \gamma(t)$ by construction so that $\mu = \nu'(0) = \bar{\partial}(\phi\bar{v})$. We compute

$$\begin{aligned} \int_{Z_0} \mu q &= \int_D q \bar{\partial}(\phi\bar{v}) = \int_{D \setminus E} q \bar{\partial}(\phi\bar{v}) = \int_{D \setminus E} \bar{\partial}(\phi q \bar{v}) \\ &= -\frac{i}{2} \int_{\partial(D \setminus E)} \phi q \bar{v} = \frac{i}{2} \int_{\partial E} q \bar{v} = -\pi \operatorname{Res}_y(q\bar{v}). \end{aligned}$$

The equality from the first line to the second is by Green's theorem and the change of sign in the next equality comes from reversing orientation on ∂E . To conclude the proof, recall that $\langle \mu, q \rangle = \operatorname{Re} \int_{Z_0} \mu q$ by definition. \square

Let Y be a finite hyperbolic surface and let $y \in Y$. If $\psi \in \mathcal{Q}(Y \setminus \{y\})$ has a simple pole at y , then there exists a tangent vector $\mathbf{v} \in T_y Y$ (unique up to rescaling) such that $\operatorname{Res}_y(\psi\mathbf{v}) < 0$. We say that \mathbf{v} is *vertical* for ψ . For example, if $\psi = \frac{1}{z} dz^2$ then $\mathbf{v} = -\frac{\partial}{\partial z}$ is vertical at 0 since

$$\operatorname{Res}_0(\psi\mathbf{v}) = \frac{1}{2\pi i} \oint \frac{-1}{z} dz = -1.$$

A vector $\bar{w} \in T_y Y$ is *mostly vertical* for ψ if $\operatorname{Re}[\operatorname{Res}_y(\psi\bar{w})] < 0$. If \mathbf{v} is vertical and $\lambda \in \mathbb{C}$ then $\lambda\mathbf{v}$ is mostly vertical if and only if $\operatorname{Re} \lambda > 0$. Thus the mostly vertical vectors are those that make an angle less than $\pi/2$ with the vertical direction (here angles are measured with respect to the Riemann surface structure on Y , not the cone metric coming from $|\psi|$).

If we push y in a mostly vertical direction, then the extremal length of $\mathcal{F}^h\psi$ will increase according to equations (3) and (4), at least for small time $|t| < \delta$. The intuition for this is that if we push in mostly vertical directions, we stretch




FIGURE 6. Pushing a simple pole of a quadratic differential in a mostly vertical direction increases the extremal length of its horizontal foliation.

the leaves of $\mathcal{F}^h\psi$ and compress them together, thereby increasing Dirichlet energy (see Figure 6).

We will need a more uniform statement where δ can be chosen independently of the direction. We can do this as long as we restrict to a compact set of angles. Given a point $y \in Y$ and a tangent vector $\vec{v} \in T_y Y$, let $\gamma_{\vec{v}} : \mathbb{R} \rightarrow Y$ be the hyperbolic geodesic such that $\gamma'_{\vec{v}}(0) = \vec{v}$. We will push y along these specific paths.

Lemma 8.3. *Let Y be a finite hyperbolic surface and let $y \in Y$. Suppose that $\psi \in \mathcal{Q}(Y \setminus \{y\})$ has a simple pole at y and let $\vec{v} \in T_y Y$ be vertical for ψ . Then for every $\theta_0 \in (0, \frac{\pi}{2})$ there exists a number $\delta > 0$ such that*

$$\text{EL}(\mathcal{F}^h\psi, Y \setminus \{\gamma_{e^{i\theta}\vec{v}}(t)\}) > \text{EL}(\mathcal{F}^h\psi, Y \setminus \{y\})$$

for every $\theta \in [-\theta_0, \theta_0]$ and every $t \in (0, \delta)$.

Proof. Suppose not and let $F = \mathcal{F}^h\psi$. Then there exist convergent sequences $t_n \searrow 0$ and $\theta_n \rightarrow \theta_\infty \in [-\theta_0, \theta_0]$ as $n \rightarrow \infty$ such that

$$\text{EL}(F, Y \setminus \{\gamma_{e^{i\theta_n}\vec{v}}(t_n)\}) \leq \text{EL}(F, Y \setminus \{y\})$$

for every $n \in \mathbb{N}$.

For $\vec{w} \in T_y Y$ let $\mu_{\vec{w}}$ be the derivative of the path $t \mapsto Y \setminus \{\gamma_{\vec{w}}(t)\}$ at $t = 0$. Then $\mu_{e^{i\theta_n}\vec{v}} \rightarrow \mu_{e^{i\theta_\infty}\vec{v}}$ as $n \rightarrow \infty$ by equation (4). Since $Z \mapsto \text{EL}(F, Z)$ is differentiable at $Z = Y \setminus \{y\}$ we have that

$$\frac{\text{EL}(F, Y \setminus \{\gamma_{e^{i\theta_n}\vec{v}}(t_n)\}) - \text{EL}(F, Y \setminus \{y\})}{t_n} \rightarrow 2\langle \mu_{e^{i\theta_\infty}\vec{v}}, \psi \rangle$$

as $n \rightarrow \infty$. But the left-hand side is non-positive for each n while

$$2\langle \mu_{e^{i\theta_\infty}\vec{v}}, \psi \rangle = -2\pi \text{Re}[\text{Res}_y(\psi e^{i\theta_\infty}\vec{v})] = -2\pi \cos(\theta_\infty) \text{Res}_y(\psi\vec{v}) > 0.$$

This is a contradiction. \square

Let us introduce some more terminology. Given $z \in \mathbb{C}$, $\mathbf{v} \in T_z\mathbb{C} \setminus \{0\}$, $\theta \in (0, \pi)$, and $\delta > 0$, we denote by $\nabla(\mathbf{v}, \theta, \delta)$ the open angular sector based at z with radius δ and angle θ on either side of \mathbf{v} . In symbols,

$$\nabla(\mathbf{v}, \theta, \delta) = \left\{ z + te^{i\phi} \frac{\mathbf{v}}{|\mathbf{v}|} : \phi \in (-\theta, \theta) \text{ and } t \in (0, \delta) \right\}.$$

Definition 8.4. Let $B \subset \mathbb{C}$ be closed. A vector $\mathbf{v} \in T_z\mathbb{C}$ with $z \in B$ is *normal to B* if $\mathbf{v} \neq 0$ and if there are angular sectors arbitrarily close to half-disks pointing in the direction of \mathbf{v} which are disjoint from B . More precisely, \mathbf{v} is normal to B if $\mathbf{v} \neq 0$ and if for every $\theta \in (0, \frac{\pi}{2})$, there exists a $\delta > 0$ such that $\nabla(\mathbf{v}, \theta, \delta) \cap B = \emptyset$.




FIGURE 7. A closed subset of the plane and some of its normal vectors. There is no normal vector at the inward corner and several normal vectors at outward corners.

Note that conformal homeomorphisms preserve normal vectors. Thus the definition extends to any Riemann surface. Furthermore, instead of Euclidean rays we can use geodesic rays for any conformal Riemannian metric (for example the hyperbolic metric if the surface is hyperbolic).

We are ready to prove that any quadratic differential realizing a point on the boundary of the blob determines a vector normal to the blob.

Proposition 8.5 (Vertical vectors are normal). *Assume hypothesis (\mathbf{H}') . Let $x \in X$, let $c \in \partial \text{Blob}(x, X, Y, h)$ and suppose that $\psi \in \mathcal{Q}^+(Y \setminus \{\pi_Y(c)\})$ realizes c . If $\mathbf{v} \in T_c \tilde{Y}$ is vertical for $\pi_Y^* \psi$, then \mathbf{v} is normal to $\text{Blob}(x, X, Y, h)$.*

Proof. Let $f \in \text{CEmb}(X, Y, h)$ be the slit mapping rel x realizing c provided by Proposition 6.9. Let $y = \pi_Y(c) = f(x)$, let $g : X \setminus \{x\} \rightarrow Y \setminus \{y\}$ be the restriction of f , let $G = \mathcal{F}^h \psi$ on $Y \setminus \{y\}$, let $\varphi = g^* \psi$ and let $F = \mathcal{F}^h \varphi$ on $X \setminus \{x\}$. Also let $\vec{w} = d\pi_Y(\vec{v})$ so that \vec{w} is vertical for ψ at y .

Fix $\theta_0 \in (0, \frac{\pi}{2})$. By Lemma 8.3, there exists a $\delta > 0$ such that for every $\theta \in [-\theta_0, \theta_0]$ and every $t \in (0, \delta)$ we have

$$\text{EL}(G, Y \setminus \{\gamma_{e^{i\theta} \mathbf{w}}(t)\}) > \text{EL}(G, Y \setminus \{y\}).$$

Since g is a slit mapping, G is measure equivalent to $g_* F$ on $Y \setminus \{y\}$ by Lemma 7.5. We also have $\text{EL}(G, Y \setminus \{y\}) = \|\psi\| = \|\varphi\| = \text{EL}(F, X \setminus \{x\})$ by Theorem 7.3, so the above inequality translates to

$$(5) \quad \text{EL}(g_* F, Y \setminus \{\gamma_{e^{i\theta} \mathbf{w}}(t)\}) > \text{EL}(F, X \setminus \{x\}).$$

Fix $\theta \in [-\theta_0, \theta_0]$ and $t \in (0, \delta)$. Let \mathcal{P} be a quasiconformal point-pushing diffeomorphism along $\gamma_{e^{i\theta} \mathbf{w}}([0, t])$. Recall that the left-hand side of (5) really stands for the extremal length of $(\mathcal{P} \circ g)_* F$ since $\mathcal{P} : Y \setminus \{y\} \rightarrow Y \setminus \{\gamma_{e^{i\theta} \mathbf{w}}(t)\}$ is our change of marking by convention (see footnote 9).

Suppose that there is a conformal embedding $M : X \setminus \{x\} \rightarrow Y \setminus \{\gamma_{e^{i\theta} \mathbf{w}}(t)\}$ homotopic to $\mathcal{P} \circ g$. Then by Corollary 5.10, either M is a slit mapping or it is homotopic to a conformal embedding which extends analytically to the ideal boundary. In either case, M is homotopic to a tame conformal embedding N . By Lemma 7.6, $N_* F$ and $(\mathcal{P} \circ g)_* F$ are measure equivalent on $Y \setminus \{\gamma_{e^{i\theta} \mathbf{w}}(t)\}$ which means that

$$\text{EL}(N_* F, Y \setminus \{\gamma_{e^{i\theta} \mathbf{w}}(t)\}) = \text{EL}((\mathcal{P} \circ g)_* F, Y \setminus \{\gamma_{e^{i\theta} \mathbf{w}}(t)\}) > \text{EL}(F, X \setminus \{x\}).$$

On the other hand, $\text{EL}(N_* F, Y \setminus \{\gamma_{e^{i\theta} \mathbf{w}}(t)\}) \leq \text{EL}(F, X \setminus \{x\})$ by Lemma 7.4. We conclude that no such M exists.

Equivalently, the concatenation of c with $\gamma_{e^{i\theta} \mathbf{w}}([0, t])$ does not belong to $\text{Blob}(x, X, Y, h)$. Since this holds for every $\theta \in [-\theta_0, \theta_0]$ and every $t \in (0, \delta)$, the open hyperbolic sector of radius $\delta \|\vec{v}\|_{\tilde{Y}}$ (the norm is with respect to the hyperbolic metric) and

angle θ_0 on either side \vec{v} in \tilde{Y} is disjoint from $\text{Blob}(x, X, Y, h)$. Thus \vec{v} is normal to $\text{Blob}(x, X, Y, h)$, as $\theta_0 \in (0, \frac{\pi}{2})$ was arbitrary. \square

As a corollary, we obtain the converse of Proposition 6.9.

Corollary 8.6. *Assume hypothesis (\mathbf{H}') and let $x \in X$. If $f : X \rightarrow Y$ is a slit mapping rel x homotopic to h , then $\text{lift}_x(f) \in \partial \text{Blob}(x, X, Y, h)$.*

Proof. Let $c = \text{lift}_x(f)$. Since f is conformal, c belongs to $\text{Blob}(x, X, Y, h)$. Let ψ be a terminal quadratic differential for f rel x . If ψ does not have a simple pole at $y = f(x)$ then f is a slit mapping, contradicting the assumption that $\text{CEmb}(X, Y, h)$ does not contain any. Thus ψ has a simple pole at y so that there exists a vector $\mathbf{v} \in T_c \tilde{Y}$ which is vertical for $\pi_Y^* \psi$. We can apply the same reasoning as in the proof of Proposition 8.5 to conclude that \vec{v} is normal to $\text{Blob}(x, X, Y, h)$. In particular, $c \in \partial \text{Blob}(x, X, Y, h)$. \square

We will also prove the converse of Proposition 8.5, namely that every vector normal to $\text{Blob}(x, X, Y, h)$ is vertical with respect to some realizing quadratic differential. The first step is to strengthen Lemma 8.3 by allowing the point $y \in Y$ and the quadratic differential ψ to vary inside a compact set. We have to be a little more careful about how to compare two surfaces $Y \setminus \{y_1\}$ and $Y \setminus \{y_2\}$ for $y_1, y_2 \in Y$ though.

Remark. Let S be a finite hyperbolic surface and let $s \in S$. Given $Y \in \mathcal{T}(S)$ and $y \in Y$, there is no canonical way to define a marking $S \setminus \{s\} \rightarrow Y \setminus \{y\}$. However, if $b \in Y$ is some basepoint, \tilde{Y} is the universal cover of Y based at b , and $f : S \setminus \{s\} \rightarrow Y \setminus \{b\}$ is a fixed marking, then point-pushing provides a continuous map $\Theta : \tilde{Y} \rightarrow \mathcal{T}(S \setminus \{s\})$ as follows. We can represent any $c \in \tilde{Y}$ by a smooth immersed curve $\gamma : [0, 1] \rightarrow Y$ with $\gamma(0) = b$. Let \mathcal{P}_γ be a quasiconformal point-pushing diffeomorphism along γ , then define $\Theta(c)$ as the marked surface $\mathcal{P}_\gamma \circ f : S \setminus \{s\} \rightarrow Y \setminus \{\pi_Y(c)\}$ in $\mathcal{T}(S \setminus \{s\})$. We use these markings implicitly below.

Lemma 8.7. *Let Y be a finite hyperbolic surface. Let $\theta_0 \in (0, \frac{\pi}{2})$ and let K be a compact set of pairs (c, ψ) where $c \in \tilde{Y}$ and $\psi \in \mathcal{Q}(Y \setminus \{\pi_Y(c)\})$ has a simple pole at $\pi_Y(c)$. There exists a number $\delta > 0$ such that $\forall (c, \psi) \in K, \forall \theta \in [-\theta_0, \theta_0]$ and $\forall t \in (0, \delta)$ we have*

$$\text{EL}(\mathcal{F}^h \psi, Y \setminus \{\gamma_{e^{i\theta} \vec{v}}(t)\}) > \text{EL}(\mathcal{F}^h \psi, Y \setminus \{\pi_Y(c)\})$$

where \vec{v} is the vertical vector for ψ at $\pi_Y(c)$ rescaled to have norm 1 with respect to the hyperbolic metric.

Proof. The idea of the proof is the same as for Lemma 8.3. If the statement fails, then there exist convergent sequences $t_n \searrow 0, (c_n, \psi_n) \rightarrow (c, \psi) \in K, \theta_n \rightarrow \theta \in [-\theta_0, \theta_0]$ such that

$$\text{EL}(\mathcal{F}^h \psi_n, Y \setminus \{\gamma_{e^{i\theta_n} \vec{v}_n}(t_n)\}) \leq \text{EL}(\mathcal{F}^h \psi_n, Y \setminus \{\pi_Y(c_n)\})$$

for every $n \in \mathbb{N}$, where \vec{v}_n is the unit vertical vector for ψ_n at $\pi_Y(c_n)$. Since $\psi_n \rightarrow \psi$ as $n \rightarrow \infty$ we have that $\vec{v}_n \rightarrow \vec{v}$ where \vec{v} is the unit vertical vector for ψ at $\pi_Y(c)$. If μ_n denotes the derivative of the path $t \mapsto Y \setminus \{\gamma_{e^{i\theta_n} \vec{v}_n}(t)\}$ at $t = 0$ and μ the derivative of $t \mapsto Y \setminus \{\gamma_{e^{i\theta} \vec{v}}(t)\}$ at $t = 0$, then $\mu_n \rightarrow \mu$ as $n \rightarrow \infty$ since $\theta_n \rightarrow \theta$ and $\vec{v}_n \rightarrow \vec{v}$. In the limit, we have

$$2\langle \mu, \psi \rangle = -2\pi \text{Re}[\text{Res}_{\pi_Y(c)}(\psi e^{i\theta} \vec{v})] = -2\pi \cos(\theta) \text{Re}[\text{Res}_{\pi_Y(c)}(\psi \vec{v})] > 0.$$

The difference with Lemma 8.3 is that the measured foliations $\mathcal{F}^h\psi_n$ are not constant, so we are looking at a sequence of functions rather than a single one.

Let $b \in Y$ be our basepoint for the universal cover $\pi_Y : \tilde{Y} \rightarrow Y$. Recall that any $c \in \tilde{Y}$ determines a marking $f_c : Y \setminus \{b\} \rightarrow Y \setminus \{\pi_Y(c)\}$. Thus the measured foliations $F_n = f_{c_n}^*(\mathcal{F}^h\psi_n)$ and $F = f_c^*(\mathcal{F}^h\psi)$ all live on the same surface $S = Y \setminus \{b\}$. Moreover $F_n \rightarrow F$ in measure since $(c_n, \psi_n) \rightarrow (c, \psi)$ as $n \rightarrow \infty$. Given a measured foliation G on S and $Z \in \mathcal{T}(S)$, let $\Phi(G, Z)$ be the quadratic differential on Z whose horizontal foliation is measure equivalent to G . By Hubbard–Masur [HM79], the map $\mathcal{MF}(S) \times \mathcal{T}(S) \rightarrow T^*\mathcal{T}(S)$ sending (G, Z) to $\Phi(G, Z)$ is continuous. Thus the function

$$Z \mapsto \Lambda_G(Z) := \text{EL}(G, Z) = \|\Phi(G, Z)\|$$

depends continuously on G . Furthermore, its derivative

$$(Z, \nu) \mapsto d_Z \Lambda_G(\nu) = 2\langle \nu, \Phi(G, Z) \rangle$$

is continuous on $T\mathcal{T}(S)$ and depends continuously on G .

Let $\{x_n\}_{n=1}^\infty$ and $\{y_n\}_{n=1}^\infty$ be sequences in \mathbb{R}^k such that $x_n \rightarrow 0$, $y_n \rightarrow 0$ and $\frac{y_n - x_n}{|y_n - x_n|} \rightarrow v$ as $n \rightarrow \infty$. Suppose that $g_n, g : \mathbb{R}^k \rightarrow \mathbb{R}$ are C^1 functions such that $g_n \rightarrow g$ and $dg_n \rightarrow dg$ uniformly on compact sets as $n \rightarrow \infty$. Then

$$\frac{g_n(y_n) - g_n(x_n)}{|y_n - x_n|} \rightarrow dg(v)$$

as $n \rightarrow \infty$. One can show this using the fundamental theorem of calculus along the line segment between x_n and y_n for instance. By the previous paragraph, this criterion applies to the extremal length functions Λ_{F_n} and Λ_F on $\mathcal{T}(S)$ (which is diffeomorphic to some \mathbb{R}^k) to conclude that

$$\frac{\text{EL}(\mathcal{F}^h\psi_n, Y \setminus \{\gamma_{e^{i\theta_n} \bar{v}_n}(t_n)\}) - \text{EL}(\mathcal{F}^h\psi_n, Y \setminus \{\pi_Y(c_n)\})}{t_n} \rightarrow 2\langle \mu, \psi \rangle$$

as $n \rightarrow \infty$. This is a contradiction, since left-hand side is non-positive for every $n \in \mathbb{N}$ while the right-hand side is positive. \square

The next observation that we need is that that any non-zero limit of a sequence of terminal quadratic differentials is a terminal quadratic differential for the limiting Teichmüller embedding.

Lemma 8.8. *Let $h : X \rightarrow Y$ be a generic embedding between finite Riemann surfaces and let $x \in X$. Let $f_n : X \rightarrow Y$ be a Teichmüller embedding rel x homotopic to h with unit norm terminal quadratic differential ψ_n . Suppose that $\text{lift}_x(f_n)$ converges to some $c \in \tilde{Y}$. Then there is a Teichmüller embedding f rel x homotopic to h with unit norm terminal quadratic differential ψ such that $f_n \rightarrow f$ and $\psi_n \rightarrow \psi$ as $n \rightarrow \infty$ after passing to a subsequence.*

Proof. Let K_n be the dilatation of f_n . By hypothesis, the marked surface $\Theta(\text{lift}_x(f_n))$ converges to $\Theta(c)$ in the Teichmüller space \mathcal{T} of Y minus a point, where $\Theta : \tilde{Y} \rightarrow \mathcal{T}$ is the map from Remark 8. In particular, the sequence K_n is bounded so we may assume it converges to some $K \geq 1$. By Lemma 5.4, there is a Teichmüller embedding f rel x homotopic to h such that $f_n \rightarrow f$ after passing to a subsequence.

Let $c_n = \text{lift}_x(f_n)$. The set $\{c\} \cup \{c_n\}_{n=1}^\infty$ is compact in \tilde{Y} and so is its image by Θ in \mathcal{T} . Since the set of unit cotangent vectors over a compact set in \mathcal{T} is compact, there is a quadratic differential $\psi \in \mathcal{Q}(Y \setminus \pi_Y(c))$ such that $\psi_n \rightarrow \psi$ after passing

to a subsequence. We have to show that ψ is a terminal quadratic differential for f .

Let φ_n be the initial quadratic differential of f_n corresponding to ψ_n . We may assume that φ_n converges to some $\varphi \in \mathcal{Q}^+(X \setminus \{x\})$ since its norm $1/K_n$ is bounded above and below. Suppose that $z_0 \in X \setminus \{x\}$ is not a zero of φ . Then there is a compact simply connected neighborhood U of z_0 on which φ does not vanish. If n is large enough, then φ_n does not have any zeros in U either. If $V = f(U)$, then ψ and ψ_n do not have zeros in V when n is large enough. We can choose square roots consistently so that $\sqrt{\varphi_n} \rightarrow \sqrt{\varphi}$ uniformly on U and $\sqrt{\psi_n} \rightarrow \sqrt{\psi}$ uniformly on V . Then for every $z \in U$ and every n we have

$$\int_{f_n(z_0)}^{f_n(z)} \sqrt{\psi_n} = K_n \operatorname{Re} \left(\int_{z_0}^z \sqrt{\varphi_n} \right) + i \operatorname{Im} \left(\int_{z_0}^z \sqrt{\varphi_n} \right)$$

since f_n restricted to $X \setminus \{x\}$ is a Teichmüller embedding of dilatation K_n with respect to φ_n and ψ_n . Taking the limit as $n \rightarrow \infty$ we get

$$\int_{f(z_0)}^{f(z)} \sqrt{\psi} = K \operatorname{Re} \left(\int_{z_0}^z \sqrt{\varphi} \right) + i \operatorname{Im} \left(\int_{z_0}^z \sqrt{\varphi} \right)$$

which means that f is locally of Teichmüller form with respect to φ and ψ . \square

Lastly, we will need the fact that the set of vectors which are vertical for some realizing quadratic differential at a given point is convex.

Lemma 8.9. *Assume hypothesis (\mathbf{H}') . Let $x \in X$, let $c \in \partial \operatorname{Blob}(x, X, Y, h)$ and suppose that $\psi_0, \psi_1 \in \mathcal{Q}^+(Y \setminus \{\pi_Y(c)\})$ realize c . If \mathbf{v}_0 and \mathbf{v}_1 are vertical for ψ_0 and ψ_1 respectively at $y = \pi_Y(c)$, then there exist $\alpha, \beta > 0$ such that $\mathbf{v}_0 + \mathbf{v}_1$ is vertical for $\alpha\psi_0 + \beta\psi_1$.*

Proof. Take $\alpha = -\frac{|\mathbf{v}_0|}{|\mathbf{v}_1|} \operatorname{Res}_y(\psi_1 \mathbf{v}_1)$ and $\beta = -\frac{|\mathbf{v}_1|}{|\mathbf{v}_0|} \operatorname{Res}_y(\psi_0 \mathbf{v}_0)$. A calculation shows that $\operatorname{Res}_y((\alpha\psi_0 + \beta\psi_1)(\mathbf{v}_0 + \mathbf{v}_1)) \leq 0$. By Lemma 6.10, the quadratic differential $\alpha\psi_0 + \beta\psi_1$ has a simple pole at y . This implies that $\mathbf{v}_0 + \mathbf{v}_1 \neq 0$ and hence that $\operatorname{Res}_y((\alpha\psi_0 + \beta\psi_1)(\mathbf{v}_0 + \mathbf{v}_1)) < 0$. \square

We are now able to show that normal vectors are vertical.

Proposition 8.10 (Normal vectors are vertical). *Assume (\mathbf{H}') and let $x \in X$. Suppose that \mathbf{v} is normal to $\operatorname{Blob}(x, X, Y, h)$ at some point c . Then there exists a quadratic differential $\psi \in \mathcal{Q}^+(Y \setminus \pi_Y(c))$ realizing c such that \mathbf{v} is vertical for $\pi_Y^* \psi$.*

Proof. Let $V_c \subset T_c \tilde{Y}$ denote the set of vectors which are vertical for the pull-back of some quadratic differential realizing c . By Lemma 6.10 and Lemma 8.9, V_c is convex. Moreover $V_c \cup \{0_c\}$ is closed by Lemma 8.8. Suppose that \mathbf{v} is not in V_c . Then there is an open half-plane H through the origin in $T_c \tilde{Y}$ containing V_c such that \mathbf{v} is not in the closure \bar{H} . Let c_n be a sequence converging to c along the geodesic ray r which is normal to \bar{H} at c . Since r makes an angle strictly less than $\pi/2$ with \mathbf{v} and since \mathbf{v} is normal to $\operatorname{Blob}(x, X, Y, h)$, we may assume that c_n is not in $\operatorname{Blob}(x, X, Y, h)$ for any n . Let f_n be the Teichmüller embedding rel x realizing c_n provided by Lemma 6.8 and let ψ_n be its terminal quadratic differential, normalized to have norm 1.

We may assume that ψ_n converges to some $\psi \in \mathcal{Q}^+(Y \setminus \pi_Y(c))$ and that f_n converges to a slit mapping f rel x with terminal quadratic differential ψ by Lemma

8.8. Indeed, the limit f is conformal since c belongs to $\text{Blob}(x, X, Y, h)$. If ψ is holomorphic at $\pi_Y(c)$, then f is a slit mapping on X , contrary to the assumption that $\text{CEmb}(X, Y, h)$ does not contain any. Therefore ψ has a simple pole at $\pi_Y(c)$, which implies that ψ_n has a simple pole at $\pi_Y(c_n)$ for all but finitely many indices. Let \mathbf{w}_n be the unit vertical vector for $\pi_Y^*\psi_n$ at c_n . Then \mathbf{w}_n converges to the unit vertical vector \mathbf{w} for $\pi_Y^*\psi$ at c . We have $\mathbf{w} \in V_c \subset H$. We will see that this yields a contradiction.

Let $\phi \in (0, \frac{\pi}{2}]$ be the angle between \mathbf{w} and the line ∂H , and let $\theta = \frac{\pi}{2} - \frac{\phi}{2}$. By Lemma 8.7, there exists a number $\delta > 0$ such that $\forall n \in \mathbb{N}$, $\forall \alpha \in [-\theta, \theta]$ and $\forall t \in (0, \delta)$ we have

$$\text{EL}(\mathcal{F}^h\psi_n, Y \setminus \{\pi_Y(\gamma_{e^{i\alpha}\bar{w}_n}(t))\}) > \text{EL}(\mathcal{F}^h\psi_n, Y \setminus \{\pi_Y(c_n)\}).$$

On the other hand, since $\mathbf{w}_n \rightarrow \mathbf{w}$, the angle between \mathbf{w}_n and the geodesic between c_n and c converges to the angle between \bar{w} and the ray opposite to r , that is, to $\frac{\pi}{2} - \phi < \theta$. Thus if n is large enough then c belongs to the open hyperbolic sector of radius δ and angle θ on either side of \bar{w}_n so that

$$\text{EL}(\mathcal{F}^h\psi_n, Y \setminus \{\pi_Y(c)\}) > \text{EL}(\mathcal{F}^h\psi_n, Y \setminus \{\pi_Y(c_n)\}).$$




FIGURE 8. The point c is eventually contained in the sector of angle θ and radius δ about the vector \mathbf{w}_n .

Let φ_n be the initial quadratic differential of f_n corresponding to ψ_n and let $K_n > 1$ be the dilatation of f_n . By Lemma 7.5, $\mathcal{F}^h\psi_n$ is measure equivalent to $(f_n)_*\mathcal{F}^h\varphi_n$ on $Y \setminus \pi_Y(c_n)$. If we let $\mathcal{P}_n : Y \setminus \{\pi_Y(c_n)\} \rightarrow Y \setminus \{\pi_Y(c)\}$ be the change of marking given by point-pushing, we get

$$\begin{aligned} \text{EL}((\mathcal{P}_n \circ f_n)_*\mathcal{F}^h\varphi_n, Y \setminus \{\pi_Y(c)\}) &= \text{EL}((\mathcal{P}_n)_*\mathcal{F}^h\psi_n, Y \setminus \{\pi_Y(c)\}) \\ &> \text{EL}(\mathcal{F}^h\psi_n, Y \setminus \{\pi_Y(c_n)\}) \\ &= K_n \text{EL}(\mathcal{F}^h\varphi_n, X \setminus \{x\}) \\ &> \text{EL}(\mathcal{F}^h\varphi_n, X \setminus \{x\}) \end{aligned}$$

when n is large enough, by the above inequality and Lemma 7.5. However, since f and $\mathcal{P}_n \circ f_n$ are homotopic tame embeddings from $X \setminus \{x\}$ to $Y \setminus \{\pi_Y(c)\}$ and f is conformal, we have

$$\begin{aligned} \text{EL}((\mathcal{P}_n \circ f_n)_*\mathcal{F}^h\varphi_n, Y \setminus \{\pi_Y(c)\}) &= \text{EL}(f_*\mathcal{F}^h\varphi_n, Y \setminus \{\pi_Y(c)\}) \\ &\leq \text{EL}(\mathcal{F}^h\varphi_n, X \setminus \{x\}) \end{aligned}$$

by Lemma 7.6 and Lemma 7.4. This is a contradiction, from which we conclude that \mathbf{v} belongs to V_c . \square

We finally come to the main result of this section, which is that the blob is semi-smooth.

Definition 8.11. A closed subset $B \subset \mathbb{C}$ is *semi-smooth* if

- for every $z \in \partial B$, the set of normal vectors to B at z is non-empty and convex;
- any non-zero limit of vectors normal to B (not necessarily at the same point) is normal to B .

For example, a 2-dimensional manifold $M \subset \mathbb{C}$ with C^1 boundary is semi-smooth, as is any convex set with non-empty interior. A line segment is not semi-smooth because at interior points the set of normal vectors is not convex.

Theorem 8.12 (The blob is semi-smooth). *Assume hypothesis (\mathbf{H}') . Then $\text{Blob}(x, X, Y, h)$ is semi-smooth for any $x \in X$.*

Proof. For every $c \in \partial \text{Blob}(x, X, Y, h)$ the set of vectors which are normal to $\text{Blob}(x, X, Y, h)$ at c coincides with the set V_c of vectors which are vertical for the pull-back of some quadratic differential realizing c , according to Propositions 8.5 and 8.10. The set V_c is non-empty by Proposition 6.9 and convex by Lemma 6.10 and Lemma 8.9. Suppose that $\{c_n\}_{n=1}^\infty \subset \partial \text{Blob}(x, X, Y, h)$ is such that $c_n \rightarrow c$, that \mathbf{v}_n is vertical for $\pi_Y^* \psi_n$ where ψ_n realizes c_n , and that $\mathbf{v}_n \rightarrow \mathbf{v} \neq 0$ as $n \rightarrow \infty$. By Lemma 5.4 we can rescale ψ_n and pass to a subsequence such that it converges to some ψ realizing c . We have $\text{Res}_c((\pi_Y^* \psi) \mathbf{v}) \leq 0$ since $\text{Res}_{c_n}((\pi_Y^* \psi_n) \mathbf{v}_n) < 0$ for every n . Moreover, ψ must have a simple pole at $\pi_Y(c)$ for otherwise $\text{CEmb}(X, Y, h)$ would contain a slit mapping. This means that $\text{Res}_c((\pi_Y^* \psi) \mathbf{v}) \neq 0$. Therefore \mathbf{v} is vertical for $\pi_Y^* \psi$, hence normal to $\text{Blob}(x, X, Y, h)$ by Proposition 8.5. \square

9. THE BLOB IS A DISK

In this section, we complete the proof of Theorem 1.7 which says that the blob is homeomorphic to a closed disk under hypothesis (\mathbf{H}') . We first prove that every semi-smooth set is a manifold.

Theorem 9.1 (The aquatic theorem). *Every closed semi-smooth subset of \mathbb{C} is a 2-dimensional manifold with boundary.*

Proof. Let B be a closed semi-smooth set. Every interior point of B has a neighborhood homeomorphic to an open subset of \mathbb{C} , namely the interior of B . Thus we only have to show that every boundary point $z \in \partial B$ has a neighborhood homeomorphic to a half-disk. By applying an isometry of the plane, we may assume that $z = 0$ and that the vector i bisects the cone N_0 of vectors normal to B at 0. Let ϕ be half the angle of N_0 , let $\alpha = \phi + \frac{\pi}{2}$ and let $\beta = \pi - \alpha$.

Since B is semi-smooth, we have $\phi < \frac{\pi}{2}$ and hence $\beta > 0$. Moreover, for every $\theta \in (0, \alpha)$ there exists a $\delta > 0$ such that the open angular sector $\nabla(i, \theta, \delta)$ is disjoint from B (notation is as in the previous section). We now show the existence of closed sectors pointing downwards contained in B .

Claim 9.2. *For every $\theta \in (0, \beta)$, there exists a $\delta > 0$ such that the closed angular sector $\overline{\nabla(-i, \theta, \delta)}$ is contained in B .*




FIGURE 9. The cone of normal vectors N_0 and the angles α and β .

Proof of Claim. Suppose not. Then there exists a $\theta \in (0, \beta)$ and a sequence $\delta_n \searrow 0$ for which the corresponding angular sector $S_n = \overline{\nabla(-i, \theta, \delta_n)}$ intersects the complement of B for every n . Let D_n be a closed disk in S_n disjoint from B . Slide the center of D_n in a straight line towards 0 until the disk first hits B , and let D_n^* be the resulting disk. The intersection points of D_n^* with B all lie on the half of ∂D_n^* which is closest to 0. Let z_n be any point in this intersection. Then z_n is on the boundary of B and the unit vector \mathbf{v}_n pointing from z_n to the center of D_n^* is normal to B . Since S_n shrinks to 0, we have $z_n \rightarrow z$. Each vector \mathbf{v}_n makes an angle at most $\theta + \frac{\pi}{2}$ with the downward direction. Therefore, the vectors \mathbf{v}_n can only accumulate onto vectors forming an angle at least $\beta - \theta$ with the cone N_0 . This contradicts the hypothesis that every limit of normal vectors is normal. \square




FIGURE 10. Bubbles floating to the surface of B .

Let $\theta^+ \in (0, \alpha)$, let $S^+ = \overline{\nabla(i, \theta^+, \delta^+)}$ be disjoint from B , let $\theta^- \in (0, \beta)$, and let $S^- = \overline{\nabla(-i, \theta^-, \delta^-)}$ be contained in B . Let $I \subset S^+$ be a compact horizontal segment symmetric about the vertical line through 0 and lying entirely above S^- . We define a map $p : I \rightarrow \partial B$ as follows. For $z \in I$, let z fall straight down until it first hits B , and let $p(z)$ be this first hitting point. Note that $p(x + iy) = x + iq(x, y)$ for some function q so that p is injective.

Claim 9.3. *The map p is continuous on some subinterval $J \subset I$ centered at the midpoint of I .*

Proof of Claim. It is easy to see that p is continuous at the midpoint $p^{-1}(0)$. This is because p keeps the x -coordinate unchanged and moreover, $p(z)$ is below S^+ and above S^- . Thus the y -coordinate of $p(z)$ converges to 0 as $z \rightarrow p^{-1}(0)$.




FIGURE 11. The vertical projection p is akin to rain falling on the surface of B . The projection is continuous (in fact Lipschitz) by the squeeze theorem.

Let $0 < \varepsilon < \beta/2$. By semi-smoothness, there exists a $\delta > 0$ such that if $w \in \partial B$ is within distance δ of 0, then every vector in N_w is within angle $\phi + \varepsilon$ of the upward direction. For every $w \in \partial B$ with $|w| < \delta$, let \mathbf{v}_w be the bisector of N_w and let ϕ_w be half the angle of N_w . For every $\theta_w^+ \in (0, \phi_w + \frac{\pi}{2})$ there is an open sector $S_w^+ = \overline{\langle \mathbf{v}_w, \theta_w^+, \delta_w^+ \rangle}$ disjoint from B by definition of N_w . Since

$$\phi + \varepsilon = \alpha - \frac{\pi}{2} + \varepsilon < \alpha + \frac{\beta - \pi}{2} = \frac{\alpha}{2} < \frac{\pi}{2} \leq \phi_w + \frac{\pi}{2},$$

we may choose θ_w^+ so that S_w^+ contains the vertical direction in its span. By the previous Claim, there is also a closed sector $S_w^- = \overline{\langle -\mathbf{v}_w, \theta_w^-, \delta_w^- \rangle}$ contained in B for every $\theta_w^- \in (0, \frac{\pi}{2} - \phi_w)$. The angle that $-\mathbf{v}_w$ makes with the downward direction is equal to the angle that \mathbf{v}_w makes with the vertical direction, which is at most $\phi + \varepsilon - \phi_w$ hence strictly less than $\frac{\pi}{2} - \phi_w$. Thus we may choose θ_w^- so that S_w^- contains the downward direction in its interior.

By continuity of p at $p^{-1}(0)$, there is a closed interval $J \subset I$ centered at $p^{-1}(0)$ such that $p(J)$ is contained in the ball of radius δ about 0. Let $z \in J$, let $w = p(z)$, and let S_w^+ and S_w^- be angular sectors as described in the previous paragraph. Also let $K \subset S_w^+$ be a compact horizontal segment crossing the vertical line through w and lying entirely above S_w^- . By construction, the vertical segment from z to w intersects B only at w . Since B is closed, we may assume that the rectangle with bottom edge K and upper edge $L \subset I$ is disjoint from B , by making K shorter if necessary. For every $\zeta \in L$, the image $p(\zeta)$ is thus squeezed between S_w^+ and S_w^- , so that p is continuous at z . \square

Thus $p(J)$ is the graph of a continuous function. Let $0 < \delta < |J|/2$. For every $z \in J$ with $|x| < \delta$, draw the open vertical segment of length 2δ centered at $p(z)$, and let U_δ be the union of those segments. The continuity of p implies that U_δ is open.

Claim 9.4. *If δ is small enough, then the component of $U_\delta \setminus p(J)$ above $p(J)$ is disjoint from B and the component below $p(J)$ is contained in B .*

Proof of Claim. If δ is small enough, then the component C^+ of $U_\delta \setminus p(J)$ above $p(J)$ lies below J itself. By definition of p , for every $z \in J$ the open vertical segment between z and $p(z)$ is disjoint from B , so that C^+ is disjoint from B .

For the component lying below $p(J)$, we use the same idea as in the proof of the first Claim. Suppose that the result does not hold. Then there is a sequence $\delta_n \searrow 0$ such that for every n , there is a closed disk D_n contained in the component of $U_{\delta_n} \setminus p(J)$ below $p(J)$. Slide the center of D_n upwards until the disk first hits B , and let D_n^* be this hitting disk. Every intersection point of D_n^* with B is on the upper half of ∂D_n^* . Let z_n be any point in that intersection. Then z_n is on the boundary of B and the unit vector \mathbf{v}_n pointing from z_n towards the center of D_n^* is normal to B . As $n \rightarrow \infty$, we have $z_n \rightarrow 0$. Moreover, the vectors \mathbf{v}_n only accumulate onto vectors forming an angle at least $\frac{\pi}{2}$ with the upwards direction at 0, hence outside N_0 . This contradicts the semi-smoothness of B at 0. \square

By this last claim, if δ is small enough then $U_\delta \cap B$ is equal to the union of $U_\delta \cap p(J)$ with the component of $U_\delta \setminus p(J)$ below $p(J)$. This neighborhood of z in B is clearly homeomorphic to the rectangle $(-\delta, \delta) \times (0, \delta]$, which in turn is homeomorphic to a half-disk. Thus B is a 2-manifold with boundary. \square

Since $\text{Blob}(x, X, Y, h)$ is semi-smooth under appropriate hypotheses, it is a 2-manifold with boundary.

Corollary 9.5. *Assume hypothesis (\mathbf{H}') . Then $\text{Blob}(x, X, Y, h)$ is a compact, connected, 2-manifold with boundary for any $x \in X$.*

Proof. $\text{Blob}(x, X, Y, h)$ is compact and connected because $\text{CEmb}(X, Y, h)$ is (Lemma 3.1 and Theorem 1.6) and lift_x is continuous (Lemma 6.3). It is semi-smooth by Theorem 8.12, hence a 2-manifold with boundary by Theorem 9.1. \square

In particular, the blob is the closure of its interior. We use this to show that the blob depends continuously on parameters. We first need to define a topology on closed subsets of a space.

Let S be a topological space and let $\text{CL}(S)$ be the hyperspace of closed subsets of S . The *Fell topology* on $\text{CL}(S)$ is the topology generated by neighborhoods of the form $N(K, \mathcal{U})$, where $K \subset S$ is compact, \mathcal{U} is a finite collection of open subsets of S , and $N(K, \mathcal{U})$ is the set of $A \in \text{CL}(S)$ such that $A \cap K = \emptyset$ and $A \cap U \neq \emptyset$ for every $U \in \mathcal{U}$.

Theorem 9.6 (Fell [Fel62]). *For any topological space S , the hyperspace $\text{CL}(S)$ is compact. If S is locally compact, then $\text{CL}(S)$ is Hausdorff.*

If S is first-countable and Hausdorff, then a sequence $\{A_n\} \subset \text{CL}(S)$ converges to $A \in \text{CL}(S)$ if and only if

- for every $a \in A$, there exist $a_n \in A_n$ such that $a_n \rightarrow a$;
- for every sequence $\{a_n\}$ with $a_n \in A_n$, if $\{a_n\}$ accumulates onto $a \in S$, then $a \in A$.

We use the Fell topology on closed subsets of $\tilde{Y} \cong \mathbb{D}$ where Y is a finite hyperbolic surface. To prove convergence, we mostly rely on Fell's compactness theorem and the above criterion for sequences.

It is fairly clear that the blob depends upper semi-continuously on parameters. The same holds for its boundary.

Lemma 9.7. *Suppose that $h : X \rightarrow Y$ is a generic embedding between finite Riemann surfaces. Then $\text{Blob}(z, Z, Y, h)$ and $\partial \text{Blob}(z, Z, Y, h)$ depend upper semi-continuously on $Z \setminus \{z\} \in \mathcal{T}(X \setminus \{x\})$. More precisely, suppose that $Z_n \setminus \{z_n\} \rightarrow Z \setminus \{z\}$ in $\mathcal{T}(X \setminus \{x\})$, that $\text{Blob}(z_n, Z_n, Y, h) \rightarrow A$ in $\text{CL}(\tilde{Y})$, and that $\partial \text{Blob}(z_n, Z_n, Y, h) \rightarrow B$ in $\text{CL}(\tilde{Y})$ as $n \rightarrow \infty$. Then $A \subset \text{Blob}(z, Z, Y, h)$ and $B \subset \partial \text{Blob}(z, Z, Y, h)$.*

Proof. Let $c \in A$. By hypothesis there exist $c_n \in \text{Blob}(z_n, Z_n, Y, h)$ such that $c_n \rightarrow c$ as $n \rightarrow \infty$. Let $f_n \in \text{CEmb}(Z_n, Y, h)$ be such that $\text{lift}_{z_n}(f_n) = c_n$. Let $\sigma_n^* : Z \setminus \{z\} \rightarrow Z_n \setminus \{z_n\}$ be the quasiconformal homeomorphism homotopic to the change of marking with minimal dilatation K_n and let $\sigma_n : Z \rightarrow Z_n$ be its extension. By hypothesis, $K_n \rightarrow 1$ as $n \rightarrow \infty$. By Lemma 3.1, we may pass to a subsequence such that $f_n \circ \sigma_n$ converges to a conformal embedding $f : Z \rightarrow Y$ homotopic to h . Then

$$\text{lift}_z(f) = \lim_{n \rightarrow \infty} \text{lift}_z(f_n \circ \sigma_n) = \lim_{n \rightarrow \infty} \text{lift}_{z_n}(f_n) = \lim_{n \rightarrow \infty} c_n = c,$$

so that $c \in \text{Blob}(z, Z, Y, h)$.

Now let $c \in B$ and let $c_n \in \partial \text{Blob}(z_n, Z_n, Y, h)$ be such that $c_n \rightarrow c$. By Proposition 6.9, there exists a slit mapping $f_n \text{ rel } z_n$ from Z_n to Y homotopic to h such that $\text{lift}_{z_n}(f_n) = c_n$. By Lemma 8.8, we can pass to a subsequence such that f_n converges to some slit mapping $f \text{ rel } z$ from Z to Y . Then $\text{lift}_z(f) = c$ so that $c \in \text{Blob}(z, Z, Y, h)$. Moreover, $y \in \partial \text{Blob}(z, Z, Y, h)$ by Corollary 8.6. \square

We do not know if the blob moves continuously in general, but it does when there are no slit mappings at the limiting parameters.

Lemma 9.8. *Assume hypothesis **(H')** and let $x \in X$. If $X_n \setminus \{x_n\} \rightarrow X \setminus \{x\}$ in $\mathcal{T}(X \setminus \{x\})$ as $n \rightarrow \infty$, then $\text{Blob}(x_n, X_n, Y, h) \rightarrow \text{Blob}(x, X, Y, h)$ and $\partial \text{Blob}(x_n, X_n, Y, h) \rightarrow \partial \text{Blob}(x, X, Y, h)$ in $\text{CL}(\tilde{Y})$ as $n \rightarrow \infty$.*

Proof. By compactness of $\text{CL}(\tilde{Y})$, it suffices to prove that if $\text{Blob}(x_n, X_n, Y, h)$ converges to some closed set A and $\partial \text{Blob}(x_n, X_n, Y, h)$ converges to some closed set B as $n \rightarrow \infty$, then $A = \text{Blob}(x, X, Y, h)$ and $B = \partial \text{Blob}(x, X, Y, h)$.

We prove convergence of the blobs first. By Lemma 9.7, the inclusion $A \subset \text{Blob}(x, X, Y, h)$ holds. We claim that the interior of $\text{Blob}(x, X, Y, h)$ is contained in A . Let c be in the interior of $\text{Blob}(x, X, Y, h)$ and suppose that there is an infinite set $J \subset \mathbb{N}$ such that c is not contained in $\text{Blob}(x_n, X_n, Y, h)$ for every $n \in J$. Then for every $n \in J$, there exists a Teichmüller embedding $f_n \text{ rel } x_n$ with $\text{lift}_{x_n}(f_n) = c$. After passing to a subsequence in J , we get that $f_n \rightarrow f$ for some Teichmüller embedding $f \text{ rel } x$ by Lemma 8.8. We have $\text{lift}_x(f) = c$ by continuity of lift_x . By Corollary 8.6, c is in the complement of the interior of $\text{Blob}(x, X, Y, h)$. This is a contradiction, which means that c is contained in $\text{Blob}(x_n, X_n, Y, h)$ for all but finitely many indices, and hence $c \in A$. Since A is closed and $\text{Blob}(x, X, Y, h)$ is the closure of its interior, we have $\text{Blob}(x, X, Y, h) \subset A$ and hence $A = \text{Blob}(x, X, Y, h)$.

By Lemma 9.7, we have $B \subset \partial \text{Blob}(x, X, Y, h)$. Let $c \in \partial \text{Blob}(x, X, Y, h)$. Let U be any connected neighborhood of c in \tilde{Y} . We claim that if n is large enough, then U intersects both the complement of $\text{Blob}(x_n, X_n, Y, h)$ and the interior of $\text{Blob}(x_n, X_n, Y, h)$. Suppose on the contrary that U is contained in $\text{Blob}(x_n, X_n, Y, h)$ for every n in an infinite set $J \subset \mathbb{N}$. Then $U \subset A = \text{Blob}(x, X, Y, h)$, which is nonsense since c is on the boundary of $\text{Blob}(x, X, Y, h)$. Similarly, suppose

that U is contained in the complement of $\text{Blob}(x_n, X_n, Y, h)$ for every n in an infinite set $J \subset \mathbb{N}$. Then for every $z \in U$ and every $n \in J$ there is a Teichmüller embedding $f_n : X_n \rightarrow Y$ rel x_n homotopic to h such that $\text{lift}_{x_n}(f) = z$. By Lemma 5.4, f_n converges to a Teichmüller embedding f rel x after passing to a subsequence. Then $\text{lift}_x(f) = z$ so that $z \in \partial \text{Blob}(x, X, Y, h)$ by Corollary 8.6. This is a contradiction, which proves the claim. Let n be large enough so that U intersects both the interior and the complement of $\text{Blob}(x_n, X_n, Y, h)$. Since U is connected, it also intersects $\partial \text{Blob}(x_n, X_n, Y, h)$. Since U can be chosen arbitrarily small, this shows that $c \in B$. \square

Similarly, nested families of blobs move continuously. In what follows, $\{X_r\}_{r \in [0, \infty]}$ is a 1-parameter family of enlargements of X as in Section 5, and \mathfrak{m} is the associated modulus of extension.

Lemma 9.9. *Let $h : X \rightarrow Y$ be a generic embedding between finite Riemann surfaces, let R be the maximum of \mathfrak{m} on $\text{CEmb}(X, Y, h)$, and let $x \in X$. Then the maps $r \mapsto \text{Blob}(x, X_r, Y, h)$ and $r \mapsto \partial \text{Blob}(x, X_r, Y, h)$ are continuous on the interval $[0, R]$.*

Proof. We may assume that $R > 0$ since otherwise there is nothing to show. If $r \in [0, R)$ and $\rho \rightarrow r$, then $\text{Blob}(x, X_\rho, Y, h) \rightarrow \text{Blob}(x, X_r, Y, h)$ and $\partial \text{Blob}(x, X_\rho, Y, h) \rightarrow \partial \text{Blob}(x, X_r, Y, h)$ by Lemma 9.8, since $\text{CEmb}(X_r, Y, h)$ does not contain any slit mapping. It remains to prove continuity at $r = R$. By compactness of the hyperspace $\text{CL}(\tilde{Y})$ and Lemma 9.7, it suffices to show that if $r_n \nearrow R$, if $\text{Blob}(x, X_{r_n}, Y, h) \rightarrow A$, and if $\partial \text{Blob}(x, X_{r_n}, Y, h) \rightarrow B$ as $n \rightarrow \infty$, then $A \supset \text{Blob}(x, X_R, Y, h)$ and $B \supset \partial \text{Blob}(x, X_R, Y, h)$. By Proposition 5.5, every element of $\text{CEmb}(X_R, Y, h)$ is a slit mapping¹⁰ so that $\text{Blob}(x, X_R, Y, h)$ is homeomorphic to a point or an interval (Proposition 6.5). Let $c \in \text{Blob}(x, X_R, Y, h) = \partial \text{Blob}(x, X_R, Y, h)$. Then $c \in \text{Blob}(x, X_{r_n}, Y, h)$ for every n . Indeed, since $r_n \leq R$ there is a canonical inclusion $X_{r_n} \subset X_R$ which means that $\text{CEmb}(X_R, Y, h) \subset \text{CEmb}(X_{r_n}, Y, h)$ and similarly for the blobs. It follows that $c \in A$. Let U be a connected neighborhood of c in \tilde{Y} . Then U intersects $\text{Blob}(x, X_{r_n}, Y, h)$ since $c \in \text{Blob}(x, X_{r_n}, Y, h)$. Thus U intersects the interior of $\text{Blob}(x, X_{r_n}, Y, h)$ because $\text{Blob}(x, X_{r_n}, Y, h)$ is the closure of its interior. Suppose that U is contained in $\text{Blob}(x, X_{r_n}, Y, h)$ for every n in an infinite set $J \subset \mathbb{N}$. Then U is contained in A and hence in $\text{Blob}(x, X_R, Y, h)$. This is absurd since $\text{Blob}(x, X_R, Y, h)$ has empty interior. Thus U intersects $\partial \text{Blob}(x, X_{r_n}, Y, h)$ for all large enough n and hence $c \in B$. \square

We use continuity to show that the blob has no holes and is thus homeomorphic to a closed disk, under hypothesis **(H')**.

Proof of Theorem 1.7. We will show that the complement of $\text{Blob}(x, X, Y, h)$ is connected, which is sufficient in view of Corollary 9.5. Let z_1 and z_2 be any two points in $\tilde{Y} \setminus \text{Blob}(x, X, Y, h)$. Let R be the maximum value of the modulus of extension \mathfrak{m} . Note that z_1 and z_2 are contained in $\tilde{Y} \setminus \text{Blob}(x, X_\rho, Y, h)$ for every $\rho \in [0, R]$ as the blobs are nested. Let r be the infimum of the set of $\rho \in [0, R]$ such that z_1 and z_2 are in the same component of $\tilde{Y} \setminus \text{Blob}(x, X_\rho, Y, h)$. The set of such ρ is

¹⁰This is unless Y is the triply-punctured sphere, in which case $R = \infty$ so that $\text{CEmb}(X_R, Y, h)$ contains only one element anyway by Lemma 2.2.

non-empty since $\text{Blob}(x, X_R, Y, h)$ is a point or a compact interval, and hence has connected complement.

Suppose that z_1 and z_2 are in different components of $\tilde{Y} \setminus \text{Blob}(x, X_r, Y, h)$. Then $r < R$. In particular, $\text{Blob}(x, X_r, Y, h)$ is a compact 2-manifold so that each boundary component of $\text{Blob}(x, X_r, Y, h)$ is a simple closed curve. Let C_1 be the component of $\partial \text{Blob}(x, X_r, Y, h)$ surrounding z_1 , let C_2 be the one surrounding z_2 , and let γ be a simple closed curve in the interior of $\text{Blob}(x, X_r, Y, h)$ separating C_1 from C_2 . For all ρ close enough to r we have that $\partial \text{Blob}(x, X_\rho, Y, h)$ is disjoint from γ . On the other hand, there is a sequence $\rho_n \searrow r$ such that z_1 and z_2 belong to the same component of $\tilde{Y} \setminus \text{Blob}(x, X_{\rho_n}, Y, h)$. Let α_n be a path in $\tilde{Y} \setminus \text{Blob}(x, X_{\rho_n}, Y, h)$ connecting z_1 and z_2 . For every n , α_n intersects γ , say at w_n . Since γ is compact, we may pass so a subsequence so that $w_n \rightarrow w$ for some $w \in \gamma$. Now w is in the interior of $\text{Blob}(x, X_r, Y, h)$. Let U be an open disk centered at w whose closure is contained in the interior of $\text{Blob}(x, X_r, Y, h)$. Since $\text{Blob}(x, X_{\rho_n}, Y, h) \rightarrow \text{Blob}(x, X_r, Y, h)$ as $n \rightarrow \infty$, the open set U must intersect $\text{Blob}(x, X_{\rho_n}, Y, h)$ for all large enough n . Since $w_n \in \gamma \setminus \text{Blob}(x, X_{\rho_n}, Y, h)$ and since $\gamma \cup U$ is connected, the intersection of $\gamma \cup U$ with $\partial \text{Blob}(x, X_{\rho_n}, Y, h)$ is non-empty. Let ζ_n be in the intersection. After passing to a subsequence, ζ_n converges to some point ζ in $\gamma \cup \bar{U}$. This is a contradiction since $\partial \text{Blob}(x, X_{\rho_n}, Y, h) \rightarrow \partial \text{Blob}(x, X_r, Y, h)$ as $n \rightarrow \infty$ but $\gamma \cup \bar{U}$ is disjoint from $\partial \text{Blob}(x, X_r, Y, h)$. Therefore z_1 and z_2 belong to the same component of $\tilde{Y} \setminus \text{Blob}(x, X_r, Y, h)$.

Suppose that $r > 0$. Let γ be a path joining z_1 to z_2 in $\tilde{Y} \setminus \text{Blob}(x, X_r, Y, h)$. Since γ is compact and $\text{Blob}(x, X_\rho, Y, h)$ depends continuously on ρ , the two are disjoint for all ρ sufficiently close to r . Then z_1 and z_2 belong to the same component of $\tilde{Y} \setminus \text{Blob}(x, X_\rho, Y, h)$ for all $\rho < r$ sufficiently close to r , which contradicts the minimality of r . We conclude that $r = 0$ and that z_1 and z_2 belong to the same component of $\tilde{Y} \setminus \text{Blob}(x, X_0, Y, h)$. Since z_1 and z_2 were arbitrary, the complement of $\text{Blob}(x, X, Y, h) = \text{Blob}(x, X_0, Y, h)$ is connected. Thus $\text{Blob}(x, X, Y, h)$ is homeomorphic to a closed disk. \square

10. THE DEFORMATION RETRACTION

The goal of this section is to prove that $\text{CEmb}(X, Y, h)$ is contractible under (\mathbf{H}') , which is the main case of Theorem 1.2. Recall that (\mathbf{H}') stands for the hypothesis that $h : X \rightarrow Y$ is a generic embedding between finite Riemann surfaces (hence X and Y are hyperbolic), X has non-empty ideal boundary, and $\text{CEmb}(X, Y, h)$ is non-empty and does not contain any slit mapping.

Fix once and for all a countable dense set $\{x_1, x_2, x_3, \dots\}$ of distinct points in X and a universal covering map $\pi_Y : \mathbb{D} \rightarrow Y$. For each $n \in \mathbb{N}$, let $b_n \in \mathbb{D}$ be such that $\pi_Y(b_n) = h(x_n)$ and define a lift $L_n = \text{lift}_{x_n} : \text{Map}(X, Y, h) \rightarrow \mathbb{D}$ as in Section 6. For each $f \in \text{Map}(X, Y, h)$, $L_n(f)$ is the endpoint of the lift of $t \mapsto H(x_n, t)$ based at b_n , where H is any homotopy from h to f . Let $\text{Blob}(x_j) = L_j(\text{CEmb}(X, Y, h))$.

Let us also fix some $F \in \text{CEmb}(X, Y, h)$ which maximizes the modulus of extension \mathfrak{m} from Section 5. Note that $\mathfrak{m}(F) > 0$ by the hypothesis that $\text{CEmb}(X, Y, h)$ does not contain any slit mapping. It follows that $L_1(F)$ does not lie on the boundary of $\text{Blob}(x_1)$ by Proposition 6.9. We will construct a (strong) deformation retraction of $\text{CEmb}(X, Y, h)$ into $\{F\}$.

Given any $f \in \text{CEmb}(X, Y, h)$, we define a sequence of paths $\gamma_n : [0, 1] \rightarrow \mathbb{D}$ inductively as follows. Let $G[1] : \mathbb{D} \rightarrow \text{Blob}(x_1)$ be the Riemann map with

$G[1](0) = L_1(F)$ and $G[1]'(0) > 0$, and let

$$\gamma_1(t) = \begin{cases} L_1(f) & \text{if } t \in [0, 1/2) \\ G[1]((2-2t)G[1]^{-1}(L_1(f))) & \text{if } t \in [1/2, 1]. \end{cases}$$

In words, γ_1 stays at $L_1(f)$ for half the time and then moves at constant speed along the conformal ray towards the ‘‘center’’ $L_1(F)$ of $\text{Blob}(x_1)$. In particular, $\gamma_1(t)$ belongs to $\text{Blob}(x_1)$ for every $t \in [0, 1]$ so that there exists some $g \in \text{CEmb}(X, Y, h)$ such that $L_1(g) = \gamma_1(t)$.

Let $n \geq 2$. Suppose that paths $\gamma_1, \dots, \gamma_{n-1}$ have been defined in such a way that

- the points $\pi_Y(\gamma_1(t)), \dots, \pi_Y(\gamma_{n-1}(t))$ are distinct for every $t \in [0, 1]$;
- γ_j is constant on the interval $[0, 2^{-j}]$ for every $j \in \{1, \dots, n-1\}$;
- $\gamma_j(0) = L_j(f)$ and $\gamma_j(1) = L_j(F)$ for every $j \in \{1, \dots, n-1\}$.

Then let

$$X[n] = X \setminus \{x_1, \dots, x_{n-1}\}, \quad Y[n, t] = Y \setminus \{\pi_Y(\gamma_1(t)), \dots, \pi_Y(\gamma_{n-1}(t))\},$$

and let $h[n, t] = \mathcal{P}[n, t] \circ h$ where $\mathcal{P}[n, t] : Y \rightarrow Y$ is a multi-point-pushing diffeomorphism chosen so that $L_j(h[n, t]) = \gamma_j(t)$ for every $j \in \{1, \dots, n-1\}$ and every $t \in [0, 1]$. The embedding $h[n, t]$ is generic for every t . Define

$$E[n, t] = \text{CEmb}(X[n], Y[n, t], h[n, t]) \quad \text{and} \quad \text{Blob}[n, t] = L_n(E[n, t]).$$

We assume that $E[n, t]$ is non-empty for every $t \in [0, 1]$ as part of the induction hypothesis. Note that $\text{Blob}[n, t]$ is either a closed disk or a point. Indeed, if $E[n, t]$ contains a slit mapping then it is homeomorphic to a point or an interval by Theorem 1.4. But since $h[n, t]$ has an essential puncture, there is at most one slit mapping in $E[n, t]$ by Remark 4. If $E[n, t]$ does not contain any slit mapping, then $\text{Blob}[n, t]$ is a closed disk by Theorem 1.7. Also, since we chose the paths $\gamma_1, \dots, \gamma_{n-1}$ to be constant on $[0, 2^{1-n}]$, the set $\text{Blob}[n, t]$ does not change for t in that interval. The next step is to choose a conformal center for $\text{Blob}[n, t]$.

Lemma 10.1. *For every $t \in [0, 1]$ there is a unique map $g[n, t]$ maximizing \mathfrak{m} within $E[n, t]$. The map $t \mapsto g[n, t]$ is continuous, constant on $[0, 2^{1-n}]$, and satisfies $g[n, 1] = F$.*

Proof. The map \mathfrak{m} is upper semi-continuous on the compact space $E[n, t]$. It thus attains its maximum at some $g[n, t]$ say with value R . By Proposition 5.5, the maximal extension of $g[n, t]$ is a slit mapping from $X_R \setminus \{x_1, \dots, x_{n-1}\}$ to $Y[n, t]$. By Remark 4, the map $g[n, t]$ is unique since it sends a puncture to a puncture. Any limit of the maximal extension of $g[n, t]$ as $t \rightarrow s$ is a slit mapping from $X_R \setminus \{x_1, \dots, x_{n-1}\}$ to $Y[n, s]$ by Lemma 5.4 and thus its restriction to $X[n]$ maximizes \mathfrak{m} in $E[n, s]$ by Lemma 5.11. Thus $g[n, t] \rightarrow g[n, s]$ as $t \rightarrow s$. The paths γ_j for $j \in \{1, \dots, n-1\}$ are all constant on $[0, 2^{1-n}]$ so $g[n, t]$ does not change on that interval. Finally, since F maximizes \mathfrak{m} on $\text{CEmb}(X, Y, h)$, it maximizes \mathfrak{m} on the subset $E[n, 1]$ as well, and we have $g[n, 1] = F$. \square

Let $G[n] : \overline{\mathbb{D}} \rightarrow \text{Blob}[n, 0]$ be the Riemann map normalized in such a way that $G[n](0) = L_n(g[n, 0])$ and $G[n]'(0) > 0$. Then let

$$\gamma_n(t) = \begin{cases} L_n(f) & \text{if } t \in [0, 2^{-n}) \\ G[n]((2-2^n t)G[n]^{-1}(L_n(f))) & \text{if } t \in [2^{-n}, 2^{1-n}) \\ L_n(g[n, t]) & \text{if } t \in [2^{1-n}, 1]. \end{cases}$$

This means that γ_n stays put at $L_n(f)$ for some time, then travels along the conformal ray towards the center of $\text{Blob}[n, 0] = \text{Blob}[n, 2^{1-n}]$, and then follows the center for the rest of the time. It is possible that $\text{Blob}[n, 0]$ is a point if $L_{n-1}(f)$ is in the boundary of $\text{Blob}[n-1, 0]$. In that case we let $G[n] : \mathbb{D} \rightarrow \text{Blob}[n, 0]$ be the constant map, i.e., we keep γ_n constant on $[0, 2^{1-n}]$. By construction we have $\pi_Y(\gamma_n(t)) \in Y[n, t]$ which means that the points $\pi_Y(\gamma_1(t)), \dots, \pi_Y(\gamma_{n-1}(t)), \pi_Y(\gamma_n(t))$ are distinct for every $t \in [0, 1]$. Moreover the path γ_n is constant on the interval $[0, 2^{-n}]$. Finally, $E[n+1, t]$ is non-empty since $\gamma_n(t) \in \text{Blob}[n, t]$ for every $t \in [0, 1]$. This finishes the induction scheme.

We now show that the paths $\{\gamma_n\}$ automatically define a path from f to F inside the space $\text{CEmb}(X, Y, h)$.

Lemma 10.2. *For every $t \in [0, 1]$, there exists a unique $f_t \in \text{CEmb}(X, Y, h)$ such that $L_n(f_t) = \gamma_n(t)$ for every $n \in \mathbb{N}$. The map $t \mapsto f_t$ is continuous and satisfies $f_0 = f$ and $f_1 = F$.*

Proof. Observe that $E[n, t]$ is a non-empty closed subset of $\text{CEmb}(X, Y, h)$ and is thus compact. Therefore, for each $t \in [0, 1]$, the nested intersection $\bigcap_{n=1}^{\infty} E[n, t]$ is non-empty. Any two functions in the intersection agree on the dense set $\{x_1, x_2, \dots\}$ and hence on all of X . Therefore, there is a unique function f_t in the intersection. Moreover, f_t varies continuously with t . Indeed, if g is any limit of any subsequence of f_t as $t \rightarrow s$, then for every $n \in \mathbb{N}$ we have

$$g(x_n) = \lim_{t \rightarrow s} f_t(x_n) = \lim_{t \rightarrow s} \pi_Y(\gamma_n(t)) = \pi_Y(\gamma_n(s)) = f_s(x_n)$$

so that $g = f_s$. It follows that $f_t \rightarrow f_s$ as $t \rightarrow s$. By construction we have $L_n(f) = \gamma_n(0)$ and $L_n(F) = \gamma_n(1)$ for every $n \in \mathbb{N}$ so that $f_0 = f$ and $f_1 = F$. \square

We thus have a map $H : \text{CEmb}(X, Y, h) \times [0, 1] \rightarrow \text{CEmb}(X, Y, h)$ defined by $H(f, t) = f_t$. This map is such that

- $t \mapsto H(f, t)$ is continuous for every $f \in \text{CEmb}(X, Y, h)$;
- $H(f, 0) = f$ and $H(f, 1) = F$ for every $f \in \text{CEmb}(X, Y, h)$;
- $H(F, t) = F$ for every $t \in [0, 1]$.

The last point holds because if $f = F$, then every path γ_n is constant and hence $f_t = F$ for every t . It remains to prove that H is continuous in both variables.

Lemma 10.3. *Suppose that for every $n \in \mathbb{N}$, the path $\gamma_n \in \text{Map}([0, 1], \mathbb{D})$ depends continuously on $f \in \text{CEmb}(X, Y, h)$, where each space is equipped with the compact-open topology. Then the map H defined above is continuous.*

Proof. If for every $n \in \mathbb{N}$ the map $(f, t) \mapsto H(f, t)(x_n)$ is continuous, then H is continuous. This is because of the compactness of $\text{CEmb}(X, Y, h)$ and the fact that $\{x_1, x_2, \dots\}$ is dense in X (the proof is the same as in Lemma 10.2). Since $H(f, t)(x_n) = f_t(x_n) = \pi_Y(\gamma_n(t))$, it thus suffices that $(f, t) \mapsto \gamma_n(t)$ be continuous. This condition is equivalent to the requirement that $f \mapsto \gamma_n$ is continuous since $[0, 1]$ is locally compact Hausdorff (see [Mun00, p.287]). \square

Since the map $f \mapsto L_1(f)$ is continuous and the Riemann map $G[1]$ is continuous, it is easy to see that $f \mapsto \gamma_1$ is continuous. We proceed by induction for the rest. Let $n \geq 2$ and suppose that the maps $f \mapsto \gamma_j$ are all continuous for $j = 1, \dots, n-1$. We will prove that the map $f \mapsto G[n]$ is continuous, which obviously implies that $f \mapsto \gamma_n$ is continuous. We use the following theorem of Radó, a proof of which is given in [Pom92, p.26].

Theorem 10.4 (Radó). *Let (D_k, w_k) and (D, w) be topological closed disks in \mathbb{C} , each with a marked point in the interior. Suppose that $w_k \rightarrow w$ and that $D_k \rightarrow D$ in the Fell topology. Suppose also that there are parametrizations $c_k : S^1 \rightarrow \partial D_k$ and $c : S^1 \rightarrow \partial D$ such that $c_k \rightarrow c$ uniformly. Then the normalized Riemann map $(\overline{\mathbb{D}}, 0) \rightarrow (D_k, w_k)$ converges uniformly on $\overline{\mathbb{D}}$ to the normalized Riemann map $(\overline{\mathbb{D}}, 0) \rightarrow (D, w)$.*

By a slight generalization¹¹ of Lemma 9.7, the maps $(f, t) \mapsto \text{Blob}[n, t]$ and $(f, t) \mapsto \partial \text{Blob}[n, t]$ are upper semi-continuous in t . Moreover, they are continuous at every (f, t) for which $E[n, t]$ does not contain any slit mapping by Lemma 9.8. But if $E[n, t]$ contains a slit mapping, then $\text{Blob}[n, t] = \partial \text{Blob}[n, t]$ is a single point and thus upper semi-continuity at (f, t) implies continuity. By Lemma 10.1, the conformal center $L_n(g[n, t])$ of $\text{Blob}[n, t]$ also depends continuously on (f, t) . The only thing that remains to be checked is that the boundary $\partial \text{Blob}[n, t]$ can be parametrized as to converge uniformly.

Definition 10.5. Let $\{c_k\}_{k=1}^\infty$ be a sequence of simple closed curves in $\widehat{\mathbb{C}}$. We say that $\{c_k\}_{k=1}^\infty$ has a *collapsing finger* if after passing to a subsequence, there exist $x_k, y_k, z_k, w_k \in S^1$ in cyclic order and $x, y \in \widehat{\mathbb{C}}$ with $x \neq y$ such that $c_k(x_k) \rightarrow x$, $c_k(y_k) \rightarrow y$, $c_k(z_k) \rightarrow x$, and $c_k(w_k) \rightarrow y$.




FIGURE 12. A piece of a curve with a finger about to collapse.

We now show that collapsing fingers are the only obstructions to uniform convergence of simple closed curves.

Lemma 10.6. *Let c_k and c be simple closed curves in $\widehat{\mathbb{C}}$ such that $c_k(S^1)$ converges to $c(S^1)$ in the Fell topology. If $\{c_k\}_{k=1}^\infty$ does not have any collapsing finger, then we can reparametrize c_k such that $c_k \rightarrow c$ uniformly.*

Proof. By the Jordan–Schoenflies Theorem, c can be extended to a homeomorphism $\hat{c} : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$. Then $\hat{c}^{-1} \circ c_k(S^1) \rightarrow S^1$ in the Fell topology and the sequence $\{\hat{c}^{-1} \circ c_k\}_{k=1}^\infty$ does not have any collapsing finger. Moreover, if $\sigma_k : S^1 \rightarrow S^1$ is a homeomorphism such that $\hat{c}^{-1} \circ c_k \circ \sigma_k$ converges uniformly to the inclusion map $S^1 \hookrightarrow \mathbb{C}$, then $c_k \circ \sigma_k$ converges uniformly to c . We may thus assume that c is the inclusion map $S^1 \hookrightarrow \mathbb{C}$.

If k is large enough, then $c_k(S^1)$ is disjoint from 0 and ∞ . We claim that if k is large enough, then the winding number of c_k around the origin is ± 1 . Since c_k is simple, its winding number is either -1 , 0 , or 1 . Suppose the claim is false. Then after passing to a subsequence, the winding number of c_k is 0 for every k . Let $\arg(c_k) = c_k/|c_k|$ and let $\text{Arg}(c_k) : S^1 \rightarrow \mathbb{R}$ be a lift of $\arg(c_k)$ under the universal covering map $\mathbb{R} \rightarrow S^1$. This lift exists because the winding number is zero. Let

¹¹The codomain $Y[n, t]$ is not fixed but depends continuously on (f, t) . The results of Section 9 generalize easily to this situation.

$[a_k, b_k]$ be the image of $\text{Arg}(c_k)$. Since $c_k(S^1)$ converges to S^1 , it follows that the image $\arg(c_k)(S^1)$ converges to S^1 as well, and hence $\liminf_{n \rightarrow \infty} b_k - a_k \geq 2\pi$. Thus if k is large enough, then $b_k - a_k > \pi$. Let x_k and z_k in S^1 be such that $\text{Arg}(c_k)(x_k) = a_k$ and $\text{Arg}(c_k)(z_k) = \min(b_k, a_k + 2\pi)$. Also let $y_k \in \overline{x_k z_k}$ and $w_k \in \overline{z_k x_k}$ be such that

$$\text{Arg}(c_k)(y_k) = \text{Arg}(c_k)(w_k) = a_k + \pi.$$

Since $c_k(S^1) \rightarrow S^1$, we may pass to a subsequence so that $c_k(x_k)$, $c_k(z_k)$, $c_k(y_k)$, and $c_k(w_k)$ converge to some x , y , z , and w in S^1 . Then $x = z$, $y = w$, and $x \neq y$, i.e. $\{c_k\}_{n=1}^\infty$ has a collapsing finger. This is a contradiction, which proves the claim.

If the winding number of c_k around the origin is -1 , then we reverse the parametrization so that it becomes $+1$. Let $\zeta_1^k < \zeta_2^k < \dots < \zeta_k^k$ be a partition of S^1 into k congruent arcs. Since c_k has winding number 1, we can find $\xi_1^k < \xi_2^k < \dots < \xi_k^k$ in S^1 such that $\arg(c_k)(\xi_j^k) = \zeta_j^k$ for every $j \in \{1, \dots, k\}$. Let $\sigma_k : S^1 \rightarrow S^1$ be any homeomorphism such that $\sigma_k(\zeta_j^k) = \xi_j^k$ for every $j \in \{1, \dots, k\}$. We claim that $c_k \circ \sigma_k$ converges uniformly to the inclusion map $c : S^1 \hookrightarrow \mathbb{C}$.

To simplify notation, we assume that c_k was parametrized correctly from the start, i.e. we assume that for every $k \gg 0$ and every $j \in \{1, \dots, k\}$, we have $\arg(c_k)(\zeta_j^k) = \zeta_j^k$. If c_k does not converge uniformly to c , then there exists an $\varepsilon > 0$ and an infinite set $J \subset \mathbb{N}$ such that for every $k \in J$, there exists a $y_k \in S^1$ such that $|c_k(y_k) - y_k| \geq \varepsilon$. Since S^1 is compact and $c_k(S^1) \rightarrow S^1$, we can pass to a subsequence such that $y_k \rightarrow x$ and $c_k(y_k) \rightarrow y$ for some x and y in S^1 . Note that $|y - x| \geq \varepsilon$ and in particular $y \neq x$. Let $j \in \{1, \dots, k\}$ be such that $\zeta_j^k \leq y_k < \zeta_{j+1}^k$, where we define $\zeta_{k+1}^k = \zeta_1^k$. Then let $x_k = \zeta_j^k$ and $z_k = \zeta_{j+1}^k$. Also let $w_k \in \{\zeta_1^k, \dots, \zeta_k^k\}$ be the closest point to y which comes after z_k but before x_k in the cyclic order on S^1 . We have $c_k(x_k) = x_k \rightarrow x$, $c_k(y_k) \rightarrow y$, $c_k(z_k) = z_k \rightarrow x$, and $c_k(w_k) = w_k \rightarrow y$. In other words, the sequence $\{c_k\}_{k=1}^\infty$ has a collapsing finger, which is a contradiction. Therefore, c_k converges uniformly to c . \square

To conclude the proof that H is continuous, we show that $\partial \text{Blob}[n, t]$ does not have any collapsing fingers. The reason for this is that the blobs $\text{Blob}[n, t]$ are uniformly semi-smooth, meaning that any non-zero limit of a sequence of vectors normal to some blob is normal to the limiting blob¹². Now if there was a collapsing finger somewhere, then we would see two normal vectors pointing opposite to each other in the limit, which is forbidden by the definition of semi-smoothness.

Lemma 10.7. *Suppose that $(f_k, t_k) \rightarrow (f, t)$ in $\text{CEmb}(X, Y, h) \times [0, 1]$. Then $\partial \text{Blob}[n, t_k]$ converges to $\partial \text{Blob}[n, t]$ without collapsing fingers.*

Proof. Let $B_k = \text{Blob}[n, t_k]$, $B = \text{Blob}[n, t]$, $c_k = \partial B_k$, and $c = \partial B$. Suppose that after passing to a subsequence we can find $x_k, y_k, z_k, w_k \in c_k$ in cyclic order and $x, y \in c$ with $x \neq y$ such that $x_k, z_k \rightarrow x$ and $y_k, w_k \rightarrow y$. Rotate and translate the picture in such a way that $x = 0$ and that the upward direction i bisects the cone N_0 of vectors normal to B at 0.

By the proof of Theorem 9.1, there exists a rectangle Q centered at 0 with sides parallel to the coordinate axes such that $Q \cap c$ is the graph of a continuous function. Since $y \neq 0$, we can shrink Q so that it does not contain y . Let $\delta > 0$ be such that

¹²The proof of this is a straightforward generalization of Theorem 8.12

the vertical δ neighborhood U_δ of $Q \cap c$ is contained in Q . Then $Q \setminus U_\delta$ is compact and disjoint from c . Let k be large enough so that x_k and z_k are in Q , y_k and w_k are not in Q , and c_k is disjoint from $Q \setminus U_\delta$. Then of the three subarcs $\overline{x_k y_k}$, $\overline{y_k z_k}$, and $\overline{z_k w_k}$ of c_k , at least two must cross the same vertical side S of Q . This implies that $S \setminus B_k$ is disconnected. Hence there is an open subinterval I of $S \setminus B_k$ whose highest point is contained in B_k . Let D be a closed round disk centered on I and contained in $\mathbb{C} \setminus B_k$. Move the center of D upwards until the boundary of the translated disk D^* first hits B_k . Any intersection point p_k of D^* with B_k is on the top half of ∂D^* . Moreover, the unit vector \mathbf{v}_k based at p_k and pointing in the direction of the center of D^* is normal to B_k .




FIGURE 13. If the sequence of blobs has a collapsing finger, then we can find a sequence of normal vectors which accumulate to a vector which is not normal to the limiting blob.

Since we can choose Q to be arbitrarily small, we can arrange so that $p_k \rightarrow 0$. Then the normal vectors \mathbf{v}_k accumulate onto vectors pointing towards the lower half-plane at 0. This is a contradiction since the cone of normal vectors N_0 is contained in the upper half-plane. \square

Thus by Lemma 10.6, the boundary of $\text{Blob}[n, t]$ can be parametrized in a way that depends uniformly continuously on (f, t) . By Theorem 10.4, this implies that the Riemann map $G[n] : \mathbb{D} \rightarrow \text{Blob}[n, t]$ depends uniformly continuously on (f, t) . Therefore the path γ_n depends uniformly continuously on f , and hence H is continuous by Lemma 10.3. This shows that $\text{CEmb}(X, Y, h)$ is contractible under hypothesis (\mathbf{H}') .

More generally, $\text{CEmb}(X, Y, h)$ is contractible whenever h is generic and $\text{CEmb}(X, Y, h)$ is non-empty. The proof for $\partial X = \emptyset$ is in Subsection 2.2 and if $\text{CEmb}(X, Y, h)$ contains a slit mapping then this follows from Theorem 1.4.

11. THE REMAINING CASES

In this section, we complete the proof of Theorem 1.2, that is, we describe the homotopy type of $\text{CEmb}(X, Y, h)$ when h is not generic, always assuming that $\text{CEmb}(X, Y, h)$ is non-empty. As before, we let $\{x_1, x_2, x_3, \dots\}$ be a dense set of distinct points in X .

11.1. h is cyclic but not parabolic. Suppose that $h : X \rightarrow Y$ is cyclic but not parabolic and that Y is hyperbolic. This case is analogous to the case where h is generic. We only provide an outline of the proof.

If $\text{CEmb}(X, Y, h)$ contains a slit mapping, then it is homeomorphic to a circle by Theorem 1.4. So we may assume that $\text{CEmb}(X, Y, h)$ does not contain any slit mapping. Form the annulus cover $\pi_A : A \rightarrow Y$ corresponding to the image of $\pi_1(h)$. We can define a continuous map $\text{lift}_{x_1} : \text{Map}(X, Y, h) \rightarrow A$ in a similar fashion as in Section 6. Let $\text{Blob}(x_1, X, Y, h)$ be the image of $\text{CEmb}(X, Y, h)$ by lift_{x_1} . This $\text{Blob}(x_1, X, Y, h)$ is compact and connected since $\text{CEmb}(X, Y, h)$ is, according to Lemma 3.1 and Theorem 1.6. All of the results from Sections 6, 8 and 9 pertaining to the local geometry of $\text{Blob}(x_1, X, Y, h)$ extend to the current setting. The conclusion is that $\text{Blob}(x_1, X, Y, h)$ is a 2-dimensional manifold with boundary.

Let $\{X_r\}_{r \in [0, \infty]}$ be a 1-parameter family of enlargements of X as in Section 5, let \mathbf{m} be the associated modulus of extension, and let R be the maximum value of \mathbf{m} . By Proposition 5.5 and Theorem 1.4, $\text{CEmb}(X_R, Y, h)$ is homeomorphic to a circle via the evaluation map. It follows that $\text{Blob}(x_1, X_R, Y, h)$ is homeomorphic to a circle so that its complement in A has two connected components. By a similar argument as in the proof of Theorem 1.7, the complement of $\text{Blob}(x_1, X, Y, h)$ has two connected components as well. Since $\text{Blob}(x_1, X, Y, h)$ is a planar 2-dimensional manifold with two boundary components, it is homeomorphic to a closed annulus.

Let D be any deformation retraction of $\text{Blob}(x_1, X, Y, h)$ into the circle $\text{Blob}(x_1, X_R, Y, h)$. The strategy after this step is the same as in Section 10. Given $f \in \text{CEmb}(X, Y, h)$, we let $\gamma_1(t) = \text{lift}_{x_1}(f)$ for $t \in [0, 1/2)$ and $\gamma_1(t) = D(\text{lift}_{x_1}(f), 2t-1)$ for $t \in [1/2, 1]$. The map $h[1, t] : X \setminus \{x_1\} \rightarrow Y \setminus \pi_A(\gamma_1(t))$ obtained by composing h with point-pushing along $\pi_A \circ \gamma_1$ is now generic so that we can construct the next paths $\gamma_2, \gamma_3, \dots$ in the same way as in the previous section. The end result is a deformation retraction of $\text{CEmb}(X, Y, h)$ into the circle $\text{CEmb}(X_R, Y, h)$.

11.2. Y is the punctured disk. Let $h : X \rightarrow \mathbb{D} \setminus \{0\}$ be a non-trivial (hence parabolic) embedding. By Montel's theorem, any sequence in $\text{CEmb}(X, \mathbb{D} \setminus \{0\}, h)$ has a subsequence converging locally uniformly to either an element in $\text{CEmb}(X, \mathbb{D} \setminus \{0\}, h)$ or to a constant map into $\overline{\mathbb{D}}$. As h is non-trivial, the only possible constant limit is 0. Thus the set of $f \in \text{CEmb}(X, \mathbb{D} \setminus \{0\}, h)$ which map x_1 outside of a fixed open neighborhood of 0 is compact. Since $\overline{\mathbb{D}} \setminus \{0\}$ acts by multiplication on $\text{CEmb}(X, \mathbb{D} \setminus \{0\}, h)$, the image $V(x_1)$ of the evaluation map at x_1 is equal to a punctured disk $r\overline{\mathbb{D}} \setminus \{0\}$ for some $r \in (0, 1)$. By a slight modification of Proposition 6.9, for every $y_1 \in \partial V(x_1)$ there is a unique $f \in \text{CEmb}(X, \mathbb{D} \setminus \{0\}, h)$ such that $f(x_1) = y_1$. We do not need to lift here: if two maps $f, g \in \text{Map}(X, \mathbb{D} \setminus \{0\}, h)$ agree at x_1 they are homotopic rel x_1 because the pure mapping class group of the twice-punctured disk is trivial. Thus the inverse image of $\partial V(x_1) = \{z \in \mathbb{D} : |z| = r\}$ in $\text{CEmb}(X, \mathbb{D} \setminus \{0\}, h)$ is homeomorphic to a circle.

Given any $f \in \text{CEmb}(X, \mathbb{D} \setminus \{0\}, h)$, let $\gamma_1 : [0, 1] \rightarrow \mathbb{D} \setminus \{0\}$ be constant equal to $f(x_1)$ on $[0, 1/2)$ followed by the radial ray from $f(x_1)$ to $rf(x_1)/|f(x_1)|$ on $[1/2, 1]$. Note that the map $h[1, t] : X \setminus \{x_1\} \rightarrow \mathbb{D} \setminus \{0, \gamma_1(t)\}$ (notation as in Section 10) is generic for every $t \in [0, 1]$, so we can construct the next paths $\gamma_2, \gamma_3, \dots$ as before. This gives a deformation retraction of $\text{CEmb}(X, \mathbb{D} \setminus \{0\}, h)$ into the circle of slit mappings rel x_1 .

11.3. Y is the disk. Suppose that Y is the unit disk \mathbb{D} . Consider the map $\mathbb{D} \rightarrow \text{Aut}_0(\mathbb{D})$ which sends $a \in \mathbb{D}$ to the automorphism $M_a(z) = \frac{z-a}{1-\bar{a}z}$. We have a homeomorphism

$$\text{CEmb}(X, \mathbb{D}, h) \rightarrow \mathbb{D} \times \text{CEmb}(X \setminus \{x_1\}, \mathbb{D} \setminus \{0\}, M_{h(x_1)} \circ h^*)$$

given by $f \mapsto (f(x_1), M_{f(x_1)} \circ f)$, where $h^* : X \setminus \{x_1\} \rightarrow \mathbb{D} \setminus \{h(x_1)\}$ is the restriction of h . Since \mathbb{D} is contractible, $\text{CEmb}(X, \mathbb{D}, h)$ is homotopy equivalent to the second factor. Note that h is trivial and $M_{h(x_1)} \circ h^*$ is parabolic, with codomain the once-punctured disk. By the previous case, $\text{CEmb}(X \setminus \{x_1\}, \mathbb{D} \setminus \{0\}, M_{h(x_1)} \circ h^*)$ is homotopy equivalent to S^1 . Thus $\text{CEmb}(X, \mathbb{D}, h)$ is homotopy equivalent to a circle, which in turn is homotopy equivalent to the unit tangent bundle of \mathbb{D} .

11.4. X is the disk. Suppose that $X = \mathbb{D}$ and that $Y \neq \mathbb{D}$ is hyperbolic. Here h is trivial so we may drop it from the notation. We first define a map from the unit tangent bundle $T^1 Y$ to $\text{CEmb}(\mathbb{D}, Y)$ as follows. Given $\mathbf{v} \in T_y^1 Y$, let $D_{\mathbf{v}} \subset Y$ be the largest embedded ball centered at y in the hyperbolic metric, and let $F_{\mathbf{v}} : \mathbb{D} \rightarrow D_{\mathbf{v}}$ be the Riemann map with $F_{\mathbf{v}}(0) = y$ and $F'_{\mathbf{v}}(0) = \lambda \mathbf{v}$ for some $\lambda > 0$. The map $\mathbf{v} \mapsto F_{\mathbf{v}}$ is an embedding from $T^1 Y$ to $\text{CEmb}(\mathbb{D}, Y)$. We will construct a deformation retraction of $\text{CEmb}(\mathbb{D}, Y)$ into the image of that map.

Given $f \in \text{CEmb}(\mathbb{D}, Y)$, let $\mathbf{v} \in T^1 Y$ be the unique vector such that $f'(0) = \lambda \mathbf{v}$ for some $\lambda > 0$. Then let $r \in (0, 1]$ be the largest number such that $f(r\mathbb{D}) \subset D_{\mathbf{v}}$ and let $f^\dagger(z) = f(rz)$. Then $F_{\mathbf{v}}^{-1} \circ f^\dagger : \mathbb{D} \rightarrow \mathbb{D}$ is a conformal embedding which fixes the origin and has positive derivative there.

Let $g : \mathbb{D} \rightarrow \mathbb{D}$ be a conformal embedding with $g(0) = 0$ and $g'(0) > 0$. For every $t \in (0, 1]$, define $\rho_t = \inf\{\rho > 0 : g(t\mathbb{D}) \subset \rho\mathbb{D}\}$ and $g_t(z) = g(tz)/\rho_t$. By Koebe's distortion theorem [Dur83, p.33] we have

$$\frac{t}{(1+t)^2} \leq \frac{\rho_t}{|g'(0)|} \leq \frac{t}{(1-t)^2}$$

and it follows that $g_t \rightarrow \text{id}$ as $t \rightarrow 0$.

We define a deformation retraction of $\text{CEmb}(\mathbb{D}, Y)$ into $\{F_{\vec{v}} : \vec{v} \in T^1 Y\}$ by the formula

$$H(f, t) = \begin{cases} z \mapsto f((1 - (1-r)2t)z) & \text{if } t \in [0, 1/2) \\ F_{\mathbf{v}} \circ g_{(2-2t)} & \text{if } t \in [1/2, 1] \end{cases}$$

where r , f^\dagger and \vec{v} are defined in terms of f as above and $g = F_{\mathbf{v}}^{-1} \circ f^\dagger$.

11.5. h is trivial. Suppose that $X \neq \mathbb{D}$, that $Y \neq \mathbb{D}$ is hyperbolic, and that h is trivial. Given $f \in \text{CEmb}(X, Y, h)$, let $D_f \subset Y$ be the smallest topological disk containing the image of f . We can define D_f by filling the holes of $f(X)$. Then let $F : \mathbb{D} \rightarrow D_f$ be the Riemann map with $F(0) = f(x_1)$ and $F^{-1}(f(x_2)) > 0$. We thus obtain an embedding

$$\text{CEmb}(X, Y, h) \rightarrow \text{CEmb}(\mathbb{D}, Y) \times W$$

defined by $f \mapsto (F, F^{-1} \circ f)$, where W is the set of all conformal embeddings from X to \mathbb{D} sending x_1 to 0 and x_2 to a positive real number. There is an obvious left inverse

$$\text{CEmb}(\mathbb{D}, Y) \times W \rightarrow \text{CEmb}(X, Y, h)$$

given by $(G, g) \mapsto G \circ g$.

By the previous subsection, there is a deformation retraction H_1 from $\text{CEmb}(\mathbb{D}, Y)$ into a subset homeomorphic to $T^1 Y$. As for W , it is homeomorphic to the quotient $\text{CEmb}(X \setminus \{x_1\}, \mathbb{D} \setminus \{0\}, g^*)/S^1$ where $g : X \rightarrow \mathbb{D}$ is any embedding with $g(x_1) = 0$, g^* is the restriction of g , and S^1 acts by multiplication. Subsection 11.2 provides a deformation retraction H_2 of W into a singleton. Therefore

$\text{CEmb}(X, Y, h)$ deformation retracts into a subset homeomorphic to T^1Y via the formula $(f, t) \mapsto H_1(F, t) \circ H_2(F^{-1} \circ f, t)$.

11.6. h is parabolic. Suppose that $h : X \rightarrow Y$ is parabolic, where Y is hyperbolic and not the once-punctured disk. Let p be the puncture around which h wraps non-trivially. Given $f \in \text{CEmb}(X, Y, h)$, we define a disk $D_f \subset Y \cup \{p\}$ by filling the holes of $f(X)$ in $Y \cup \{p\}$. Then we define $F : \mathbb{D} \rightarrow D_f$ to be the Riemann map with $F(0) = p$ and $F^{-1}(f(x_1)) > 0$. This yields an embedding

$$\text{CEmb}(X, Y, h) \rightarrow \text{CEmb}(\mathbb{D} \setminus \{0\}, Y, G^*) \times W$$

defined by $f \mapsto (F, F^{-1} \circ f)$, where G^* is the restriction of some embedding $G : \mathbb{D} \rightarrow Y \cup \{p\}$ satisfying $G(0) = p$ and W is the set of all conformal embeddings $g : X \rightarrow \mathbb{D} \setminus \{0\}$ such that $g(x_1) > 0$ which are homotopic to $F_0^{-1} \circ f_0$ for any $f_0 \in \text{CEmb}(X, Y, h)$. The first factor deformation retracts into a circle by Subsection 11.4 whereas the second factor deformation retracts into a point by Subsection 11.2. By applying the left inverse of the above embedding (the composition map), we get a deformation retraction of $\text{CEmb}(X, Y, h)$ into a circle.

It remains to treat the cases where Y is not hyperbolic. In those cases, we can quotient $\text{CEmb}(X, Y, h)$ by the action of $\text{Aut}_0(Y)$ to reduce to the hyperbolic case.

11.7. Y is a torus. Suppose that Y is a torus. Then $\text{CEmb}(X, Y, h)$ is homeomorphic to

$$\text{Aut}_0(Y) \times \text{CEmb}(X \setminus \{x_1\}, Y \setminus \{h(x_1)\}, h^*)$$

where $h^* : X \setminus \{x_1\} \rightarrow Y \setminus \{h(x_1)\}$ is the restriction of h . Recall that $\text{Aut}_0(Y)$ is homeomorphic to Y itself.

If h is trivial, then h^* is parabolic and its codomain is a hyperbolic surface. Subsection 11.6 shows that $\text{CEmb}(X \setminus \{x_1\}, Y \setminus \{h(x_1)\}, h^*)$ is then homotopy equivalent to S^1 . This means that $\text{CEmb}(X, Y, h)$ is homotopy equivalent to a 3-dimensional torus, or the unit tangent bundle of Y .

If h is non-trivial, then h^* is generic so that $\text{CEmb}(X \setminus \{x_1\}, Y \setminus \{h(x_1)\}, h^*)$ is contractible. Thus $\text{CEmb}(X, Y, h)$ is homotopy equivalent to $\text{Aut}_0(Y) \approx Y$.

11.8. Y is the sphere with at most 2 punctures. Suppose that Y is the Riemann sphere $\widehat{\mathbb{C}}$. Let $P = \{x_1, x_2, x_3\}$. Then

$$\text{CEmb}(X, Y, h) \approx \text{Aut}_0(Y) \times \text{CEmb}(X \setminus P, Y \setminus h(P), h^*)$$

where $h^* : X \setminus P \rightarrow Y \setminus h(P)$ is the restriction of h . Observe that h is automatically trivial and that h^* is generic. Thus $\text{CEmb}(X, Y, h)$ is homotopy equivalent to $\text{Aut}_0(Y)$. We stated in Subsection 2.2 that $\text{Aut}_0(\widehat{\mathbb{C}})$ is homotopy equivalent to the unit tangent bundle of $\widehat{\mathbb{C}}$. Here is a proof. First, $\text{Aut}_0(\widehat{\mathbb{C}})$ is homeomorphic to the set of triples (a, \mathbf{v}, b) where $a, b \in \widehat{\mathbb{C}}$ are distinct and $\mathbf{v} \in T_a \widehat{\mathbb{C}}$ is non-zero. The homeomorphism is given by $f \mapsto (f(0), f'(0), f(\infty))$. This set of triples deformation retracts into $T\widehat{\mathbb{C}} \setminus \vec{0}$ (the complement of the zero section in the tangent bundle of $\widehat{\mathbb{C}}$) by moving the point b along the spherical geodesic to the antipode of a . Now $T\widehat{\mathbb{C}} \setminus \vec{0}$ clearly deformation retracts into the unit tangent bundle $T^1\widehat{\mathbb{C}}$.

Suppose that Y is the Riemann sphere minus a point. Then we can repeat the same trick with $P = \{x_1, x_2\}$ instead. Again, h is trivial and its restriction $h^* : X \setminus P \rightarrow Y \setminus h(P)$ is generic. Thus $\text{CEmb}(X, Y, h)$ is homotopy equivalent to $\text{Aut}_0(\mathbb{C})$. The latter is homeomorphic to the complement of the zero section in $T\mathbb{C}$.

This deformation retracts onto the unit tangent bundle of \mathbb{C} (which deformation retracts further into a circle).

Suppose that Y is the Riemann sphere minus two points. Then we can puncture at one point to factor out the action of $\text{Aut}_0(Y)$. That group is homeomorphic to $S^1 \times \mathbb{R}$, hence homotopy equivalent to S^1 . If h is trivial, then its restriction $h^* : X \setminus \{x_1\} \rightarrow Y \setminus \{h(x_1)\}$ is a parabolic embedding into a hyperbolic surface and $\text{CEmb}(X \setminus \{x_1\}, Y \setminus \{h(x_1)\}, h^*)$ is homotopy equivalent to S^1 by Subsection 11.6. Thus $\text{CEmb}(X, Y, h)$ is homotopy equivalent to $S^1 \times S^1$, which is in turn homotopy equivalent to the unit tangent bundle of Y . If h is non-trivial, then it is cyclic. In this case, h^* is generic so that $\text{CEmb}(X \setminus \{x_1\}, Y \setminus \{h(x_1)\}, h^*)$ is contractible and $\text{CEmb}(X, Y, h)$ is homotopy equivalent to S^1 .

The reader can check that we have exhausted all possibilities for the embedding $h : X \rightarrow Y$, which concludes the proof of Theorem 1.2. The latter obviously implies Theorem 1.1.

REFERENCES

- [Ahl53] L.V. Ahlfors, *On quasiconformal mappings*, J. Anal. Math. **3** (1953), no. 1, 1–58.
- [Ahl06] ———, *Lectures on Quasiconformal Mappings: Second Edition*, vol. 38, American Mathematical Society, 2006, with additional chapters by C.J. Earle and I. Kra, M. Shishikura, J.H. Hubbard.
- [Ber58] L. Bers, *Quasiconformal mappings and Teichmüller's theorem*, New York University, Institute of Mathematical Sciences, 1958.
- [BH09] M.R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Springer, 2009.
- [Dur83] P.L. Duren, *Univalent functions*, Springer-Verlag, 1983.
- [Ear64] C.J. Earle, *Teichmüller spaces of groups of the second kind*, Acta Math. **112** (1964), no. 1, 91–97.
- [Ear67] ———, *Reduced Teichmüller spaces*, Trans. Amer. Math. Soc. **126** (1967), no. 1, 54–63.
- [EL02] C.J. Earle and N. Lakic, *Variability sets on Riemann surfaces and forgetful maps between Teichmüller spaces*, Ann. Acad. Sci. Fenn. Math. **27** (2002), no. 2, 307–324.
- [EM78] C. Earle and A. Marden, *Conformal embeddings of Riemann surfaces*, J. Anal. Math. **34** (1978), 194–203.
- [EM02] Y. Eliashberg and N. Mishachev, *Introduction to the h-Principle*, American Mathematical Society, 2002.
- [Fel62] J.M.G. Fell, *A Hausdorff topology for the closed subsets of a locally compact non-Hausdorff space*, Proc. Amer. Math. Soc. **13** (1962), no. 3, 472–476.
- [FLP12] A. Fathi, F. Laudenbach, and V. Poénaru, *Thurston's work on surfaces*, Mathematical Notes 48, Princeton University Press, 2012.
- [FM11] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton University Press, 2011.
- [Gar84] F.P. Gardiner, *Measured foliations and the minimal norm property for quadratic differentials*, Acta Math. **152** (1984), 57–76.
- [GG01] A.A. Golubev and S.Yu. Graf, *Extremal problems on classes of quasiconformal embeddings of Riemann surfaces*, Siberian Adv. Math. **11** (2001), no. 4, 47–64.
- [GL10a] F.P. Gardiner and N. Lakic, *The Dirichlet principle for measured foliations*, preprint, 2010.
- [GL10b] ———, *A synopsis of the Dirichlet principle for measured foliations*, RIMS Kôkyûroku Bessatsu **B17** (2010), 55–64.
- [Hat01] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2001.
- [HM79] J. Hubbard and H. Masur, *Quadratic differentials and foliations*, Acta Math. **142** (1979), 221–274.
- [Iof75] M.S. Ioffe, *Extremal quasiconformal embeddings of Riemann surfaces*, Sib. Math. J. **16** (1975), 398–411.
- [Iof78] ———, *Some problems in the calculus of variation “in the large” for conformal and quasiconformal imbeddings of Riemann surfaces*, Sib. Math. J. **19** (1978), 412–424.

- [Jen58] J.A. Jenkins, *Univalent functions and conformal mapping*, Springer-Verlag Berlin Heidelberg, 1958.
- [Jen59] ———, *A counter-example to a statement on conformal mapping of Riemann surfaces*, Proc. Amer. Math. Soc. **10** (1959), 423–424.
- [KPT15] J. Kahn, K.M. Pilgrim, and D.P. Thurston, *Conformal surface embeddings and extremal length*, preprint, arXiv:1507.05294, 2015.
- [Kru05] S.L. Krushkal, *Variational principles in the theory of quasiconformal maps*, Handbook of Complex Analysis: Geometric Function Theory, vol. 2, Elsevier, 2005, pp. 31–98.
- [Mas75] H. Masur, *On a class of geodesics in Teichmüller space*, Ann. of Math. (2) **102** (1975), no. 2, 205–221.
- [McM13] C. McMullen, *Navigating moduli space with complex twists*, J. Eur. Math. Soc. (JEMS) **15** (2013), 1223–1243.
- [Mil97] J.W. Milnor, *Topology from the differentiable viewpoint*, Princeton University Press, 1997.
- [Mil06] J. Milnor, *Dynamics in one complex variable*, third ed., Princeton University Press, 2006.
- [MS84] A. Marden and K. Strebel, *The Heights theorem for quadratic differentials on Riemann surfaces*, Acta Math. **153** (1984), 153–211.
- [Mun00] J.R. Munkres, *Topology*, second ed., Pearson, 2000.
- [Pom92] C. Pommerenke, *Boundary behavior of conformal maps*, Springer, 1992.
- [RF10] H.L. Royden and P.M. Fitzpatrick, *Real analysis*, fourth ed., Pearson, 2010.
- [Roy54a] H.L. Royden, *The conformal rigidity of certain subdomains on a Riemann surface*, Trans. Amer. Math. Soc. **76** (1954), 14–25.
- [Roy54b] ———, *The interpolation problem for schlicht functions*, Ann. of Math. (2) **60** (1954), no. 2, 326–344.
- [Shi04] M. Shiba, *Conformal mapping of Riemann surfaces and the classical theory of univalent functions*, Publ. Inst. Math. (Beograd) (N.S.) **75** (2004), no. 89, 217–232.
- [Slo91] Z. Słodkowski, *Holomorphic motions and polynomial hulls*, Proc. Amer. Math. Soc. **11** (1991), no. 2, 347–355.
- [SS54] M.M. Schiffer and D.C. Spencer, *Functionals of finite Riemann surfaces*, Princeton University Press, 1954.
- [Str84] K. Strebel, *Quadratic differentials*, Springer-Verlag, 1984.
- [Str98] ———, *Point shift differentials and extremal quasiconformal mappings*, Ann. Acad. Sci. Fenn. Math. **23** (1998), 475–494.
- [Thu86] W.P. Thurston, *Zippers and univalent functions*, The Bieberbach Conjecture: Proceedings of the Symposium on the Occasion of the Proof, American Mathematical Society, 1986, pp. 185–197.
- [Thu16] D.P. Thurston, *From rubber bands to rational maps: a research report*, Res. Math. Sci. **3:15** (2016).

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

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