Geometry & Topology 22 (2018) 4259–4306

Rigidity of Teichmüller space

Alex Eskin
Howard Masur
Kasra Rafi

We prove that every quasi-isometry of Teichmüller space equipped with the Teichmüller metric is a bounded distance from an isometry of Teichmüller space. That is, Teichmüller space is quasi-isometrically rigid.

32G15

1 Introduction and statement of the theorem

In this paper we continue our study of the coarse geometry of Teichmüller space begun in Eskin, Masur and Rafi [11]. Our goal, as part of Gromov's broad program to understand spaces and groups by their coarse or quasi-isometric geometry, is to carry this out in the context of Teichmüller space equipped with the Teichmüller metric. To state the main theorem, let S be a connected surface of finite hyperbolic type. Define the complexity of S to be

$$\xi(S) = 3\mathsf{g}(S) + \mathsf{p}(S) - 3,$$

where $\mathsf{g}(S)$ is the genus of S and $\mathsf{p}(S)$ is the number of punctures. Let $\mathcal{T}(S)$ denote the Teichmüller space of S equipped with the Teichmüller metric $d_{\mathcal{T}(S)}$.

Theorem 1.1 *Assume $\xi(S) \geq 2$. Then, for every $\mathsf{K}_S, \mathsf{C}_S > 0$ there is a constant $\mathsf{D}_S > 0$ such that if*

$$f_S\colon \mathcal{T}(S) \to \mathcal{T}(S)$$

is a $(\mathsf{K}_S, \mathsf{C}_S)$–quasi-isometry then there is an isometry

$$\Psi_S\colon \mathcal{T}(S) \to \mathcal{T}(S)$$

such that, for $x \in \mathcal{T}(S)$,

$$d_{\mathcal{T}(S)}(f(x), \Psi(x)) \leq \mathsf{D}_S.$$

Published: 6 December 2018 DOI: 10.2140/gt.2018.22.4259

Recall that a map $f\colon \mathcal{X} \to \mathcal{Y}$ from a metric space $(\mathcal{X}, d_{\mathcal{X}})$ to a metric space $(\mathcal{Y}, d_{\mathcal{Y}})$ is a (K, C)–quasi-isometry if it is C–coarsely onto and, for $x_1, x_2 \in \mathcal{X}$,

$$\frac{1}{K} d_{\mathcal{X}}(x_1, x_2) - C \le d_{\mathcal{Y}}(f(x_1), f(x_2)) \le K d_{\mathcal{X}}(x_1, x_2) + C. \tag{1}$$

If equation (1) holds and the map is not assumed to be onto, then f is called a *quasi-isometric embedding*. One defines, for a metric space $(\mathcal{X}, d_{\mathcal{X}})$, the group $\mathrm{QI}(\mathcal{X})$ as the equivalence classes of quasi-isometries from $\mathcal{X}$ to itself, with two quasi-isometries being equivalent if they are a bounded distance apart. When the natural homomorphism $\mathrm{Isom}(\mathcal{X}) \to \mathrm{QI}(\mathcal{X})$ is an isomorphism, we say $\mathcal{X}$ is quasi-isometrically rigid. Then Theorem 1.1 restated is that $\mathcal{T}(S)$ is quasi-isometrically rigid.

Note also that, by Royden's theorem [24], $\mathrm{Isom}(\mathcal{T}(S))$ is essentially the mapping class group (the exceptional cases are the twice-punctured torus and the closed surface of genus 2 where the two groups differ by a finite index). Hence, except for the lower complexity cases, $\mathrm{QI}(\mathcal{T}(S))$ is isomorphic to the mapping class group. This theorem has also been proven by Brian Bowditch [3] by a different method.

History and related results

There is a fairly long history that involves the study of the group $\mathrm{QI}(\mathcal{X})$ in different contexts. Among these, symmetric spaces are the closest to our setting.

In the case when $\mathcal{X}$ is $\mathbb{R}^n$, $\mathbb{H}^n$ or $\mathbb{CH}^n$, the quasi-isometry group is complicated and much larger than the isometry group. In fact, if a self map of $\mathbb{H}^n$ is a bounded distance from an isometry then it induces a conformal map on S^{n-1}. But every quasiconformal homeomorphism from $S^{n-1} \to S^{n-1}$ extends to a quasi-isometry of $\mathbb{H}^n$. This, in particular, shows why the condition $\xi(S) \ge 2$ in Theorem 1.1 is necessary. When $\xi(S) = 1$, the Teichmüller space $\mathcal{T}(S)$ is isometric (up to a factor of 2) to the hyperbolic plane $\mathbb{H}$ and, as mentioned above, $\mathbb{H}$ is not rigid.

Pansu [19] proved that other rank 1 symmetric spaces of noncompact type such as quaternionic hyperbolic space $\mathbb{HH}^n$ and the Cayley plane $\mathbb{P}^2(\mathbb{O})$ are rigid. In contrast, higher-rank irreducible symmetric spaces are rigid. This was proven by Kleiner and Leeb [16]; see also Eskin and Farb [7] for a different proof. In our setting, when $\xi(S) \ge 2$, the space $\mathcal{T}(S)$ is analogous to a higher-rank symmetric space; see Eskin, Masur and Rafi [11] and Bowditch [3] for discussion about flats in $\mathcal{T}(S)$. The curve complex $\mathcal{C}(S)$, which plays a prominent role in Teichmüller theory, is analogous to the Tits boundary of symmetric space.

Continuing the above analogy, the action of the mapping class group on Teichmüller space is analogous to the action of a nonuniform lattice on a symmetric space. Quasi-isometric rigidity was shown for nonuniform lattices in rank 1 groups other than $\mathrm{SL}(2,\mathbb{R})$ and for some higher-rank lattices by Schwartz [25; 26], and in general by Eskin [6]. The quasi-isometric rigidity of the mapping class group $\mathrm{Mod}(S)$ was shown by Behrstock, Kleiner, Minsky and Mosher [1], by Hamenstädt [13] and later by Bowditch [4]. More generally, Bowditch in that same paper showed that if S and S' are closed surfaces with $\xi(S) = \xi(S') \geq 4$ and ϕ is a quasi-isometric embedding of $\mathrm{Mod}(S)$ in $\mathrm{Mod}(S')$, then $S = S'$ and ϕ is a bounded distance from an isometry. He also shows quasi-isometric rigidity for Teichmüller space with the Weil–Petersson metric [2].

Inductive step

We prove this theorem inductively. To apply induction, we need to consider non-connected surfaces. Let Σ be a possibly disconnected surface of finite hyperbolic type. We always assume that Σ does not have a component that is a sphere, an annulus, a pair of pants or a torus. We define the complexity of Σ to be

$$\xi(\Sigma) = 3\mathsf{g}(\Sigma) + \mathsf{p}(\Sigma) - 3\mathsf{c}(\Sigma),$$

where $\mathsf{c}(\Sigma)$ is the number of connected components of Σ. For a point $x \in \mathcal{T}(\Sigma)$, let P_x be the short pants decomposition at x and for a curve $\gamma \in P_x$ define

$$\tau_x(\gamma) \simeq \log \frac{1}{\mathrm{Ext}_x(\gamma)},$$

where $\mathrm{Ext}_x(\gamma)$ denotes the extremal length of the curve γ on the Riemann surface X. (See Section 2.2 for exact definitions.) For a constant $L > 0$, consider the sets

$$\mathcal{T}(\Sigma, L) = \{x \in \mathcal{T}(\Sigma) \mid \tau_x(\gamma) \leq L \text{ for every curve } \gamma\},$$
$$\partial_L(\Sigma, L) = \{x \in \mathcal{T}(\Sigma, L) \mid \tau_x(\gamma) = L \text{ for all } \gamma \in P_x\}.$$

Thinking of L as a very large number, we say a quasi-isometry

$$f_\Sigma \colon \mathcal{T}(\Sigma, L) \to \mathcal{T}(\Sigma, L)$$

is C_Σ–anchored if the restriction of f_Σ to $\partial_L(\Sigma)$ is nearly the identity. That is, for every $x \in \partial_L(\Sigma, L)$,

$$d_{\mathcal{T}(\Sigma)}(x, f_\Sigma(x)) \leq \mathsf{C}_\Sigma.$$

Our induction step is the following.

Theorem 1.2 *Let Σ be a surface of finite hyperbolic type. For every K_Σ and C_Σ, there are L_Σ and D_Σ such that for $L \ge \mathsf{L}_\Sigma$, if*

$$f_\Sigma \colon \mathcal{T}(\Sigma, L) \to \mathcal{T}(\Sigma, L)$$

is a $(\mathsf{K}_\Sigma, \mathsf{C}_\Sigma)$–quasi-isometry that is C_Σ–anchored, then for every $x \in \mathcal{T}(\Sigma)$ we have

$$d_{\mathcal{T}(\Sigma)}(x, f_\Sigma(x)) \le \mathsf{D}_\Sigma.$$

Note that D_Σ depends on the topology of Σ and the constants K_Σ, C_Σ and L_Σ, but is independent of L.

Outline of the proof

The overall strategy is to take the quasi-isometry and prove it preserves more and more of the structure of Teichmüller space. Section 2 is devoted to establishing notation and background material. In Section 3 we define the rank of a point as the number of short curves plus the number of complementary components that are not pairs of pants. A point has maximal rank if all complementary components W have $\xi(W) \le 1$. We show in Proposition 3.8 that a quasi-isometry preserves points with maximal rank. The proof uses the ideas of coarse differentiation, previously developed in the context of Teichmüller space in [11], which in turn is based on the work of Eskin, Fisher and Whyte [9; 10]. Coarse differentiation was also previously used by Peng [20; 21] to study quasiflats in solvable Lie groups. The important property of maximal rank is that near such a point of maximal rank, Teichmüller space is close to being isometric to a product of copies of $\mathbb{H}$, with the supremum metric; see Minsky [18].

In Section 4 we prove a local version of the splitting theorem shown in Kleiner and Leeb [16] in the context of symmetric spaces and later in Eskin and Farb [8] for products of hyperbolic planes. There it is proved in Theorem 4.1 that a quasi-isometric embedding from a large ball in $\prod \mathbb{H}$ to $\prod \mathbb{H}$ can be restricted to a smaller ball where it factors, up to a fixed additive error. This local factoring is applied in Section 5 to give a bijective association $f_x^\star$ between factors at x and at $f(x)$. We also prove a notion of analytic continuation, namely, we examine how $f_x^\star$ and $f_{x'}^\star$ are related when local factors around points x and x' overlap.

We use this to show (Proposition 6.1) that maximal cusps are preserved by the quasi-isometry. A maximal cusp is the subset of maximal rank consisting of points where there is a maximal set of short curves all about the same length. There Teichmüller space looks like a cone in a product of horoballs inside the product of $\mathbb{H}$. The set of maximal cusps is disconnected; there is a component associated to every pants decomposition. Thus, f induces a bijection on the set of pants decompositions. The next step is then to show this map is induced by automorphism of the curve complex, and hence by Ivanov's theorem, it is associated to an isometry of Teichmüller space. Composing f by the inverse, we can assume that f sends every component of the set of maximal cusps to itself. An immediate consequence of this is that f restricted to the thick part of Teichmüller space is a bounded distance away from the identity (Proposition 6.6). From this fact and again applying Propositions 5.1 and 5.3, we then show, in Proposition 6.9, that for any point in the maximal rank set, the shortest curves are preserved and in fact, for any shortest curve α at points x, one has $f_x^\star(\alpha) = \alpha$.

This allows one to cut along the shortest curve, induce a quasi-isometry on Teichmüller space of lower complexity and proceed by induction. Most of the discussion above also applies for the disconnected subsurfaces. Hence, much of Sections 5 and 6 is written in a way to apply to both $\mathcal{T}(S)$ and $\mathcal{T}(\Sigma, L)$ settings. The induction step is carried out in Section 7.

Acknowledgements

We would like to thank Brian Bowditch for helpful comments on an earlier version of this paper. We also thank the GEAR network for supporting this project.

2 Background

The purpose of this section is to establish notation and recall some statements from the literature. We refer the reader to [14; 12] for basic background on Teichmüller theory and to [18; 23] for some background on the geometry of the Teichmüller metric.

For much of this paper, the arguments are meant to apply to both $\mathcal{T}(S)$, which is the Teichmüller space of a connected surface, and to $\mathcal{T}(\Sigma, L)$, where Σ is disconnected and the space is truncated. In such situations, we use the notation $\mathcal{T}$ to refer to either case and use the full notation where the discussion is specific to one case or the other. Similarly, f denotes either a quasi-isometry f_S of $\mathcal{T}(S)$ or a quasi-isometry f_Σ of $\mathcal{T}(\Sigma, L)$. A similar convention is also applied to other notation as we suppress symbols S or Σ to unify the discussion in the two cases. For example, K and C could refer to K_S and C_S or to K_Σ and C_Σ, and ξ could refer to $\xi(S)$ or $\xi(\Sigma)$.

By a *curve*, we mean the free isotopy class of a nontrivial, nonperipheral simple closed curve in either S or Σ. Also, by a subsurface we mean a free isotopy class of a subsurface U, where the inclusion map induces an injection between the fundamental groups. We always assume that U is not a pair of pants. When we say γ is a curve in U, we always assume that it is not peripheral in U (not just in S). We write $\alpha \subset \partial U$ to indicate that the curve α is a boundary component of U.

2.1 Product regions

We often examine a point $x \in \mathcal{T}$ from the point of view of its subsurfaces. For every subsurface U, there is a projection map

$$\psi_U \colon \mathcal{T} \to \mathcal{T}(U)$$

defined using the Fenchel–Nielsen coordinates; see [18] for details. When the boundary of U is not short, these maps are not well-behaved. Hence, these maps should be applied only when there is an upper bound on the length of ∂U; see below.

When U is an annulus, or when $\xi(U) = 1$, the space $\mathcal{T}(U)$ can be identified with the hyperbolic plane $\mathbb{H}$; however the Teichmüller metric differs from the usual hyperbolic metric by a factor of 2. We always assume that $\mathbb{H}$ is equipped with this metric, which has constant curvature -4. We denote the point $\psi_U(x)$ simply by x_U.

A *decomposition* of S is a set $\mathcal{U}$ of pairwise disjoint subsurfaces of S that fill S. Subsurfaces in $\mathcal{U}$ are allowed to be annuli (but not pairs of pants). In this context, *filling* means that every curve in S either intersects or is contained in some $U \in \mathcal{U}$. In particular, for every $U \in \mathcal{U}$, the annulus associated to every boundary curve of U is also included in $\mathcal{U}$. The same discussion holds for Σ.

For a decomposition $\mathcal{U}$ and $\ell_0 > 0$ define

$$\mathcal{T}_{\mathcal{U}} = \{x \in \mathcal{T} \mid \operatorname{Ext}_x(\alpha) \le \ell_0 \text{ for all } U \in \mathcal{U} \text{ and } \alpha \subset \partial U\}.$$

Theorem 2.1 (product regions theorem [18]) *For ℓ_0 small enough,*

$$\psi_{\mathcal{U}} = \prod_{U \in \mathcal{U}} \psi_U \colon \mathcal{T}_{\mathcal{U}} \to \prod_{U \in \mathcal{U}} \mathcal{T}(U)$$

is an isometry up to a uniform additive error D_{pr}. Here, the product on the right-hand side is equipped with the sup metric. That is, for x^1 and x^2,

$$(2) \qquad d_{\mathcal{T}}(x^1, x^2) - \mathsf{D}_{\mathrm{pr}} \le \sup_{U \in \mathcal{U}} d_{\mathcal{T}(U)}(x^1_U, x^2_U) \le d_{\mathcal{T}}(x^1, x^2) + \mathsf{D}_{\mathrm{pr}}.$$

For the rest of the paper, we fix a value for ℓ_0 that makes this statement hold. We also assume that two curves of length less than ℓ_0 do not intersect. We refer to $\mathcal{T}_{\mathcal{U}}$ as *the product region associated to* $\mathcal{U}$.

A point $x \in \mathcal{T}_{\mathcal{U}}$ can be coarsely described by its projection $\{x_U\}_{U \in \mathcal{U}}$. We often say "let $x \in \mathcal{T}_{\mathcal{U}}$ be a point whose projection to every $U \in \mathcal{U}$ is a given point in $\mathcal{T}(U)$". Such a description does not determine x uniquely, but it does up to a distance of at most D_{pr}.

We say $\mathcal{U}$ is *maximal* if every $U \in \mathcal{U}$ is either an annulus, or $\xi(U) = 1$ (recall that pairs of pants are always excluded). That is, if $\mathcal{U}$ is maximal then the associated product region is isometric, up to an additive error D_{pr}, to a subset of a product space $\prod_{i=1}^{\xi} \mathbb{H}$ equipped with the sup metric. For the rest of the paper, we always assume the product $\prod_{i=1}^{\xi} \mathbb{H}$ is equipped with the sup metric. For points $x, y \in \mathcal{T}(S)$, we say *x and y are in the same maximal product region* if there is a maximal decomposition $\mathcal{U}$ with $x, y \in \mathcal{T}_{\mathcal{U}}$. Note that such a $\mathcal{U}$ is not unique. For example, let P be a pants decomposition and let x be a point such that the length in x of every curve in P is less than ℓ_0. Then there are many decompositions $\mathcal{U}$ where every $U \in \mathcal{U}$ is either a punctured torus or a four-times-punctured sphere with $\partial U \subset P$ or an annulus whose core curve is in P. The point x belongs to $\mathcal{T}_{\mathcal{U}}$ for every such decomposition $\mathcal{U}$.

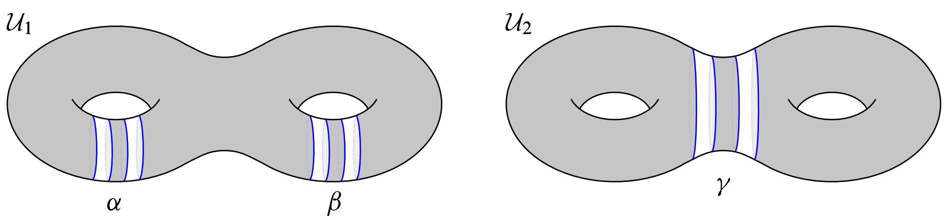

Figure 1: Two maximal decompositions $\mathcal{U}_1$ and $\mathcal{U}_2$ are depicted above. Note that $\mathcal{U}_1 \cap \mathcal{U}_2 = \varnothing$. However, for the pants decomposition $P = \{\alpha, \beta, \gamma\}$, any point $x \in \mathcal{T}$ for which all curves in P have a length less than ℓ_0 is contained in $\mathcal{T}_{\mathcal{U}_1} \cap \mathcal{T}_{\mathcal{U}_2}$.

2.2 Short curves on a surface

The thick part $\mathcal{T}_{\mathrm{thick}}$ of $\mathcal{T}$ is the set of points x for which $\mathrm{Ext}_X(\gamma) \ge \ell_0$ for every curve γ. There is a constant B (the Bers constant) such that for any point $z \in \mathcal{T}_{\mathrm{thick}}$, the set of curves of extremal length at most B *fills* the surface. That is, every curve intersects a curve of length at most B. Note that every $x \in \mathcal{T}$ contains a curve of length at most B. In fact, we choose B large enough so that every $x \in \mathcal{T}$ has a pants decomposition of length at most B. We call such a pants decomposition the *short pants decomposition at* x and denote it by P_x. We also assume that ℓ_0 is small enough that if $\mathrm{Ext}_x(\alpha) \le \ell_0$ (α is ℓ_0–short) then α does not intersect any B–short curve. Hence, P_x contains every ℓ_0–short curve.

It is often more convenient to work with the logarithm of length. For $x \in T$ and $\alpha \in P_x$, define $\tau_x(\alpha)$ to be the largest number such that if $d_{\mathcal{T}}(x, x') \le \tau_x(\alpha)$ then $\mathrm{Ext}_{x'}(\alpha) \le \ell_0$. From the product regions theorem, we have

$$\left| \tau_x(\alpha) - \log \frac{1}{\mathrm{Ext}_x(\alpha)} \right| = O(1),$$

where the constant on the right-hand side depends on the value of D_{pr} and $\log 1/\ell_0$. We often need to *pinch* a curve. Let $x \in \mathcal{T}$ and τ be given and let $\alpha \in P_x$ with $\tau_x(\alpha) = O(1)$. Let x' be a point with $d_{\mathcal{T}}(x, x') = O(1)$ and for which the length of α is ℓ_0. Let $U = S - \alpha$ and let x'' be the point such that

$$\tau_{x''}(\alpha) = \tau, \quad x'_U = x''_U, \quad \Re(x'_\alpha) = \Re(x''_\alpha).$$

The last condition means x' and x'' have no relative twisting around α. We then say *x'' is a point obtained from x by pinching α*. There is a constant $\mathsf{d}_{\mathrm{pinch}}$ such that

$$\tau \le d_{\mathcal{T}}(x, x'') \le \tau + \mathsf{d}_{\mathrm{pinch}}.$$

2.3 Subsurface projection

Let U be a subsurface of S with $\xi(U) \ge 1$. Let $\mathcal{C}(U)$ denote the curve graph of U; that is, a graph where a vertex represents a curve in U, and an edge represents a pair of disjoint curves. When $\xi(U) = 1$, the subsurface U does not contain disjoint curves. Here an edge is a pair of curves that intersect minimally; once in the punctured-torus case and twice in the four-times-punctured sphere case. In the case that U is an annulus with core curve α, in place of the curve complex we use the subset $H_\alpha \subset \mathcal{T}(U)$ of all points for which the extremal length of α is at most ℓ_0. This is a horoball in $\mathbb{H} = \mathcal{T}(U)$. Depending on context, we use the notation $\mathcal{C}(U)$ or H_α.

There is a projection map

$$\pi_U \colon \mathcal{T}(U) \to \mathcal{C}(U)$$

that sends a point $z \in \mathcal{T}(U)$ to a curve γ in U with $\mathrm{Ext}_z(\gamma) \le \mathsf{B}$. This is not unique but the image has a uniformly bounded diameter and hence the map is coarsely well-defined.

When U is an annulus, $\pi_U(z)$ is the same as $\psi_{\mathcal{U}}(z)$ if $\mathrm{Ext}_z(\alpha) \le \ell_0$ and, otherwise, is the point on the boundary of H_α where the real value is twisting of z around α. (See [23, Section 2] for the definition and discussion of twisting.)

We can also define a projection $\pi_U(\gamma)$, where γ is any curve that intersects U nontrivially. If $\gamma \subset U$ then choose the projection to be γ. If γ is not contained in U then $\gamma \cap U$ is a collection of arcs with endpoints on ∂U. Choose one such arc and perform a surgery using this arc and a subarc of ∂U to find a point in $\mathcal{C}(U)$. The choice of different arcs or different choices of intersecting pants curves determines a set of diameter 2 in $\mathcal{C}(U)$; hence the projection is coarsely defined. Note that this is not defined when γ is disjoint from U. We also define a projection $\mathcal{T}(S) \to \mathcal{C}(U)$ to be $\pi_U \circ \pi_S$, however we still denote it by π_U. For $x, y \in \mathcal{T}$, we define

$$d_U(x,y) := d_{\mathcal{C}(U)}(\pi_U(x), \pi_U(y)).$$

For curves α and β, $d_U(\alpha,\beta)$ is similarly defined.

In fact, the subsurface projections can be used to estimate the distance between two points in $\mathcal{T}$; see [22]. There is a threshold T such that

$$d_{\mathcal{T}}(x,y) \stackrel{*}{\asymp} \sum_{W \in \mathcal{W}_{\mathsf{T}}} d_W(x,y), \tag{3}$$

where $\mathcal{W}_{\mathsf{T}}$ is the set of subsurfaces W where $d_W(x,y) \ge \mathsf{T}$.

Definition 2.2 We say a pair of points $x, y \in \mathcal{T}(S)$ are M–cobounded relative to a subsurface $U \subset S$ if ∂U is ℓ_0–short in x and y and if $d_V(x,y) \le M$ for every surface $V \ne U$. If $U = S$, we simply say x and y are M–cobounded.

Similarly, we say a pair of curves α, β in U are M–cobounded relative to U if for every $V \subsetneq U$, we have $d_V(\alpha,\beta) \le M$ when defined. If $U = S$, we simply say α and β are M–cobounded.

A path g in $\mathcal{T}(S)$ or in $\mathcal{C}(U)$ is M–cobounded relative to U if every pair of points in g are M–cobounded relative to U. Once and for all, we choose a constant M so that through every point $x \in \mathcal{T}(S)$ and for every U whose boundary length is at most ℓ_0 in x, there is a bi-infinite path in $\mathcal{T}(S)$ passing through x that is M–cobounded relative to U. One can, for example, take an axis of a pseudo-Anosov element in $\mathcal{T}(U)$ and then use the product regions theorem to elevate that to a path in $\mathcal{T}(S)$ whose projections to disjoint subsurfaces are constant. When we say a geodesic g in $\mathcal{T}(S)$ is cobounded relative to U, we always mean that it is M–cobounded relative to U. From equation (3) (the distance formula) we have, for x and y along such g,

$$d_{\mathcal{T}}(x,y) \stackrel{*}{\asymp} d_U(x,y). \tag{4}$$

See below for the definition of the notation $\stackrel{*}{\asymp}$.

3 Rank is preserved

In this section, we recall some results from [11] and we develop them further to show that the set of points in Teichmüller space with maximum rank is coarsely preserved; see Proposition 3.8 below. Since the notation $\stackrel{*}{\asymp}$ and $\stackrel{*}{\succ}$ were used in [11], we continue to use them in this section. Recall from [11] that $\boldsymbol{A} \stackrel{*}{\asymp} \boldsymbol{B}$ means there is a constant C, depending only on the topology of S or Σ, such that $\frac{1}{C}\boldsymbol{A} \le \boldsymbol{B} \le C\boldsymbol{A}$. We say $\boldsymbol{A}$ and $\boldsymbol{B}$ are *comparable*. Similarly, $\boldsymbol{A} \stackrel{+}{\asymp} \boldsymbol{B}$ means there is a constant C, depending only on the topology of S or Σ, such that $\boldsymbol{A} - C \le \boldsymbol{B} \le \boldsymbol{A} + C$. We say $\boldsymbol{A}$ and $\boldsymbol{B}$ are *the same up to an additive error*. For the rest of the paper, we will name our constants explicitly, as we did with D_{pr} and $\mathsf{d}_{\mathrm{pinch}}$. The only constant from this section that is used later is the constant d_0 from Proposition 3.8.

Coarse differentiation and preferred paths

A path $g\colon [a,b] \to \mathcal{T}$ is called *a preferred path* if, for every subsurface U, the image of $\pi_U \circ g$ is a reparametrized quasigeodesic in $\mathcal{C}(U)$. We use preferred paths as coarse analogues of straight lines in $\mathcal{T}$. Teichmüller geodesics are examples of preferred paths, but the latter are much more general. One explanation of this greater generality is that given a pair of disjoint domains, a preferred path may move in the curve complex of each essentially arbitrarily as long as they are quasigeodesics, while Teichmüller geodesics move in a product region in a manner determined by the endpoints of the geodesic.

Definition 3.1 A *box* in $\mathbb{R}^n$ is a product of intervals; namely $B = \prod_{i=1}^n I_i$, where I_i is an interval in $\mathbb{R}$. We say a box B is *of size* R if $|I_i| \stackrel{*}{\succ} R$ for every i and if the diameter of B is less than R. Note that if B is of size R and of size R', then $R \stackrel{*}{\asymp} R'$. The box in $\mathbb{R}^n$ is always assumed to be equipped with the usual Euclidean metric.

For points $a, b \in B$, we often treat the geodesic segment $[a,b]$ in B as an interval of times parametrized by t. A map $f\colon B \to \mathcal{T}$ from a box of size R in $\mathbb{R}^n$ to $\mathcal{T}$ is called

ϵ*–efficient* if, for any pair of points $a, b \in B$, there is a preferred path $g\colon [a,b] \to \mathcal{T}$ such that, for $t \in [a,b]$,

$$d_{\mathcal{T}}(f(t), g(t)) \le \epsilon R.$$

Let B be a box of size L in $\mathbb{R}^n$ and let $\underline{B}$ be a central sub-box of B with comparable diameter (say half). For any constant $0 < R \le \frac{1}{3}L$, let $\mathcal{B}_R$ be a subdivision of $\underline{B}$ into boxes of size R. That is,

(1) boxes in $\mathcal{B}_R$ are of size R,

(2) they are contained in $\underline{B}$ and hence their distance to the boundary of B is comparable to L,

(3) they have disjoint interiors and

(4) $|\mathcal{B}_R| \stackrel{*}{\asymp} (L/R)^n$.

The following combines Theorem 2.5 and Theorem 4.9 in [11].

Theorem 3.2 (coarse differentiation) *For every K, C, ϵ, θ and R_0 there is an L_0 such that the following holds. For $L \ge L_0$, let $f\colon B \to \mathcal{T}$ be a (K,C)–quasi-Lipschitz map, where B is a box of size L in $\mathbb{R}^n$. Then, there is a size $\frac{1}{3}L \ge R \ge R_0$ such that the proportion of boxes $B' \in \mathcal{B}_R$ where $f|_{B'}$ is ϵ–efficient is at least $(1-\theta)$.*

Even though we have no control over the distribution of efficient boxes, the following lemma says we can still connect every two points in $\underline{B}$ with a path that does not intersect too many nonefficient boxes.

Lemma 3.3 *Let L, R, $\mathcal{B}_R$, ϵ and θ be as above. Then, for any pair of points $a, b \in \underline{B}$, there is a path γ in $\underline{B}$ connecting them so that γ is covered by at most $O(L/R)$ boxes and the number of boxes in the covering that are not ϵ–efficient is at most $O(\sqrt[n]{\theta}L/R)$.*

Proof Let $N = \sqrt[n]{\theta}L/R$. First assume that the distance between a and b to the boundary of $\underline{B}$ is at least NR. Consider the geodesic segment $[a,b]$. Take $(n-1)$–dimensional totally geodesic boxes Q_a and Q_b containing a and b respectively that are perpendicular to $[a,b]$, parallel to each other and have a diameter NR. Choose an R–net of points $p_1, \dots, p_k$ in Q_a and $q_1, \dots, q_k$ in Q_b so that $[p_i, q_i]$ is parallel to $[a,b]$. We have, for $1 \le i, j \le k$,

$$d_{\mathbb{R}^n}([p_i, q_i], [p_j, q_j]) > R \quad \text{and} \quad k \stackrel{*}{\asymp} N^{n-1}.$$

For $1 \le i \le k$, let γ_i be the path that is a concatenation of geodesic segments $[a, p_i]$, $[p_i, q_i]$ and $[q_i, b]$. We claim one of these paths satisfies the conclusions of the lemma.

Assume, for the sake of contradiction, that the number of nonefficient boxes along each γ_i is larger than cN for some large $c > 0$. Then the total number of nonefficient boxes is at least

$$kcN = cN^n = c\,\theta\Big(\frac{L}{R}\Big)^n.$$

But this is not possible for large enough values of c; see property (4) of $\mathcal{B}_R$ above and Theorem 3.2. Hence, there is a $c = O(1)$ and i such that $[p_i, q_i]$ intersects at most cN nonefficient boxes.

Note that the segments $[a, p_i]$ and $[q_i, b]$ intersect at most N boxes each. Hence, the number of inefficient boxes intersecting γ_i is at most $(c+2)N$. In the case that a or b are close to the boundary, we choose points a' and b' nearby (distance NR) and apply the above argument to find an appropriate path between a' and b' and then concatenate this path with segments $[a, a']$ and $[b, b']$. The total number of inefficient boxes along this path is at most $(c+4)N$. This finishes the proof. □

Efficient quasi-isometric embeddings

In this section, we examine efficient maps that are also assumed to be quasi-isometric embeddings. We will show that they have *maximal rank*; they make small progress in any subsurface W with $\xi(W) \ge 2$.

Definition 3.4 Let $\mathcal{U}$ be a decomposition of S. For every $U \in \mathcal{U}$, let $g_U\colon I_U \to \mathcal{T}(U)$ be a preferred path. Consider the box $B = \prod_U I_U \subset \mathbb{R}^m$, where m is the number of elements in $\mathcal{U}$. Consider the map

$$F\colon B \to \mathcal{T}_{\mathcal{U}} = \prod_{U \in \mathcal{U}} \mathcal{T}(U), \quad \text{where } F = \prod_{U \in \mathcal{U}} g_U.$$

Then F is a quasi-isometric embedding because each g_U is a quasigeodesic. We call this map a *standard flat* in $\mathcal{T}_{\mathcal{U}}$. The dimension of a standard flat is defined to be the number of elements in $\mathcal{U}$. The maximal dimension is ξ.

The map $\bar{f}$ below will be a modified version of our map f from Theorem 1.1. The next theorem is a basic tool. It says that if a large box of maximal dimension is quasi-isometrically embedded and mapped efficiently, then the projection to a cobounded geodesic in any subsurface must have small diameter.

Theorem 3.5 *For every K, C and M, there is an ϵ and R_0 such that the following holds. Let $\overline{f}\colon B \to \mathcal{T}(S)$ be an ϵ–efficient (K, C)–quasi-isometric embedding defined on a box $B \subset \mathbb{R}^\xi$ of size $R \geq R_0$, let ω_0 be a M–cobounded geodesic in $\mathcal{C}(W)$, where W is a subsurface with $\xi(W) \geq 2$, and let π_{ω_0} be the closest-point projection map from $\mathcal{C}(W)$ to ω_0. Define*

$$\pi = \pi_{\omega_0} \circ \pi_W \circ \overline{f}.$$

Then, for all points $a, b \in B$, we have

$$d_W(\pi(a), \pi(b)) \leq \sqrt{\epsilon} R.$$

Proof Assume by way of contradiction that for all large R and all small ϵ there is a subsurface W, $\xi(W) \geq 2$, an M–cobounded geodesic ω_0 in $\mathcal{C}(W)$, a box B of size R in $\mathbb{R}^m$, an ϵ–efficient map $\overline{f}\colon B \to \mathcal{T}$ and a pair of points $a, b \in B$ such that

$$d_W(\pi(a), \pi(b)) \geq \sqrt{\epsilon} R. \tag{5}$$

Let $g\colon [a, b] \to \mathcal{T}$ be the preferred path joining $\overline{f}(a)$ and $\overline{f}(b)$ coming from the efficiency assumption. Since g is a preferred path, if ω is a geodesic in $\mathcal{C}(W)$ joining $\pi_W(\overline{f}(a))$ and $\pi_W(\overline{f}(b))$, then ω can be reparametrized so that

$$d_W(g(t), \omega(t)) = O(1).$$

As stated in equation (5), we are assuming that the projection of the geodesic ω to ω_0 has a length of at least $\sqrt{\epsilon} R$.

The hyperbolicity of $\mathcal{C}(W)$ implies that ω and therefore $\pi_W \circ g$ lie in a uniformly bounded neighborhood of ω_0 along a segment of g of length $\stackrel{*}{\succ} \sqrt{\epsilon} R$. Divide this piece of ω_0 into three segments. Let $\alpha_1, \alpha_2, \alpha_3, \alpha_4 \in \mathcal{C}(W)$ be the corresponding endpoints of these segments, which for $i = 1, 2, 3$ satisfy

$$d_W(\alpha_i, \alpha_{i+1}) \stackrel{*}{\succ} \sqrt{\epsilon} R. \tag{6}$$

Let $\omega_{\mathrm{mid}} = [\alpha_2, \alpha_3]$ be the middle segment, let $[c, d] \subset [a, b]$ be the associated time interval and let $g_{\mathrm{mid}} = g|_{[c,d]}$.

Step 1 We claim that for every V such that $V \cap W \neq \varnothing$ (that is, either $V \subset W$ or $V \pitchfork W$), the image of the projection of g_{mid} to $\mathcal{C}(V)$ has a bounded diameter. Note that, since g is a preferred path, it is enough to prove either $d_V(\alpha_2, \alpha_3)$ or $d_V(\alpha_1, \alpha_4)$ is uniformly bounded, assuming those curves intersect V.

We argue in two cases. If every curve in $[\alpha_2,\alpha_3]$ intersects V, the claim follows from the bounded geodesic image theorem [17, Theorem 3.1]. Otherwise, ∂V is close in $\mathcal{C}(W)$ to this segment and hence it is far in $\mathcal{C}(W)$ from curves α_1 and α_4. Let $\overline{\alpha}_1$ and $\overline{\alpha}_4$ be curves on ω_0 that are close to α_1 and α_4, respectively. Then, ∂V intersects every curve in $[\alpha_1,\overline{\alpha}_1]$ and $[\alpha_4,\overline{\alpha}_4]$ and, by the bounded geodesic image theorem, the projections of these segments to $\mathcal{C}(V)$ have bounded diameters. But ω_0 is cobounded. Hence, $d_V(\overline{\alpha}_1,\overline{\alpha}_4)=O(1)$ and therefore $d_V(\alpha_1,\alpha_4)=O(1)$. This proves the claim.

Step 2 To obtain a contradiction, we will find a large sub-box of B that maps near a standard flat F of maximal rank.

Note that the map π above is quasi-Lipschitz. Choose a constant D which is large compared to the quasi-Lipschitz constant of π and the hyperbolicity constant of $\mathcal{C}(W)$. Let a',b' be points in B in a neighborhood of a,b, respectively, such that

$$\|a-a'\| \le \frac{\sqrt{\epsilon}R}{D} \quad \text{and} \quad \|b-b'\| \le \frac{\sqrt{\epsilon}R}{D}. \tag{7}$$

Let g' be the preferred path joining $\overline{f}(a')$, $\overline{f}(b')$ and ω' be the geodesic in $\mathcal{C}(W)$ connecting $\pi_W(\overline{f}(a'))$ to $\pi_W(\overline{f}(b'))$. Consider the quadrilateral

$$\beta=\pi_W(\overline{f}(a)),\quad \gamma=\pi_W(\overline{f}(b)),\quad \beta'=\pi_W(\overline{f}(a')),\quad \gamma'=\pi_W(\overline{f}(b'))$$

in $\mathcal{C}(W)$. From the assumption on D, the edges $[\beta,\beta']$ and $[\gamma,\gamma']$ are short compared to $[\beta,\gamma]$. From the hyperbolicity of $\mathcal{C}(W)$, we conclude that ω' has a subsegment ω'_{mid} that has a bounded Hausdorff distance to ω_{mid}. Let g'_{mid} be the associated subsegment of g' (see previous step). As we argued in the previous step, the projection of g'_{mid} to $\mathcal{C}(V)$ has a bounded diameter for every $V\cap W\neq\varnothing$. In fact, it is close to the projection of g_{mid} to $\mathcal{C}(V)$.

The union of subsegments of type $[a',b']$ fill a $(\sqrt{\epsilon}R/D)$–neighborhood of $[c,d]$ and

$$|d-c| \overset{*}{\asymp} \sqrt{\epsilon}R.$$

Therefore, there is a subbox $B'\subset B$ of size $R'\overset{*}{\asymp}\sqrt{\epsilon}R$ such that $(\pi_V\circ\overline{f})(B')$ has bounded diameter for every $V\cap W\neq\varnothing$. For ϵ_0 small to be chosen later (it will depend only on universal constants), set $\epsilon_\xi=\epsilon_0^{6^\xi}$ and assume ϵ is small enough so that

$$\sqrt{\epsilon}<\epsilon_\xi\epsilon_0.$$

Since f is ϵ–efficient and $\epsilon<\epsilon_\xi$ it is ϵ_ξ–efficient. By Theorem 7.2 of [11] (which can be applied if R is large enough) there is a sub-box $B''\subset B'$ of size $R''\ge\epsilon_\xi R'$ such that $\bar{f}(B'')$ is within $O(\epsilon_0 R'')$ of a standard flat F. The implied constants depend only on K, C and ξ.

We show this is impossible for ϵ_0 sufficiently small. Note that ξ is the maximum dimension of any standard flat. Since B'' is a box of dimension ξ and $\bar{f}$ is a quasi-isometric embedding, the standard flat F must have dimension ξ as well. Let $\mathcal{V}$ be the decomposition of S with $|\mathcal{V}| = \xi$, and let $F_V \colon I_V \to \mathcal{T}(V)$ be the preferred paths, where

$$F \colon \prod_{V \in \mathcal{V}} I_V \to \mathcal{T}.$$

Then $\bar{f}(B'')$ is contained in the $O(\epsilon_0 R'')$–neighborhood of the image of F. We assume I_V is the smallest possible interval for which this holds. Then, for $V \in \mathcal{V}$, $F_V(I_V)$ has a diameter comparable to R'', which is the size of B''.

Since $|\mathcal{V}| = \xi$, every $V \in \mathcal{V}$ is either an annulus or $\xi(V) = 1$. Hence, they cannot equal W and, for at least one $V \in \mathcal{V}$, we have $V \cap W \neq \varnothing$. In fact, we can assume V is an annulus, because $\mathcal{V}$ is maximal and if a subsurface is in $\mathcal{V}$ the annuli associated to its boundary curves are also in $\mathcal{V}$.

From the assumption of the minimality of lengths of I_V, we know that every $t_V \in I_V$ can be completed to a vector in $\prod_{V \in \mathcal{V}} I_V$ whose image is in the $O(\epsilon_0 R'')$–neighborhood of $\bar{f}(B'') \subset \bar{f}(B')$. In addition, since $\bar{f}$ is ϵ–efficient, any point in $\bar{f}(B')$ is ϵR–close to some g'_{mid}. We already know that for any such V the projection of g'_{mid} to $\mathcal{C}(V)$ is $O(1)$. Combining these statements we find that the projection of the image of F to $\mathcal{C}(V)$ of any such V has a diameter $O(\epsilon R + \epsilon_0 R'')$. This means the same bound also holds for the diameter of the projection to $\mathcal{T}(V)$; in the case where V is an annulus and ∂V is short, the two distances are the same. We have shown

$$R'' \stackrel{*}{\prec} \operatorname{diam}_{\mathcal{T}(V)}(F_V(I_V)) \stackrel{*}{\prec} \epsilon R + \epsilon_0 R''.$$

Therefore, $R'' \stackrel{*}{\prec} \epsilon R$. But

$$R'' \geq \epsilon_\xi R' \stackrel{*}{\succ} \epsilon_\xi \sqrt{\epsilon} R \geq \frac{\epsilon R}{\epsilon_0}.$$

For ϵ_0 sufficiently small, this is a contradiction. That is, the theorem holds for appropriate values of ϵ and R_0. □

Maximal rank is preserved

Recall that, for $x \in \mathcal{T}(S)$ and a curve α, $\tau_x(\alpha)$ is the largest number such that if $d(x, x') \leq \tau_x(\alpha)$ then $\operatorname{Ext}_{x'}(\alpha) \leq \ell_0$.

For a point x in $\mathcal{T}(S)$, let $\mathcal{S}_x = \mathcal{S}_x(\ell_0)$ be the set of curves α such that $\mathrm{Ext}_x(\alpha) \le \ell_0$.

Definition 3.6 A point $x \in \mathcal{T}(S)$ has *maximal rank* when the components of $S \setminus \mathcal{S}_x$ are all either a pair of pants, a once-punctured torus or a four-times-punctured sphere. Denote the set of points with maximal rank by $\mathcal{T}_{\mathrm{MR}}$. Let $\mathcal{T}_{\mathrm{LR}}$ be its complement; the set with lower rank.

Definition 3.7 Suppose $x \in \mathcal{T}_{\mathrm{MR}}$. We say a curve $\alpha \in \mathcal{S}_x$ is *isolated* if by increasing its length to ℓ_0 while keeping all other lengths the same one leaves $\mathcal{T}_{\mathrm{MR}}$. A pair of curves in $\mathcal{S}_x$ are called *adjacent* if increasing both of their lengths to ℓ_0 one leaves $\mathcal{T}_{\mathrm{MR}}$.

The importance of this definition is that if $d(x, \mathcal{T}_{\mathrm{LR}}) \ge d$, then every isolated curve α satisfies $\tau_x(\alpha) \ge d$, and for any pair of adjacent curves α_1, α_2, at least one of them satisfies $\tau_x(\alpha_i) \ge d$.

Proposition 3.8 *There exists a* $\mathsf{d}_0 > 0$ *such that given* $x \in \mathcal{T}$, *if* $d(x, \mathcal{T}_{\mathrm{LR}}) \ge \mathsf{d}_0$ *then* $f(x) \in \mathcal{T}_{\mathrm{MR}}$.

Proof Suppose by way of contradiction that, for large d_0, we have a point x such that $d(x, \mathcal{T}_{\mathrm{LR}}) \ge \mathsf{d}_0$ but $f(x) \in \mathcal{T}_{\mathrm{LR}}$. Then there is a surface W (possibly the whole surface) with $\xi(W) \ge 2$ such that the boundary curves of W are ℓ_0–short on $f(x)$, but no curve in W is shorter than ℓ_0 on $f(x)$.

Let g_0 be a path passing through $f(x)$ that is M–cobounded relative to W; see the discussion after Definition 2.2. Let $f(y)$ be a point in g_0 such that the distance in $\mathcal{T}$ between $f(x)$ and $f(y)$ is $L = \mathsf{d}_0/2\mathsf{K} - \mathsf{C}$. Since g_0 is M–cobounded, from equation (4) we have

$$d_W(f(x), f(y)) \stackrel{*}{\succ} L. \tag{8}$$

Since f is a (K, C)–quasi-isometry, solving equation (1) for $d(x, y)$ we get

$$\frac{\mathsf{d}_0}{2\mathsf{K}^2} - \frac{2\mathsf{C}}{\mathsf{K}} \le d(x, y) \le \frac{\mathsf{d}_0}{2}.$$

In particular, $y \in \mathcal{T}_{\mathrm{MR}}$.

Note that, in addition, for $\mathsf{d}_0 \ge 2\,\mathrm{Log}(1/\ell_0)$, there is a maximal product region $\mathcal{T}_{\mathcal{U}}$ containing both x and y. For $d_{\mathcal{T}}(x, \mathcal{T}_{\mathrm{LR}}) \ge \mathsf{d}_0$ implies that there is a set of curves $\boldsymbol{\alpha}$ such that $\tau_x(\alpha) \ge \mathsf{d}_0$ for $\alpha \in \boldsymbol{\alpha}$, and such that the complementary regions have complexity at most one. For all these curves, we have $\tau_y(\alpha) \geq \frac{1}{2}\mathsf{d}_0$; in particular they are at least ℓ_0 short in y. This means both x and y are in $\mathcal{T}_{\mathcal{U}}$.

In fact, there is a box

$$B = \prod_{U \in \mathcal{U}} I_U \subset \mathbb{R}^\xi$$

of size L and a quasi-isometry

$$Q = \prod_{U \in \mathcal{U}} Q_U \colon B \to \mathcal{T}$$

such that each $Q_U \colon I_U \to \mathcal{T}(U)$ is a geodesic and x and y are contained in $Q(\underline{B})$ (recall that $\underline{B}$ is the central sub-box of half the diameter). The map Q is a quasi-isometry because B is equipped with the Euclidean metric and $T_{\mathcal{U}}$ is equipped with the sup metric up to an additive error of D_{pr}. Define

$$\bar{f} \colon B \to \mathcal{T} \quad \text{by} \quad \bar{f} = f \circ Q.$$

Then $\bar{f}$ is a (K, C)–quasi-isometric embedding, where K and C depend on K, C and the complexity $\xi = |\mathcal{U}|$. (D_{pr} depends only on these constants.)

Let ω_0 be the geodesic in $\mathcal{C}(U)$ that fellow-travels the projection $\pi_U(g_0)$. Then ω_0 is M–cobounded with M slightly larger than M. Define

$$\pi = \pi_{\omega_0} \circ \pi_U \circ \bar{f},$$

and let $l_\pi \overset{*}{\asymp} \mathsf{K}$ be the Lipschitz constant of π.

Let ϵ and R_0 be constants from Theorem 3.5 associated to K, C and M and choose θ so that $\sqrt[\xi]{\theta}\, l_\pi$ is small (see below). Then, let L_0 be the constant given by Theorem 3.2 (the dimension n equals ξ). Choose d_0 large enough that

$$L = d(x, y) \geq \frac{\mathsf{d}_0}{2\mathsf{K}^2} - \frac{2\mathsf{C}}{\mathsf{K}} \geq L_0.$$

Applying Theorem 3.2 to B, we conclude that there is a scale R and a decomposition $\mathcal{B}_R$ of $\underline{B}$ to boxes of size R so that a proportion at least $(1-\theta)$ of boxes in $\mathcal{B}_R$ are ϵ–efficient.

By Lemma 3.3, there exists a path γ joining x to y that is covered by at most $O(L/R)$ boxes in $\mathcal{B}_R$ of which at most $O(\sqrt[\xi]{\theta} L/R)$ are not ϵ–efficient.

Assume γ intersects boxes $B_1, \dots, B_k$ and let γ_i be the subinterval of γ associated to B_i. By the triangle inequality, the sum of the diameters of $\pi(\gamma_i)$ is larger than $d_W(f(x), f(y))$, which is $\stackrel{*}{\succ} L$ by equation (8). However, by Theorem 3.5,

$$\sum_{B_i \text{ is efficient}} \operatorname{diam}_{\mathcal{C}(W)} \pi(\gamma_i) \stackrel{*}{\prec} \frac{L}{R}\sqrt{\epsilon} R = \sqrt{\epsilon} L. \tag{9}$$

The assumption on the number of nonefficient boxes gives

$$\sum_{B_i \text{ is not efficient}} \operatorname{diam}_{\mathcal{C}(W)} \pi(\gamma_i) \stackrel{*}{\prec} \left(\sqrt[\xi]{\theta} \frac{L}{R} \right) l_\pi \, R = \sqrt[\xi]{\theta}\, l_\pi \, L. \tag{10}$$

For ϵ and θ small enough, equations (9) and (10) contradict the fact that the sum of the diameters is $\stackrel{*}{\succ} L$. This finishes the proof. □

4 Local splitting theorem

In this section, we prove a local version of a splitting theorem proven by Kleiner and Leeb [16] and Eskin and Farb [8].

Theorem 4.1 *For every K, C, $\bar{\rho}$ there are constants R_0, D and ρ such that for all*

$$z = (z_1, \dots, z_m) \in \prod_{i=1}^m \mathbb{H}_i$$

and $R \geq R_0$, the following holds. Let $B_R(z)$ be a ball of radius R centered at z in $\prod_{i=1}^m \mathbb{H}_i$ and let $\bar{f}\colon B_R(z) \to \prod_{i=1}^m \mathbb{H}_i$ be a (K, C)–quasi-isometric embedding whose image coarsely contains a ball of radius $\bar{\rho} R$ about $\bar{f}(z)$. Then there is a smaller ball $B_{\rho R}(z)$, a permutation $\sigma\colon \{1, 2, \dots, m\} \to \{1, 2, \dots, m\}$ and (K, C)–quasi-isometric embeddings

$$\phi_i\colon B_{\rho R}(z_i) \to \mathbb{H}_{\sigma(i)}$$

such that the restriction of $\bar{f}$ to $B_{\rho R}(z)$ is D–close to

$$\phi_1 \times \dots \times \phi_m\colon B_{\rho R}(z) \to \prod_{i=1}^m \mathbb{H}_{\sigma(i)}.$$

Remark 4.2 When we use this theorem in Section 5, we need to equip $\prod \mathbb{H}$ with the L^∞–metric. However, it is more convenient to use the L^2–metric for the proof. Note that if $\bar{f}$ is a quasi-isometry with respect to one metric, it is also a quasi-isometry with respect to the other. For the rest of this section, we assume $\prod \mathbb{H}$ is equipped with the L^2–metric. To simplify notation, we use $d_{\mathbb{H}}$ to denote the distance in both in $\mathbb{H}$ and in $\prod \mathbb{H}$.

For the proof, we will use the notion of an asymptotic cone. Our brief discussion is taken from [16]. A nonprincipal ultrafilter is a finitely additive probability measure ω on the subsets of the natural numbers $\mathbb{N}$ such that

- $\omega(S) = 0$ or 1 for every $S \subset \mathbb{N}$,
- $\omega(S) = 0$ for every finite subset $S \subset \mathbb{N}$.

Given a bounded sequence $\{a_n\}$ in $\mathbb{R}$, there is a unique limit point $a_\omega \in \mathbb{R}$ such that for every neighborhood U of a_ω, the set $\{n \mid a_n \in U\}$ has full ω measure. We write $a_\omega = \omega\text{-}\lim a_n$.

Let $(\mathcal{X}_n, d_n, *_n)$ be a sequence of metric spaces with basepoints. Consider

$$\mathcal{X}_\infty = \bigl\{\vec{x} = (x_1, x_2, \dots) \in \prod \mathcal{X}_i \mid d_i(x_i, *_i) \text{ is bounded}\bigr\}.$$

Define $\bar{d}_\omega\colon \mathcal{X}_\infty \times \mathcal{X}_\infty \to \mathbb{R}$ by

$$\bar{d}_\omega(\vec{x}, \vec{y}) = \omega\text{-}\lim d_i(x_i, y_i).$$

Now $\bar{d}_\omega$ is a pseudodistance. Define the ultralimit of the sequence $(\mathcal{X}_n, d_n, *_n)$ to be the quotient metric space $(\mathcal{X}_\omega, d_\omega)$ identifying the points of distance zero.

Let $\mathcal{X}$ be a metric space and $*$ be a basepoint. The asymptotic cone $\mathrm{Cone}(\mathcal{X})$ of $\mathcal{X}$, with respect to the nonprincipal ultrafilter ω and the sequence λ_n of scale factors with $\omega\text{-}\lim \lambda_n = \infty$ and the basepoint $*$, is defined to be the ultralimit of the sequence of rescaled spaces $(\mathcal{X}_n, d_n, *_n) := (\mathcal{X}, (1/\lambda_n) d_n, *)$. The asymptotic cone is independent of the basepoint.

In the case of $\mathbb{H}$, the asymptotic cone $\mathbb{H}_\omega$ is a metric tree which branches at every point, and the asymptotic cone $\bigl(\prod_{i=1}^m \mathbb{H}\bigr)_\omega$ of the product of hyperbolic planes is $\prod_{i=1}^m \mathbb{H}_\omega$, the product of the asymptotic cones. A flat in $\prod_{i=1}^m \mathbb{H}$ is a product $\prod_{i=1}^m g_i$, where g_i is a geodesic in the i^{th} factor.

We first prove a version of Theorem 4.1 with small linear error term. We then show that, by taking an even smaller ball, the error term can be made to be uniform additive.

Proposition 4.3 *Given $K, C, \bar{\rho}$ there exist $\rho' > 0$ and D_0 such that for all sufficiently small $\epsilon > 0$, there exists an R_0 such that if $R \geq R_0$ and $\bar{f}$ is a (K, C)–quasi-isometric embedding defined on $B_R(z)$ with $\bar{f}(B_R(z))$ C–coarsely containing $B_{\bar{\rho}R}(\bar{f}(z))$, then:*

- *There is a permutation σ and, for $1 \le i \le m$, there is a quasi-isometric embedding $\phi_i^{\boldsymbol{z}}\colon B_{\rho' R}(z_i) \to \mathbb{H}_{\sigma(i)}$ such that, for*

$$\phi^{\boldsymbol{z}} = \phi_1^{\boldsymbol{z}} \times \cdots \times \phi_m^{\boldsymbol{z}}\colon B_{\rho' R}(\boldsymbol{z}) \to \prod_{i=1}^{m} \mathbb{H}$$

and for $\boldsymbol{x} \in B_{\rho' R}(\boldsymbol{z})$, we have

$$d_{\mathbb{H}}(\bar{f}(\boldsymbol{x}), \phi^{\boldsymbol{z}}(\boldsymbol{x})) \le \epsilon d_{\mathbb{H}}(\boldsymbol{z}, \boldsymbol{x}) + D_0. \tag{11}$$

- *For any $\boldsymbol{x} \in B_{\rho' R}(\boldsymbol{z})$ and any flat $F_{\boldsymbol{x}}$ through $\boldsymbol{x}$, there is a flat $F'_{\boldsymbol{x}}$ such that for $\boldsymbol{p} \in N_{\rho' R}(\boldsymbol{x}) \cap F_{\boldsymbol{x}}$,*

$$d_{\mathbb{H}}(\bar{f}(\boldsymbol{p}), F'_{\boldsymbol{x}}) \le m\epsilon d_{\mathbb{H}}(\boldsymbol{x}, \boldsymbol{p}) + D_0.$$

Proof We begin with a claim.

Claim Assume, for given ϵ and D_0, that there is a permutation σ and a constant ρ' so that, if $\boldsymbol{x}, \boldsymbol{y} \in B_{\rho' R}(\boldsymbol{z})$ differ only in the i^{th} factor, then $\bar{f}(\boldsymbol{x})$ and $\bar{f}(\boldsymbol{y})$ differ in all factors besides the $\sigma(i)^{\text{th}}$ factor by at most $\frac{1}{m}(\epsilon d_{\mathbb{H}}(\boldsymbol{y}, \boldsymbol{x}) + D_0)$. Then the first conclusion holds.

Proof of claim For $x \in B_{\rho' R}(z_i)$ and an index i define $\boldsymbol{z}_x^i$ to be a point whose i^{th} coordinate is x and whose other coordinates are the same as the coordinates of $\boldsymbol{z}$. We then define the map $\phi_i^{\boldsymbol{z}}$ by letting $\phi_i^{\boldsymbol{z}}(x)$ be the $\sigma(i)^{\text{th}}$ coordinate of $\bar{f}(\boldsymbol{z}_x^i)$. We show that $\phi_i^{\boldsymbol{z}}$ is a quasi-isometric embedding. For $x, x' \in B_{\rho' R}(z_i)$,

$$d_{\mathbb{H}}(\phi_i^{\boldsymbol{z}}(x), \phi_i^{\boldsymbol{z}}(x')) \le d_{\mathbb{H}}(\bar{f}(\boldsymbol{z}_x^i), \bar{f}(\boldsymbol{z}_{x'}^i)) \le K d_{\mathbb{H}}(\boldsymbol{z}_x^i, \boldsymbol{z}_{x'}^i) + C \le K d_{\mathbb{H}}(x, x') + C.$$

In addition, since $\boldsymbol{z}_x^i$ and $\boldsymbol{z}_{x'}^i$ differ in only one factor, for $\epsilon \le \frac{1}{2K}$,

$$\begin{aligned} d_{\mathbb{H}}(\phi_i^{\boldsymbol{z}}(x), \phi_i^{\boldsymbol{z}}(x')) &\ge d_{\mathbb{H}}(\bar{f}(\boldsymbol{z}_x^i), \bar{f}(\boldsymbol{z}_{x'}^i)) - (m-1)\frac{\epsilon d_{\mathbb{H}}(\boldsymbol{z}_x^i, \boldsymbol{z}_{x'}^i) + D_0}{m} \\ &\ge \frac{d_{\mathbb{H}}(\boldsymbol{z}_x^i, \boldsymbol{z}_{x'}^i)}{K} - C - \epsilon d_{\mathbb{H}}(\boldsymbol{z}_x^i, \boldsymbol{z}_{x'}^i) - D_0 \\ &\ge \frac{d_{\mathbb{H}}(x, x')}{2K} - (C + D_0). \end{aligned}$$

Hence, $\phi_i^{\boldsymbol{z}}$ is a $(2K, C+D_0)$–quasi-isometry. Equation (11) follows from applying the triangle inequality m times. ◁

Now, suppose the first conclusion is false. Then there exist $K, C, \epsilon > 0$, sequences $\rho_n \to 0$ and $D_n \to \infty$, and a sequence $\bar{f}_n$ of (K, C)–quasi-isometric embeddings defined on the balls $B_{R_n}(z_n)$, with $R_n \to \infty$, such that the restriction of $\bar{f}_n$ to $B_{\rho_n R_n}(z_n)$ does not factor as above. Then, by the above claim, there exist points x_n, y_n in $B_{\rho_n R_n}(z_n)$ which differ in one factor only, and are such that $\bar{f}_n(x_n)$ and $\bar{f}_n(y_n)$ differ in at least two factors by an amount that is at least $(\epsilon d_{\mathbb{H}}(y_n, x_n) + D_n)/m$ in each. We can assume $d_{\mathbb{H}}(y_n, x_n) \to \infty$, for otherwise $d_{\mathbb{H}}(\bar{f}(y_n), \bar{f}(x_n))$ is bounded.

Let $\lambda_n = 1/d_{\mathbb{H}}(y_n, x_n)$ and scale the metric on $B_{R_n}(z_n)$ with basepoint z_n by λ_n. Since $\rho_n \to 0$, we have $\lambda_n R_n \geq 1/2\rho_n \to \infty$. That is, the radius of $B_{R_n}(z_n)$ in the scaled metric still goes to ∞. However, in the scaled metric, the distance between x_n and y_n equals 1. Let $\prod_{i=1}^m \mathbb{H}_\omega$ be the asymptotic cone of $\prod_{i=1}^m \mathbb{H}$ with basepoint z_n and metric $d_n = \lambda_n d_{\mathbb{H}}$. For any $(u_1, u_2, \dots) \in \prod_{i=1}^m \mathbb{H}_\omega$ we define

$$\bar{f}_\omega \colon \prod_{i=1}^m \mathbb{H}_\omega \to \prod_{i=1}^m \mathbb{H}_\omega, \quad \bar{f}_\omega(u_1, u_2, \dots) = (\bar{f}_1(u_1), \bar{f}_2(u_2), \dots).$$

By definition $\lambda_n d_{\mathbb{H}}(u_n, z_n)$ is bounded, so for n large enough, $u_n \in B_{R_n}(z_n)$ and $f_n(u_n)$ is defined. It is clear that f_ω is bi-Lipschitz. We show $\bar{f}_\omega$ is onto.

By assumption $\bar{f}_n(B_{R_n}(z_n))$ C–coarsely contains $B_{\bar{\rho} R_n}(\bar{f}_n(z_n))$. Consider a point

$$(w_1, w_2, \dots) \in \prod_{i=1}^m \mathbb{H}_\omega.$$

Since $\lambda_n d_{\mathbb{H}}(\bar{f}_n(z_n), w_n)$ is bounded, for n large enough, $w_n \in B_{\bar{\rho} R_n}(\bar{f}_n(z_n))$. This means there is a sequence $u_n \in B_{R_n}(z_n)$ such that

$$d_{\mathbb{H}}(\bar{f}_n(u_n), w_n) \leq C.$$

But $\lambda_n \to \infty$. Thus,

$$(w_1, w_2, \dots) = \bar{f}_\omega(u_1, u_2, \dots),$$

and so $\bar{f}_\omega$ is onto and hence a homeomorphism.

By the argument in Step 3 of Section 9 in [16], the map $\bar{f}_\omega$ factors. The ω–limit points of x_n and y_n give a pair of points x_ω and y_ω in $\prod_{i=1}^m \mathbb{H}_\omega$ that have the same coordinate in every factor but one, and satisfy

$$d_{\mathbb{H}_\omega}(x_\omega, y_\omega) = 1.$$

But $\bar{f}_\omega(x_\omega)$ and $\bar{f}_\omega(y_\omega)$ differ in at least two coordinates by at least ϵ/m. This contradicts the fact that $\bar{f}_\omega$ factors. This contradiction proves the first conclusion.

We now use the first conclusion to prove the second conclusion. Let $r = \rho' R$. Consider a flat $F_{\boldsymbol{x}}$ through $\boldsymbol{x}$ and let $g_i = [a_i, b_i]$ be a geodesic in the i^{th} factor such that

$$F_{\boldsymbol{x}} \cap B_r(\boldsymbol{x}) \subset \prod_{i=1}^{m} [a_i, b_i].$$

Let g_i' be the geodesic joining $\phi_i^{\boldsymbol{z}}(a_i)$ to $\phi_i^{\boldsymbol{z}}(b_i)$. Since $\phi_i^{\boldsymbol{z}}$ is a quasi-isometric embedding,

$$d_{\mathbb{H}}(\phi_i^{\boldsymbol{z}}(g_i), g_i') = O(1),$$

where the bound depends on K, C, D_0. Let $F_{\boldsymbol{x}}'$ be the flat determined by the g_i'. For a point $\boldsymbol{p} \in F \cap B_r(\boldsymbol{x})$, the i^{th} coordinate of $\boldsymbol{p}$ lies on g_i. Therefore, the i^{th} coordinate of $\bar{f}(\boldsymbol{p})$ is distance at most $\epsilon d(\boldsymbol{x}, \boldsymbol{p})$ from a point whose i^{th} coordinate lies on $\phi_i^{\boldsymbol{z}}(g_i)$ and that in turn is distance $O(1)$ from g_i'. Since this is true for each i, the triangle inequality implies that

$$d_{\mathbb{H}}(\bar{f}(\boldsymbol{p}), F') \leq m\epsilon d_{\mathbb{H}}(\boldsymbol{x}, \boldsymbol{p}) + O(1). \qquad \square$$

Lemma 4.4 *Fix a constant $\rho'' < 1$. There exists a D'' such that for all sufficiently small ϵ and large R' the following holds. Suppose F, F' are flats, $\boldsymbol{p} \in F$ and*

$$B_{R'}(\boldsymbol{p}) \cap F \subset \mathcal{N}_{\epsilon R'}(F').$$

Then, for $r \leq \rho'' R'$,

$$B_r(\boldsymbol{p}) \cap F \subset \mathcal{N}_{D''}(F').$$

Proof We can assume $\epsilon R'$ is larger than the hyperbolicity constant for $\mathbb{H}$. Consider any geodesic $\gamma \subset F$ whose projection to each factor $\mathbb{H}$ intersects the disc of radius $r < \rho'' R'$ centered at the projection of $\boldsymbol{p}$ to that factor. Extend the geodesic so that its endpoints lie further than $\epsilon R'$ from the disc. Choose a pair of points in F' within $\epsilon R'$ of the endpoints of γ and let γ' be the geodesic in F' joining these points. Then γ' lies within Hausdorff distance $\epsilon R'$ of γ. Form the quadrilateral with two additional segments joining the endpoints of γ and γ'. Since the endpoints are within $\epsilon R'$ of each other, the segments joining the endpoints do not enter the disc of radius r. The quadrilateral is 2δ–thin, where δ is the hyperbolicity constant for $\mathbb{H}$. Therefore γ and γ' are within Hausdorff distance $O(2\delta)$ of each other on the disc of radius r. $\square$

Proof of Theorem 4.1 By the second conclusion of Proposition 4.3 there exist ρ' and D_0 such that for all small ϵ and large R, for any $\boldsymbol{x} \in B_{\rho' R}(\boldsymbol{z})$ and any flat F

through $\boldsymbol{x}$, there is a flat $F'_{\boldsymbol{x}}$ such that for $\boldsymbol{p} \in F \cap B_{\rho' R}(\boldsymbol{x})$,

$$d_{\mathbb{H}}(\bar{f}(\boldsymbol{p}), F'_{\boldsymbol{x}}) \leq m\epsilon d(\boldsymbol{x}, \boldsymbol{p}) + D_0.$$

Now clearly, $d(\bar{f}(\boldsymbol{x}), F'_{\boldsymbol{x}}) \leq D_0$ for all $\boldsymbol{x}$. By Lemma 4.4 there are $\bar{D} = \bar{D}(D_0, D'')$ and ρ'' such that for ϵ sufficiently small and R large, and any pair of points

$$\boldsymbol{x}, \boldsymbol{p} \in F \cap B_{\rho''\rho' R}(\boldsymbol{z}),$$

the flats $F'_{\boldsymbol{p}}$ and $F'_{\boldsymbol{x}}$, which correspond to $\boldsymbol{p}$ and $\boldsymbol{x}$ satisfying the above inequality, are within $\bar{D}$ of each other. That is, given F there is a single flat F' such that for all $\boldsymbol{x}$ in $F \cap B_{\rho''\rho' R}(\boldsymbol{z})$,

$$d_{\mathbb{H}}(\bar{f}(\boldsymbol{x}), F') \leq \bar{D}.$$

Consider any geodesic g_i in the i^{th} factor of $\prod_{i=1}^m \mathbb{H}$ that intersects $B_{\rho''\rho' R}(z_i)$ and fix the other coordinates so that we have a geodesic in $\prod_{i=1}^m \mathbb{H}$ that intersects $B_{\rho''\rho' R}(\boldsymbol{z})$. Again denote it by g_i. Choose a pair of flats F_1 and F_2 that intersect exactly along g_i. Then, as we have seen, there are flats F'_1 and F'_2 such that for $j = 1, 2$,

$$\bar{f}(F_j \cap B_{\rho''\rho' R}(\boldsymbol{z}_0)) \subset \mathcal{N}_{\bar{D}}(F'_j)$$

and thus, for both $j = 1, 2$,

$$\bar{f}(g_i \cap B_{\rho''\rho' R}(\boldsymbol{z}_0)) \subset \mathcal{N}_{\bar{D}}(F'_j).$$

Since $\bar{f}$ is a quasi-isometric embedding the pair of flats F'_1 and F'_2 must come $O(\bar{D})$–close along a single geodesic of length comparable to R in one factor in each. Thus we can assume that $\bar{f}$ factors along g_i and sends its intersection with $B_{\rho''\rho' R}(\boldsymbol{z}_0)$ to within $O(\bar{D})$ of a geodesic g'_j in a factor j.

Now let $\rho''' < 1$ and set $\rho = \rho'''\rho''\rho'$. Now consider *any* geodesic g in the i^{th} factor that intersects the smaller ball $B_{\rho R}(\boldsymbol{z}_0)$. We have that $\bar{f}$ factors in the bigger ball $B_{\rho''\rho' R}(\boldsymbol{z}_0)$ along $g \cap B_{\rho''\rho' R}(\boldsymbol{z}_0)$. We claim that it also sends it to the same j^{th} factor. Suppose not, so that it sends it to the $k \neq j$ factor. Choose a geodesic ℓ that comes close to both g_i and g, possibly in the bigger ball $B_{\rho''\rho' R}(\boldsymbol{z}_0)$. Its image must change from the j^{th} to the k^{th} factor, which is impossible since the image of $\ell \cap B_{\rho''\rho' R}(\boldsymbol{z}_0)$ lies in a single factor up to bounded error. Thus the map is factor-preserving. □

5 Local factors in Teichmüller space

In this section we apply Theorem 4.1 to balls in Teichmüller space.

For a given $R > 0$ and $x \in \mathcal{T}$, define the R–decomposition at x to be the decomposition $\mathcal{U}$ that contains a curve α if and only if $\tau_x(\alpha) \geq R$. That is, elements of $\mathcal{U}$ are either such curves or their complementary components. By convention, if $\mathcal{T} = \mathcal{T}(\Sigma, L)$ and Σ has a component U with $\xi(U) = 1$, then U (not any annulus in U) is always included in any R–decomposition of Σ. This is because $\mathcal{T}(U)$ is already a copy of $\mathbb{H}$. An R–decomposition is maximal if there are no complementary components W with $\xi(W) > 1$. We always assume $f\colon \mathcal{T} \to \mathcal{T}$ is a (K, C)–quasi-isometry but, unless specified, it is not always assumed that f is anchored.

Proposition 5.1 *For K and C as before, there are constants $0 < \rho_1 < 1$, d_1, C_1 and D_1 such that the following holds. For $x \in \mathcal{T}_{\mathrm{MR}}$ let R be such that*

$$d(x, \mathcal{T}_{\mathrm{LR}}) \geq R \geq \mathsf{d}_1.$$

Let $\mathcal{U}$ be the R–decomposition at x, let $\mathcal{V}$ be the $(R/2\mathsf{K})$–decomposition at $f(x)$ and let $r = \rho_1 R$. We have:

(1) *The decompositions $\mathcal{U}$ and $\mathcal{V}$ are maximal. For $x' \in B_r(x)$, we have $x' \in \mathcal{T}_{\mathcal{U}}$ and $f(x') \in \mathcal{T}_{\mathcal{V}}$.*

(2) *There is a bijection $f_x^\star\colon \mathcal{U} \to \mathcal{V}$ and, for every $U \in \mathcal{U}$ and $V = f_x^\star(U)$, there is a $(\mathsf{K}, \mathsf{C}_1)$–quasi-isometry*

$$\phi_x^U\colon B_r(x_U) \to \mathcal{T}(V)$$

such that for all $x' \in B_r(x)$,

$$d_{\mathcal{T}(V)}(\phi_x^U(x'_U), f(x')_V) \leq \mathsf{D}_1.$$

Remark 5.2 Property (2) above states that the map f restricted to $B_r(x)$ is close to the product map $\prod_{U \in \mathcal{U}} \phi_x^U$. Note also that $\mathcal{U}$ depends on R, so $f_x^\star$ too depends on R.

Proof Let $\boldsymbol{\alpha}$ be the set of curves α with $\tau_X(\alpha) \geq R$. Since $d_{\mathcal{T}}(x, \mathcal{T}_{\mathrm{LR}}) \geq R$ any complementary component U of $\boldsymbol{\alpha}$ must satisfy $\xi(U) \leq 1$. Otherwise there would be a pair of adjacent curves $\gamma_1, \gamma_2 \subset U$ with $\tau_x(\gamma_i) < R$, which means that $d_{\mathcal{T}}(x, \mathcal{T}_{\mathrm{LR}}) < R$. We conclude that $\mathcal{U}$ is a maximal decomposition.

Let $\rho_0 = 1/5\mathsf{K}^2$. For $x' \in B_{\rho_0 R}(x)$, we have

$$\tau_{x'}(\alpha) \geq \tau_x(\alpha) - \rho_0 R \geq (1 - \rho_0) R.$$

This, if d_1 is large enough, implies that the curves $\alpha \in \boldsymbol{\alpha}$ are ℓ_0–short in x' and so $x' \in \mathcal{T}_{\mathcal{U}}$.

Let $y \in \mathcal{T}_{\mathrm{LR}}$ be a point such that

$$d_{\mathcal{T}}(f(x), y) = d_{\mathcal{T}}(f(x), \mathcal{T}_{\mathrm{LR}}).$$

By Proposition 3.8,

$$d_{\mathcal{T}}(f^{-1}(y), \mathcal{T}_{\mathrm{LR}}) \le \mathsf{d}_0.$$

Using first the triangle inequality and then the fact that for some C', f^{-1} is a (coarse) $(\mathsf{K}, \mathsf{C}')$–quasi-isometry, we get

$$R = d_{\mathcal{T}}(x, \mathcal{T}_{\mathrm{LR}}) \le d_{\mathcal{T}}(x, f^{-1}(y)) + \mathsf{d}_0 \le \mathsf{K} d_{\mathcal{T}}(f(x), \mathcal{T}_{\mathrm{LR}}) + \mathsf{C}' + \mathsf{d}_0.$$

By picking d_1 large enough in terms of $\mathsf{d}_0, \mathsf{K}, \mathsf{C}'$, we have

$$d_{\mathcal{T}}(f(x), \mathcal{T}_{\mathrm{LR}}) \ge \frac{R}{2\mathsf{K}}. \tag{12}$$

Thus, as argued above, if $\boldsymbol{\beta}$ is the collection of curves β with $\tau_{f(x)}(\beta) \ge R/2\mathsf{K}$, then any complementary component V satisfies $\xi(V) \le 1$. Hence, $\mathcal{V}$ is a maximal decomposition.

Again, since f is a (K, C)–quasi-isometry, for all $\beta \in \boldsymbol{\beta}$ we have

$$\begin{aligned} |\tau_{f(x)}(\beta) - \tau_{f(x')}(\beta)| &\le d(f(x), f(x')) \le \mathsf{K} d(x, x') + \mathsf{C} \\ &\le \mathsf{K}(\rho_0 R) + \mathsf{C} \le \frac{R}{5\mathsf{K}} + \mathsf{C} \le \frac{R}{4\mathsf{K}}. \end{aligned}$$

The last inequality holds for d_1 large enough. Thus, for $x' \in B_{\rho_0 R}(x)$ and $\beta \in \boldsymbol{\beta}$,

$$\tau_{f(x')}(\beta) \ge \tau_{f(x)}(\beta) - \frac{R}{4\mathsf{K}} \ge \frac{R}{2\mathsf{K}} - \frac{R}{4\mathsf{K}}.$$

Again, for d_1 large enough, this means β is ℓ_0–short and thus $f(x') \in \mathcal{T}_{\mathcal{V}}$.

We have shown that $x' \in B_{\rho_0 R}(x)$ implies $x' \in \mathcal{T}_{\mathcal{U}}$ and $f(x') \in \mathcal{T}_{\mathcal{V}}$. By the Minsky product region theorem (Theorem 2.1) the maps $\psi_{\mathcal{U}}$ and $\psi_{\mathcal{V}}$ are distance D_{pr} from an isometry. Define

$$\bar{f}\colon \prod_{U \in \mathcal{U}} B_{\rho_0 R}(x_U) \to \prod_{V \in \mathcal{V}} \mathcal{T}(V), \quad \bar{f} = \psi_{\mathcal{V}} \circ f \circ \psi_{\mathcal{U}}^{-1}.$$

Then, if we set $\mathsf{C}_1 = 2\mathsf{D}_{\mathrm{pr}} + \mathsf{C}$, the map $\bar{f}$ is a $(\mathsf{K}, \mathsf{C}_1)$–quasi-isometry.

Because the map f has an inverse, $\bar{f}(B_{\rho_0 R}(\boldsymbol{x}))$ contains a ball of comparable radius about $\bar{f}(\boldsymbol{x})$. Now Theorem 4.1 applied to $K = \mathsf{K}$ and $C = \mathsf{C}_1$ says that there are constants R_0, ρ, D and a bijection $f_x^*\colon \mathcal{U} \to \mathcal{V}$ such that for $r = \rho\rho_0 R$, the following holds. Assume $\mathsf{d}_1 \ge R_0$. Then, for each $U \in \mathcal{U}$ and $V = f_x^*(U)$, there is a $(\mathsf{K}, \mathsf{C}_1)$–quasi-isometry

$$\phi_x^U\colon B_r(x_U) \to \mathcal{T}(V)$$

such that for $x' \in B_r(x)$,

$$d_{\mathcal{T}_\mathcal{V}}\Big(\bar{f}(\psi_\mathcal{U}(x')), \prod_{U \in \mathcal{U}} \phi_x^U(x'_U)\Big) \le D.$$

The distance in $\mathcal{T}_\mathcal{V}$ is the sup metric, each factor of which is either a copy of $\mathbb{H}$ or a horosphere $H_\beta \subset \mathbb{H}$. Hence, the inequality holds for every factor:

$$d_{\mathcal{T}(V)}\big(\bar{f}(\psi_\mathcal{U}(x'))_V, \phi_x^U(x'_U)\big) \le D.$$

Therefore, for $\rho_1 = \rho\rho_0$, $\mathsf{D}_1 = D$ and d_1 large enough, the proposition holds. □

Proposition 5.3 *Choose R so that*

$$r = \rho_1 R \ge \max(\rho_1 \mathsf{d}_1, 4\mathsf{K}(4\mathsf{D}_1 + \mathsf{C}_1)).$$

Let $x^1, x^2 \in \mathcal{T}$ be points such that

$$\mathrm{d}_\mathcal{T}(x^1, \mathcal{T}_{\mathrm{LR}}) \ge R, \quad \mathrm{d}_\mathcal{T}(x^2, \mathcal{T}_{\mathrm{LR}}) \ge R, \quad d_\mathcal{T}(x^1, x^2) \le r.$$

Assume every point $x \in B_r(x^1) \cup B_r(x^2)$ has the same R–decomposition $\mathcal{U}$ and the $(R/2\mathsf{K})$–decomposition at $f(x)$ always contains some subsurface V. Then, there is a $U \in \mathcal{U}$ such that

$$f^\star_{x^1}(U) = f^\star_{x^2}(U) = V,$$

and, for every $u \in B_r(x^1_U) \cap B_r(x^2_U)$,

$$d_{\mathcal{T}(V)}(\phi^U_{x^1}(u), \phi^U_{x^2}(u)) \le 2\mathsf{D}_1.$$

Remark 5.4 Since $f^*_{x^i}$ is a bijection there must be some U that is mapped to V. The content of the first conclusion is that the same U works at both points.

Proof By assumption, for $x \in B_r(x^1) \cap B_r(x^2)$, the domain of $f^\star_x$ is $\mathcal{U}$. We start by proving the following claim.

Claim For $z^1, z^2 \in B_r(x^1) \cap B_r(x^2)$ and $U \in \mathcal{U}$, suppose $f^\star_{z^1}(U) = V$. Also, assume either

$$d_{\mathcal{T}(U)}(z^1_U, z^2_U) \ge \frac{r}{4} \quad \text{and} \quad z^1_W = z^2_W \text{ for all } W \in \mathcal{U} - \{U\}$$

or

$$d_{\mathcal{T}(W)}(z_W^1, z_W^2) \geq \frac{r}{4} \text{ for all } W \in \mathcal{U} - \{U\} \quad \text{and} \quad z_U^1 = z_U^2.$$

That is, z^1 and z^2 either differ by $\frac{r}{4}$ in only one factor, or all but one factor. Then $f_{z^2}^{\star}(U) = V$.

Proof of claim Assume the first case holds. Since $\phi_{z^1}^U$ is a $(\mathsf{K}, \mathsf{C}_1)$–quasi-isometry (Proposition 5.1), we have

$$d_{\mathcal{T}(V)}(\phi_{z^1}^U(z^1), \phi_{z^1}^U(z^2)) \geq \frac{r}{4\mathsf{K}} - \mathsf{C}_1.$$

By Proposition 5.1 and the triangle inequality applied twice, we get

$$d_{\mathcal{T}(V)}(f(z^1)_V, f(z^2)_V) \geq \frac{r}{4\mathsf{K}} - \mathsf{C}_1 - 2\mathsf{D}_1.$$

Now suppose $f_{z^2}^{\star}(W) = V$ for $W \neq U$. Since $z_W^1 = z_W^2$, Proposition 5.1 and the triangle inequality then imply

$$d_{\mathcal{T}(V)}(f(z^1)_V, f(z^2)_V) \leq 2\mathsf{D}_1.$$

These two inequalities contradict the choice of R in the statement of the proposition, proving the claim. The proof of the claim under the second assumption is similar. $\triangleleft$

We now prove the proposition. In each $W \in \mathcal{U}$ choose z_W so that

$$d_{\mathcal{T}(W)}(z_W, x_W^1) = d_{\mathcal{T}(W)}(z_W, x_W^2) = \frac{r}{4}.$$

Let z^1, $z^{1,2}$ and z^2 be points in $B_r(x^1) \cap B_r(x^2)$ such that for $i = 1, 2$,

$$z_U^i = z_U \qquad \text{and} \qquad z_W^i = x_W^i \quad \text{for all } W \in \mathcal{U} - \{U\},$$

and such that

$$z_W^{1,2} = z_W \quad \text{for all } W \in \mathcal{U}.$$

Note that the claim can be applied to pairs (x^1, z^1), $(z^1, z^{1,2})$, $(z^{1,2}, z^2)$ and (z^2, x^2), concluding that

$$V = f_{x^1}^{\star}(U) = f_{z^1}^{\star}(U) = f_{z^{1,2}}^{\star}(U) = f_{z^2}^{\star}(U) = f_{x^2}^{\star}(U).$$

Now, consider $u \in B_r(x_U^1) \cap B_r(x_U^2)$ and let $z \in B_r(x^1) \cap B_r(x^2)$ be such that $z_U = u$. We know

$$d_{\mathcal{T}(V)}(f(z)_V, \phi_{x^1}^U(u)) \leq \mathsf{D}_1 \quad \text{and} \quad d_{\mathcal{T}(V)}(f(z)_V, \phi_{x^2}^U(u)) \leq \mathsf{D}_1.$$

Therefore,

$$d_{\mathcal{T}(V)}(\phi^U_{x^1}(u), \phi^U_{x^2}(u)) \le 2\mathsf{D}_1.$$

This finishes the proof of the proposition. □

Corollary 5.5 *If $\mathcal{T} = \mathcal{T}(\Sigma, L)$ and U is a component of Σ with $\xi(U) = 1$, then for every $x \in \mathcal{T}(\Sigma, L)$ where $f_x^\star$ is defined, $f_x^\star(U) = U$.*

Proof Note that, by definition, the subsurface U is always included in any R–decomposition $\mathcal{U}$ of Σ. Hence, Proposition 5.3 applies. That is, for a large R as in Proposition 5.3, if $f_x^\star(U) = U$ then the same holds for points in an r–neighborhood of x. But the set of points where $f_x^\star$ is defined is connected. Hence, it is enough to show $f_x^\star(U) = U$ for just one point x.

Let $\partial_L(U)$ be the projection of the boundary of $\partial_L(\Sigma)$ to $\mathcal{T}(U)$. Consider a geodesic g_U in $\mathcal{T}(U)$ connecting a pair of points a and a' where

$$d_{\mathcal{T}(U)}(a, \partial_L(U)) = d_{\mathcal{T}(U)}(a', \partial_L(U)) = R,$$
$$d_{\mathcal{T}(U)}(g, \partial_L(U)) \ge R, \quad d_{\mathcal{T}(U)}(a, a') \ge 4\mathsf{K}R.$$

For example, we can choose a geodesic connecting two L–horoballs in $\mathcal{T}(U)$ that otherwise stays in the thick part of $\mathcal{T}(U)$ and then we can cut off a subsegment of length R from each end.

Let $W = \Sigma - U$. Choose $b \in \mathcal{T}(W)$ to have distance R from $\partial_L(W)$ and let g be a path in $\mathcal{T}(\Sigma, L)$ that has constant projection to W and projects to g_U in U. That is,

$$g(t)_W = b \quad \text{and} \quad g(t)_U = g_U(t).$$

Then g connects a point $x \in \mathcal{T}(\Sigma, L)$ to a point $x' \in \mathcal{T}(\Sigma, L)$ where $x_U = a$, $x'_U = a'$ and $x_W = x'_W = b$.

Let $y = f(x)$, $y' = f(x')$ and let z and z' be points on $\partial_L(\Sigma)$ that are distance R to x and x' respectively. Since f is anchored, we have

$$d_{\mathcal{T}}(f(z), z) \le \mathsf{C} \quad \text{and} \quad d_{\mathcal{T}}(f(z'), z') \le \mathsf{C},$$

and hence

$$d_{\mathcal{T}}(y, z) \le d_{\mathcal{T}}(f(x), f(z)) + d_{\mathcal{T}}(f(z), z) \le (\mathsf{K}R + \mathsf{C}) + \mathsf{C}.$$

The same holds for $d_{\mathcal{T}}(y', z')$. Also,

$$d_{\mathcal{T}(U)}(z_U, z'_U) \geq d_{\mathcal{T}(U)}(x_U, x'_U) - d_{\mathcal{T}}(z, x) - d_{\mathcal{T}}(z', x') \geq d_{\mathcal{T}(U)}(x_U, x'_U) - 2R.$$

Therefore,

$$\begin{aligned} d_{\mathcal{T}(U)}(y_U, y'_U) &\geq d_{\mathcal{T}(U)}(z_U, z'_U) - d_{\mathcal{T}(U)}(y_U, z_U) - d_{\mathcal{T}(U)}(y'_U, z'_U) \\ &\geq d_{\mathcal{T}(U)}(x_U, x'_U) - 2R - 2(\mathsf{K}R + 2\mathsf{C}) \\ &\geq \tfrac{1}{4} d_{\mathcal{T}(U)}(x_U, x'_U). \end{aligned}$$

The last inequality holds since $d_{\mathcal{T}(U)}(x_U, x'_U) \geq 4\mathsf{K}R$ and $2R + 2(\mathsf{K}R + 2\mathsf{C}) \leq 3\mathsf{K}R$. Now choose points

$$x = x_0, \dots, x_N = x'$$

along g so that $d_{\mathcal{T}}(x_i, x_{i+1}) \leq r$ and

$$N = \frac{d_{\mathcal{T}(U)}(x_U, x'_U)}{r} = \frac{d_{\mathcal{T}(U)}(x_U, x'_U)}{\rho_1 R}.$$

As mentioned before, we already know $f^\star_x(U) = f^\star_{x_i}(U)$. If $f^\star_x(U) \neq U$, then $f^\star_{x_i}(U) \neq U$ and hence, by Proposition 5.1,

$$d_{\mathcal{T}(U)}(f(x_i)_U, f(x_{i+1})_U) \leq 2\mathsf{D}_1.$$

Therefore,

$$d_{\mathcal{T}(U)}(y_U, y'_U) \leq 2N\mathsf{D}_1 \leq 2\mathsf{D}_1 \frac{d_{\mathcal{T}(U)}(x_U, x'_U)}{\rho_1 R}.$$

For R large enough, this contradicts equation (13). Hence, $f^\star_x(U) = U$. □

6 Nearly shortest curves

The goal of this section is to prove Proposition 6.9. Essentially, this states that if α is one of the shortest curves in x, then α is also short in $f(x)$ and, furthermore, $f^\star_x(\alpha) = \alpha$.

Throughout this section we always assume, if $\mathcal{T} = \mathcal{T}(\Sigma, L)$, that Σ does not have any components U with $\chi(U) = 1$. Hence, Proposition 6.9 is the complementary statement to Corollary 5.5. The effect of this assumption is that every curve has an adjacent curve. The arguments are conceptually very elementary. However, we need to keep careful track of constants.

Cone is preserved

Recall that $\mathcal{S}_x$ is the set of curves whose extremal length on x is smaller than some fixed ℓ_0. For $\eta > 1$, define the η–maximal cone $\mathcal{T}_{\mathcal{MC}}(\eta)$ to be the set of points $x \in \mathcal{T}$ such that:

- $\mathcal{S}_x$ is a pants decomposition, that is, $\mathcal{S}_x = P_x$.
- For $\alpha, \beta \in P_x$,
$$\frac{\tau_x(\alpha)}{\tau_x(\beta)} \le \eta.$$

Let $\tau_x = \max_{\gamma \in P_x} \tau_x(\gamma)$. The value τ_x gives the maximum distance to travel to make the shortest curve have length ℓ_0. To say x lies in the cone says that the distance traveled to make any short curve have length ℓ_0 are comparable up to factor η. For a point z to be in $\mathcal{T}_{\mathrm{LR}}$, it has to contain two adjacent curves that have lengths larger than ℓ_0. Hence, to reach lower rank costs at most τ_x and at least τ_x/η. That is, for $x \in \mathcal{T}_{\mathcal{MC}}(\eta)$, we have

$$\frac{\tau_x}{\eta} \le d_{\mathcal{T}}(x, \mathcal{T}_{\mathrm{LR}}) \le \tau_x. \tag{13}$$

Proposition 6.1 *For every η_0, there is a τ_0 such that if $x \in \mathcal{T}_{\mathcal{MC}}(\eta_0)$ and $\tau_x \ge \tau_0$, then*
$$f(x) \in \mathcal{T}_{\mathcal{MC}}(16\mathsf{K}^2\eta_0).$$
In fact, for every $\gamma \in P_{f(x)}$, we have
$$\frac{\tau_x}{4\mathsf{K}\eta_0} \le \tau_{f(x)}(\gamma) \le 4\mathsf{K}\tau_x.$$

Proof Choose $R \ge \mathsf{d}_1$ large enough so that if $d(x, \mathcal{T}_{\mathrm{LR}}) \ge R$ then, similarly to equation (12), we have

$$\frac{d(x, \mathcal{T}_{\mathrm{LR}})}{2\mathsf{K}} \le d(f(x), \mathcal{T}_{\mathrm{LR}}) \le 2\mathsf{K} d(x, \mathcal{T}_{\mathrm{LR}}), \tag{14}$$

and so that

$$r = \rho_1 R \ge 32\mathsf{K}\,\eta_0\,\mathsf{D}_1. \tag{15}$$

Let $\tau_0 = 16\mathsf{K}^2\eta_0 R$.

Claim For adjacent curves $\beta, \beta' \in P_{f(X)}$ we have
$$\tau_{f(x)}(\beta) \le \frac{\tau_x}{4\mathsf{K}\eta_0} \implies \tau_{f(x)}(\beta') \ge 4\mathsf{K}\tau_x.$$

We will later prove that in fact $\tau_{f(x)}(\beta') < 4\mathsf{K}\tau_x$. Assuming this inequality and the claim we can finish the proof of the proposition. Indeed the claim says that for β adjacent to β', one has $\tau_x/(4\mathsf{K}\eta_0) \le \tau_{f(x)}(\beta)$.

Proof of claim We prove the claim by contradiction. Assume that the first inequality in the claim holds and the second is false. Note that we still have, by equations (13) and (14), that

$$\max(\tau_{f(x)}(\beta'), \tau_{f(x)}(\beta)) \ge d_{\mathcal{T}}(f(x), \mathcal{T}_{\mathrm{LR}}) \ge \frac{d_{\mathcal{T}}(x, \mathcal{T}_{\mathrm{LR}})}{2\mathsf{K}} \ge \frac{\tau_x}{2\eta_0\mathsf{K}}.$$

To summarize, we have two adjacent curves β and β' with

$$\tau_{f(x)}(\beta) \le \frac{\tau_x}{4\mathsf{K}\eta_0} \quad \text{and} \quad \frac{\tau_x}{2\mathsf{K}\eta_0} \le \tau_{f(x)}(\beta') \le 4\mathsf{K}\tau_x.$$

Consider a geodesic g moving only in the β' factor that increases the length of β' and connects $f(x)$ to a point $f(x')$ with

$$\tau_{f(x')}(\beta') = \frac{\tau_x}{4\mathsf{K}\eta_0} \ge \frac{\tau_0}{4\mathsf{K}\eta_0} \ge R.$$

We have

$$d(f(x'), \mathcal{T}_{\mathrm{LR}}) \le \frac{\tau_x}{4\mathsf{K}\eta_0}. \tag{16}$$

Take a sequence of points

$$f(x) = y_1, y_2, \dots, y_N = f(x')$$

along g with $d_{\mathcal{T}}(y_i, y_{i+1}) \le r$ and

$$N = \frac{4\mathsf{K}\tau_x}{r}.$$

Let $h = f^{-1}$, let $x^i = h(y_i)$, let $\mathcal{V}_i$ be the R–decomposition at y_i and $\mathcal{U}_i$ be the $(R/2\mathsf{K})$–decomposition at x_i. By assumption, $\tau_{y_i}(\beta') \ge R$ and hence $\beta' \in \mathcal{V}_i$ for every i and, since the length of no other curve is changing along g, all $\mathcal{V}_i$ are in fact the same decomposition (which we denote by $\mathcal{V}$). Also $\mathcal{U}_1 = P_x$. Let $\alpha' = h^\star_{y_1}(\beta')$.

Let $\alpha \in P_x$ with $\alpha \ne \alpha'$. Let $V \in \mathcal{V}$ be the component such that $h^\star_{y_1}(V) = \alpha$. We show by induction on i that $\alpha \in \mathcal{U}_i$ and $h^\star_{y_i}(V) = \alpha$. Assume this for $1 \le i < j$. Since $y_{i+1} \in B_r(y_i)$, and y_i and y_{i+1} have the same projection to $\mathcal{T}(V)$, Proposition 5.1 implies

$$d_{\mathcal{T}(\alpha)}(x^i_\alpha, x^{i+1}_\alpha) \le 2\mathsf{D}_1.$$

Therefore,

$$d_{\mathcal{T}(\alpha)}(x_\alpha^1, x_\alpha^j) \le 2j\,\mathsf{D}_1 \le 2N\,\mathsf{D}_1 \le \frac{8\mathsf{K}\tau_x}{r}\mathsf{D}_1 \le \frac{\tau_x}{4\eta_0}.$$

The last inequality is from the assumption on r. Hence,

$$\tau_{x^j}(\alpha) \ge \frac{\tau_x}{\eta_0} - \frac{\tau_x}{4\eta_0} \ge \frac{3\tau_x}{4\eta_0}.$$

This means in particular that α is short enough ($\tau_{x^j}(\alpha) \ge R/2\mathsf{K}$) and is included in $\mathcal{U}_j$. Moreover, by Proposition 5.3, $h^\star_{y_j}(V) = \alpha$, completing the induction step. Continuing this way, we conclude that, for every $\alpha \in P_x$, $\alpha \ne \alpha'$, we have

$$\tau_{x'}(\alpha) \ge \frac{3\tau_x}{4\eta_0}.$$

In particular, the above inequality holds for every pair of adjacent curves in x'. But

$$d_{\mathcal{T}}(x', \mathcal{T}_{\mathrm{LR}}) \le 2\mathsf{K} d_{\mathcal{T}}(f(x'), \mathcal{T}_{\mathrm{LR}}) \le \frac{\tau_x}{2\eta_0}.$$

This is a contradiction, which proves the claim. $\lhd$

We now show that all curves $\beta \in P_{f(x)}$ must in fact satisfy

$$\tau_{f(x)}(\beta) < 4\mathsf{K}\tau_x.$$

The argument is similar to the one above, so we skip some of the details. Suppose this is false for some curve β. Let α be a curve such that $f_x^*(\alpha) = \beta$, and α' be a curve adjacent to α. Let g be a geodesic that moves only in the α and α' factors, increasing their lengths, and that connects x to a point x' where

$$\tau_x(\alpha) = \tau_x(\alpha') = R.$$

We find a collection of points

$$x = x_1, \dots, x_N = x'$$

on g such that $d(x, x') \le r$ and $N = \tau_x/r$. Let $\mathcal{U}_i$ be the R–decomposition at x_i and $\mathcal{V}_i$ be the $(R/2\mathsf{K})$–decomposition at $y_i = f(x_i)$. Then, as before, α and α' are in every $\mathcal{U}_i$ (in fact, all $\mathcal{U}_i$ are the same decompositions which we denote by $\mathcal{U}$). Let $f_x^\star(\alpha') = V'$. By an argument as above, the total movement in any other factor (besides β and V') is at most $2N\,\mathsf{D}_1$.

Since $d_{\mathcal{T}}(x_N, \mathcal{T}_{\mathrm{LR}}) = R$, we have $d_{\mathcal{T}}(y_N, \mathcal{T}_{\mathrm{LR}}) \le 2\mathsf{K}R$. That is, there are two adjacent curves γ and γ' with

$$\tau_{y_N}(\gamma), \tau_{y_N}(\gamma') \le 2\mathsf{K}R. \tag{17}$$

But, for R large enough, we have

$$\tau_{y_N}(\beta) \ge \tau_{y_1}(\beta) - N(\mathsf{K}r + \mathsf{C}) \ge 4\mathsf{K}\tau_x - \frac{\tau_x}{r}(\mathsf{K}r - \mathsf{C}) \ge \tfrac{5}{2}\mathsf{K}\tau_x.$$

Hence, neither γ nor γ' equals β. And one, say γ', is not contained in V'. Therefore,

$$\begin{aligned}\tau_{y_1}(\gamma') &\le \tau_{y_N}(\gamma') + 2N\mathsf{D}_1 \le 2\mathsf{K}R + 2\frac{\tau_x}{r}\mathsf{D}_1 \\ &\le \frac{2\mathsf{K}R}{\tau_0}\tau_x + \frac{2\mathsf{D}_1}{r}\tau_x \le \left(\frac{1}{8\mathsf{K}\eta_0} + \frac{1}{16\mathsf{K}\eta_0}\right)\tau_x < \frac{\tau_x}{4\mathsf{K}\eta_0}.\end{aligned}$$

Now, γ is adjacent to γ', so $\tau_{y_1}(\gamma) \ge 4\mathsf{K}\tau_x$ by the claim. However, again as above,

$$\tau_{y_N}(\gamma') \ge \tau_{y_1}(\gamma') - N(\mathsf{K}r + \mathsf{C}) \ge 4\mathsf{K}\tau_x - \frac{\tau_x}{r}(\mathsf{K}r + \mathsf{C}) \ge \tfrac{5}{2}\mathsf{K}\tau_x.$$

This contradicts equation (17). We are done. $\square$

Induced map on the curve complex is cellular

Let $\mathcal{P}(S)$ be the complex of pants decompositions of S. For a pants decomposition $P \in \mathcal{P}(S)$, define $\mathcal{T}_{\mathcal{MC}}(P, \eta_0)$ to be the set of points $x \in \mathcal{T}_{\mathcal{MC}}(\eta_0)$ where $P_x = P$. Note that $\mathcal{T}_{\mathcal{MC}}(\eta_0)$ is not connected and its connected components are parametrized by $\mathcal{P}(S)$:

$$\mathcal{T}_{\mathcal{MC}}(\eta_0) = \bigsqcup_{P \in \mathcal{P}(S)} \mathcal{T}_{\mathcal{MC}}(P, \eta_0).$$

The same is true for $\mathcal{T}_{\mathcal{MC}}(16\mathsf{K}^2\eta_0)$ and, by Proposition 6.1,

$$f(\mathcal{T}_{\mathcal{MC}}(\eta_0)) \subset \mathcal{T}_{\mathcal{MC}}(16\mathsf{K}^2\eta_0).$$

Hence, we can define a bijection $f_{\mathcal{P}}^\star \colon \mathcal{P}(S) \to \mathcal{P}(S)$ so that

$$f(\mathcal{T}_{\mathcal{MC}}(P, \eta_0)) \subset \mathcal{T}_{\mathcal{MC}}(f_{\mathcal{P}}^\star(P), 16\mathsf{K}^2\eta_0).$$

We will show that $f_{\mathcal{P}}^\star$ is induced by a simplicial automorphism $f_{\mathcal{C}}^\star$ of the curve complex. Note that the definition of $f_{\mathcal{P}}^\star$ depends on the choice of η_0.

Proposition 6.2 *Assuming η_0 is large enough, there exists a simplicial automorphism $f_{\mathcal{C}}^\star\colon \mathcal{C}(S) \to \mathcal{C}(S)$ such that, for a pants decomposition $P = \{\alpha_1, \dots, \alpha_n\}$, we have*

$$f_{\mathcal{P}}^\star(P) = \{f_{\mathcal{C}}^\star(\alpha_1), \dots, f_{\mathcal{C}}^\star(\alpha_n)\}.$$

Proof Recall we have constants d_1 and ρ_1 from Proposition 5.1. Choose R, η_0 and d so that

$$\eta_0 d \geq R \geq \mathsf{d}_1 \quad \text{and} \quad r = \rho_1 R \geq 2d.$$

Let $\boldsymbol{\alpha}$ be a multicurve containing $\xi(S) - 1$ curves and let

$$P = \boldsymbol{\alpha} \cup \{\beta\} \quad \text{and} \quad P' = \boldsymbol{\alpha} \cup \{\beta'\}$$

be two extensions of $\boldsymbol{\alpha}$ to pants decompositions with $\mathrm{i}(\beta, \beta') \leq 2$. Let x be a point in $\mathcal{T}_{\mathcal{MC}}(P, \eta_0)$ such that $\tau_x(\beta) = d$ and $\tau_x(\alpha) = \eta_0 d$ for $\alpha \in \boldsymbol{\alpha}$. Similarly, define $x' \in \mathcal{T}_{\mathcal{MC}}(P', \eta_0)$ so that

$$x'_\alpha = x_\alpha \quad \text{for all } \alpha \in \boldsymbol{\alpha}, \quad \tau_{x'}(\beta') = d \quad \text{and} \quad d_{\mathcal{T}}(x, x') \leq 2d.$$

Let $\mathcal{U}$ be the decomposition consisting of $\boldsymbol{\alpha}$ and the complementary subsurface U. Then $x, x' \in T_{\mathcal{U}}$. Let g be a path connecting x to x' whose projection is constant in every $\mathcal{T}(\alpha)$ for $\alpha \in \boldsymbol{\alpha}$, is a geodesic connecting x_U to x'_U in $\mathcal{T}(U)$, and has length $2d$. We have

$$d_{\mathcal{T}}(g, \mathcal{T}_{\mathrm{LR}}) \geq \eta_0 d \geq R.$$

Hence, Propositions 5.1 and 5.3 apply to x and x'. This means $f(x)$ and $f(x')$ are contained in the same product region, say $T_{\mathcal{V}}$. Let $V = f_x^\star(U)$. Then, for $W \in \mathcal{V}$ with $W \neq V$,

$$d_{\mathcal{T}(W)}(f(x)_W, f(x')_W) \leq 2\mathsf{D}_1.$$

But

$$f(x) \in \mathcal{T}_{\mathcal{MC}}(f_{\mathcal{P}}^\star(P), 16\mathsf{K}^2\eta_0) \quad \text{and} \quad f(x') \in \mathcal{T}_{\mathcal{MC}}(f_{\mathcal{P}}^\star(P'), 16\mathsf{K}^2\eta_0).$$

Hence, $\mathcal{V}$ contains a multicurve $\boldsymbol{\gamma}$ with $\xi(S) - 1$ curves, and $f_{\mathcal{P}}^\star(P)$ and $f_{\mathcal{P}}^\star(P')$ share all but one curve. That is, there are curves δ and δ' such that

$$f_{\mathcal{P}}^\star(P) = \boldsymbol{\gamma} \cup \{\delta\} \quad \text{and} \quad f_{\mathcal{P}}^\star(P') = \boldsymbol{\gamma} \cup \{\delta'\}.$$

However, for the moment, we do not have a good bound on the intersection number $\mathrm{i}(\delta, \delta')$.

We first show that the multicurve $\boldsymbol{\gamma}$ does not depend on the choice of β'. Assume β'' is another curve with $i(\beta, \beta'') \leq 2$, and let $P'' = \boldsymbol{\alpha} \cup \{\beta''\}$. If $\mathrm{i}(\beta,' \beta'') \leq 2$ then

$f_{\mathcal{P}}^{\star}(P)$, $f_{\mathcal{P}}^{\star}(P')$ and $f_{\mathcal{P}}^{\star}(P'')$ share $\xi(S)-1$ curves. Hence, these have to be the same multicurve $\boldsymbol{\gamma}$. If $\mathrm{i}(\beta,' \beta'')$ is large, then we can find a sequence

$$\beta' = \beta_1, \dots, \beta_m = \beta''$$

of curves that are disjoint from $\boldsymbol{\alpha}$, with $\mathrm{i}(\beta, \beta_i) \leq 2$ and $\mathrm{i}(\beta_i, \beta_{i+1}) \leq 2$. Define $P_i = \boldsymbol{\alpha} \cup \{\beta_i\}$. Then arguing as above shows that all $f_{\mathcal{P}}^{\star}(P_i)$ share the same $\xi(S)-1$ curves. That is, the same curve changes from $f_{\mathcal{P}}^{\star}(P)$ to $f_{\mathcal{P}}^{\star}(P'')$ as it did from $f_{\mathcal{P}}^{\star}(P)$ to $f_{\mathcal{P}}^{\star}(P')$ and $\boldsymbol{\gamma} \subset f_{\mathcal{P}}^{\star}(P'')$.

Note that we have shown that there is an association between curves in P and $f_{\mathcal{P}}^{\star}(P)$. Namely let $\beta \in P$ and $\boldsymbol{\alpha}$ be the set of complementary $\xi-1$ curves in P. Let P' any pants decomposition that contains $\boldsymbol{\alpha}$ but not β. Then $\beta \in P$ is associated to the curve $P(\beta) \in f_{\mathcal{P}}^{\star}(P)$ that is not contained in $f_{\mathcal{P}}^{\star}(P')$. We will show $P(\beta)$ is the same for every P; that is, it does not depend on $\boldsymbol{\alpha}$.

Let $\boldsymbol{\alpha}$ be a multicurve with $\xi(S)-2$ curves and let $P = \boldsymbol{\alpha} \cup \{\beta_1, \beta_2\}$ be a pants decomposition. Let

$$P_1 = \boldsymbol{\alpha} \cup \{\beta_1', \beta_2\}, \quad P_2 = \boldsymbol{\alpha} \cup \{\beta_1, \beta_2'\}, \quad P_{12} = \boldsymbol{\alpha} \cup \{\beta_1', \beta_2'\},$$

with $\mathrm{i}(\beta_i, \beta_i') \leq 2$ for $i = 1, 2$. Denote

$$Q = f_{\mathcal{P}}^{\star}(P), \quad Q_1 = f_{\mathcal{P}}^{\star}(P_1), \quad Q_2 = f_{\mathcal{P}}^{\star}(P_2), \quad Q_{12} = f_{\mathcal{P}}^{\star}(P_{12}).$$

Let $\boldsymbol{\gamma}$ be a multicurve with $\xi(S)-2$ curves such that $Q = \boldsymbol{\gamma} \cup \{P(\beta_1), P(\beta_2)\}$. From the discussion above, we know that there are curves δ_1 and δ_2 such that

$$Q_1 = \boldsymbol{\gamma} \cup \{\delta_1, P(\beta_2)\} \quad \text{and} \quad Q_2 = \boldsymbol{\gamma} \cup \{P(\beta_1), \delta_2\}.$$

Since $P(\beta_1)$ and $P(\beta_2)$ are disjoint, the curves δ_1 and δ_2 must be different. Also Q_{12} shares $\xi(S)-1$ curves with both Q_1 and Q_2. But the map $f_{\mathcal{P}}^{\star}$ is a bijection and $Q_{12} \neq Q$. Moreover, since $\delta_1 \neq \delta_2$ and $P(\beta_1)$ and $P(\beta_2)$ are disjoint, Q_{12} cannot contain both $P(\beta_1)$ and δ_1 and similarly it cannot contain both $P(\beta_2)$ and δ_2. Therefore,

$$Q_{12} = \boldsymbol{\gamma} \cup \{\delta_1, \delta_2\}.$$

That means $P_2(\beta_1) = P(\beta_1)$ because it is the curve that changes from Q_2 to Q_{12}. Similarly, $P_1(\beta_2) = P(\beta_2)$.

W have shown the association $\beta \to P(\beta)$ is the same for adjacent pants decompositions in $\mathcal{P}(S)$. Hence, it does not depend on P and we can define $f_{\mathcal{C}}^{\star}(\beta) = P(\beta)$ for any pants decomposition P containing β.

Every multicurve is contained in a pants decomposition. Hence $f_{\mathcal{C}}^{\star}$ sends disjoint curves to disjoint curves and in fact sends simplices in $\mathcal{C}(S)$ to simplices. Since $f_{\mathcal{P}}^{\star}$ is onto, $f_{\mathcal{C}}^{\star}$ is also onto. We show $f_{\mathcal{C}}^{\star}$ is one-to-one. Assume for the sake of contradiction that $f_{\mathcal{C}}^{\star}(\alpha) = f_{\mathcal{C}}^{\star}(\beta)$, and let P_α and P_β be pants decompositions that contain α and β, respectively, but have no curves in common. Then $f_{\mathcal{P}}^{\star}(P_\alpha)$ and $f_{\mathcal{P}}^{\star}(P_\beta)$ share a curve δ. Let

$$f_{\mathcal{P}}^{\star}(P_\alpha) = Q_1, \dots, Q_m = f_{\mathcal{P}}^{\star}(P_\beta)$$

be a sequence of adjacent pants decompositions all containing δ and let $P_i = {f_{\mathcal{P}}^{\star}}^{-1}(Q_i)$. Then ${f_{\mathcal{C}}^{\star}}^{-1}(\delta)$ is contained in every P_i. This contradicts the assumption that P_α and P_β have no common curves.

We have shown $f_{\mathcal{C}}^{\star}$ is simplicial and it is a bijection; that is, it is a simplicial automorphism of $\mathcal{C}(S)$. □

Remark 6.3 In the case $\mathcal{T} = \mathcal{T}(S)$, by a theorem of Ivanov [15], $f_{\mathcal{C}}^{\star}$ is induced by an isometry $f^{\star}$ of Teichmüller space $\mathcal{T}(S)$. Hence, after applying the inverse of this isometry to f, we can assume that $f_{\mathcal{C}}^{\star}$ is the identity map. In the case that $\mathcal{T} = \mathcal{T}(\Sigma, L)$ and f is anchored, we know that $f_{\mathcal{P}}^{\star}$ is the identity map. Thus, so is $f_{\mathcal{C}}^{\star}$. Indeed, for the remainder of the paper, we assume that in these cases $f_{\mathcal{C}}^{\star}$ is the identity.

Remark 6.4 We point out that $f_{\mathcal{C}}^{\star}$ and $f_x^{\star}$ are different maps. We now show that on each η_0 cone the latter map is the identity as well.

Corollary 6.5 *Assume either $f = f_S$ or $f = f_\Sigma$ is anchored. Let $x \in \mathcal{T}_{\mathcal{MC}}(P, \eta_0)$ be a point with*

$$\tau_x \ge \max(\tau_0, 8\eta_0^2 \mathsf{d}_1).$$

Then, for $\alpha \in P$, we have $f_x^{\star}(\alpha) = \alpha$.

Proof Let $R = \tau_x / 2\eta_0$. Then the R–decomposition at x is P_x and the $(R/2\mathsf{K})$–decomposition of $f(x)$ is also P_x. This is because, by Remark 6.3 and Proposition 6.1 we have, for $\alpha \in P_x$,

$$\tau_{f(x)}(\alpha) \ge \frac{\tau_x}{4\mathsf{K}\eta_0} = \frac{R}{2\mathsf{K}}.$$

Hence, $f_x^{\star}(\alpha)$ is some curve in P_x.

Take $\gamma \in P$, $\gamma \neq \alpha$, and let γ' be a curve intersecting γ once or twice and that is disjoint from other curves in P, and let

$$P' = (P - \{\gamma\}) \cup \{\gamma'\}.$$

Let $x' \in \mathcal{T}_{\mathcal{MC}}(P', \eta_0)$ be a point such that

$$x_\beta = x'_\beta \quad \text{for all } \beta \in P - \{\gamma\} \quad \text{and} \quad \tau_x(\gamma) = \tau_{x'}(\gamma').$$

Let g be a path connecting x to x' that is constant in all other factors. The length of γ changes by at least $2\tau_x/\eta_0$. As argued before, since α remains short along g, the length of $f_x^\star(\alpha)$ changes by at most $2N\mathsf{D}_1$ where $N = \tau_x/\rho_1 R$. The choice of τ_x says this is less than the change in the length of γ. Hence, $f_x^\star(\alpha) \neq \gamma$. Since we can use this argument for every curve $\gamma \neq \alpha$, we conclude $f_x^\star(\alpha) = \alpha$. □

The restriction to the thick part

Proposition 6.6 *Assume $f_{\mathcal{P}}^\star$ is the identity. Then there is constant $\mathsf{D}_{\text{thick}}$ such that if x is ℓ_0–thick, then*

$$d_{\mathcal{T}}(f(x), x) \le \mathsf{D}_{\text{thick}}.$$

Proof Let $\mathsf{B} > 0$ be the Bers constant, which has the property that the set of curves of length at most B fill x. To find an upper-bound for $d(x, f(x))$ it is enough to show that α has bounded length in $f(x)$ for each α with $\operatorname{Ext}_x(\alpha) \le \mathsf{B}$. (This follows, for example, from [5, Theorem B] and the fact that extremal length and hyperbolic lengths are comparable in a thick surface x.)

Let η_0 be as in Proposition 6.2 and τ_0 be as in Proposition 6.1. For α as above, let P be a pants decomposition containing α, and $z \in \mathcal{T}_{\mathcal{MC}}(P, \eta_0)$ be such that $\tau_z = \tau_0$ and $d_{\mathcal{T}}(x, z) \overset{+}{\prec} \tau_0$. Then, by Proposition 6.1,

$$\tau_{f(z)}(\alpha) \ge \frac{\tau_0}{4\mathsf{K}\eta_0}.$$

Now we have

$$\mathsf{K}\tau_0 \overset{+}{\succ} d_{\mathcal{T}}(f(x), f(z)) \overset{+}{\succ} \log \frac{\operatorname{Ext}_{f(x)}(\alpha)}{\operatorname{Ext}_{f(z)}(\alpha)} \overset{+}{\succ} \log \operatorname{Ext}_{f(x)}(\alpha).$$

Hence,

$$\log \operatorname{Ext}_{f(x)}(\alpha) \overset{+}{\prec} \mathsf{K}\tau_0.$$

That is, the length of all such α in $f(x)$ is uniformly bounded and hence $d_{\mathcal{T}}(x, f(x))$ is uniformly bounded as well. □

Grouping of sizes

Definition 6.7 Fix once and for all

$$\mu = 64\mathsf{K}^4\mathsf{D}_1^2.$$

Suppose we have a set $\{d_1, d_2, \dots, d_m\}$ of positive numbers. We say that η *is an admissible scale for this set* if there is a decomposition $\mathcal{E}$ of this set such that, if $d_i, d_j \in E$ for $E \in \mathcal{E}$, then $d_j/d_i \le \eta$, and if they are in different subsets, then $\max(d_i/d_j, d_j/d_i) \ge \mu\,\eta$. We refer to $\mathcal{E}$ as *the partition associated to* η. We call the set $E \in \mathcal{E}$ containing the largest elements the *top group*.

Lemma 6.8 *Given $\eta_0 > 1$, there are scales $\eta_0 < \eta_1 < \cdots < \eta_m$ such that for any set $\{d_1, d_2, \dots, d_m\}$ of positive numbers, some η_i is an admissible scale for this set. In fact, we can define η_i recursively as*

$$\eta_{i+1} = \mu\eta_i^3.$$

Proof For $1 \le i \le m$, let $\mathcal{E}_i$ be the partition of $\{d_1, d_2, \dots, d_m\}$ containing the smallest number of subsets such that, for $E \in \mathcal{E}_i$ and $d, d' \in E$, we have $d/d' \le \eta_i$. If we also have $d/d' \ge \mu\eta_i$ for $d \ge d'$ in different sets, then η_i is an admissible scale and we are done. Otherwise, for every i, there are two distinct sets $E, E' \in \mathcal{E}_i$, with $d \in E$ and $d' \in E'$, where $d \ge d'$ and $d/d' \le \eta_i\mu$. Then, for any $c \in E$ and $c' \in E'$ we have

$$\frac{c}{c'} \le \frac{c}{d} \cdot \frac{d}{d'} \cdot \frac{d'}{c'} \le \mu\eta_i^3 = \eta_{i+1}.$$

This means E and E' fit in one group in $\mathcal{E}_{i+1}$ and

$$|\mathcal{E}_{i+1}| \le |\mathcal{E}_i| - 1.$$

If this holds for all i, then some $\mathcal{E}_j$ has size 1 and η_j would be an admissible scale. □

The shortest curves are preserved

Proposition 6.9 *Assume either $f = f_S$ or $f = f_\Sigma$ is anchored. For τ_1 large enough, the following holds. Let $x \in \mathcal{T}$ be a point such that $\tau_x(\alpha) \ge \tau_1$ for $\alpha \in P_x$, and let η be an admissible scale for the set $\{\tau_x(\alpha)\}_{\alpha \in P_x}$ with the associated partition $\mathcal{E}$. Let E be the top group in $\mathcal{E}$ and let $\boldsymbol{\alpha}$ be the set of curves α where $\tau_x(\alpha) \in E$. Then, for every $\alpha \in \boldsymbol{\alpha}$, we have*

$$\frac{\tau_x}{\sqrt{\mu}\eta} \le \tau_{f(x)}(\alpha) \le 2\mathsf{K}\tau_x,$$

and $f_x^\star(\alpha) = \alpha$.

Proof For τ_1 large enough,

$$\tau_{f(x)}(\alpha) \le d_{\mathcal{T}}(f(x), \mathcal{T}_{\text{thick}}) \le \mathsf{K} d_{\mathcal{T}}(x, \mathcal{T}_{\text{thick}}) + \mathsf{C} + \mathsf{D}_{\text{thick}} \le \mathsf{K}\tau_x + \mathsf{C} + \mathsf{D}_{\text{thick}} \le 2\mathsf{K}\tau_x.$$

Hence, we have the upper-bound. We also require that

$$\tau_1 > \frac{128\mathsf{D}_1^2\mathsf{K}\eta_m}{\rho_1},$$

where $\eta_m \ge \eta$ is from Lemma 6.8. Let $r = \rho_1\tau_1$ and let $y = f(x)$. Let $\boldsymbol{\beta}$ be the set of curves β with $\tau_y(\beta) \ge \tau_x/\sqrt{\mu}\eta$.

Claim $$|\boldsymbol{\beta}| \le |\boldsymbol{\alpha}|.$$

Proof of claim The proof is essentially the same as the proof of Proposition 6.1. Choose a path g that changes the lengths of curves in $\boldsymbol{\alpha}$ only connecting x to a point x' such that $\tau_{x'}(\alpha) = \tau_x/\eta\mu$ for $\alpha \in \boldsymbol{\alpha}$, and $d_{\mathcal{T}}(x, x') \le \tau_x$. Since all other curves in x are already shorter that $\tau_x/\eta\mu$, we have

$$\tau_{x'} = \frac{\tau_x}{\eta\mu}.$$

We can cover g with points

$$x = x_1, \dots, x_N = x'$$

so that $d_{\mathcal{T}}(x_i, x_{i+1}) \le r$ and $N = \tau_x/r$. Let $k = |\boldsymbol{\alpha}|$ and $y' = f(x')$. Then, as in the proof of Proposition 6.1, only the lengths of k curves can change substantially from y to y'. More precisely, if $\boldsymbol{\beta}$ has more than k curves, then there is a $\beta \in \boldsymbol{\beta}$ such that

$$|\tau_y(\beta) - \tau_{y'}(\beta)| \le 2N\mathsf{D}_1.$$

Also, $\tau_{y'} \le 2\mathsf{K}\tau_{x'} = 2\mathsf{K}\tau_x/\eta\mu$. Hence,

$$\tau_y(\beta) \le \tau_{y'}(\beta) + 2N\mathsf{D}_1 \le \frac{2\mathsf{K}\tau_x}{\eta\mu} + \frac{2\mathsf{D}_1\tau_x}{r} \le \frac{\tau_x}{3\eta\sqrt{\mu}}.$$

The last inequality comes from our choice of μ, τ_1 and the fact that $r = \rho_1\tau_1$. We have a contradiction to the assumption $\beta \in \boldsymbol{\beta}$. This proves the claim. ◁

We now show that, in fact, $\boldsymbol{\alpha} = \boldsymbol{\beta}$. Assume there is a curve $\alpha \in \boldsymbol{\alpha} - \boldsymbol{\beta}$. Let w be a point such that

$$w_\beta = y_\beta \quad \text{for all } \beta \in \boldsymbol{\beta}, \qquad \tau_w(\beta) = 0 \quad \text{for all } \beta \in P_y - \boldsymbol{\beta}, \qquad d_{\mathcal{T}}(y, w) \le \frac{\tau_x}{\sqrt{\mu}\eta}.$$

Let P be a pants decomposition containing $\boldsymbol{\beta}$ such that

$$\mathrm{i}(P, \alpha) \ne 0 \qquad \text{and} \qquad \mathrm{Ext}_w(\beta') \le \mathsf{B} \quad \text{for all } \beta' \in P - \boldsymbol{\beta}.$$

Let $y' \in \mathcal{T}_{\mathcal{MC}}(P, \sqrt{\mu}\eta)$ be a point obtained from w by pinching curves $\beta' \in P - \boldsymbol{\beta}$ until $\tau_{y'}(\beta') = \tau_y/\sqrt{\mu}\eta$. Then

$$d_{\mathcal{T}}(w, y') \le \frac{\tau_y}{\sqrt{\mu}\eta} + \mathsf{d}_{\mathrm{pinch}} \le \frac{(2\mathsf{K}+1)\tau_x}{\sqrt{\mu}\eta} \quad \text{and hence} \quad d_{\mathcal{T}}(y, y') \le \frac{(2\mathsf{K}+2)}{\sqrt{\mu}\eta}\tau_x.$$

By Proposition 6.1, all curves in P are still short in $x' = f^{-1}(y')$. This means α is not short, and hence

$$d_{\mathcal{T}}(x, x') \ge \tau_x(\alpha) \ge \frac{\tau_x}{\eta}.$$

But we also have

$$d_{\mathcal{T}}(x, x') \le \mathsf{K} d_{\mathcal{T}}(y, y') + \mathsf{C} \le \frac{(2\mathsf{K}^2 + 2\mathsf{K})}{\sqrt{\mu}\eta}\tau_x + \mathsf{C}.$$

The choice of μ gives a contradiction between the last two inequalities. This proves $\boldsymbol{\alpha} \subset \boldsymbol{\beta}$. This and the claim imply $\boldsymbol{\alpha} = \boldsymbol{\beta}$.

We now show $f_x^\star(\alpha) = \alpha$ for $\alpha \in \boldsymbol{\alpha}$. This essentially follows from Corollary 6.5. Let g be a path connecting x to a point $x' \in \mathcal{T}_{\mathcal{MC}}(P_x, \eta_0)$ that is constant in the α coordinate (that is, changes the length of all curves until they are comparable to α). Note that $\tau_{x'}$ is large enough that Corollary 6.5 applies. But, by Corollary 6.5, $f_{x'}^\star(\alpha) = \alpha$. And, as we have argued in the proof of claim in Proposition 6.1, as we move along g from x' to x, the curve α remains short both in $g(t)$ and $f(g(t))$. Hence, Proposition 5.3 applies to all points along this path and $f_x^\star(\alpha) = \alpha$ as well. We are done. □

7 Applying induction

We start by proving the base case of induction (Theorem 1.2). Note that when $\xi(\Sigma) = 1$, the surface Σ is connected and is either a punctured torus or a four-times-punctured sphere and $\mathcal{T}(\Sigma) = \mathbb{H}$.

Proposition 7.1 *Assume $\xi(\Sigma) = 1$ and K_Σ and C_Σ are given. Then Theorem 1.2 holds for $\mathsf{L}_\Sigma = 0$ and some constant D_Σ.*

Proof Let $L \ge \mathsf{L}_\Sigma = 0$ be given. Consider a point $z \in \mathcal{T}(\Sigma, L)$. For $i = 1, 2$, let $g_i \colon [a_i, b_i] \to \mathcal{T}(\Sigma, L)$ be a geodesic such that $g_i(0) = z$ and $g_i(a_i)$ and $g_i(b_i)$ lie on distinct L–horocycles, and such that g_1 and g_2 have a definite angle between them; see Figure 2. Define $\overline{g}_i = f_\Sigma \circ g_i$. Then $\overline{g}_i$ is a (K_Σ, C_Σ)–quasigeodesic. Since f_Σ is anchored, the endpoints of $\overline{g}_i$ are uniformly bounded distance from the endpoints of g_i. This implies $\overline{g}_i$ is contained in a uniformly bounded neighborhood of g_i. This

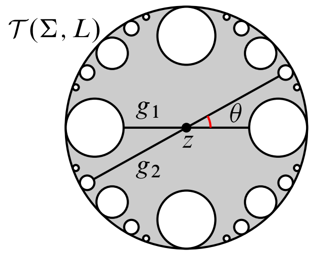

Figure 2: The geodesics g_1 and g_2 pass through z, have their end points on $\partial_L(\Sigma)$ and the angle θ between them is of a definite size.

means $f(z) = \overline{g}_i(0)$ is contained in a uniformly bounded neighborhood of both g_1 and g_2. But, since g_1 and g_2 have a definite size angle between them, the diameter of this set is uniformly bounded. Hence, $d_{\mathbb{H}}(z, f(z))$ is uniformly bounded. □

Our plan is to apply induction by removing from Σ all components of complexity 1 and by cutting along the shortest curves.

Proposition 7.2 *Assume that Σ has a component W with $\xi(W) = 1$. Suppose that $f\colon \mathcal{T}(\Sigma, L) \to \mathcal{T}(\Sigma, L)$ is a $(\mathsf{K}_\Sigma, \mathsf{C}_\Sigma)$–quasi-isometry that is C_Σ–anchored. Pick a large R so that the statements in Section 5 apply. Then, there is a constant D_W such that the following holds. Let $z \in \mathcal{T} = \mathcal{T}(\Sigma, L)$ be a point such that*

$$d_{\mathcal{T}}(z, \mathcal{T}_{\mathrm{LR}}) \geq R \quad \text{and} \quad d_{\mathcal{T}}(z, \partial_L(\Sigma)) \geq R. \tag{18}$$

Then

$$d_{\mathcal{T}(W)}(z_W, f(z)_W) \leq \mathsf{D}_W.$$

Proof Let $\Sigma' = \Sigma - W$. We denote a point $x \in \mathcal{T}(\Sigma)$ as a tuple $(x_W, x_{\Sigma'})$. Let $r = \rho_1 R$.

Claim 1 Let $g\colon [a,b] \to \mathcal{T}(W)$ be a geodesic that stays a distance at least R from $\partial_L(W)$ and let

$$x_t = (g(t), z_{\Sigma'}) \in \mathcal{T}(\Sigma, L).$$

Define

$$\overline{g}(t) = f(x_t)_W.$$

Then $\overline{g}$ is a quasigeodesic.

Proof of Claim 1 To see the upper bound, we note that, for times s and t,

$$\begin{aligned} d_{\mathcal{T}(W)}(\overline{g}(t), \overline{g}(s)) = d_{\mathcal{T}(W)}(f(x_t)_W, f(x_s)_W) &\leq d_{\mathcal{T}(\Sigma)}(f(x_t), f(x_s)) \\ &\leq \mathsf{K}\, d_{\mathcal{T}(\Sigma)}(x_t, x_s) + \mathsf{C} = \mathsf{K}\,|t-s| + \mathsf{C}. \end{aligned}$$

We now check the lower bound. Pick a sequence of points in $\mathcal{T}(\Sigma, L)$

$$x_s = x_1, \ldots, x_N = x_t$$

so that $d_{\mathcal{T}(\Sigma)}(x_i, x_{i+1}) \le r$ and $N \le |t-s|/r + 1$. Since W is always a factor in any decomposition, we know from Proposition 5.1 that

$$d_{\mathcal{T}(\Sigma')}(f(x_i)_{\Sigma'}, f(x_{i+1})_{\Sigma'}) \le 2\mathsf{D}_1.$$

Therefore (assuming $r \ge 4\mathsf{K}\mathsf{D}_1$),

$$d_{\mathcal{T}(\Sigma')}(f(x_s)_W, f(x_t)_W) \le 2N\mathsf{D}_1 \le \frac{|t-s|}{2\mathsf{K}} + 2\mathsf{D}_1.$$

Now the desired lower bound in the claim follows:

$$\begin{aligned} d_{\mathcal{T}(W)}(\overline{g}(s), \overline{g}(t)) &= d_{\mathcal{T}(W)}(f(x_s)_W, f(x_t)_W) \\ &\ge d_{\mathcal{T}(\Sigma)}(f(x_s), f(x_t)) - d_{\mathcal{T}(\Sigma')}(f(x_s)_{\Sigma'}, f(x_t)_{\Sigma'}) \\ &\ge \frac{1}{\mathsf{K}}\, d_{\mathcal{T}(\Sigma)}(f(x_s), f(x_t)) - \mathsf{C} - \frac{|t-s|}{2\mathsf{K}} - 2\mathsf{D}_1 \\ &= \frac{|t-s|}{2\mathsf{K}} - (\mathsf{C} + 2\mathsf{D}_1). \end{aligned} \qquad \triangleleft$$

Next, we show that if the endpoints of g are close to $\partial_L(W)$, then the endpoints of $\overline{g}$ are close to the endpoints of g, which would imply, exactly as in the proof of Proposition 7.1, that $\overline{g}$ stays near g. That is, the reader should think of w below as an endpoint of g.

Claim 2 Let $w \in \mathcal{T}(W)$ be a point such that $d_{\mathcal{T}(W)}(w, \partial_L(W)) = R$. Then

$$d_{\mathcal{T}(W)}(f(w, z_{\Sigma'})_W, w) \le \frac{2L\mathsf{D}_1}{r} + (\mathsf{K}+1)R + 2\mathsf{C}. \tag{19}$$

Proof of Claim 2 We choose a sequence of points

$$z_{\Sigma'} = u_1, \ldots, u_N$$

in $\mathcal{T}(\Sigma')$ satisfying

$$d_{\mathcal{T}(\Sigma')}(\partial_L(\Sigma'), u_N) = R \quad \text{and} \quad d_{\mathcal{T}(\Sigma')}(u_N, z_{\Sigma'}) \le L,$$

and also

$$d_{\mathcal{T}(\Sigma')}(u_i, u_{i+1}) \le r \quad \text{and} \quad N \le \frac{L}{r}.$$

Let $z_i = (w, u_i)$. Note that z_1 is the point of interest. By Proposition 5.1, we have

$$d_{\mathcal{T}(W)}(f(z_i)_W, f(z_{i+1})_W) \le 2\mathsf{D}_1.$$

Hence,

$$d_{\mathcal{T}(W)}(f(z_1)_W, f(z_N)_W) \le 2N\mathsf{D}_1 \le \frac{2L\mathsf{D}_1}{r}.$$

But z_N is distance R from some point $z_L \in \partial_L(\Sigma)$. Hence

$$\begin{aligned} d_{\mathcal{T}(W)}(w, f(z_N)_W) &\le d_{\mathcal{T}}(z_N, f(z_N)) \\ &\le d_{\mathcal{T}}(z_N, z_L) + d_{\mathcal{T}}(z_L, f(z_L)) + d_{\mathcal{T}}(f(z_L), f(z_N)) \\ &\le R + \mathsf{C} + \mathsf{K}R + \mathsf{C} \le (\mathsf{K}+1)R + 2\mathsf{C}. \end{aligned}$$

The claim follows from the last two inequalities by the triangle inequality. $\triangleleft$

We now prove the proposition. For $i = 1, 2$, consider geodesic segments

$$g_i \colon [a_i, b_i] \to \mathcal{T}(W)$$

such that $g_i(0) = z_W$, $g_i(a_i)$ and $g_i(b_i)$ are distance R from $\partial_L(W)$, and g_1 and g_2 make a definite size angle at z_W. By Claim 1, the paths $\overline{g}_i$ are quasigeodesics, and Claim 2 provides a bound for the distance between $g_i(a_i)$ and $\overline{g}_i(a_i)$ and also between $g_i(b_i)$ and $\overline{g}_i(b_i)$. We can make g_1 and g_2 as long as needed. If the lengths of g_i are long enough compared with the right-hand side of equation (19), this implies that $\overline{g}_i$ stays in a uniform neighborhood of g_i and, in particular, $f(z)_W = \overline{g}_i(0)$ is near g_i for $i = 1, 2$. But g_1 and g_2 make a definite size angle. Hence, $f(z)_W$ is near z_W. We are done. $\square$

Proposition 7.3 *Assume either $\mathcal{T} = \mathcal{T}(S)$, or $\mathcal{T} = \mathcal{T}(\Sigma, L)$ and Σ has no component with complexity one. Let $f \colon \mathcal{T} \to \mathcal{T}$ be a (K, C)–quasi-isometry such that the restriction of f to the thick part is $\mathsf{D}_{\text{thick}}$–close to the identity. Pick a large R so that the statements in Section 5 apply. Then there is a constant D_{Top} such that the following holds. Let $z \in \mathcal{T}$ be a point such that*

$$d_{\mathcal{T}}(z, \mathcal{T}_{\text{LR}}) \ge \mu\eta R \quad \text{and} \quad d_{\mathcal{T}}(z, \partial_L(\Sigma)) \ge R$$

(the second condition applies only when $\mathcal{T} = \mathcal{T}(\Sigma, L)$), *and let α be the shortest curve in z, so $\tau_z = \tau_z(\alpha)$. Then*

$$d_{\mathcal{T}(\alpha)}(z_\alpha, f(z)_\alpha) \le \mathsf{D}_{\text{Top}}.$$

Proof The proof is essentially the same as the proof of Proposition 7.2. Let $\Sigma' = S - \alpha$ or $\Sigma - \alpha$. For any geodesic $g \colon [a, b] \to \mathcal{T}(\alpha)$ with $g(0) = z_\alpha$ that stays τ_z/η_0 away from $\mathcal{T}_{\text{LR}}$ and $\partial_L(\alpha)$, we let x^t be a point in $\mathcal{T}$ that projects to $g(t)$ in $\mathcal{T}(\alpha)$ and has the same projection to Σ' as z. Then α is still in the top group of x^t. Define

$$\overline{g}(t) = f(x^t)_\alpha.$$

Assuming L_Σ is large enough, such geodesics exist. Also, by Proposition 6.9, we have $f^\star_{x^t}(\alpha) = \alpha$. As in Claim 1 in Proposition 7.2, the map $\overline{g}$ is a quasigeodesic with uniform constants. We choose the end points of g so that $\tau_{x^a}(\alpha) = \tau_{x^b}(\alpha) = \tau_z/\eta_0$.

Choose a sequence of points in $\mathcal{T}(\Sigma')$,

$$z_{\Sigma'} = u_1, \dots, u_N, \quad d_{\mathcal{T}(\Sigma')}(u_N, \mathcal{T}_{\text{thick}}(\Sigma')) = \frac{\tau_z}{\eta_0},$$

similar to Claim 2 in Proposition 7.2, and let $z^i \in \mathcal{T}(S)$ be a point such that

$$z^i_\alpha = g(a) \quad \text{and} \quad z^i_\Sigma = u_i.$$

As before, we have (here τ_z plays the role of L in Claim 2 above)

$$d_{\mathcal{T}(\alpha)}(f(z^1)_\alpha, f(z^N)_\alpha) \le \frac{2\tau_z \mathsf{D}_1}{r}.$$

Also, using the fact that points in the thick part move by at most $\mathsf{D}_{\text{thick}}$ and z^N has a distance τ_z/η_0 to the thick part, as in Claim 2 of Proposition 7.2 we have

$$d_{\mathcal{T}(\alpha)}(f(z^N)_\alpha, g(a)) \le d_{\mathcal{T}}(f(z^N), z^N) \le (\mathsf{K}+1)\frac{\tau_z}{\eta_0} + \mathsf{C} + \mathsf{D}_{\text{thick}}.$$

Noting that $f(z^1_\alpha) = \overline{g}(a)$, by the triangle inequality we have

$$d_{\mathcal{T}(\alpha)}(\overline{g}(a), g(a)) \le \left(\frac{2\mathsf{D}_1}{r} + \frac{\mathsf{K}+1}{\eta_0}\right)\tau_z + (\mathsf{C} + \mathsf{D}_{\text{thick}}). \tag{20}$$

The same holds for $\overline{g}(b)$ and $g(b)$. Note that $a < 0 < b$ and $|a|$ and b are larger than $\frac{1}{2}\tau_z$. Assuming η_0 is large compared with K, and τ_z is large compared with the additive error in equation (20), we have that the distance between the endpoints of $\overline{g}$ and g is much less than the length of g, and hence, $\overline{g}(0)$ is uniformly close to g.

Taking two such geodesics passing through z_α that make a definite angle with each other, we have that $f(z)_\alpha$ is uniformly close to z_α. □

Proof of Theorem 1.2 Let $z \in \mathcal{T}(\Sigma, L)$ be a point satisfying equation (18). Let Σ' be the surface obtained from Σ after removing all subsurfaces of complexity one. Let α be the shortest curve in $z_{\Sigma'}$. Let $\Sigma'' = \Sigma' - \alpha$.

By Proposition 7.2, for any component W of Σ with $\xi(W) = 1$ we have

$$d_{\mathcal{T}(W)}(z_W, f(z)_W) \le \mathsf{D}_W. \tag{21}$$

For $x' \in \mathcal{T}(\Sigma', L)$, the point $z' = (z_W, x')$ is a point in $\mathcal{T}(\Sigma, L)$. Define, for some $\overline{L}$,

$$f_{\Sigma'} \colon \mathcal{T}(\Sigma', L) \to \mathcal{T}(\Sigma', \overline{L}), \quad f_{\Sigma'}(x') = f(z')_{\Sigma'}.$$

We will show that $f_{\Sigma'}$ is a quasi-isometry. It will also turn out that $|\bar{L}-L|=O(1)$. We first show that $f_{\Sigma'}$ is coarsely onto. This follows from equation (21), which says that under the map f the slice $\{(z_W,x') : x' \in \mathcal{T}(\Sigma',L)\}$ is mapped a bounded distance from itself. The same is true for f^{-1} and the coarse onto statement follows.

For $x',x'' \in \mathcal{T}(\Sigma',L)$, let $z'=(x',z_W)$ and $z''=(x'',z_W)$. Then

$$\begin{aligned}
d_{\mathcal{T}(\Sigma')}(f_{\Sigma'}(x'), f_{\Sigma'}(x'')) &= d_{\mathcal{T}(\Sigma')}(f(z')_{\Sigma'}, f(z'')_{\Sigma'}) \\
&\le d_{\mathcal{T}(\Sigma)}(f(z'), f(z'')) \\
&\le \mathsf{K}\, d_{\mathcal{T}(\Sigma)}(z',z'') + \mathsf{C} = \mathsf{K}\, d_{\mathcal{T}(\Sigma')}(x',x'') + \mathsf{C}.
\end{aligned}$$

Hence, we only need to check the lower bound. It is enough to prove this for points x' and x'' that have a distance of at least R from $\mathcal{T}_{\mathrm{LR}}(\Sigma')$ and $\partial_L(\Sigma')$. Assuming L_Σ is large enough, such points exist and form a connected subset of $\mathcal{T}(\Sigma',L)$. But we have shown that

$$d_{T(W)}(f(z')_W, z'_W) \le \mathsf{D}_W \quad \text{and} \quad d_{T(W)}(f(z'')_W, z''_W) \le \mathsf{D}_W.$$

Hence, since $z'_W = z''_W = z_W$,

$$\begin{aligned}
d_{\mathcal{T}(\Sigma')}(f_{\Sigma'}(x'), f_{\Sigma'}(x'')) &= d_{\mathcal{T}(\Sigma')}(f(z')_{\Sigma'}, f(z'')_{\Sigma'}) \\
&\ge d_{\mathcal{T}}(f(z'), f(z'')) - d_{\mathcal{T}(W)}(f(z')_W, f(z'')_W) \\
&\ge \frac{1}{\mathsf{K}} d_{\mathcal{T}}(z',z'') - \mathsf{C} - 2\mathsf{D}_W \\
&\ge \frac{1}{\mathsf{K}} d_{\mathcal{T}(\Sigma')}(x',x'') - (\mathsf{C} + 2\mathsf{D}_W).
\end{aligned}$$

That is, $f_{\Sigma'}$ is a $(\mathsf{K}_{\Sigma'}, \mathsf{C}_{\Sigma'})$–quasi-isometry for $\mathsf{K}_{\Sigma'} = \mathsf{K}$ and $\mathsf{C}_{\Sigma'} = \mathsf{C} + 2\mathsf{D}_W$.

Now, Propositions 6.1 and 6.2 apply. In fact, since f is anchored, we can conclude that $f^\star_{\mathcal{P}}$ is the identity. Hence, Proposition 6.6 implies that $f_{\Sigma'}$ is in fact $\mathsf{D}_{\mathrm{thick}}$–close to the identity in the thick part of $\mathcal{T}(\Sigma',L)$. Therefore, by Proposition 7.3,

$$d_{\mathcal{T}(\alpha)}(z_\alpha, f(z)_\alpha) \le \mathsf{D}_{\mathrm{Top}}. \tag{22}$$

Let $L'' = \tau_{z_{\Sigma'}} = \tau_z(\alpha)$. Now, for any $u \in \mathcal{T}(\Sigma'',L'')$, let $z'_u \in \mathcal{T}(\Sigma',L'')$ be the point that projects to z_α in $\mathcal{T}(\alpha)$ and projects to u in Σ''. Call the set of points z'_u which project to z_α the *slice* through z_α. At each point on the slice, α is the shortest curve. For some L''', we can define

$$f_{\Sigma''}\colon \mathcal{T}(\Sigma'',L'') \to \mathcal{T}(\Sigma'',L'''), \quad f_{\Sigma''}(u) = f_{\Sigma'}(z'_u)_{\Sigma''}.$$

A similar argument to the one above shows that $f_{\Sigma''}$ is a quasi-isometry. Again by equation (22), both f and f^{-1} preserve the slice up to bounded error, which says the map $f_{\Sigma''}$ is coarsely onto. The upper and lower bounds in the definition of quasi-isometry go as before.

We next show that up to bounded additive error for any z'_u in the slice, α is the shortest curve on $f(z'_u)$. For let β be the shortest curve. Applying (22) to f^{-1} we see that

$$\tau_{f(z'_u)}(\alpha) \le \tau_{f(z'_u)}(\beta) \le \tau_{z'_u}(\beta) + \mathsf{D}_{\mathrm{Top}} \le \tau_{z'_u}(\alpha) + \mathsf{D}_{\mathrm{Top}} \le \tau_{f(z'_u)}(\alpha) + 2\mathsf{D}_{\mathrm{Top}}.$$

This says that $|L'''-L''| \le 2\mathsf{D}_{\mathrm{Top}}$ and so by introducing a slightly larger additive error in the constants of quasi-isometry we can assume the image of our map is in $\mathcal{T}(\Sigma'', L'')$.

We now show that $f_{\Sigma''}$ is anchored. This is because if $u \in \partial_{L''}(\Sigma'')$, then for every $\beta \in P_u$ we have $\tau_u(\beta) = L''$ and every curve β is the shortest curve in z'_u. Then as in the proof of equation (22), the projection of $f_{\Sigma''}(u) = f_{\Sigma'}(z'_u)$ to every $\mathcal{T}(\beta)$ is also close to the projection of z'_u to $\mathcal{T}(\beta)$ which in turn is the same as the projection of u to $\mathcal{T}(\beta)$. That is, $f_{\Sigma''}(u)$ is close to u.

By induction (Theorem 1.2 applied to Σ''), $f_{\Sigma''}$ is $\mathsf{D}_{\Sigma''}$–close to the identity. We have shown that the projections of z to $\mathcal{T}(W)$, $\mathcal{T}(\alpha)$ and $\mathcal{T}(\Sigma'')$ are close to the projections of $f(z)$ to the same. That is, $f(z)$ is close to z. But the set of points satisfying equation (18) is R–dense in $\mathcal{T}(\Sigma, L)$. Thus, for an appropriate value of D_Σ, the theorem holds. □

Proof of Theorem 1.1 The proof is the same as above. From Proposition 6.2 we have that there is an isometry of $\mathcal{T}(S)$ such that if we precompose f with this isometry, then $f_{\mathcal{P}}^{\star}$ is the identity. Assuming this is done, Proposition 6.6 implies that the restriction of f to $\mathcal{T}_{\mathrm{thick}}$ is $\mathsf{D}_{\mathrm{thick}}$–close to the identity.

Now consider a point $z \in \mathcal{T}(S)$ and let α be the shortest curve in z. Choose R large enough that the statements in Section 6 and Proposition 7.3 apply and so that $R \ge \mathsf{L}_\Sigma$ for any subsurface Σ of S. If $d_{\mathcal{T}}(z, \mathcal{T}_{\mathrm{LR}}) \ge \mu\eta R$, then, applying Proposition 7.3,

$$d_{\mathcal{T}(\alpha)}(z_\alpha, f(z)_\alpha) \le \mathsf{D}_{\mathrm{Top}}. \tag{23}$$

Let $\Sigma = S - \alpha$, let $L = \tau_z$ and, as before, for some $\overline{L}$, define a map

$$f_\Sigma \colon \mathcal{T}(\Sigma, L) \to \mathcal{T}(\Sigma, \overline{L})$$

as follows: For $u \in \mathcal{T}(\Sigma, L)$, let $x \in \mathcal{T}(S)$ be a point such that

$$x_\Sigma = u \quad \text{and} \quad x_\alpha = z_\alpha.$$

Now, define

$$f_\Sigma(u) = f(x)_\Sigma.$$

As we argued in the proof of Theorem 1.2, this map is a quasi-isometry with uniform constants and it is anchored. Hence, by Theorem 1.2, we have

$$d_{\mathcal{T}(\Sigma)}(z_\Sigma, f(z)_\Sigma) \le D_\Sigma. \tag{24}$$

The theorem follows from equations (23) and (24). $\square$

References

[1] **J Behrstock**, **B Kleiner**, **Y Minsky**, **L Mosher**, *Geometry and rigidity of mapping class groups*, Geom. Topol. 16 (2012) 781–888 MR

[2] **B H Bowditch**, *Large-scale rank and rigidity of the Weil–Petersson metric*, preprint (2015) Available at `http://tinyurl.com/bowditchWP`

[3] **B H Bowditch**, *Large-scale rank and rigidity of the Teichmüller metric*, J. Topol. 9 (2016) 985–1020 MR

[4] **B H Bowditch**, *Large-scale rigidity properties of the mapping class groups*, Pacific J. Math. 293 (2018) 1–73 MR

[5] **Y-E Choi**, **K Rafi**, **C Series**, *Lines of minima and Teichmüller geodesics*, Geom. Funct. Anal. 18 (2008) 698–754 MR

[6] **A Eskin**, *Quasi-isometric rigidity of nonuniform lattices in higher rank symmetric spaces*, J. Amer. Math. Soc. 11 (1998) 321–361 MR

[7] **A Eskin**, **B Farb**, *Quasi-flats and rigidity in higher rank symmetric spaces*, J. Amer. Math. Soc. 10 (1997) 653–692 MR

[8] **A Eskin**, **B Farb**, *Quasi-flats in* $\mathbb{H}^2 \times \mathbb{H}^2$, from "Lie groups and ergodic theory" (S G Dani, editor), Tata Inst. Fund. Res. Stud. Math. 14, Tata Inst. Fund. Res., Bombay (1998) 75–103 MR

[9] **A Eskin**, **D Fisher**, **K Whyte**, *Coarse differentiation of quasi-isometries, I: Spaces not quasi-isometric to Cayley graphs*, Ann. of Math. 176 (2012) 221–260 MR

[10] **A Eskin**, **D Fisher**, **K Whyte**, *Coarse differentiation of quasi-isometries, II: Rigidity for Sol and lamplighter groups*, Ann. of Math. 177 (2013) 869–910 MR

[11] **A Eskin**, **H Masur**, **K Rafi**, *Large-scale rank of Teichmüller space*, Duke Math. J. 166 (2017) 1517–1572 MR

[12] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Mathematical Series 49, Princeton Univ. Press (2012) MR

[13] **U Hamenstädt**, *Geometry of the mapping class groups, III: Quasi-isometric rigidity*, preprint (2007) arXiv

[14] **J H Hubbard**, *Teichmüller theory and applications to geometry, topology, and dynamics, I*, Matrix Editions, Ithaca, NY (2006) MR

[15] **N V Ivanov**, *Automorphism of complexes of curves and of Teichmüller spaces*, Int. Math. Res. Not. 1997 (1997) 651–666 MR

[16] **B Kleiner**, **B Leeb**, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, Inst. Hautes Études Sci. Publ. Math. 86 (1997) 115–197 MR

[17] **H A Masur**, **Y N Minsky**, *Geometry of the complex of curves, II: Hierarchical structure*, Geom. Funct. Anal. 10 (2000) 902–974 MR

[18] **Y N Minsky**, *Extremal length estimates and product regions in Teichmüller space*, Duke Math. J. 83 (1996) 249–286 MR

[19] **P Pansu**, *Métriques de Carnot–Carathéodory et quasiisométries des espaces symétriques de rang un*, Ann. of Math. 129 (1989) 1–60 MR

[20] **I Peng**, *Coarse differentiation and quasi-isometries of a class of solvable Lie groups, I*, Geom. Topol. 15 (2011) 1883–1925 MR

[21] **I Peng**, *Coarse differentiation and quasi-isometries of a class of solvable Lie groups, II*, Geom. Topol. 15 (2011) 1927–1981 MR

[22] **K Rafi**, *A combinatorial model for the Teichmüller metric*, Geom. Funct. Anal. 17 (2007) 936–959 MR

[23] **K Rafi**, *Hyperbolicity in Teichmüller space*, Geom. Topol. 18 (2014) 3025–3053 MR

[24] **H L Royden**, *Automorphisms and isometries of Teichmüller space*, from "Advances in the theory of Riemann surfaces" (L V Ahlfors, L Bers, H M Farkas, R C Gunning, I Kra, H E Rauch, editors), Ann. of Math. Studies 66, Princeton Univ. Press (1971) 369–383 MR

[25] **R E Schwartz**, *The quasi-isometry classification of rank one lattices*, Inst. Hautes Études Sci. Publ. Math. 82 (1995) 133–168 MR

[26] **R E Schwartz**, *Quasi-isometric rigidity and Diophantine approximation*, Acta Math. 177 (1996) 75–112 MR

Department of Mathematics, University of Chicago
Chicago, IL, United States

Department of Mathematics, University of Chicago
Chicago, IL, United States

Department of Mathematics, University of Toronto
Toronto, ON, Canada

eskin@math.uchicago.edu, masur@math.uchicago.edu,
rafi@math.toronto.edu

Proposed: Benson Farb — Received: 6 October 2017
Seconded: Tobias H Colding, Jean-Pierre Otal — Accepted: 23 April 2018