Algebraic & Geometric Topology

Volume 24 (2024)

A qualitative description of the horoboundary of the Teichmüller metric

Aitor Azemar

msp

Algebraic & Geometric Topology **24**:7 (2024) 3919–3984
DOI: 10.2140/agt.2024.24.3919
Published: 9 December 2024

A qualitative description of the horoboundary of the Teichmüller metric

AITOR AZEMAR

Two commonly studied compactifications of Teichmüller spaces of finite type surfaces with respect to the Teichmüller metric are the horofunction and visual compactifications. We show that these two compactifications are related, by proving that the horofunction compactification is finer than the visual compactification. This allows us to use the straightforwardness of the visual compactification to obtain topological properties of the horofunction compactification. Among other things, we show that Busemann points of the Teichmüller metric are not dense within the horoboundary, answering a question of Liu and Su. We also show that the horoboundary of Teichmüller space is path connected, determine for which surfaces the horofunction compactification is isomorphic to the visual one and show that some horocycles diverge in the visual compactification based at some point. As an ingredient in one of the proofs we show that extremal length is not C^2 along some paths that are smooth with respect to the piecewise linear structure on measured foliations.

30F60, 32G15; 51F30

1 Introduction

The horofunction compactification of a metric space is defined in terms of the metric, so its properties are well aligned for studying the metric properties of the space. For example, all geodesic rays converge to points and isometries of the space can be extended to homeomorphisms of the compactification. This compactification was first introduced by Gromov [16] as a natural, general compactification, based on

previous ideas of Busemann. The horofunction compactification has since found many applications; it was used to obtain asymptotic properties of random walks on weakly hyperbolic spaces by Maher and Tiozzo [29], to determine the isometry group of some Hilbert geometries by Lemmens and Walsh [25] and to obtain properties of quantum metric spaces by Rieffel [39]. The compactification is obtained by embedding the metric space X into the space $C(X)$ of continuous functions on X via the map $h : X \hookrightarrow C(X)$ defined by

$$h(p)(\cdot) = d(p, \cdot) - d(p, b), \tag{1}$$

where $b \in X$ is an arbitrarily chosen basepoint. As explained, for example, by Walsh [42, Section 2], if the space X is proper then h is an embedding, the closure of $h(X)$ is compact and the *horofunction compactification* of X is defined as the pair $(h, \overline{h(X)})$. By considering two functions equivalent if they differ by a constant one can show that the compactification does not depend on the basepoint b. While this compactification has been rather useful, it is sometimes hard to visualize, and there are not that many examples where the horofunction boundary is explicitly known. Some cases where the horofunction compactification is understood include Hadamard manifolds and some of their quotients, by Dal'bo, Peigné, and Sambusetti [8], as well as the Heisenberg group with the Carnot–Carathéodory metric, by Klein and Nicas [24], and Hilbert geometries, by Walsh [43].

On the other hand, for a proper, uniquely geodesic, straight metric space X (see Section 2 for definitions) the visual compactification based at some point $b \in X$ is defined by pasting the set of geodesic rays exiting b, denoted D_b, to the space X in such a way that a sequence $(x_n) \subset X$ converges to some ray $\gamma \in D_b$ if the distance $d(b, x_n)$ goes to infinity as $n \to \infty$, and the geodesic ray between b and x_n converges uniformly on compacts to γ. This compactification was introduced by Eberlein and O'Neil [10] as a generalization of the Poincaré disk model, and we give a brief introduction in Section 2. This compactification may depend on the basepoint b, which restricts its usefulness. It can even happen that isometries of X that move the basepoint cannot be extended continuously to the compactification, as Kerckhoff showed for Teichmüller spaces [23]. However, the visual compactification usually has a simple geometric interpretation. For example, for a Hadamard manifold, as well as for a Teichmüller space with the Teichmüller metric, this compactification is homeomorphic to a closed ball of the same dimension as the space, where the boundary of that ball is the space of geodesic rays based at b. In the context of Teichmüller spaces with the Teichmüller metric, the visual compactification is often called the Teichmüller compactification.

1.1 Horoboundary of proper, uniquely geodesic, straight metric spaces

To make this work as general as possible, we begin our analysis by using the aforementioned metric properties of the Teichmüller metric. The relationship between the horofunction compactification and the visual compactification is established by observing that, for such a metric space, a sequence converging to a point in the horofunction compactification also converges in the visual compactification. This allows us to build a continuous map Π_b from the horofunction compactification $\overline{h(X)}$ to the visual compactification

$X \cup D_b$, showing that the former is finer than the latter. In the context of Teichmüller spaces without boundary, the map Π_b coincides with the one defined by Liu and Shi [27, Definition 3.3]. We may denote this map as simply Π when the basepoint is not relevant to the discussion.

Given a geodesic γ, the path $\gamma(t)$ converges, as $t \to \infty$, to the *Busemann point* associated to γ in the horofunction compactification, which we denote by B_γ. As the map Π is defined in terms of sequences it follows that $\Pi(B_\gamma) = \gamma$. The existence of the map Π shows a strong relation between the horofunction and the visual compactification, which we state in the following result.

Theorem 1.1 *Let (X, d) be a proper, uniquely geodesic, straight metric space. For any basepoint $b \in X$, there is a continuous surjection Π from the horofunction compactification to the visual compactification based at b such that $\Pi(B_\gamma) = \gamma$ for every ray γ starting at b and $\Pi(h(p)) = p$ for every $p \in X$.*

In particular, the horofunction compactification of X is finer than the visual compactification of X based at any point.

Most of the subsequent results in the paper follow as applications of this theorem.

It is not the first time that a map such as Π appears in the literature. Similar maps have been found for δ–hyperbolic spaces by Webster and Winchester [45]. Walsh [43] defined such a map for Hilbert geometries, which satisfy the hypothesis of the theorem whenever there are no coplanar noncollinear segments in the boundary of the convex set, as shown by de la Harpe [17, Proposition 2].

The map Π does not induce a fiber bundle, as its fibers $\Pi^{-1}(\gamma)$ vary from points to higher dimensional sets (see Theorem 6.10). Still, Theorem 1.1 characterizes the horoboundary as the disjoint union of all the fibers $\Pi^{-1}(\gamma)$. Furthermore, our analysis of the topology of these fibers shows that they are path connected (see Proposition 3.11), which gives the following characterization of the connectivity of the horoboundary.

Proposition 1.2 *The horoboundary of a proper, uniquely geodesic straight metric space is connected if and only if its visual boundary based at some point (and hence, any) is connected.*

The *Busemann map* B from the visual compactification $X \cup D_b$ to the horofunction compactification is defined by setting $B(\gamma) = B_\gamma$ for each geodesic ray $\gamma \in D_b$ and $B(p) = h(p)$ for each $p \in X$. With this definition, the map satisfies $\Pi \circ B = \mathrm{id}$. As the next result shows, the continuity of this map is related with the topology of the horofunction compactification.

Proposition 1.3 *The visual compactification of a proper, uniquely geodesic, straight metric space based at some point is isomorphic to its horofunction compactification if and only if the Busemann map is continuous.*

The Busemann map is essentially the identity inside X, so the only possible points of discontinuity are at the boundary. It is therefore of interest to find a criterion for the continuity of B at the boundary, which turns out to give a criterion for when the fibers $\Pi^{-1}(\gamma)$ are singletons.

Proposition 1.4 *Let X be a proper, uniquely geodesic, straight metric space, $b \in X$ a basepoint and B the corresponding Busemann map. Furthermore, let γ be a geodesic ray based at b. Then the following three statements are equivalent*:

(1) *The Busemann map B restricted to the boundary is continuous at γ.*

(2) *The fiber $\Pi^{-1}(\gamma)$ is a singleton.*

(3) *The Busemann map B is continuous at γ.*

In other words, we have reduced the continuity of B to the continuity restricted to the boundary. This result can then be applied to different settings to obtain a more precise characterization. In the case of Teichmüller spaces, Proposition 1.4 can be used to get an explicit criterion for the continuity of the Busemann map in terms of the quadratic differentials associated to the geodesic rays, giving us a characterization of the fibers that are singletons.

1.2 Horoboundary of the Teichmüller metric

Many compactifications have been defined for Teichmüller space, such as Thurston's compactification, the visual compactification (also known as the Teichmüller compactification) and the Gardiner–Masur compactification. These compactifications play an important role in the study of mapping class groups and asymptotic aspects of Teichmüller space. See for example the articles by Thurston [41], Kerckhoff [23] or Ohshika [36]. The main reason multiple compactifications have been introduced is that each one has been designed with a certain application in mind.

Thurston's compactification takes the rather simple shape of a ball, upon which the mapping class groups acts as homeomorphisms. This facts make this compactification well suited for studying properties of the mapping class group. Indeed, Thurston's classification of the elements of the mapping class group relied on this compactification [41]. However, the Teichmüller metric is not directly related to the compactification, which results in some quirks when trying to use it to study the asymptotic geometry. For example, Lenzhen, Modami, and Rafi [26] prove that there exist geodesic rays with high-dimensional limit sets.

The visual compactification is defined directly using the metric, and takes the shape of a sphere where each point in the boundary has a clear geometric interpretation. This makes the compactification a good tool to interpret asymptotic geometric results. For example, Walsh [44, Theorem 7] has proven that all geodesic rays converge to points in the visual boundary. However, as proved by Kerckhoff [23], the action of the mapping class group does not extend continuously to this compactification, which implies that the compactification depends on the choice of basepoint. This fact limits the usability of the visual compactification.

The Gardiner–Masur compactification was initially defined to study the asymptotic properties of extremal lengths, following an analogous construction to that of the Thurston's compactification. It was later proved by Liu and Su [28] that this compactification is isomorphic to the horofunction compactification

with respect to the Teichmüller metric, giving it a geometric meaning. Furthermore, the mapping class group extends continuously to the compactification. These two properties make the Gardiner–Masur compactification a good candidate to study asymptotic properties of the Teichmüller metric. However, as noted by Miyachi [32] and Liu and Su [28], there is a lack of information on the shape of this compactification. In this paper, we start working towards an understanding of the shape of this boundary.

Let S be a compact surface with (possibly empty) boundary and finitely many marked points, where we allow marked points to be on the boundary. Denote by $\mathcal{T}(S)$ its Teichmüller space equipped with the Teichmüller metric. Furthermore, for any quadratic differential q based at some basepoint $b \in \mathcal{T}(S)$, denote by $R(q;\cdot)$ the geodesic ray in $\mathcal{T}(S)$ starting at b in the direction q, and $V(q)$ the vertical foliation associated to q, see Section 4 for a quick introduction or the book by Farb and Margalit [11] for a more in-depth explanation of these concepts. Recall that a measured foliation is *indecomposable* if it is either a thickened curve, or a component with a transverse measure that cannot be expressed as the sum of two projectively distinct nonzero transverse measures. Furthermore, each measured foliation can be decomposed uniquely into finitely many indecomposable components (see Section 4.1 for detailed definitions). Walsh has shown the following characterization of the convergence of Busemann points in terms of the convergence of the associated quadratic differentials.

Theorem 1.5 (Walsh [44, Theorem 10]) *Let (q_n) be a sequence of unit area quadratic differentials based at $b \in \mathcal{T}(S)$. Then, $B_{R(q_n;\cdot)}$ converges to $B_{R(q;\cdot)}$ if and only if both of the following hold:*

(1) *(q_n) converges to q with respect to the L^1 norm on $T_b^*\mathcal{T}(S)$.*

(2) *For every subsequence $(G^n)_n$ of indecomposable measured foliations such that, for each $n \in \mathbb{N}$, G^n is a component of $V(q_n)$, we have that every limit point of G^n is indecomposable.*

While Walsh's proof is done in the context of surfaces without boundary, it can be easily extended to our setting. In view of this theorem, we say that a sequence of quadratic differentials (q_n) *converges strongly* to q if it satisfies the two conditions of Theorem 1.5. Furthermore, we say that q is *infusible* if every sequence of quadratic differentials converging to q converges strongly. By Proposition 1.4, a quadratic differential q is infusible if and only if the Busemann map is continuous at $R(q;\cdot)$. In Theorem 5.4, we derive a topological characterization of the vertical foliations of infusible quadratic differentials. This allows us to determine precisely which surfaces only admit infusible quadratic differentials, yielding the following result.

Theorem 1.6 *Let S be a compact surface of genus g with b_m and b_u boundary components with and without marked points respectively and p interior marked points. Then the horofunction compactification of $\mathcal{T}(S)$ is isomorphic to the visual compactification if and only if $3g + 2b_m + b_u + p \leq 4$.*

This result had been previously proven by Miyachi [32] for surfaces without boundary, that is, when $b_m = b_u = 0$. For the cases where we do not have an isomorphism Miyachi found non-Busemann points in the boundary. These points are in the closure of Busemann points, which prompted Liu and Su to ask the following question:

Question 1.7 (Liu and Su [28, Question 1.4.2]) *Is the set of Busemann points dense in the horofunction boundary?*

We give a negative answer to this question, summed up in the following result.

Theorem 1.8 *Let S be a closed surface of genus g with p marked points. Then the Busemann points are not dense in the horofunction boundary of $\mathcal{T}(S)$ whenever $3g + p \geq 5$.*

To achieve this result we use Liu and Su's [28] and Walsh's [44] characterization of the horofunction compactification as the Gardiner–Masur compactification. We use an equivalent but slightly different definition than usual for the Gardiner–Masur compactification, as the definition we use is more well suited for our computations, and more easily extendable to surfaces with boundary (see Section 4.4 for the precise definition). For each point in the horofunction compactification there is an associated real-valued function on the set of measured foliations. We show that the functions associated to elements in the closure of Busemann points are polynomials of degree 2 with respect to some variables (see Proposition 6.2 for the precise statement). We then show that the elements of the Gardiner–Masur boundary found by Fortier Bourque [13] do not satisfy that condition. The main ingredient for this last part of the reasoning is the following result, which shows that extremal length is not C^2 along certain smooth paths in $\mathcal{MF}$.

Theorem 1.9 *Let S be a closed surface of genus g with p marked points and empty boundary satisfying $3g + p \geq 5$. Then there is a point $X \in \mathcal{T}(S)$ and a path G_t, $t \in [0, t_0]$, in the space of measured foliations on X, smooth with respect to the canonical piecewise linear structure of the space of measured foliations, such that $\mathrm{Ext}(G_t)$ is not C^2.*

The canonical piecewise linear structure of the space of measured foliations was developed by Bonahon [3; 4; 5]. The first derivative of the extremal length along such a path was determined by Miyachi [33], so our proof is based on finding cases where Miyachi's expression is not C^1. This follows from an explicit computation, whose complication is greatly reduced by using previous estimates established by Markovic [30].

The relation between the Thurston compactification and the horofunction compactification was studied by Miyachi [34]. He proves that, while neither Thurston's nor the horofunction compactification is finer than the other, there is a bicontinuous map from the union of $\mathcal{T}(S)$ and uniquely ergodic foliations in Thurston's boundary to a subset of the horofunction compactification. Masur showed [31] that this result can be interpreted to say that these two compactifications are the same almost everywhere according to the Lebesgue measure on Thurston's boundary. The image of uniquely ergodic foliations by the bicontinuous map is the set of Busemann points associated to uniquely ergodic foliations. As we show in Theorem 7.5, this set is nowhere dense within the horoboundary. Hence the map defined by Miyachi does not show that these two are the same almost everywhere according to any strictly positive measure on the horoboundary. In fact, any attempt to extend the identity map from the interior of the Thurston compactification to the interior of the horoboundary compactification to a set of full measure within the Thurston compactification results in the same problem.

Corollary 1.10 *Let ν be any finite strictly positive measure on the horoboundary and let μ be the Lebesgue measure on the Thurston boundary. Furthermore, let ϕ be a map from the Thurston compactification to the horofunction compactification satisfying $\phi|_{\mathcal{T}(S)} = h$, where h is as in* (1)*. Then there is no subset U of the Thurston boundary with full μ–measure such that ϕ is continuous at every point in U and $\phi(U)$ has full ν–measure.*

Under some smoothness assumptions satisfied by Teichmüller metric, we are able to use the maps Π_b to give an alternative definition of the horofunction compactification based on geometric notions. This definition characterizes the horofunction compactification as the reachable subset of the product of all visual compactifications obtained by choosing different basepoints (see Section 3.3 for details). Hence, the horofunction compactification can be interpreted as a collection of the asymptotic information provided by all visual compactifications. As a straightforward result of this alternative definition we get the following characterization of converging sequences in the horofunction compactification.

Corollary 1.11 *A sequence $(x_n) \subset \mathcal{T}(S)$ converges in the horofunction compactification if and only if the sequence converges in all the visual compactifications.*

Considering the horocycles diverging in the horofunction compactification found by Fortier Bourque [13] we get that there is some visual compactification in which these horocycles do not converge.

Corollary 1.12 *Let S be a closed surface of genus g with p marked points, such that $3g + p \geq 5$. There is a basepoint such that a horocycle diverges in the visual compactification based at that point.*

This contrasts with the behavior of Teichmüller rays, which converge in all visual compactifications (see [44, Theorem 7] by Walsh).

The structure of the horoboundary provided by Theorem 1.1, as well as the path-connectivity of the fibers, allows us to prove the following connectivity result.

Theorem 1.13 *The horoboundary of any Teichmüller space of real dimension at least* 2 *is path connected.*

Furthermore, we also prove that whenever the surface has empty boundary the map Π restricted to the horoboundary admits a section, while it only admits a section for surfaces of low complexity if the boundary is nonempty (see Theorem 8.1 for details).

Figure 1 shows a sketch of what we think the horoboundary looks like based on the results of this paper. The outer circle represents the section given by Theorem 8.1. Each line perpendicular to the sphere represents one of the fibers induced by the map Π, so it is associated with a unique Teichmüller ray starting at b. Note that while by Proposition 3.11 the fibers are path connected, by Theorem 6.10 they are bigger than segments in some cases. Furthermore, a priori they might not be contractible.

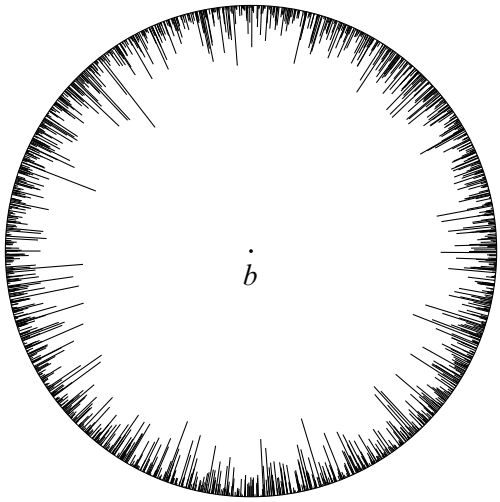

Figure 1: The shape of the horoboundary of the Teichmüller metric for surfaces without boundary.

The nearest point to the basepoint b of each fiber represents the Busemann point associated to the geodesic joining b to the fiber. This point could indeed be considered the nearest point to b from the fiber, as one can access it in a straight way, through a geodesic exiting b. On the other hand, the points in the outer circle represents the points associated to the section alluded to earlier. These can be accessed through a sequence of Busemann points whose associated fiber is a point, which can be considered as the most tangentially possible way to reach points in the boundary. Following a result by Masur [31], with respect to the measure on the fibers induced by the Lebesgue measure on the set of Teichmüller rays exiting b, almost all the fibers are actually points. As we shall see in Theorem 7.5 these points are nowhere dense in the boundary.

Note that there exist paths within the horoboundary connecting the fibers without passing through the section, and a priori there may be paths not represented in the sketch along which the fibers vary continuously. For surfaces for which the map Π does not admit a global section, a similar sketch could be drawn, although there would be no continuous global section in some cases. Hence, the outer circle would be broken at some places.

Finally, Liu and Su's and Walsh's characterization of the horofunction compactification as the Gardiner–Masur compactification can be used to translate some of these findings to results regarding the asymptotic value of extremal length functions. For example, we get the following estimate.

Theorem 1.14 *Let (q_n) be a sequence of unit quadratic differentials converging strongly to a unit quadratic differential q. Denote by G_j the components of the vertical foliation associated to q, and $H(q)$ the horizontal foliation. Then, for any $F \in \mathcal{MF}$ and sequence (t_n) of real values converging to positive infinity we have*

$$\lim_{n\to\infty} e^{-2t_n} \operatorname{Ext}_{R(q_n;t_n)}(F) = \sum_j \frac{i(G_j, F)^2}{i(G_j, H(q))}.$$

This generalizes a previous result by Walsh [44, Theorem 1], where the same is shown for q_n constant.

1.3 Outline of the paper and a note for the reader interested in surfaces without boundary

The paper is structured as follows. In Section 2 we introduce the necessary metric notions used in the paper. We follow in Section 3 by proving the results related to the more general metric setting, such as showing that the horofunction compactification is finer than the visual one. In Section 4 we give a short review of the necessary background on Teichmüller spaces. In Section 5 we determine which quadratic differentials are infusible, and find which surfaces admit infusible quadratic differentials, getting a proof of Theorem 1.6. In Section 6 we characterize the points in the closure of Busemann points, and get some bounds on the dimension of the fibers of the map Π. In Section 7 we show that Busemann points are not dense. In Section 8 we determine which surfaces result in the map Π having a section, and prove that the horoboundary is path connected. Finally, in Section 9 we use the previous results to obtain estimates regarding asymptotic values of extremal lengths.

Some of the most dense parts of this paper are due to the added complexity of considering surfaces with boundary. As such, the reader focused on surfaces with empty boundary might want to omit the corresponding sections on a first reading. One of the largest related parts starts after the remark following Theorem 5.4 and ends before the start of Section 5.2. The other sizable part starts with Proposition 8.3 and ends at the start of the proof of Theorem 1.13, where we note that the proof is significantly simpler in the case of surfaces without boundary.

Acknowledgments

The author would like to thank Maxime Fortier Bourque and Vaibhav Gadre for many helpful discussions and corrections.

2 Metric definitions

2.1 Compactifications

A compactification of a space serves, among other things, as a way of characterizing convergence to infinity. Formally, a *compactification* of a topological space X is a pair $(f, \overline{X})$, where $\overline{X}$ is a compact topological space and $f: X \to \overline{X}$ is an embedding such that $f(X)$ is dense in $\overline{X}$. The boundary of a compactification $\partial\overline{X} = \overline{X} - X$ describes the different ways of converging to infinity provided by that compactification. We shall usually identify the points in X with the ones in $\overline{X}$ via the map f, and say that a sequence $(x_n) \subset X$ converges in $\overline{X}$ if $f(x_n)$ converges.

A compactification (f_1, X_1) is *finer* than another one (f_2, X_2) if there exists a continuous map $\bar{f_2}: X_1 \to X_2$ such that $\bar{f_2} \circ f_1 = f_2$. Since $f_2(X)$ is dense in X_2, the continuous extension $\bar{f_2}$ is surjective. Furthermore, we can restrict the map $\bar{f_2}$ to the boundary to get a surjective map $\bar{f_2}|_{\partial X_1}: \partial X_1 \to \partial X_2$, which can be seen as a projection. Having a compactification finer than another ones means, from an intuitive point of view, that the finer compactification catalogs more ways of converging to infinity than the other one. Namely, any sequence in X converging in the finer compactification converges also in the coarser one, while the opposite may not be true.

We say that two compactifications are isomorphic if each one is finer than the other one. The following lemma found in [44, Lemma 17] coincides with the intuitive notion of finer compactifications.

Lemma 2.1 *Let (f_1, X_1) and (f_2, X_2) be two compactifications of a Hausdorff topological space X such that f_2 extends continuously to an injective map $\bar{f_2}\colon X_1 \to X_2$. Then the two compactifications are isomorphic.*

We will usually refer to the space $\overline{X}$ as the compactification when the embedding is clear from the context. Since the images of X by the embedding are dense, the extensions we get to compare the compactifications are unique. That is, we have the following result:

Lemma 2.2 *Let (f_1, X_1) and (f_2, X_2) be two compactifications of a Hausdorff topological space X such that X_1 is finer than X_2. Then the extension $\bar{f_2}\colon X_1 \to X_2$ is unique.*

Proof For any $x \in X$ we have $\bar{f_2}(f_1(x)) = f_2(x)$. Hence, the image of $\bar{f_2}$ is determined on a dense subset of X_1, so by continuity it is determined on X_1. □

2.2 Visual compactification of proper, uniquely geodesic, straight spaces

Let (X, d) be a metric space. We shall say that a map γ from an interval $I \subset \mathbb{R}$ to X is a *geodesic* if it is an isometric embedding, that is, if $d(\gamma(t), \gamma(s)) = |t - s|$. We shall consider two geodesics to be equal if their image is equal and have the same orientation. A space is *uniquely geodesic* if for any two distinct points $a, b \in X$ there is a unique geodesic starting at a and ending at b.

Furthermore, we say that the space is *proper* if the closed balls $D(x, r) = \{p \in X \mid d(p, x) \leq r\}$ are compact.

If geodesic segments can be extended uniquely, that is, if for any geodesic segment γ_1 there is a unique biinfinite geodesic γ_2 such that $\gamma_1 \cap \gamma_2 = \gamma_1$, we say that the space is *straight*.

Let then X be a proper, uniquely geodesic, straight space and let D_b be the set of infinite geodesic rays starting at b, with the topology given by uniform convergence on compact sets. Furthermore, denote by $S_b^1 = \{x \in X \mid d(x, b) = 1\}$ the sphere of radius 1 around b.

Lemma 2.3 *The map from D_b to S_b^1 defined by sending $\gamma \in D_b$ to $\gamma(1)$ is a homeomorphism.*

Proof Since the topology on D_b is given by uniform convergence on compact sets, the point $\gamma(1)$ varies continuously with respect to γ.

On the other hand, since the space is straight and has unique geodesics, given any point $a \in S_b^1$ there is a unique geodesic ray starting at b and passing through a. This is the inverse to the map obtained by evaluating the geodesics. To see that the relation is continuous we consider a sequence $(a_n) \subset S_b^1$ converging to some a, and denote by (γ_n) and γ the associated geodesics. Assume γ_n does not converge to γ. Then we have a subsequence without γ as an accumulation point. For any $t > 0$, the geodesic

segments $\gamma|_{[0,t]}$ are contained in the ball of radius t, which is compact, as X is proper. As these are geodesics we have equicontinuity, so by Arzelà–Ascoli we can take a subsequence converging uniformly to some path γ'. Since the distance function is continuous, γ' is a geodesic. Furthermore, $\gamma'(1) = \lim_{n\to\infty} \gamma_n(1) = \lim_{n\to\infty} a_n = a$. By uniqueness of geodesics, γ' and γ are equal when restricted to $[0, 1]$, which by straightness implies they are equal. Hence, γ_n converges to γ uniformly on the compact $[0, t]$. □

Following a similar reasoning it is possible to show the following, still under the same hypotheses on X.

Lemma 2.4 *The space X is homeomorphic to $D_b \times [0,\infty)/D_b \times \{0\}$.*

Proof We define the map $C : D_b \times [0,\infty)/D_b \times \{0\} \to X$ given by $C(\theta, r) = \theta(r)$. This is well defined, as $C(\theta, 0) = b$ for any $\theta \in D_b$. Furthermore, this is a bijection, since for every $x \in X - \{b\}$ there is a unique geodesic ray from b to x. The map is continuous, as the topology on D_b is given by uniform convergence on compact sets. To see that the inverse is continuous consider a sequence $a_n \in X$ converging to some $a \in X$. If $a = b$, then $d(a_n, b) \to 0$, so we have continuity. Otherwise we let $r_n = d(a_n, b)$ and $r = d(a, b)$. We have $r_n \to r$, so denoting (γ_n) and γ the unique geodesic in D_b such that $\gamma_n(r_n) = a_n$ and $\gamma(r) = a$ and applying Arzelà–Ascoli's theorem in the same way as in Lemma 2.3, we have that γ_n converges to γ. □

The space $D_b \times [0,\infty)/D_b \times \{0\}$ can be included into the compact space $D_b \times [0,\infty]/D_b \times \{0\}$, which can be written as $(D_b \times [0,\infty)/D_b \times \{0\}) \cup D_b \times \{\infty\}$. Using the homeomorphism from Lemma 2.4, we can use this inclusion to give a compact topology on the space $X \cup D_b$. The *visual compactification* is defined as the pair $(i, X \cup D_b)$, where i is the inclusion $i : X \to X \cup D_b$ and the topology on the space $X \cup D_b$ is the one we just defined. We shall denote $X \cup D_b$ as $\overline{X}^v_b$, or $\overline{X}^v$ when the basepoint is not relevant to the discussion.

2.3 Horofunction compactification

The second compactification that will play a part in this paper is slightly more involved and difficult to visualize.

Let X be a proper, uniquely geodesic, straight metric space. Given a basepoint $b \in X$, one can embed X into the space of continuous functions from X to $\mathbb{R}$ via the map $h : X \to C(X)$ defined by

$$h(x)(\cdot) := d(x, \cdot) - d(x, b).$$

The topology given to $C(X)$ is that of uniform convergence on compact sets. The map h is indeed continuous, as the distance function is continuous. Furthermore, h is injective, as $h(x)$ has a strict global minimum at x. It can also be proven that since X is proper, h is an embedding. For more details about this construction see [42, Section 2]. Furthermore, the properness of X implies it is second countable, so the closure of $h(X)$ is compact, Hausdorff and second countable. We shall denote the closure of $h(X)$ on $C(X)$ as $\overline{X}^h$. The *horofunction compactification* is defined as the pair $(h, \overline{X}^h)$. We call the set $\partial\overline{X}^h = \overline{X}^h - X$ the *horofunction boundary* or *horoboundary*, and we call its members

horofunctions. If we want to specify the chosen basepoint we write $\overline{X}^h_b$. However, it is possible to see that quotienting the compactification by letting $f \sim g$ whenever the difference is constant we get an isomorphic compactification, showing that the horofunction compactification does not depend on the basepoint.

Usually the easier points to identify in the horoboundary are the Busemann points. These are the ones that can be reached as a limit along almost geodesics, which is a slight weakening of the notion of geodesic by allowing an additive constant approaching 0. That is, a path $\gamma\colon [0,\infty) \to X$ is an *almost geodesic* if for each $\varepsilon > 0$,

$$\left| d(\gamma(0),\gamma(s)) + d(\gamma(s),\gamma(t)) - t \right| < \varepsilon$$

for all s and t large enough, with $s \le t$. Rieffel [39] proved that every almost geodesic converges to a limit in $\partial\overline{X}^h$. A horofunction is called a *Busemann point* if there exists an almost geodesic converging to it. We shall denote the Busemann point associated in this way to the almost geodesic γ by B_γ.

3 Horofunction compactification of proper, uniquely geodesic, straight metric spaces

3.1 The relation between the horofunction compactification and the visual compactification

Fix a uniquely geodesic, proper and straight metric space (X,d) and a basepoint $b \in X$. We will assume X satisfies these hypotheses through this section. For each geodesic ray $\gamma \in \partial\overline{X}^v$ starting at b there is an associated Busemann point $B_\gamma \in \partial\overline{X}^h$. We can extend this map to all the visual compactification by setting it as the identification with the map h on X given by the horofunction compactification. That is, we define the *Busemann map* $B\colon \overline{X}^v \to \overline{X}^h$ by setting $B(\gamma) = B_\gamma$ for $\gamma \in \partial\overline{X}^v$ and $B(x) = h(x)$ for $x \in X$. The relevance of this map can be seen with the following result.

Lemma 3.1 *The visual compactification $(i,\overline{X}^v)$ is finer than the horofunction compactification $(h,\overline{X}^h)$ if and only if the map B is continuous.*

Proof We have that $B(i(x)) = h(x)$, so B is an extension of h to $\overline{X}^v$. Hence, if B is continuous, then the visual compactification is finer than the horofunction compactification.

On the other hand, if the visual compactification is finer than the horofunction compactification, then we have a continuous map $f\colon \overline{X}^v \to \overline{X}^h$. For every $x \in X$, we have $f(i(x)) = h(x) = B(i(x))$. Furthermore, for any ray γ starting at the basepoint we have $f(\gamma) = \lim_{t\to\infty} f(i(\gamma(t)) = \lim h(\gamma(t)) = B(\gamma)$. Hence, $B = f$, and B is continuous. □

In general, the Busemann map may not be surjective nor continuous. However, we have the following.

Proposition 3.2 *For a proper, uniquely geodesic, straight metric space (X,d) the Busemann map is injective.*

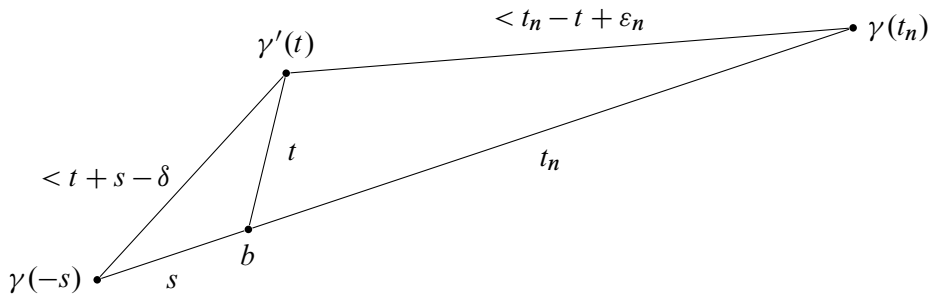

Figure 2: The triangles involved in the proof of Proposition 3.2.

Proof For each $x \in X$, the associated function $h(x)$ has a global minimum at x, while B_γ is unbounded below for every $\gamma \in \partial \overline{X}^v$. Hence, in the interior of $\overline{X}^v$ the map is injective and $B(X) \cap B(\partial \overline{X}^v) = \varnothing$. Assume we have $\gamma, \gamma' \in \partial \overline{X}^v$ such that $\gamma \neq \gamma'$ and $B(\gamma) = B(\gamma') = \xi$. Then, for a given sequence $t_n \to \infty$ we have $\lim_{n\to\infty} h(\gamma(t_n)) = \lim_{n\to\infty} h(\gamma'(t_n)) = \xi$. For any $t \in \mathbb{R}$ and any n such that $t_n > t$ we have

$$h(\gamma(t_n))(\gamma(t)) = d(\gamma(t_n), \gamma(t)) - d(\gamma(t_n), \gamma(0)) = t_n - t - t_n = -t,$$

and similarly for γ'. Hence $\xi(\gamma(t)) = \xi(\gamma'(t)) = -t$ for all t.

Fix now a $t > 0$. We have

$$-t = \xi(\gamma'(t)) = \lim_{n\to\infty} \big(d(\gamma'(t), \gamma(t_n)) - d(b, \gamma(t_n))\big) = \lim_{n\to\infty} \big(d(\gamma'(t), \gamma(t_n)) - t_n\big).$$

That is, there is a sequence ε_n with $\varepsilon_n \to 0$ such that

$$t_n - t + \varepsilon_n \geq d(\gamma'(t), \gamma(t_n)) \geq t_n - t - \varepsilon_n.$$

for every n.

By straightness we can extend γ in the negative direction towards $\gamma(-s)$ for some $s > 0$. We shall now show that the geodesic γ does not minimize the distance between $\gamma(-s)$ and $\gamma(t_n)$ for n big enough. Since the space is straight, the geodesic segment between $\gamma(-s)$ and b can be extended uniquely, so concatenating it with the segment between b and $\gamma'(t)$ does not result in a geodesic. Hence, the distance between $\gamma'(t)$ and $\gamma(-s)$ is strictly smaller than $s + t$. That is, there is some $\delta > 0$ such that $d(\gamma(-s), \gamma'(t)) < t + s - \delta$. As shown in Figure 2 we get a path going from $\gamma(-s)$, to $\gamma(t_n)$, passing through $\gamma'(t)$ that has length less than $t + s - \delta + t_n - t + \varepsilon_n = t_n + s - \delta + \varepsilon_n$. Hence, taking n big enough so that $\varepsilon_n < \delta$ we get that the geodesic segment between $\gamma(-s)$ and $\gamma(t_n)$ is not minimizing. This is a contradiction, from which we conclude that $\gamma = \gamma'$. Therefore, B is injective. □

Hence, given a Busemann point ξ in $B(\partial \overline{X}^v)$ we have a unique associated geodesic ray $\gamma \in \partial \overline{X}^v$ such that $\xi(\gamma(t)) = -t$ for all t. Our next aim is to build a similar relation for all other horofunctions. Our approach is similar to the one used by Walsh in [44, Section 7].

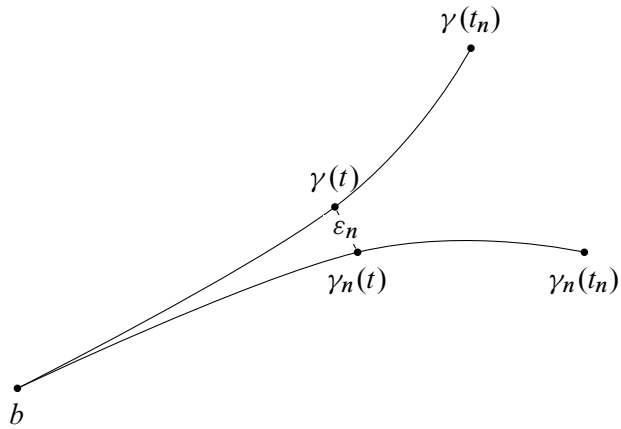

Figure 3: In the proof of Lemma 3.3, γ_n converges to γ, so $\gamma_n(t)$ converges to $\gamma(t)$, and hence the distance between $\gamma_n(t_n)$ and $\gamma_n(t)$ gets arbitrarily close to the distance between $\gamma_n(t_n)$ and $\gamma(t)$.

We say that a geodesic γ is an *optimal geodesic* for a certain horofunction $\xi \in \overline{X}^h$ if $\xi(\gamma(t)) - \xi(\gamma(0)) = -t$ for all $t \in \mathbb{R}$. We shall now see that each function in the horoboundary has at least one optimal geodesic.

Lemma 3.3 *Let X be a proper, uniquely geodesic, straight metric space and let $\xi \in \partial\overline{X}^h$ be a horofunction. Suppose that $(x_n) \subset X$ converges to ξ, with $x_n = \gamma_n(t_n)$, $\gamma_n \in \partial\overline{X}^v$ and (γ_n) converging to γ as $n \to \infty$. Then $\xi(\gamma(t)) = -t$ for every $t \in \mathbb{R}$. That is, $\gamma(t)$ is an optimal geodesic for ξ.*

Proof Fix t. We have that

$$\xi(\gamma(t)) = \lim_{n\to\infty} \big(d(\gamma(t), \gamma_n(t_n)) - d(b, \gamma_n(t_n))\big) = \lim_{n\to\infty} \big(d(\gamma(t), \gamma_n(t_n)) - t_n\big).$$

As n goes to infinity, γ_n converges to γ. Hence by the given topology on the visual boundary, the maps $\gamma_n(\cdot)$ converge uniformly on compact sets to the geodesic $\gamma(\cdot)$. In particular, denoting $d(\gamma(t), \gamma_n(t)) = \varepsilon_n$ we have $\varepsilon_n \to 0$. We get then Figure 3, so by the triangle inequality,

$$|d(\gamma(t), \gamma_n(t_n)) - (t_n - t)| = \big|d(\gamma(t), \gamma_n(t_n)) - d(\gamma_n(t), \gamma_n(t_n))\big| \le \varepsilon_n,$$

and so $\xi(\gamma(t)) = -t$. □

Since $\partial\overline{X}^v$ is compact, for any horofunction $\xi \in \partial\overline{X}^h$ and sequence $(x_n) \subset X$ converging to ξ we can take a subsequence such that the hypotheses of Lemma 3.3 are satisfied, so each $\xi \in \partial\overline{X}^v$ does have at least one optimal geodesic.

If ξ has another optimal geodesic γ' with $\gamma'(0) = \gamma(0)$ we have at least two geodesics along which $\xi(\gamma(t)) = \xi(\gamma'(t)) = -t$ for all t. Following a reasoning similar to the one in the proof of Proposition 3.2, we get a contradiction. This time, however, we have to be a bit more careful about the distances, as instead of two fixed rays we have a fixed ray and a sequence converging to a distinct fixed ray.

Proposition 3.4 *Let $\xi \in \partial \overline{X}^h$ and $b \in X$. Then there is a unique optimal geodesic for ξ passing through b.*

Proof Let $(x_n) = (\gamma_n(t_n))$ be a sequence converging to ξ, with $(\gamma_n) \subset \partial \overline{X}^v$, and take a subsequence such that γ_n converges to some geodesic γ. By Lemma 3.3, γ is an optimal geodesic. Assume that we have a different optimal geodesic γ' passing through b.

Using that $h(\gamma_n(t_n))$ converges pointwise to ξ, we have

$$-t = \xi(\gamma'(t)) = \lim_{n\to\infty} \big(d(\gamma'(t), \gamma_n(t_n)) - d(b, \gamma_n(t_n))\big) = \lim_{n\to\infty} \big(d(\gamma'(t), \gamma_n(t_n)) - t_n\big).$$

Hence, there is a sequence ε_n with $\varepsilon_n \to 0$ such that

$$t_n - t + \varepsilon_n \geq d(\gamma'(t), \gamma_n(t_n)) \geq t_n - t - \varepsilon_n.$$

We proceed by showing that for n big enough there is some $s > 0$ such that the geodesic γ_n does not minimize the distance between $\gamma_n(-s)$ and $\gamma_n(t_n)$. As in the proof of Proposition 3.2, by applying the triangle inequality between $\gamma'(t), \gamma(-s)$ and b we have $d(\gamma'(t), \gamma(-s)) < s + t$. Fix $s > 0$ and pick $\delta > 0$ such that $d(\gamma'(t), \gamma(-s)) < t + s - \delta$. Since γ_n converges to γ uniformly on compact sets, $\gamma_n(-s)$ converges to $\gamma(-s)$. Hence, $d(\gamma'(t), \gamma_n(-s))$ converges to $d(\gamma'(t), \gamma(-s))$. Then for n big enough we have $d(\gamma'(t), \gamma_n(-s)) < t + s - \delta$. Consider then n big enough so that $\varepsilon_n \leq \delta/2$ as well. The triangle between $\gamma'(t), \gamma_n(-s)$ and $\gamma_n(t_n)$ gives

$$d(\gamma_n(-s), \gamma_n(t_n)) \leq d(\gamma_n(-s), \gamma'(t)) + d(\gamma'(t), \gamma_n(t_n)) < (t + s - \delta) + (t_n - t + \varepsilon_n) < t_n + s.$$

This is a contradiction, which proves the uniqueness of γ. □

Given a basepoint $b \in X$ we can now define a map $\Pi_b \colon \overline{X}^h \to \overline{X}^v_b$ by sending any $\xi \in \partial \overline{X}^h$ to the unique optimal geodesic γ of ξ with $\gamma(0) = b$, and by sending $h(x)$ to x for any $x \in X$. This map is indeed an extension of the relation we had established for Busemann points in $\mathcal{B}(\partial \overline{X}^v)$, since if $\xi = B(\gamma)$ for $\gamma \in D_b$ then γ is an optimal geodesic of ξ, giving us $\Pi_b(B(\gamma)) = \gamma$.

We will often write Π instead of Π_b whenever the basepoint is not relevant to the discussion. To prove that Π is continuous, we first have to see the following result.

Proposition 3.5 *Let $(x_n) \subset X$ be a sequence converging to $\xi \in \partial \overline{X}^h$. Then, (x_n) has a unique accumulation point in the visual compactification. Further, this accumulation point depends only on ξ.*

Proof Since $\partial \overline{X}^v$ is compact, (x_n) has accumulation points in the visual compactification. If (x_n) has two accumulation points we can take two subsequences converging to two different geodesics, which by Lemma 3.3 are optimal geodesics, contradicting Proposition 3.4.

If there is another sequence (y_n) converging to ξ with a different accumulation point the result follows by merging both sequences and repeating the reasoning. □

Hence, Π can be alternatively defined by sending any $\xi \in \partial \overline{X}^h$ to the unique accumulation point in $\overline{X}^v$ of the sequences converging to ξ in $\overline{X}^h$, and by sending $h(x)$ to x for any $x \in X$. By Proposition 3.5, this definition is equivalent to the previous one.

By this second definition of the map Π, we see how it is mostly related to the convergence of sequences, so using a diagonal sequence argument we can prove its continuity.

Proposition 3.6 *The map Π is continuous.*

Proof Take a sequence $(\xi_n) \subset \overline{X}^h$ converging to ξ. If $\xi \in h(X)$ we have that, as $h(X)$ is open, $\xi_n \in h(X)$ for n big enough. Hence, $\Pi(\xi_n) = h^{-1}(\xi_n)$, which converges to $h^{-1}(\xi)$, as h is a homeomorphism with its image.

If $\xi \in \partial\overline{X}^h$ we split the sequence into two subsequences, one contained in $h(X)$ and one contained in $\partial\overline{X}^h$. The one contained in $h(X)$ converges to ξ, so by definition of Π and we have $\Pi(\xi) = \lim_{n\to\infty} h^{-1}(\xi_n)$.

Assume then that $(\xi_n) \subset \partial\overline{X}^h$ converges to ξ. We want to see that $\gamma_n = \Pi(\xi_n)$ converges to $\gamma = \Pi(\xi)$. For each ξ_n we can take a sequence $\big(h(\gamma_n^m(t_n^m))\big)_m$ converging, as $m \to \infty$ to ξ_n. By Proposition 3.5 the sequence $\gamma_n^m(t_n^m)$ converges to γ_n. Let γ' be an accumulation point of γ_n. Take a convergent subsequence of γ_n converging to γ', and relabel it as γ_n. Let (V_n) be a nested sequence of open neighborhoods of ξ in $\overline{X}^h$ such that $\xi_n \in V_n$ and $\bigcap_n V_n = \{\xi\}$ and let (W_n) be a nested sequence of open neighborhoods of γ' in $\overline{X}^v$ such that $\gamma_n \in W_n$ and $\bigcap_n W_n = \{\gamma'\}$. We can take such sequences of sets, as both spaces are metrizable.

For each n, there exists $m(n)$ big enough so that $\gamma_n^{m(n)} \in W_n$ and $h(\gamma_n^{m(n)}(t_n^{m(n)})) \in V_n$. By the first condition on $m(n)$, we have that $\gamma_n^{m(n)}$ converges to γ'. By the second condition, $h(\gamma_n^{m(n)}(t_n^{m(n)}))$ converges to ξ, so by the definition of Π and Proposition 3.5 the sequence $\gamma_n^{m(n)}$ converges to $\Pi(\xi) = \gamma$. Hence, $\gamma = \gamma'$, so the only accumulation point of (γ_n) is γ and by compactness of $\partial\overline{X}^v$ the sequence (γ_n) converges to γ. □

By combining Propositions 3.5 and 3.6 we get that Π is the map announced at the introduction, giving us a proof of Theorem 1.1. As mentioned in the introduction, this map shows that the horofunction compactification is finer than the visual compactification. By using the Busemann map to insert the visual boundary inside the horoboundary, we can consider the map Π as a projection.

One straightforward consequence of the continuity of Π_b is as follows.

Corollary 3.7 *Let γ be a geodesic ray, not necessarily starting at the basepoint $b \in X$. Then, γ converges in the visual compactification of X based at b.*

Proof The ray γ converges in the horofunction compactification to B_γ. Since Π_b is continuous, the ray also converges in the visual compactification based at b to $\Pi_b(B_\gamma)$. □

For Teichmüller spaces with the Teichmüller metric this result was first proved by Walsh [44, Theorem 7].

By Lemma 3.1, the visual compactification is finer than the horofunction compactification if and only if the Busemann map is continuous. Hence, since the horofunction compactification is always finer than the visual compactification, we obtain an isomorphism whenever this is the case, resulting in Proposition 1.3.

3.2 The fiber structure

To get a better picture of the shape of the horoboundary we shall study the shape of the preimages of the projection Π restricted to the boundary. That is, for a given point γ in the visual boundary we are interested in finding out information about the fiber $\Pi^{-1}(\gamma)$. We first prove the following lemma, which we will use to get bounds on the values of $\Pi^{-1}(\gamma)$.

Lemma 3.8 *Fix a geodesic ray $\gamma \in \partial\overline{X}^v$ and $p \in X$ not in the biinfinite extension of the geodesic ray γ. Then, the function $h(\gamma(\cdot))(p)$, with domain $[0,\infty)$, is strictly decreasing.*

Proof Take $t, s \geq 0$ with $s < t$. By the triangle inequality we have

$$d(\gamma(t), p) \leq d(\gamma(s), p) + d(\gamma(t), \gamma(s)) = d(\gamma(s), p) + t - s.$$

Further, we have strict inequality, as equality would give us two different paths with the same length between $\gamma(t)$ and p, with one of them being geodesic. Hence,

$$h(\gamma(t))(p) = d(\gamma(t), p) - d(\gamma(t), b) < d(\gamma(s), p) + t - s - t = h(\gamma(s))(p). \qquad \square$$

The set $C(X)$ can be partially ordered by saying that $f \geq g$ whenever $f(x) \geq g(x)$ for all $x \in X$. If $f \geq g$ and $f \neq g$ then we write $f > g$. If $p = \gamma(r)$ for some r and $s < t$ we have $h(\gamma(s))(p) = h(\gamma(t))(p) = -r$ for $r \leq s$ and $-s = h(\gamma(s))(p) > h(\gamma(t))(p) = -\min(r, t)$ otherwise. Hence, adding the previous lemma we have $h(\gamma(s)) > h(\gamma(t))$ whenever $s < t$. By attempting to extend this relation to the horofunction boundary we get that Busemann points are maximal in their fibers.

Proposition 3.9 *Let $\gamma \in \partial\overline{X}^v$ and $\xi \in \Pi^{-1}(\gamma)$. Then, $\xi \leq B(\gamma)$.*

Proof Choose any sequence $(x_n) \subset X$ such that $h(x_n)$ converges to ξ. Since $\xi \in \Pi^{-1}(\gamma)$ the sequence (x_n) converges to γ in $\overline{X}^v$, so we can write $x_n = \gamma_n(t_n)$ with t_n converging to infinity and γ_n converging to γ.

Fix $p \in X$ and let $\varepsilon > 0$. Denote $s_n = \sup\{t : d(\gamma(t), \gamma_n(t)) < \varepsilon \text{ and } t < t_n\}$. The geodesics γ_n converge to γ uniformly on compact sets, so $s_n \to \infty$ as $n \to \infty$. Hence, by definition of the Busemann point and since $d(\gamma_n(s_n), \gamma(s_n)) < \varepsilon$,

$$B_\gamma(p) = \lim_{n\to\infty} h(\gamma(s_n))(p) \geq \limsup_{n\to\infty} h(\gamma_n(s_n))(p) - 2\varepsilon.$$

Furthermore, $s_n \leq t_n$, so by Lemma 3.8,

$$\xi(p) = \lim_{n\to\infty} h(\gamma_n(t_n))(p) \leq \limsup_{n\to\infty} h(\gamma_n(s_n))(p) \leq B_\gamma(p) + 2\varepsilon.$$

Since ε can be arbitrarily small we get the proposition. $\square$

While it might not be possible to get a similar unique minimum in each fiber, we can get the following result.

Proposition 3.10 *Let $\gamma \in \partial \overline{X}^v$ and $\xi \in \Pi^{-1}(\gamma)$. Furthermore, let $(x_n) \subset X$ be a sequence converging to ξ with $x_n = \gamma_n(t_n)$. For any p, define $\eta(p) = \liminf_{n\to\infty} B(\gamma_n)(p)$. Then $\xi \geq \eta$.*

Proof The proof follows a similar reasoning as the last one.

Fix $p \in X$, choose a subsequence so $B(\gamma_n)(p)$ converges to $\eta(p)$ and let (ε_m) be a sequence of positive numbers converging to 0. For each ε_m, take $n(m)$ big enough so that $B(\gamma_{n(m)})(p) \geq \eta(p) - \varepsilon_m$. Further, take s_m bigger than $t_{n(m)}$, and big enough so that

$$h(\gamma_{n(m)}(s_m))(p) \geq B(\gamma_{n(m)})(p) - \varepsilon_m.$$

Such an s_m always exists by the definition of $B(\gamma_{n(m)})$. In particular, we have that

$$\liminf_{m\to\infty} h(\gamma_{n(m)}(s_m))(p) \geq \eta(p).$$

By Lemma 3.8 we have

$$\xi(p) = \lim_{m\to\infty} h(\gamma_{n(m)}(t_{n(m)}))(p) \geq \liminf_{m\to\infty} h(\gamma_{n(m)}(s_m))(p) \geq \eta(p). \qquad \square$$

The intuition one might get from these propositions is that approaching γ "through the boundary", that is, through the furthest way possible from the interior of X, gives a lower bound on the possible values of approaching through other angles, and approaching γ in a straight way, that is, through the geodesic, gives an upper bound. Hence, when these two ways of approaching γ are the same, every other possible angle of approach should also yield the same limit. Following this reasoning we get our next result, announced in the introduction.

Proposition 1.4 *Let X be a proper, uniquely geodesic, straight metric space, $b \in X$ a basepoint and B the corresponding Busemann map. Furthermore, let γ be a geodesic ray based at b. Then the following three statements are equivalent:*

(1) *The Busemann map B restricted to the boundary is continuous at γ.*

(2) *The fiber $\Pi^{-1}(\gamma)$ is a singleton.*

(3) *The Busemann map B is continuous at γ.*

Proof (1) $\Longrightarrow$ (2) Take $\xi \in \Pi^{-1}(\gamma)$. By Proposition 3.9 we have $\xi \leq B(\gamma)$. Since B is continuous at γ when restricted to the boundary we have that for any $\gamma_n \to \gamma$ the horofunctions $B(\gamma_n)$ converge to $B(\gamma)$. Hence, by Proposition 3.10, $\xi \geq B(\gamma)$, so $\xi = B(\gamma)$ and we have (2).

(2) $\Longrightarrow$ (3) Take then any $(x_n) \subset \overline{X}^v$ converging to γ, consider the sequence $(B(x_n)) \subset \overline{X}^h$ and let η be an accumulation point. By the definition of Π we have $\eta \in \Pi^{-1}(\gamma)$, so $\eta = B(\gamma)$ since we assumed that $\Pi^{-1}(\gamma)$ is a singleton. This shows that B is continuous at γ.

Finally, it is clear that (3) $\Longrightarrow$ (1). $\square$

The relation obtained in Lemma 3.8 can be exploited further. Indeed, trying to carry it to the boundary in a more delicate manner we can see that the fibers are path connected.

Proposition 3.11 *Let $\gamma \in \partial \overline{X}^v$. For any $\xi \in \Pi^{-1}(\gamma)$ there exists a continuous path from $B(\gamma)$ to ξ contained in $\Pi^{-1}(\gamma)$.*

Proof Take a sequence $(x_n) \subset X$ converging to ξ in the horofunction compactification, and write $x_n = \gamma_n(u_n)$. As we have seen in the proof of Proposition 3.10, we can take a sequence $(l_n) \subset \mathbb{R}$ with $\gamma_n(l_n)$ converging to B_γ such that $l_n < u_n$ for all n. For each n we have a path $\tilde{\alpha}^n(t)$ connecting $\gamma_n(l_n)$ and $\gamma_n(u_n)$ by setting $\tilde{\alpha}^n(t) = \gamma_n(tu_n + (1-t)l_n)$ for $t \in [0,1]$. We would like to carry this path to the limit, getting a path between ξ and $B(q)$. However, directly taking such a limit might result in some discontinuities, so we have to choose a parametrization carefully.

To find a good parametrization we shall use a certain functional as a control. We want the functional to carry discontinuities and strict increases in the path of functions to discontinuities and strict increases in the value of the functional. Since X is proper, it is separable, so let $(p_i)_{i\in\mathbb{N}}$ be a countable dense set in X. We define the functional $I\colon \overline{X}^h \to \mathbb{R}$ given by

$$I(f) = \sum_{i\in\mathbb{N}} \frac{f(p_i)}{2^i d(b, p_i)}.$$

Since $|f(x)| \le d(b,x)$ for all $f \in \overline{X}^h$, the summation in the definition of $I(f)$ is absolutely convergent, so $I(f)$ is defined, finite, continuous with respect to f, and for any two $f, g \in \overline{X}^h$ we have $I(f+g) = I(f) + I(g)$. Furthermore, since (p_n) is dense and we are taking continuous functions, we have that the functional translates strict inequalities. That is, $f > g$ implies $I(f) > I(g)$. Hence, if $I(f) = 0$ and $f \ge 0$ we have $f = 0$.

We define then the function $F_n(t) = I(h(\gamma_n(t))$. By continuity of I this function is continuous, and by Lemma 3.8 it is strictly decreasing with respect to t. That is, we have continuous strictly decreasing functions $F_n\colon [l_n, u_n] \to [F_n(u_n), F_n(l_n)]$. Hence, we can define implicitly the continuous parametrizations $s_n\colon [0,1] \to [l_n, u_n]$ by taking the unique value $s_n(t)$ such that

$$F_n(s_n(t)) = (1-t)F_n(l_n) + tF_n(u_n).$$

Denote the $F_n(s_n(t))$ as $E_n(t)$. By the continuity of I we have that $E_n(t)$ converges to $(1-t)I(B_\gamma) + tI(\xi)$ as $n \to \infty$, which we denote by $E(t)$.

Take now a countable dense set $(t^k)_{k\in\mathbb{N}} \subset [0,1]$ containing 0 and 1. We are now ready to start defining the path $\alpha\colon [0,1] \to \Pi^{-1}(\gamma)$, and we begin defining it for the dense set (t^k). For $k = 1$ we define $\alpha(t^1)$ as an accumulation point of $h\big(\gamma_n(s_n(t^1))\big)$. Denote $(\gamma_{m^1(n)})$ the subsequence of γ_n such that $h\big(\gamma_{m^1(n)}(s_{m^1(n)}(t^1))\big)$ converges to $\alpha(t^1)$. Define inductively $\alpha(t^k)$ and $(\gamma_{m^k(n)})$ by taking an accumulation point and a corresponding converging subsequence of $h\big(\gamma_{m^{k-1}(n)}(s_{m^{k-1}(n)}(t^k))\big)$. By the continuity of I we have

$$I(\alpha(t^k)) = \lim_{n\to\infty} \big(F_{m^k(n)}(s_{m^k(n)}(t^k))\big) = E(t^k).$$

For each pair $i > j$ we have that $m^i(n)$ is a subsequence of $m^j(n)$, so $h\big(\gamma_{m^i(n)}(s_{m^i(n)}(t^j))\big)$ converges to $\alpha(t^j)$. Assume $t^i > t^j$. By Lemma 3.8 we have that $h\big(\gamma_{m^i(n)}(s_{m^i(n)}(t^i))\big) < h\big(\gamma_{m^i(n)}(s_{m^i(n)}(t^j))\big)$, so $\alpha(t^i) \leq \alpha(t^j)$.

We now have to prove that the definition we have given for α on (t^k) can be extended continuously to $[0,1]$. Fix any $t \notin (t^k)$ and take a subsequence of t^k, labeled t^{k_n}, such that $t^{k_n} \to t$. We shall now see that $\alpha(t^{k_n})$ converges to a function which does not depend on the chosen subsequence, and define $\alpha(t)$ as that limit. We can split and reorder the sequence (t^{k_n}) into (t_n^+) and (t_n^-) satisfying $t_n^+ > t_{n+1}^+ > t > t_{n+1}^- > t_n^-$. The associated $\alpha(t_n^\pm)$ are ordered, so for any $p \in X$ the sequence $\alpha(t_n^\pm)(p)$ is an increasing (or decreasing) sequence of values in $\mathbb{R}$, bounded above (or below) by $\alpha(0)(p)$ (or $\alpha(1)(p)$). Hence, both sequences converge pointwise, which implies uniform convergence on compact sets, as these functions are 1–Lipschitz. Furthermore, these limits do not depend on the chosen sequence, since if we had any other we could intercalate them and the sequences would still converge. Denote then α^+ the limit associated to t_n^+, and α^- the limit associated to t_n^-. Since $\alpha(t_n^+) < \alpha(t_m^-)$ for all n, m we have $\alpha^+ \leq \alpha^-$. For each $\alpha(t^k)$ we have $I(\alpha(t^k)) = E(t^k)$. Hence by the continuity of I we have that

$$I(\alpha^+) = E(t) = I(\alpha^-).$$

That is, we have

$$I(\alpha^- - \alpha^+) = 0.$$

Since α^- and α^+ are continuous and $\alpha^- - \alpha^+ \geq 0$ we have $\alpha^- = \alpha^+$. We thus define $\alpha(t)$ to be either one. The same reasoning shows that α is continuous. □

We would like to remark that several choices where made in the proof of the previous lemma, and the obtained path may not be unique.

We can use the previous result to observe that the horoboundary is connected if and only if the visual boundary is connected.

Proof of Proposition 1.2 Assume that the visual boundary is not connected. Then we have $U, V \subset \partial\overline{X}^v$ nonempty and open such that $U \cap V = \varnothing$ and $U \cup V = \partial\overline{X}^v$. As Π is continuous, the sets $\Pi^{-1}(U)$ and $\Pi^{-1}(V)$ are open, so the horoboundary is not connected.

For the other implication, assume that the visual boundary is connected while the horoboundary is not connected. Then we have $U, V \subset \partial\overline{X}^h$ nonempty and open such that $U \cap V = \varnothing$ and $U \cup V = \partial\overline{X}^h$. Since fibers are path connected, each of them is contained in only one of U or V, so $\Pi(U)$ and $\Pi(V)$ are disjoint. Since $U \cup V = \partial\overline{X}^h$ we have $\Pi(U) \cup \Pi(V) = \partial\overline{X}^v$, and since both U and V are nonempty, so are the images. Hence, both images cannot be open at the same time, as $\partial\overline{X}^v$ is connected. Therefore, these sets cannot be both closed. Assume $\Pi(U)$ is not closed. We then have a sequence $(\gamma_n) \subset \Pi(U)$ converging to a point in $\Pi(V)$. Again, since $U \cup V = \partial\overline{X}^h$, we have that $U = \Pi^{-1}\Pi(U)$ and $V = \Pi^{-1}\Pi(V)$. Hence, any lift of the sequence (γ_n) to $\Pi^{-1}\Pi(U)$ is contained in U and, since $\partial\overline{X}^h$ is compact, has accumulation points which, by the continuity of the projection map, are be contained in $\Pi^{-1}\Pi(V) = V$. Hence, U is not closed and we get a contradiction. □

3.3 An alternative definition of the horofunction compactification

Under what a priori seem to be more restrictive hypotheses on the space X it is possible to characterize the horofunction compactification as a subset of the product of all of its visual compactifications. We detail the construction in this section.

The new extra hypotheses are both related to the differentiability of the distance function. We say a that a uniquely geodesic metric space X is C^1 *along geodesics* if given a point $p \in X$ and a geodesic segment γ that does not intersect p, the distance function $d(\gamma(t), p)$ is first differentiable and the value of the derivative depends continuously on both t and p. Furthermore, the space X has *constant distance variation* if for any two distinct geodesics γ, η with $\gamma(0) = \eta(0)$ we have either

$$\frac{d}{dt} d(\gamma(t), \eta(s))\Big|_{t=0} = \frac{d}{dt} d(\gamma(t), \eta(1))\Big|_{t=0} \tag{2}$$

for all $s > 0$, or $\frac{d}{dt} d(\gamma(t), \eta(s))\big|_{t=0}$ does not exist for any $s > 0$.

Many commonly studied metric spaces have constant distance variation. For example, spaces with bounded curvature, either above of below, have constant distance variation, as explained in the book by Burago, Burago, and Ivanov [6, Section 4]. Importantly to our case, Teichmüller spaces with the Teichmüller distance satisfy both hypotheses. Earle [9] proved that the distance function is C^1 by providing a formula for its derivative. Applying the formula to (2) we get that the derivative depends only on the tangential vector to γ at 0 and the unit area quadratic differential associated to η at 0, so we also have constant distance variation. Furthermore, Teichmüller spaces with the Teichmüller distance are also straight and proper, so the results from this section can be applied to them.

Consider the product of all the possible visual compactifications obtained by changing the basepoint,

$$E = \prod_{b \in X} \overline{X}^v_b,$$

with the usual product topology. See the book by Munkres [35, Chapters 2.19 and 5.37] for some background on infinite products of topological spaces. Denote by π_b the projection from E to $\overline{X}^v_b$. By definition of the product topology, the diagonal inclusion $i\colon X \hookrightarrow E$ such that by $\pi_b(i(x)) = x$ for every $x, b \in X$ is continuous, and has continuous inverse restricted to $i(X)$ given by π_b. Hence, $i(X)$ is homeomorphic to X. That is, i is an embedding. Furthermore, by Tychonoff's theorem the product is compact, as each factor of the product is compact. Hence the closure $\overline{i(X)}$, which we shall denote by $\overline{X}^V$, is compact. The pair $(i, \overline{X}^V)$ is then a compactification of X, which tracks the information given by the visual boundary at each point. That is, a sequence in X converges in the topology of $\overline{X}^V$ if and only if it converges for every possible visual compactification $\overline{X}^v_b$. The main interest of this compactification comes from the following result.

Theorem 3.12 *Let X be a proper, uniquely geodesic, straight metric space which is C^1 along geodesics and has constant distance variation. Then $(i, \overline{X}^V)$ is isomorphic to $(h, \overline{X}^h)$.*

Denote by Π_b the continuous map from $\overline{X}^h$ to $\overline{X}^v_b$ given by Theorem 1.1. The isomorphism between $\overline{X}^h$ and $\overline{X}^V$ is defined by recording the value of each possible Π_b within $\overline{X}^V$. That is, we define $\widetilde{\Pi}\colon \overline{X}^h \to \overline{X}^V$ in such a way that $\pi_b \circ \widetilde{\Pi} := \Pi_b$ for each $b \in X$. The only property required to prove that $\widetilde{\Pi}$ is an isomorphism not following directly from previous results is the injectivity. By Proposition 3.4 we know that if $f \in \Pi_b^{-1}(\gamma)$ then γ is an optimal geodesic of f. That is, $f(\gamma(t)) - f(\gamma(s)) = -(t-s)$. Hence, if $f, g \in \Pi_b^{-1}(\gamma)$, then they differ by a constant along the geodesic γ. If f and g are horofunctions in the preimage of a point by $\widetilde{\Pi}$, then they differ by a constant along infinitely many geodesics, which cover X. However, the constant might depend on the geodesic, so we need a way to connect these constants. We proceed by strengthening Proposition 3.4 to show that any two functions in $\Pi_b^{-1}(\gamma)$ also have the same directional derivatives at points in γ, which allows us to connect the geodesics. Precisely, we prove the following.

Proposition 3.13 *Let X be a proper, uniquely geodesic, straight metric space which is C^1 along geodesics and has constant distance variation. Furthermore, let γ be a geodesic ray starting at b, and let α be a geodesic starting at some point on γ. Then, $\frac{d}{dt} f \circ \alpha(t)\big|_{t=0}$ exists and its value is the same for all $f \in \Pi_b^{-1}(\gamma)$.*

Proof For any $b' \in \gamma$ we have that γ is an optimal geodesic of f passing through b'. Denoting $\gamma_{b'}$ the geodesic ray starting at b' with the same biinfinite extension as γ we have that $f \in \Pi_{b'}^{-1}(\gamma_{b'})$, by Proposition 3.4. Hence, we can assume that $\alpha(0) = b$ by changing the basepoint if necessary. Let x_n be a sequence converging to f. Furthermore, let η_t^n be the geodesic from $\alpha(t)$ to x_n and $g_n(t)$ be the value of $\frac{d}{ds} h(x_n) \circ \alpha(s)\big|_{s=t}$. By the definition of the map h we have $g_n(t) = \frac{d}{ds} d(\alpha(s), x_n)\big|_{s=t}$. By the constant distant variation we have $g_n(t) = \frac{d}{ds} d(\alpha(s), \eta_t^n(1))\big|_{s=t}$, which since X is C^1 along geodesics depends continuously on $\eta_t^n(1)$ and t.

By Proposition 3.5 the geodesics η_t^n converge as $n \to \infty$ to some geodesics η_t, so $\eta_t^n(1)$ converges to $\eta_t(1)$. Since the space is C^1 along geodesics, the value of $\frac{d}{ds} d(\alpha(s), \eta_t^n(1))\big|_{s=t}$ depends continuously on $\eta_t^n(1)$, and so g_n converges pointwise to $g(t) = \frac{d}{ds} d(\alpha(s), \eta_t(1))\big|_{s=t}$.

Take some $\delta > 0$ and assume the convergence is not uniform on $[-\delta, \delta]$. Then there is some $\varepsilon > 0$ such that for each n there is at least one $t_n \in [-\delta, \delta]$ such that $|g_n(t_n) - g(t_n)| > \varepsilon$. Since $[-\delta, \delta]$ is compact we can take a converging subsequence such that t_n converges to some $T \in [-\delta, \delta]$. Hence, the point $\eta_{t_n}^n(1)$ does not converge to $\eta_T(1)$, so by properness of X we can take a subsequence such that $\eta_{t_n}^n(1)$ converges to some $p \in X$ different from $\eta_T(1)$. Let β be the geodesic starting at $\alpha(T)$ passing through p. The geodesics $\eta_{t_n}^n$ converge uniformly to β, and $\beta \neq \eta_T$. For any fixed $t > 0$ we have, following the same reasoning than in the proof of Proposition 3.5,

$$f(\beta(t)) - f(\beta(0)) = \lim_{n\to\infty} d(x_n, \beta(t)) - d(x_n, \beta(0)) = -t.$$

Hence, β is an optimal geodesic of f passing through $\alpha(T)$. However, $f \in \Pi_{\alpha(T)}^{-1}(\eta_T)$, so η_T is also an optimal geodesic passing through $\alpha(T)$, contradicting Proposition 3.4.

Hence, the convergence of $(h(x_n)\circ\alpha)' = g_n$ to g is uniform on $[-\delta,\delta]$. Therefore, f is differentiable and $f'(0) = g(0) = \frac{d}{ds}d(\alpha(s),\gamma(1))\big|_{s=0}$, which is the same for all $f \in \Pi^{-1}(\gamma)$. □

Proof of Theorem 3.12 Each Π_b is continuous, so by the definition of the product topology the map $\widetilde{\Pi}$ is continuous. Hence, by Lemma 2.1 to see that $\widetilde{\Pi}$ is an isomorphism it is enough to show that $\widetilde{\Pi}$ is injective.

Let $f, g \in \overline{X}^h$ be such that $\widetilde{\Pi}(f) = \widetilde{\Pi}(g)$. If there is some $b \in X$ such that $\pi_b \circ \widetilde{\Pi}(f) \in X$ then $f = h(\pi_b \circ \widetilde{\Pi}(f)) = g$. Assume then $\pi_b \circ \widetilde{\Pi}(f) \in \partial\overline{X}^v_b$ for all $b \in X$. By Proposition 3.13 they have the same directional derivatives at every point. Let α be a geodesic from a fixed basepoint b to any other point. We have $(f\circ\alpha)' = (g\circ\alpha)'$, so $f - g$ is constant along α, and hence everywhere, since any point can be connected to b by a geodesic. Hence, f and g are the same horofunctions. □

By the definition of the convergence in the product topology, this characterization gives us the following equivalence for the convergence to points in the horoboundary.

Corollary 3.14 *Let X be a proper, uniquely geodesic, straight metric space, C^1 along geodesics and with constant distance variation. A sequence $(x_n) \subset X$ converges in the horofunction compactification if and only if the sequence converges in all the visual compactifications.*

Restricting the result to the Teichmüller metric we get Corollary 1.11 announced in the introduction.

4 Background on Teichmüller spaces

A *surface with marked points* S is a pair (Σ, P), where Σ is a compact, orientable surface with possibly empty boundary, and $P \subset \Sigma$ is a finite, possibly empty, set of points, where we allow points to be on the boundary. The *Teichmüller space* $\mathcal{T}(S)$ is the set of equivalence classes of pairs (X, f), where X is a Riemann surface and $f\colon \Sigma \to X$ is an orientation-preserving homeomorphism. Two pairs (X, f) and (Y, g) are equivalent if there is a conformal diffeomorphism $h\colon X \to Y$ such that $g^{-1}\circ h\circ f$ is isotopic to identity rel P.

The *Teichmüller distance* between two points $[(X, f)], [(Y, g)] \in \mathcal{T}(S)$ is defined as the value $\frac{1}{2}\log\inf K$, where the infimum is taken over all $K \geq 1$ such that there exists a K–quasiconformal homeomorphism $h\colon X \to Y$ with $g^{-1}\circ h\circ f$ isotopic to identity rel P. Together with the smooth structure provided by the Fenchel–Nielsen coordinates $\mathcal{T}(S)$ satisfies all the metric properties discussed in the previous section. That is, $\mathcal{T}(S)$ with the Teichmüller distance is a proper, uniquely geodesic and straight metric space which is C^1 along geodesics and has constant distance variation. See [11, Part 2] for some background on the Teichmüller metric and the Fenchel–Nielsen coordinates.

A *quadratic differential* on a Riemann surface X is a map $q\colon TX \to \mathbb{C}$ such that $q(\lambda v) = \lambda^2 q(v)$ for every $\lambda \in \mathbb{C}$ and $v \in TX$. Considering only holomorphic quadratic differentials with finite area $\int_X |q|$ we get a

characterization of the cotangent space to the Teichmüller space based at $[(X, f)]$. Given a point $p \in \mathcal{T}(S)$ and a quadratic differential $q \in T_p^*\mathcal{T}(S)$ there is a unique geodesic γ such that $\gamma(0) = p$ and $\gamma'(0) = |q|/q$. We shall denote such a geodesic as $R(q; \cdot)$ and denote the associated Busemann points as $B(q)$ or B_q.

4.1 Measured foliations

A *multicurve* on S is an embedded 1–dimensional submanifold of $\Sigma \setminus P$ with boundary in $\partial\Sigma \setminus P$ such that

- no circle component bounds a disk with at most 1 marked point;
- no arc component bounds a disk with no interior marked points and at most 1 marked point on $\partial\Sigma$; and
- no two components are isotopic to each other in Σ rel P.

Each of the components is called *curve.* A *weighted multicurve* is a multicurve together with a positive weight associated to each curve. We shall consider (weighted) multicurves up to isotopy rel P. If a simple curve is a circle we shall call it a *closed curve*, and a *proper arc* otherwise.

A *measured foliation* on S is a foliation with isolated prong singularities, where we allow 1–prong singularities at marked points, equipped with an invariant transverse measure μ_F [12, exposé 5]. Denoting α_i and w_i the components and the weights of α respectively, the intersection number $i(\alpha, F)$ is defined as $\inf \sum_i w_i \int_{\alpha_i} |\mu_F| \, d\alpha_i$, where the infimum is taken over all representatives of α. Two measured foliations F and G are *equivalent* if $i(\alpha, F) = i(\alpha, G)$ for every multicurve α. We shall always consider measured foliations up to this equivalence relation. The set of measured foliations is usually denoted as $\mathcal{MF}$, and its topology is defined in such a way that a sequence $(F_n) \subset \mathcal{MF}$ converges to F if and only if $i(\alpha, F_n)$ converges to $i(\alpha, F)$ for every multicurve α.

Given a quadratic differential one can define the *vertical foliation* as the union of *vertical trajectories*, that is, maximal smooth paths γ such that $q(\gamma'(t)) < 0$ for every t in the interior of the domain. This foliation can be equipped with the transverse measure given by $|\operatorname{Re} \sqrt{q}|$. This measured foliation is called the *vertical measured foliation* of q, and shall be denoted as $V(q)$. This map is actually a homeomorphism. As such, given a measured foliation F and a complex structure X, there is a unique quadratic differential $q_{F,X}$ on X such that $V(q_{F,X}) = F$. We call this quadratic differential the *Hubbard–Masur* differential associated to F on X [19]. Furthermore, for each $\lambda > 0$ we have $q_{\lambda F,X} = \lambda q_{F,X}$. Similarly, the *horizontal foliation* $H(q)$ can be defined as the union of maximal smooth paths γ such that $q(\gamma'(t)) > 0$, with the transverse measure $|\operatorname{Im} \sqrt{q}|$.

It is possible to associate a measured foliation to each weighted multicurve by thickening each proper arc and closed curve to a rectangle or cylinder respectively with width equal to the weight of the curve, and then collapsing the rest of the surface. The intersection numbers are maintained by this construction. This association is injective, and hence we shall consider the set of weighted multicurves as a subset of the measured foliations, and use both expressions of weighted multicurve indistinctly.

By removing the critical graph, a measured foliation is decomposed into a finite number of connected components, each of which is either a thickened curve, or a minimal component which does not intersect the boundary, in which every leaf is dense [40, Chapter 24.3]. Each transverse measure within the minimal components can be further decomposed into a sum of finitely many projectively distinct ergodic measures. A foliation F' is an indecomposable component of F if it is either a thickened curve or a minimal component with a transverse measure that cannot be decomposed as a sum of more than one projectively distinct ergodic measure. Every foliation can be decomposed uniquely into a union of indecomposable foliations. For a surface of genus g with no boundaries nor marked points Papadopoulos shows [37] that the maximum number of indecomposable components for any foliation is $3g-3$. It is possible to get an upper bound for foliations on surfaces with boundary and marked points by swapping the marked points for boundaries and using the doubling trick we will explain in Section 4.3.

It was shown by Thurston that for surfaces without boundary it is possible to achieve a dense subset by restricting to simple closed curves, see Fathi, Laudenbach, and Poénaru [12] for a reference. When there are boundaries the picture gets slightly more complicated, but it has been shown by Kahn, Pilgrim and Thurston [21, Proposition 2.12] that multicurves can be seen as a dense subset. More precisely, they show the following.

Proposition 4.1 (Kahn–Pilgrim–Thurston) *Let F be a measured foliation in S not containing proper arcs. Then there exists a sequence of multicurves composed solely of closed curves approaching F.*

The result can be extended to any foliation by cutting along the proper arcs and approaching the foliation in the resulting surfaces by multicurves. Then, joining the multicurves from the proposition with the proper arcs and the adequate weights we get a sequence of multicurves converging to our original foliation.

4.2 Extremal length

Given a marked conformal structure on S, that is, a point $X \in \mathcal{T}$, the *extremal length* of F on X is defined as

$$\mathrm{Ext}_X(F) := \int_X |q_{F,X}|.$$

The map $\mathrm{Ext}\colon \mathcal{MF}(S) \times \mathcal{T}(S) \to \mathbb{R}$ is continuous and homogeneous of degree 2 in the first variable.

Given two points $x, y \in \mathcal{T}(S)$ we can define the function

$$K_{x,y} := \sup_{F \in P_b} \frac{\mathrm{Ext}_x(F)}{\mathrm{Ext}_y(F)},$$

where P_b is the set of measured foliations F satisfying $\mathrm{Ext}_b(F) = 1$. As revealed by Kerckhoff's formula [23], the value $\frac{1}{2}\log K_{x,y}$ coincides with the usual definition of the Teichmüller distance $d(x, y)$.

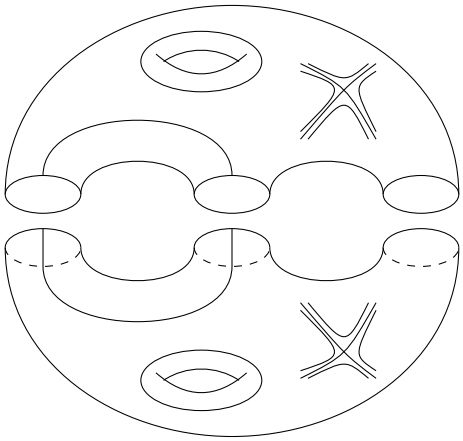

Figure 4: Visual representation of the doubling trick.

4.3 The doubling trick

Let X be a Riemann surface with nonempty boundary. Denote by $\overline{X}$ the mirror surface, obtained by composing each atlas of X with the complex conjugation. Gluing X to $\overline{X}$ along the corresponding boundary components we obtain the *conformal double* $X^d = X \cup \overline{X}/\sim$ of X. Note that X^d has empty boundary. See Figure 4. Given a foliation F or a quadratic differential q on X, we can repeat the same process, obtaining the corresponding conformal doubles F^d and q^d on X^d. For a more detailed treatment of this argument see [1, Section II.1.5].

The main interest of the conformal doubles is that these are surfaces without boundary, so most of the results relating to Teichmüller theory of surfaces without boundary can be translated to surfaces with boundary. We have the following.

Proposition 4.2 *Let X be a Riemann surface with boundary, and F be a foliation on X. Then,*

$$\operatorname{Ext}_{X^d}(F^d) = 2\operatorname{Ext}_X(F).$$

Proof We have $q_{F^d,X^d} = q^d_{F,X}$, so the result follows, as $\int_{X^d} |q^d_{F,X}| = 2\int_X |q_{F^d,X^d}|$. □

4.4 The Gardiner–Masur compactification

For a surface S with marked points and empty boundary we can embed $\mathcal{T}(S)$ into the space of continuous functions from the set $\mathcal{S}$ of simple closed curves on S to $\mathbb{R}$ via the map $\phi\colon \mathcal{T}(S) \to P(\mathbb{R}^{\mathcal{S}})$ defined by

$$\phi(X) = [\operatorname{Ext}_X(\alpha)^{1/2}]_{\alpha\in\mathcal{S}},$$

where the square brackets indicate a projective vector. Gardiner and Masur show [14] that this map is indeed an embedding, and that $\phi(\mathcal{T}(S))$ is precompact. The Gardiner–Masur compactification of a surface without boundary is then defined as the pair $\big(\phi, \overline{\phi(\mathcal{T}(S))}\big)$.

Alternatively, after choosing a basepoint $b \in \mathcal{T}(S)$, it is also possible to consider the map

$$\mathcal{E}\colon \mathcal{T}(S) \to C(\mathcal{MF}) \quad \text{defined by} \quad \mathcal{E}(X)(\cdot) := \left(\frac{\mathrm{Ext}_X(\cdot)}{K_{b,X}}\right)^{1/2},$$

This map is quite similar to the original map ϕ, the differences being that $\mathcal{E}$ considers all measured foliations instead of just the closed curves, and normalizes instead of projectivizing. Walsh proves [44] that, for surfaces without boundary, the map $\mathcal{E}$ defines a compactification in the same way that ϕ does, and in fact this compactification is isomorphic to the one defined by ϕ.

The compactification defined by $\mathcal{E}$ fits better our goal, so we shall define the Gardiner–Masur compactification of Teichmüller spaces of surfaces with boundary as the one obtained by using $\mathcal{E}$. With this in mind, we first need the following result.

Proposition 4.3 *Let S be a compact surface with possibly boundary and marked points. Then the map $\mathcal{E}\colon \mathcal{T}(S) \to C(\mathcal{MF})$ is injective.*

Proof Assume we have $x, y \in \mathcal{T}(S)$ with $\mathcal{E}(x)(F) = \mathcal{E}(y)(F)$ for all $F \in \mathcal{MF}$. Then,

$$K_{x,y} = \sup_{F \in P_b} \frac{\mathrm{Ext}_x(F)}{\mathrm{Ext}_y(F)} = \frac{K_{b,x}}{K_{b,y}} \quad \text{and} \quad K_{y,x} = \sup_{F \in P_b} \frac{\mathrm{Ext}_y(F)}{\mathrm{Ext}_x(F)} = \frac{K_{b,y}}{K_{b,x}} = K_{x,y}^{-1}.$$

However, $K_{y,x} = K_{x,y}$, since the Teichmüller distance is symmetric. Hence, $K_{x,y} = 1$ and, by Kerckhoff's formula, $d(x, y) = \frac{1}{2} \log K_{x,y} = 0$. □

Miyachi shows [32] that the set $E(S) := \{\mathcal{E}(X) \mid X \in \mathcal{T}(S)\}$ is precompact when S is a surface without boundary. Given a surface with boundary S, denote by $\mathcal{MF}^d(S)$ the set of measured foliations on S^d obtained by doubling the foliations $\mathcal{MF}(S)$. The set $E(S^d)|_{\mathcal{MF}^d(S)} = \{\mathcal{E}(X)|_{\mathcal{MF}^d(S)} \mid X \in \mathcal{T}(S^d)\}$, obtained by restricting the functions in $E(S^d)$ to $\mathcal{MF}^d$, is precompact. Furthermore, we can embed $E(S)$ into $E(S^d)|_{\mathcal{MF}^d(S)}$ by sending $f \in E(S)$ to $f^d \in E(S^d)|_{\mathcal{MF}^d(S)}$ defined by $f^d(F^d) = f(F)$. Hence, $E(S)$ is precompact.

We define the *Gardiner–Masur compactification* for a surface with boundary as the closure $\overline{E}$ of $E(S)$, together with the map $\mathcal{E}$. We shall be using the same characterization for surfaces without boundary.

One of the relevant features of the Gardiner–Masur compactification is that it coincides with the horofunction compactification. Indeed, Liu and Su [28] and Walsh [44] prove that for surfaces without boundary these two compactifications are isomorphic. In the following, we shall extend the relevant results to surfaces with boundary. We begin with the driving theorem from Walsh's paper.

Theorem 4.4 (extension of [44, Theorem 1] to surfaces with boundary) *Let $R(q;\cdot)\colon \mathbb{R}_+ \to \mathcal{T}(S)$ be the Teichmüller ray with initial unit-area quadratic differential q, and let F be a measured foliation. Then,*

$$\lim_{t\to\infty} e^{-2t}\,\mathrm{Ext}_{R(q;t)}(F) = \sum_j \frac{i(G_j, F)^2}{i(G_j, H(q))},$$

where the $\{G_j\}$ are the indecomposable components of the vertical foliation $V(q)$, and $H(q)$ is the horizontal foliation.

Proof If S does not have boundary the result follows from Walsh's paper. Assume then that S has boundary. Let p be the number of proper arcs of $V(q)$, and reorder the components so G_j is a proper arc for $j \le p$. The conformal double G_j^d is indecomposable whenever G_j is a proper arc, and decomposes into two components otherwise, as it is not incident to the boundary of S. Denote G_j^1 and G_j^2 the two components of G_j for $j > p$. We have

$$2 \lim_{t\to\infty} e^{-2t} \operatorname{Ext}_{R(q;t)}(F) = \lim_{t\to\infty} e^{-2t} \operatorname{Ext}_{R(q^d;t)}(F^d) = \sum_{j\le p} \frac{i(G_j^d, F^d)^2}{i(G_j^d, H(q)^d)} + \sum_{i\in\{1,2\}} \sum_{j>p} \frac{i(G_j^i, F^d)^2}{i(G_j^i, H(q)^d)}.$$

For foliations $G, F \in \mathcal{MF}(S)$ we have $i(G^d, F^d) = 2i(G, F)$. Hence, $i(G_j^d, F^d) = 2i(G_j, F)$. Using the symmetry, $i(G_j^1, F^d) = i(G_j^2, F^d)$, so for $j > p$ we have $i(G_j^1, F^d) = i(G_j, F)$. Using these identities we get the result. □

Following the same reasoning we can extend as well the next result.

Lemma 4.5 (extension of [44, Lemma 3] to surfaces with boundary) *Let q be a unit area quadratic differential. Then,*

$$e^{-2t} \operatorname{Ext}_{R(q;t)}(F) \ge \sum_j \frac{i(G_j, F)^2}{i(G_j, H(q))},$$

where $t \in \mathbb{R}_+$ and $\{G_j\}$ are the indecomposable components of the vertical foliation $V(q)$.

Most of the results in Walsh's paper use the previous theorem. In particular, we have the following.

Corollary 4.6 (extension of [44, Corollary 1] to surfaces with boundary) *Let q be a quadratic differential and denote by G_j the indecomposable components of its vertical foliation. Then, the Teichmüller ray $R(q;\cdot)$ converges in the Gardiner–Masur compactification to*

$$\left(\sum_j \frac{i(G_j, \cdot)^2}{i(G_j, H(q))}\right)^{1/2}.$$

The relation between the Gardiner–Masur compactification and the horoboundary compactification is given by the map $\Xi\colon \overline{E} \to \overline{\mathcal{T}(S)}^h$ defined by

$$\Xi(f)(x) := \tfrac{1}{2} \log \sup_{F\in\mathcal{P}} \frac{f(F)^2}{\operatorname{Ext}_x(F)}.$$

The following result can be extended to surfaces with boundary by repeating the proof found in Walsh's paper in this context.

Theorem 4.7 (extension of [44, Lemma 21] to surfaces with boundary) *The map Ξ is an isomorphism between the compactifications $(\mathcal{E}, \overline{E})$ and $(h, \overline{\mathcal{T}(S)}^h)$.*

Directly from the definition of Ξ we have the following.

Corollary 4.8 *Let $f, g \in \overline{E}$. If $f \ge g$, then $\Xi(f) \ge \Xi(g)$.*

We shall denote the representation of the Busemann point $B(q)$ in the Gardiner–Masur compactification as $\mathcal{E}(q)$. By Corollary 4.6 we have an explicit representation of $\mathcal{E}(q)$. As we have seen in Propositions 1.3 and 1.4, the continuity of the Busemann map has some interesting implications, and it is enough to look for continuity of the map restricted to the boundary. Related to this question we have the following result, which can also be derived by the same proof found in Walsh's paper, applied to this context.

Theorem 4.9 (extension of [44, Theorem 10] to surfaces with boundary) *Let (q_n) be a sequence of quadratic differentials based at $b \in \mathcal{T}(S)$. Then $B(q_n)$ converges to $B(q)$ if and only if both of the following hold*:

(1) *(q_n) converges to q.*

(2) *For every subsequence $(G^n)_n$ of indecomposable elements of $\mathcal{MF}$ such that, for each $n \in \mathbb{N}$, G^n is a component of $V(q_n)$, we have that every limit point of G^n is indecomposable.*

In view of this theorem, we say that a sequence of quadratic differentials (q_n) converges *strongly* to q if it does so in the sense described by the theorem.

Finally, while the following result may be extendable to surfaces with boundary, we only use it in the context of surfaces without boundary, so we shall not be working on finding an extension.

Theorem 4.10 [44, Theorem 3] *For the Teichmüller space of a surface without boundary with the Teichmüller metric, for any basepoint $X \in \mathcal{T}(S)$, all Busemann points can be expressed as $B(q)$ for some quadratic differential q based at X.*

5 Horoboundary convergence for Teichmüller spaces

5.1 Continuity of the Busemann map

We begin by using Proposition 1.4 to determine when the Busemann map is continuous. Recall that a sequence (q_n) converges to q strongly if and only if the sequence satisfies the conditions of Theorem 4.9. That is, a sequence (q_n) converges to q strongly if and only if the associated Busemann points $B(q_n)$ converge to $B(q)$. With this in mind we introduce the following notion.

Definition 5.1 Let q be a quadratic differential. We say that q is *infusible* if any sequence of quadratic differentials converging to q converges strongly. We say that q is *fusible* if it is not infusible.

In other words, we say that q is fusible when it can be approached by a sequence of quadratic differentials (q_n) such that there is some sequence (G^n) of measured foliations with each G^n being an indecomposable component of $V(q_n)$, with (G^n) having at least one decomposable accumulation point. The following statement follows directly from this definition, Proposition 1.4 and Walsh's result.

Proposition 5.2 *Let q be a unit area quadratic differential. The Busemann map B is continuous at q if and only if q is infusible.*

Proof If q is fusible then we have a sequence converging to q but not strongly. Hence, by Theorem 4.9 the sequence $(B(q_n))$ does not converge to $B(q)$, and so the Busemann map is not continuous at q.

If q is infusible we have that any sequence (q_n) converging to q does so strongly, and so we have that $B(q_n)$ converges to $B(q)$, so B is continuous at q when restricted to the boundary. By Proposition 1.4 this implies that B is continuous at q. □

We shall now find a criterion on the vertical foliation to determine when a unit area quadratic differential is infusible.

Definition 5.3 Let F be a measured foliation on a surface S and let G be one of its indecomposable components. We say that G is a *boundary annulus* if it is an annulus parallel to a boundary with no marked points, and a *boundary component* if it is a boundary annulus or a proper arc. If G is not a boundary component, we shall call it an *interior component*. Each of the connected components of the surface obtained after removing the proper arcs shall be called *interior part*. If each of these interior parts has at most one interior component, then we say that F is *internally indecomposable*. If F is not internally indecomposable we say that it is *internally decomposable*.

For surfaces without boundary, a foliation F is internally indecomposable if and only if it is indecomposable, as we do not have boundary components. Given these definitions we can state our main result of this section.

Theorem 5.4 *Let q be a quadratic differential. Then q is infusible if and only if its vertical foliation $V(q)$ is internally indecomposable.*

This result is somewhat straightforward whenever S does not have boundary, as in order to have a sequence (q_n) that converges to q but not strongly we need a sequence of components of $V(q_n)$ converging to a decomposable component of $V(q)$, but if S is closed and $V(q)$ is internally indecomposable, then $V(q)$ only has one indecomposable component. Conversely, if $V(q)$ has more than one indecomposable component, as S does not have boundary $V(q)$ can be approached by a sequence of simple closed curves, so the associated sequence of quadratic differentials converges to q but not strongly.

For surfaces with boundary the proof is more involved, as simple closed curves are no longer dense. However, the density of multicurves from Proposition 4.1 allows us to follow a slightly similar reasoning. We begin by proving some results regarding the shape that foliations have to take when approaching a foliation with boundary components, namely, boundary components have to be eventually included in the approaching foliations.

Proposition 5.5 *Let (F_n) be a sequence of measured foliations converging to a measured foliation F, let G be the union of the boundary components of F and let H be such that $F = H + G$. Then, for n big enough, $F_n = H_n + a_n G$, with a_n converging to 1 and H_n converging to H.*

In particular, the proper arcs of the limiting foliation have to be included in the approaching foliations. Hence, we will be able to separate the surface along these proper arcs into the interior parts of the limiting foliation, and study the convergence in each of these parts.

We say that a subset of a boundary component is a *boundary arc* if it is homeomorphic to an open interval or a circle, does not contain marked points and, if it is homeomorphic to an open interval, it is delimited by marked points.

Repeating the argument by Chen, Chernov, Flores, Fortier Bourque, Lee, and Yang [7] to a more general setting we get the following characterization of foliations on simple surfaces, which we shall use to solve the simpler cases.

Lemma 5.6 *Let S be a sphere with one boundary component possibly containing boundary marked points and one interior marked point. Then every indecomposable foliation on S is a proper arc and there are finitely many distinct proper arcs.*

Proof Assuming that there is some foliation F with a recurrent leaf to some part of S we get a contradiction, as explained in the proof of [7, Lemma 4.1]. Hence, each indecomposable foliation is a curve. Any closed curve in S is contractible to the marked point. Hence, a each indecomposable foliation is a proper arc.

A proper arc in S must have two endpoints, which must be contained in the boundary arcs in the boundary component of S. Denote b_1 and b_2 these two boundary arcs, which might be the same. We aim to show that there are at most two classes of arcs with endpoints in b_1 and b_2. Fix three proper arcs with endpoints on b_1 and b_2. Any intersection between these arcs can be removed by doing isotopies moving the endpoints along the arcs b_1 or b_2. Hence, these arcs can be isotoped to not intersect each other. Since there is only one interior marked point, two of these arcs delimit a rectangle with no marked interior marked points, so are isotopic. Hence, there are at most two different proper arcs between b_1 and b_2. There are finitely marked points in the boundary component, so there are finitely many boundary arcs. Therefore, there are finitely many pairs of boundary arcs, and since we have at most two proper arcs per pair, there are also finitely many different proper arcs. □

We shall first see the proposition for the case where G contains a proper arc and we are approaching with a sequence of indecomposable foliations.

Lemma 5.7 *Let S be a surface and let (F_n) be a sequence of indecomposable foliations on S converging to a measured foliation G. Then G is either a multiple of a proper arc γ, in which case F_n is also a multiple of γ for n big enough, or G does not contain a proper arc.*

Proof Assume G contains a proper arc γ with weight $w > 0$ and denote by b one of the boundary arcs where γ is incident.

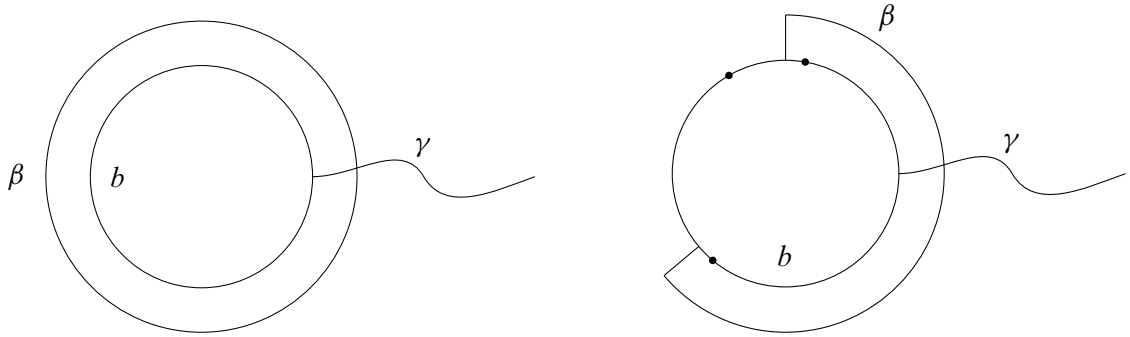

Figure 5: Sample curves used in the proof of Lemma 5.7.

Our first step is seeing that, for n big enough, F_n intersects b. We shall do this by finding different test curves β depending on the shape of b. If the boundary component containing b has at most one marked point, we consider β to be a curve parallel to that boundary component as in Figure 5, left. Otherwise we consider β to be the curve defined by taking a small arc starting at the boundary arc next to b, concatenating with a curve parallel to b, and concatenating another segment with endpoint in the boundary arc after b, as shown in Figure 5, right.

If the curve β is contractible then S is a sphere with one boundary component and at most one interior marked point, so by Lemma 5.6 the result follows. Assume then that β is not contractible. We have $i(\gamma, \beta) > 0$, so $i(G, \beta) > 0$ and hence $i(F_n, \beta) > 0$ for n big enough, which implies that F_n intersects b. Hence, since F_n is indecomposable, it is a weighted proper arc, which we denote by $w_n\gamma_n$, where $w_n > 0$ is the weight at γ_n is a proper arc.

Denote b_1 and b_2 the boundary arcs where γ has its endpoints, and denote by β_1 and β_2 the associated test curves shown in Figure 5. If both endpoints are in the same boundary arc we set b_2 and β_2 as null curves. We shall now find a multicurve A surrounding γ, b_1 and b_2 such that any leaf of G intersecting A but not γ has an endpoint in either b_1 or b_2. The multicurve A is chosen so that, together with the boundaries where γ has its endpoints, delimits the smallest surface containing γ. The precise shape of A depends on whether the endpoints of γ are in the same boundary component or not, and the distribution of marked points in these boundaries.

If both endpoints of γ are in different boundary components we proceed differently according to the distribution of marked points at these boundaries. If each of the boundaries contains at most one marked point then we define A as the curve shown in Figure 6, left. If one of the boundary components has two or more marked points, but the other has at most one marked point we define A as the arc shown in Figure 6, middle. Finally, if each of the boundaries contains at least two marked points we define A as the multicurve formed by the curves A_1 and A_2 as shown in Figure 6, right.

If both endpoints γ are in the same boundary we also proceed differently according to the distribution of marked points. In all cases A is defined as a multicurve formed by two curves. If each possible segment

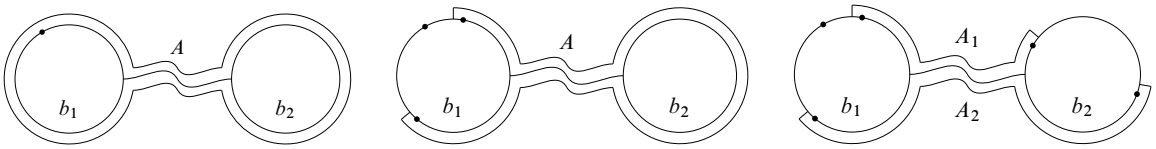

Figure 6: Construction of the curves A_1 and A_2 whenever γ has endpoints in different boundary components in the proof of Lemma 5.7.

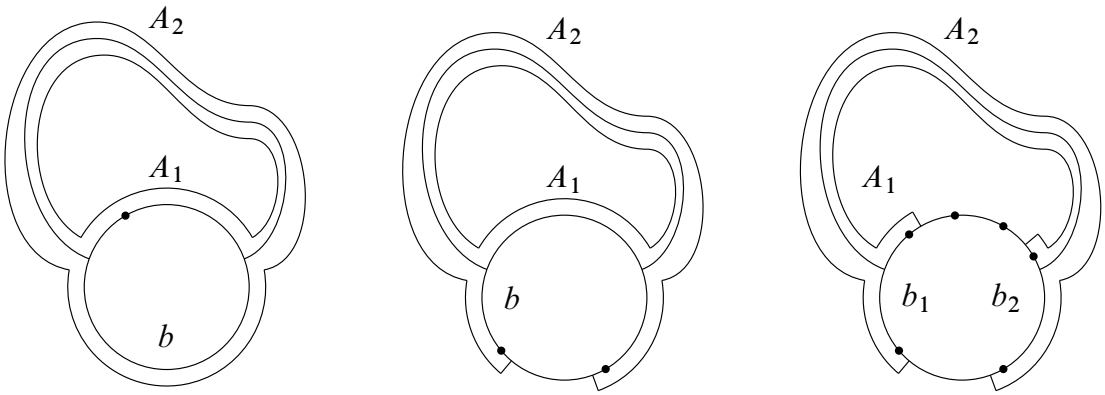

Figure 7: Construction of the curves A_1 and A_2 whenever γ has endpoints in the same boundary component in the proof of Lemma 5.7.

within the boundary component joining the two endpoints has at most one marked points we proceed as in Figure 7, left. If one of these segments has two or more marked points, while the other has at most one we proceed as in Figure 7, middle. Finally, if both of these segments have two or more marked points we proceed as in Figure 7, right.

In any of the cases above if a component of A is nonessential we remove it from A. The following argument also applies whenever A is a null curve. Put A and G in minimal position and denote by P the surface containing γ, delimited by A and the boundary components where γ has its endpoints. Let α be a connected component of a noncritical leaf of G restricted to P intersecting A. Since G contains γ the proper arc α cannot intersect γ. Furthermore, by observing the possible configurations, if α has one endpoint in A_1, the other one cannot be in A_2, as whenever we have both A_1 and A_2, these are separated within P by the proper arc γ. Furthermore, if both endpoints are in A_1 then α can be isotoped to not intersect A. Therefore, the other endpoint of α is in either b_1 or b_2. Hence, $i(G,\beta_1)+i(G,\beta_2)\geq i(G,A)+w\,i(\gamma,\beta_1)+w\,i(\gamma,\beta_2)>i(G,A)$. Since $w_n\gamma_n$ converges to G, this last inequality implies that for n big enough,

$$i(\gamma_n,\beta_1)+i(\gamma_n,\beta_2)>i(\gamma_n,A).$$

Fix n such that γ_n satisfies the previous inequality. Assume γ_n has just one endpoint inside P. Then, $i(\gamma_n,\beta_1)+i(\gamma_n,\beta_2)=1$, so $i(\gamma_n,A)=0$ and γ_n cannot leave P. If γ_n has both endpoints in P then

$i(\gamma_n, \beta_1) + i(\gamma_n, \beta_2) = 2$. Furthermore, if γ_n leaves P, then it has to reenter at some point, resulting in $i(\gamma_n, A_1 + A_2) = 2$. Hence, γ_n stays inside P.

The weights w_n do not converge to 0, as $w_n i(\gamma_n, \beta)$ converges to $i(G, \beta)$, but $i(\gamma_n, \beta) \leq 2$. Since γ is contained in G we have $i(G, \gamma) = 0$. Therefore, for any $\varepsilon > 0$ and n big enough we have $w_n i(\gamma_n, \gamma) < \varepsilon$, so for n big enough $i(\gamma_n, \gamma) = 0$. Since γ_n does not intersect γ and stays inside P, γ_n can be isotoped to stay inside one of the components obtained after removing γ from P. Denote such a component by C. The component C has either one or two boundary components and no interior marked points or one boundary component and one interior marked point. By Lemma 5.6 the only case where we do not have finitely many different proper arcs is when C has two boundary components. However, in that case one of the boundary components is associated to a curve in A, so γ_n does not intersect it and that boundary can be treated as a marked point. Hence, in all cases there are finitely many possible proper arcs, and so γ_n is a multiple of γ for n big enough. □

When the boundary component is an annulus we have to be a bit more careful, so we start by proving it for approaching curves.

Lemma 5.8 *Let S be a surface and let $(w_n \gamma_n)$ be a sequence of weighted curves on S converging to a foliation G, where (w_n) are the weights and (γ_n) are the curves. Then G is either a multiple of a boundary annulus γ, in which case γ_n is γ for n big enough, or G does not contain a boundary annulus.*

Proof If S is a polygon with at most one interior marked point, then G cannot contain a boundary annulus. If S is a cylinder then, since we have a boundary annulus, at least one of the boundaries must not contain marked points. Hence, the number of curves is finite, as there is only one possible closed curve, and for counting the proper arcs we can consider the boundary without marked points as a marked point and apply Lemma 5.6. In that case, the conclusion follows.

Assume then that S is neither a disk with at most one interior marked point nor a cylinder with no interior marked points. Then there is a pair of pants P in S containing γ where each boundary component of P is either noncontractible or contractible to a marked point. Denote by B_1 the boundary component parallel to γ and B_2 and B_3 the other two boundary components of P. Furthermore, assume that G contains γ with weight w.

Begin by assuming that B_2 and B_3 are not contractible to marked points. Let C be the proper arc contained in P with both endpoints in B_1. Put B_2, B_3 and C in a minimal position with respect to G, and consider a connected component of a noncritical leaf of G intersecting C restricted to P. This noncritical leaf either is isotopic to γ, or to the curves F, E and D shown in Figure 8. Since the leaves of G do not intersect, there cannot be leaves isotopic to E and leaves isotopic to D at the same time. Breaking symmetry, assume there are no leaves isotopic to D. Then, $i(C, G) = i(C, \gamma) + i(B_3, G) > i(B_3, G) \geq i(B_2, G)$. Doing the same reasoning assuming that there are no leaves isotopic to E we get $i(C, G) > \max(i(B_2, G), i(B_3, G))$.

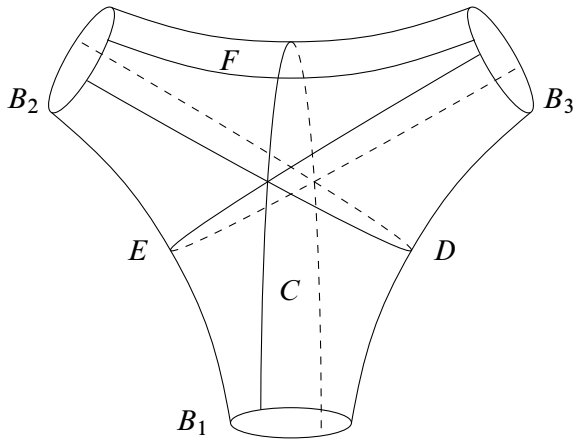

Figure 8: Curve labeling for the proof of Lemma 5.8.

Hence, since $w_n\gamma_n$ converges to G, γ_n has to satisfy

$$i(C,\gamma_n) > \max(i(B_2,\gamma_n), i(B_3,\gamma_n))$$

for n big enough.

For each n put B_3, B_2 and C in a minimal position with respect to γ_n, and consider the restriction of γ_n to P. Assume γ_n is not γ. Then, the curves on the restriction of γ_n to P intersecting C are isotopic to either E, F and D, but not γ. As before, this restriction cannot contain curves isotopic to E and curves isotopic to D for the same n, so assuming there are no curves isotopic to D we have $i(C,\gamma_n) = i(B_3,\gamma_n)$ which is a contradiction. Doing the same reasoning assuming that there are no curves isotopic to E also gives a contradiction. Hence, γ_n is γ for n big enough.

If B_2 or B_3 are contractible to marked points we have $i(G,B_2)$ or $i(G,B_3)$ is 0, and a similar reasoning yields the same result. □

Proof of Proposition 5.5 Let (F_n) be a sequence of measured foliations converging to F. As pointed out before, Proposition 4.1 can be extended to get sequences of weighted multicurves $(\gamma_n^m)_m$ converging to each F_n. Denote by $\gamma_{n,1}^m, \gamma_{n,2}^m, \dots, \gamma_{n,k(n,m)}^m$ the weighted curves of γ_n^m. For each n we take a subsequence such that $k(n,m)$ is constant with respect to m, and $\gamma_{n,i}^m$ converges for each i as $m\to\infty$. Denoting $F_{n,i}$ the limit of $\gamma_{n,i}^m$ as $m\to\infty$, we can write $F_n = \sum F_{n,i}$.

Denote by β_j the boundary components of F. That is, $\sum \beta_j = G$. Furthermore, denote by $b_{n,j}$ and $b_{n,j}^m$ the weights of β_j on F_n and γ_n^m, where we set the weight to be 0 if β_j is not contained in the foliation. It is clear that if $b_{n,j}=0$ then $b_{n,j}^m\to 0$, as we must have $b_{n,j}\geq \liminf_{m\to\infty} b_{n,j}^m$. If $b_{n,j}>0$ for some n, then $F_{n,i}$ contains β_j for some i. Hence, by Lemmas 5.7 and 5.8 we have $F_{n,i}$ and $\gamma_{n,i}^m$ are both multiples of β_j for m big enough. Then, since each of the multicurves in γ_n^m has to be different, β_j is not contained in any other foliation $F_{n,i}$ for that given n, so $F_{n_i} = b_{n,j}\beta_j$ and $\gamma_{n,i}^m$ can be written as $b_{n,i}^m\beta_j$ for m big enough, with $b_{n,i}^m$ converging to $b_{n,j}$ as $m\to\infty$.

Assume for some j we have $b_{n,j}$ not converging to 1. We can then take a subsequence such that $b_{n,j}$ converges to some $\lambda \neq 1$. Denote $\delta = |1-\lambda|/2$. For each n, there exists some $m_0(n)$ big enough so that $|1-b^m_{n,j}| > \delta$ for all $m \geq m_0(n)$. We can then take a diagonal sequence $\gamma_n^{m(n)}$ converging to F with $m(n) \geq m_0(n)$. However, following the previous reasoning we get that $\gamma_n^{m(n)}$ should contain β_j for n big enough, and the weight should converge to the weight in G, that is, to 1. However, $|1-b^{m(n)}_{n,j}| > \delta$, giving us a contradiction. Hence, $b_{n,j}$ converges to 1 for all j. Let then $a_n = \min_j(b_{n,j})$. Since $b_{n,j} \geq a_n$ we can define $H_n = F_n - a_n G$ and we have $F_n = H_n + a_n G$. Finally, $a_n \to 1$ as $n \to \infty$, so the proposition is proved. □

Proposition 5.9 *Let q be a unit area quadratic differential such that $V(q)$ is internally indecomposable. Then q is infusible.*

Proof Assume q is fusible, that is, we have a sequence of quadratic differentials (q_n) converging to q but not strongly. Let F_i^n be the indecomposable components of $V(q_n)$. To have nonstrong convergence we must have at least one sequence of indecomposable components converging to a decomposable component G, which we assume is $(F_1^n)_n$. Let β be a boundary component of $V(q)$. By Proposition 5.5 for n big enough a multiple of β must be contained in $V(q_n)$. Furthermore, β cannot be contained in G. Since G cannot contain boundary components, it must contain at least two interior components. On the other hand, since $V(q)$ is internally indecomposable, each interior part obtained by removing the proper arcs contains at most one interior component. Hence, for n big enough F_1^n must intersect at least two interior parts, that is, F_1^n must cross at least one proper arc. However, for each proper arc γ there is some n big enough such that γ is contained in the foliation $V(q_n)$, so F_1^n, a component of $V(q_n)$, intersects the foliation $V(q_n)$, giving us a contradiction. □

To prove the other direction we shall first see the following lemma.

Lemma 5.10 *Let S be a compact surface with possibly nonempty boundary and finitely many marked points, let $k \geq 2$ and let $\alpha = \{\alpha_1, \alpha_2, \dots, \alpha_k\}$ be a collection of nonintersecting closed curves on S. Furthermore, let p be the number of curves in α parallel to a boundary. Then there exists a collection of* $\max(\lceil (p/2) \rceil), 1)$ *nonintersecting curves intersecting each α_i.*

Our main interest in the lemma is that the amount of curves needed is strictly smaller than the amount of closed curves in α. This will allow us, by doing Dehn twists along the closed curves in α, to create a sequence of foliations converging to a foliation with strictly more components, which can be translated to a sequence of quadratic differentials that converge but not strongly. The proof of this lemma is based on a reasoning found in [11, Proposition 3.5].

Proof We start by replacing all boundaries of S without parallel curves in α by marked points. Let then α' be a completion of α to a pair of pants decomposition. Glue the remaining boundaries pairwise until we have at most one left. After cutting the surface along the closed curves that were not parallel to boundaries we get a collection of $\lceil p/2 \rceil$ tori with one boundary component and some spheres with b

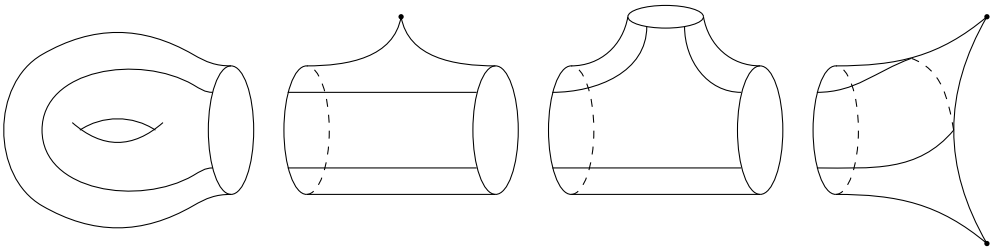

Figure 9: Laying out of curve segments for the proof of Lemma 5.10.

boundary components and n marked points, with $b+n=3$ and $b \geq 1$. If p is odd, one of these spheres has a boundary of S as a boundary. We join the boundaries of each of these surfaces with nonintersecting arcs, as shown in Figure 9, that is, in such a way that each boundary component has two arcs incident to it. We can then paste these surfaces back together in order to obtain a collection $\beta_1, \beta_2, \dots, \beta_l$ of pairwise disjoint curves in S. If p is odd this collection contains precisely one proper arc, as we only have two endpoints coming from the boundary we did not paste. If p is even the collection does not contain any proper arc. By the bigon criterion each β_j is in minimal position with respect to each α_i, and each α_i intersects either one or two of the β_j. Furthermore, since we did not cut along the original boundaries we pasted from S, each α_i parallel to a boundary of S intersects precisely one of the β_j. Suppose we have β_j and $\beta_{j'}$ intersecting a curve $\kappa \in \alpha'$ and that β_j and $\beta_{j'}$ are distinct. Since we have at most one proper arc, at least one of β_j and $\beta_{j'}$ is a closed curve. Hence, doing a half twist about κ, β_j and $\beta_{j'}$ become a single curve. Since this process does not create any bigons, the resulting collection is still in minimal position with α. Continuing this way we obtain a single curve γ intersecting each curve in κ. Furthermore, γ intersects each pasted boundary once. Cutting along the pasted boundaries, we get the curves from the lemma. If p is odd, β is a proper arc, so each cut along a pasted boundary increases the curve count by one, totaling $(p+1)/2$ curves. If p is even, β is a closed curve, so the first cut transforms it into a proper arc, and the following ones increase the curve count by one, giving a total of $\max(p/2, 1)$ curves. □

Proposition 5.11 *Let F be an internally decomposable measured foliation. Then, F can be approached by a sequence of weighted multicurves with fewer components than F.*

Proof By the extension to Proposition 4.1, we have a sequence of weighted multicurves γ^n converging to F, with the only proper arcs being the ones contained in F. Cutting the surface along the proper arcs of γ^n and quotienting these proper arcs to points we get k many surfaces $Z_1, Z_2, \dots, Z_k$ with boundary. Let γ_i^n be the restriction of γ^n to Z_i, and let F_i be the limit of γ_i^n. The foliation F is the union of the foliations F_i and the proper arcs.

Fix some i such that F_i is nonempty, and let $\alpha_1, \dots, \alpha_b$ be the closed curves parallel to the boundaries of Z_i. Let $a_1^n, \dots, a_b^n$ be the weights of $\alpha_1, \dots, \alpha_b$ in γ_i^n. We can take a subsequence such that a_j^n converges for each j to some a_j. If $a_j > 0$, the closed curve α_j is contained in F_i. If $a_j = 0$, then the weights a_j^n can be set to 0 on the multicurves γ_i^n while leaving the limit intact. Hence, we can assume

that $a_j^n = 0$ for all j such that $a_j = 0$. Let p and u be the number of closed curves with $a_j > 0$ parallel to boundaries with or without marked points respectively. Since we have removed all the closed curves with $a_j = 0$, the multicurve γ_i^n contains precisely p and u closed curves parallel to boundaries with or without marked points for n big enough. Denote by B the set of closed curves parallel to boundary components without marked points. Applying Lemma 5.10 to the multicurve γ_i^n minus B we get $\max(\lceil (p/2) \rceil), 1)$ curves β_i^n intersecting all closed curves in γ_i^n except the ones parallel to boundaries without marked points. Doing the appropriate Dehn twists along the closed curves of γ_i^n and rescaling to the curves β_i^n, and adding with the corresponding weights the curves in B, we get a sequence converging to γ_i^n with $\max(\lceil (p/2) \rceil), 1) + u$ many components. As such, taking a diagonal sequence we can get a sequence of multicurves converging to F_i with each multicurve containing $\max(\lceil (p/2) \rceil), 1) + u$ components.

Finally, since F is internally decomposable, there is at least one F_i with at least 2 interior components, so one of these multicurves has strictly less components than the limiting foliations, and we have nonstrong convergence. □

Theorem 5.4 follows by combining Propositions 5.9 and 5.11.

We do not need S to have a lot of topology to find internally decomposable foliations. In fact, determining which surfaces do not support internally decomposable foliations we get the following result.

Proposition 5.12 *Let $S_{g,b_m,b_u,p}$ be a surface of genus g with b_m and b_u boundaries with and without marked points respectively and p interior marked points. Then the Busemann map is continuous if and only if $3g + 2b_m + b_u + p \leq 4$.*

We shall split the proof in the following two lemmas

Lemma 5.13 *Let $S_{g,b_m,b_u,p}$ be a surface with $3g + 2b_m + b_u + p > 4$. Then it admits an internally decomposable foliation.*

Proof A multicurve consisting of two interior closed curves generates an internally decomposable foliation, so we just have to find such a pair for each possible surface satisfying the hypothesis. If S has genus at least 2 we can take a multicurve consisting of 2 nonseparating closed curves. If S is a torus with at least 2 boundaries or marked points, or a boundary with marked points, we can take a nonseparating closed curve and a separating closed curve around 2 boundaries or marked points, or around a boundary with marked points. If S is a sphere with at least 5 marked points or boundaries, we can take a closed curve around two interior points or boundaries, and a closed curve around two different interior points or boundaries. If S is a sphere with 1 boundary with marked points and at least 3 other boundaries or interior points we can take a closed curve around the boundary with marked points, and a closed curve around two other interior points or boundaries. Lastly, if S is a sphere with 2 boundaries with marked points and another interior marked point or boundary we take a closed curve around each boundary with marked points. □

Lemma 5.14 *Let $S_{g,b_m,b_u,p}$ be a surface with $3g + 2b_m + b_u + p \leq 4$. Then every foliation on S is internally indecomposable.*

Proof Assume we have an internally decomposable foliation on $S_{g,b_m,b_u,p}$. Then we can get an internally decomposable foliation on $S_{g,0,0,b_u+p+2b_m}$ by removing the boundary components, replacing the boundaries without marked points with marked points and each boundary with marked points for 2 marked points. Furthermore, if we have at least one marked point, we can get an internally decomposable foliation in $S_{g,0,0,b_u+p+2b_m+k}$, $k \in \mathbb{N}$, by replacing a marked point with a $k+1$ marked points.

Hence, we only need to prove that a torus with one marked point and a sphere with 4 marked points do not admit internally decomposable foliations. However, since these do not have boundaries, a foliation being internally decomposable translates to a foliation having at least two indecomposable components.

Assume the torus with one marked point admits a foliation with two indecomposable components. We can replace the marked point with a boundary, and add to the foliation a boundary component parallel to that boundary. Considering the doubled surface explained in Section 4.3 we get a closed surface of genus 2 without boundaries nor marked points, with at least 5 indecomposable components. Recall that the maximum number of indecomposable components for a foliation on a surface of genus g is $3g-3$, so for genus 2 the maximum is 3, giving us a contradiction. A similar process applies for the sphere with 4 marked points. □

Proof of Proposition 5.12 The Busemann map is continuous at every point in the interior of Teichmüller space, as it is the identity when restricted in there and $\partial\overline{X}^v$ is closed. Hence, we only need to prove continuity or discontinuity at the points on the boundary. By Lemma 5.13 if $3g+2b_m+b_u+p>4$ then S admits an internally decomposable foliation F, so by Theorem 5.4 the Hubbard–Masur quadratic differential associated to F at the basepoint X is fusible and hence the Busemann map is not continuous at that point. On the other hand, if $3g+2b_m+b_u+p \leq 4$ then by Lemma 5.14 for any quadratic differential q, the vertical foliation $V(q)$ is internally indecomposable, so again by Theorem 5.4 every quadratic differential is infusible an B is continuous at every boundary point. □

By combining Proposition 5.12 with Proposition 1.3, we get the precise classification of surfaces with horofunction compactification isomorphic to visual compactification announced in Theorem 1.6 from the introduction.

Proof of Theorem 1.6 As shown in Proposition 1.3, the visual compactification and the horofunction compactification are isomorphic if and only if the Busemann map is continuous, so the theorem follows by applying Proposition 5.12. □

5.2 Criteria for convergence

One straightforward consequence of the horofunction compactification being finer than the visual compactification is the following criterion regarding the convergence of sequences in the horofunction compactification.

Corollary 5.15 *Let $(x_n) \subset \mathcal{T}(S)$ be a sequence. If (x_n) converges to a quadratic differential q in the visual compactification, then all accumulation points of (x_n) in the horofunction compactification are contained in $\Pi^{-1}(q)$. In particular, if $V(q)$ is internally indecomposable, then (x_n) converges in the horofunction compactification.*

Furthermore, if (x_n) does not converge in the visual compactification, then it does not converge in the horofunction compactification.

Proof If x_n converges in the visual compactification to a quadratic differential q then by the continuity of Π all its accumulation points are in $\Pi^{-1}(q)$. If $V(q)$ is internally indecomposable, then by Theorem 5.4 the quadratic differential q is infusible, so the Busemann map is continuous at q and by Proposition 1.4 the fiber $\Pi^{-1}(q)$ is a singleton. Therefore x_n converges to $\Pi^{-1}(q)$, as that is the only accumulation point of x_n and the horofunction compactification is compact.

On the other hand, if x_n converges to ξ in the horofunction compactification, by continuity of Π, x_n converges to $\Pi(\xi)$ in the visual compactification. □

A frequent topic in the study of compactifications of Teichmüller spaces is the convergence of certain measure-preserving paths. We shall see now how the previous results can be applied in that study.

Let $X \in \mathcal{T}(S)$ be a point in Teichmüller space and q be a unit quadratic differential based at X. It is a well known fact that there exists a unique orientation-preserving isometric embedding $\iota\colon \mathbb{H} \to \mathcal{T}(S)$ from the hyperbolic plane $\mathbb{H}$ to the Teichmüller space such that $\iota(i) = X$ and $\iota^*(q) = i$, see the work of Herrlich and Schmithüsen [18] for a detailed explanation. The path $\iota(i+t)$ for $t \in \mathbb{R}_+$ is called the *horocycle* generated by q. Since ι is an isometric embedding, $h(X)(p) = d(\iota^{-1}X, \iota^{-1}p) - d(\iota^{-1}X, \iota^{-1}b)$ for $X, b, p \in \iota(\mathbb{H})$. That is, if we restrict the evaluations of horofunctions to the image of the Teichmüller disc, the value coincides with the values in the hyperbolic plane. Hence, since the path $i+t$ is a horocycle of the Busemann point obtained by moving along the geodesic $e^t i$ along the hyperbolic plane, the path $\iota(i+t)$ is also a horocycle of the corresponding Busemann point $B(q)$, obtained by moving along the geodesic $\iota(e^t i)$.

Since ι is an isometric embedding, the geodesic between X and $\iota(i+t)$ is contained in $\iota(\mathbb{H})$. Furthermore, the pushforward and pullback maps are continuous, so denoting q_t the unit quadratic differential spawning the geodesic between X and $\iota(i+t)$, we have $\lim_{t\to\infty} \iota^*(q_t) = i$, and $\iota_*(i) = q$, so $\lim_{t\to\infty} q_t = q$. The distance between $\iota(i+t)$ and X grows to infinity, so any horocycle path generated by some q based at X converges to q in the visual compactification based at X. Hence, horocycles generated by infusible quadratic differentials converge in the horofunction compactification, which had been previously shown by Jiang and Su [20] and Alberge [2] in the context of surfaces without boundary.

Corollary 5.16 *Let S be a compact surface with possibly nonempty boundary and finitely many marked points and let q be an infusible quadratic differential based at any $X \in \mathcal{T}(S)$. Then the horocycle generated by q converges in the horofunction compactification.*

Proof The horocycle path converges to q in the visual compactification based at X, so by Corollary 5.15 all accumulation points in the horofunction compactification are contained in $\Pi_X^{-1}(q)$. Furthermore, since q is infusible, $\Pi_X^{-1}(q)$ is a singleton, so the horocycle path has a unique accumulation point in the horofunction compactification, and hence it converges. □

On the other hand, Fortier Bourque found some diverging horocycles in the horofunction compactification.

Theorem 5.17 (Fortier Bourque [13, Theorem 1.1]) *Let S be a closed surface of genus g with p marked points, such that $3g + p \geq 5$. Then there is some fusible quadratic differential q based at some basepoint $X \in \mathcal{T}(S)$ such that the associated horocycle path does not converge in the horofunction compactification.*

Corollary 5.15 gives an upper limit on the set of accumulation points, as it has to be contained in $\Pi_X^{-1}(q)$. Furthermore, by Corollary 3.14 we have that a path converges in the horofunction compactification if and only if it converges in each visual compactification. Hence, such a divergent horocycle also diverges in some visual compactification. That is, we get Corollary 1.12. This contrasts with the behavior of Teichmüller rays, which by Corollary 3.7 or [44, Theorem 7] converge in all visual compactifications.

6 Dimension of the fibers

Our first approach in determining the shape of the fibers is looking at the limits of Busemann points, which by Proposition 3.10 give us bounds on the elements of $\Pi^{-1}(q)$. For a given quadratic differential q and a foliation G we define $\mathcal{W}^q(G)$ as the map from measured foliations to $\mathbb{R}$ given by

$$\mathcal{W}^q(G) = \frac{i(G,\cdot)^2}{i(G,H(q))}$$

if $i(G,H(q)) > 0$, and $\mathcal{W}^q(G) = 0$ otherwise. By the extension of Walsh's Corollary 4.6 describing Busemann points in the Gardiner–Masur compactification, we see that the element $\mathcal{E}_q = \Xi^{-1}B_q$ has the form $\sqrt{\sum_i \mathcal{W}^q(V_i)}$, where V_i are the indecomposable components of $V(q)$. Hence, a reasonable path to follow for understanding the limits of Busemann points is understanding the limits of $\mathcal{W}^q$ as q varies.

Lemma 6.1 *Let q_n be a sequence of quadratic differentials on X converging to q, and let V_j^n, $1 \leq j \leq c(n)$ be the indecomposable components of $V(q_n)$. Let G^n be a sequence of nonzero measured foliations of the form $\sum \alpha_j^n V_j^n$, converging to a measured foliation G. Then*

$$\lim_{n\to\infty} \mathcal{W}^{q_n}(G^n) = \mathcal{W}^q(G)$$

if G is nonzero and $\lim_{n\to\infty} \mathcal{W}^{q_n}(G^n) = 0$ if G is zero, where the convergence is pointwise in both cases.

Proof For any measured foliation F we have

$$\mathcal{W}^{q_n}(G^n)(F) = \frac{i(G^n,F)^2}{i(G^n,H(q_n))},$$

so if G is nonzero the lemma follows by continuity of the intersection number.

If G is zero the result follows from applying the same proof than in [44, Lemma 27]. □

Denote $\mathcal{B}$ the set of Busemann points, $\overline{\mathcal{B}}$ its closure and $\overline{\mathcal{B}}(q)$ the intersection $\overline{\mathcal{B}} \cap \Pi^{-1}(q)$. We can use the previous lemma to show that the elements of $\overline{\mathcal{B}}(q)$ satisfy certain properties.

Proposition 6.2 *Let S be a closed surface with possibly marked points, $\xi \in \overline{\mathcal{B}}(q)$ and V_i, $i \in \{1, \dots, k\}$ be the indecomposable components of $V(q)$. Let $x_i = i(V_i, \cdot)/i(V_i, H(q))$. Then, the square of the representation of ξ in the Gardiner–Masur compactification, $(\Xi^{-1}\xi)^2$, is a homogeneous polynomial of degree 2 in the variables x_i, whose coefficients sum to 1.*

Recall that we are using a normalized version of the Gardiner–Masur compactification. Under the projectivized version the sum of the coefficients cannot have any fixed value.

Proof Since the surface does not have boundary, all Busemann points are of the form $B(q')$ for some quadratic differential of unit area q'. Consider a sequence (q_n) such that $B(q_n)$ converges to ξ and q_n converges to q. Let $c(n)$ be the number of indecomposable vertical components of $V(q_n)$, and let V_j^n, $0 < j \le c(n)$ be those components. We know that $c(n)$ is bounded by some number depending on the topology of the surface. Take a subsequence such that $c(n)$ is equal to some constant c and V_j^n converges for each j. The sum $\sum_{j=1}^c V_j^n$ converges as $n \to \infty$ to $\sum_{i=1}^k V_i$, so the limit of each V_j^n has to be of the form $\sum_{i=1}^k \alpha_j^i V_i$. Furthermore, $\sum_{j=1}^c \alpha_j^i = 1$, since

$$\sum_{i=1}^k V_i = V(q) = \lim_{n\to\infty} V(q_n) = \lim_{n\to\infty} \sum_{j=1}^c V_j^n = \sum_{j=1}^c \sum_{i=1}^k \alpha_j^i V_i = \sum_{i=1}^k \Big(\sum_{j=1}^c \alpha_j^i \Big) V_i.$$

The element associated to the Busemann point $B(q_n)$ in the Gardiner–Masur compactification satisfies

$$\mathcal{E}_{q_n}^2 = \sum_{j=1}^c \mathcal{W}^{q_n}(V_j^n).$$

Hence, applying Lemma 6.1 we get the following expressions for the square of the limit of Busemann points:

$$(\Xi^{-1}\xi)^2 = \sum_{j=1}^c \mathcal{W}^q \Big(\sum_{i=1}^k \alpha_j^i V_i \Big) = \sum_{j=1}^c \frac{\big(\sum_{i=1}^k \alpha_j^i i(V_i, H(q)) x_i\big)^2}{\sum_{i=1}^k \alpha_j^i i(V_i, H(q))}.$$

That is, we get a homogeneous polynomial of degree 2 in the variables x_i. Since q has unit area, the sum of the coefficients is

$$\sum_{j=1}^c \sum_{i=1}^k \alpha_j^i i(V_i, H(q)) = \sum_{i=1}^k i(V_i, H(q)) = 1,$$

which completes our claim. □

By Proposition 3.9, the Busemann point $B(q)$ gives an upper bound on all functions in $\Pi^{-1}(q)$. While Proposition 3.10 does not give us a lower bound directly, we can use Lemma 2.1 to get one. For a unit area quadratic differential q, let Z_j be the interior parts of $V(q)$, and denote by G_j the union of interior indecomposable components within Z_j. Further, let P_i be the boundary components of $V(q)$. We define the *minimal point* at q as

$$M(q) = \Xi\left(\sum_i \mathcal{W}^q(P_i) + \sum_j \mathcal{W}^q(G_j)\right)^{1/2}.$$

Proposition 6.3 *Let q be a quadratic differential. Then, for any $\xi \in \Pi^{-1}(q)$, we have*

$$\Xi^{-1}\xi \geq \Xi^{-1}M(q)$$

in the Gardiner–Masur compactification. Furthermore, $M(q) \in \Pi^{-1}(q)$ whenever each G_j has at most two annuli parallel to the boundaries of Z_j with marked points.

In the context of surfaces without boundary the previous result has been also proven by Liu and Shi [27, Lemma 3.10]. In such context we have $M(q) = \Xi i(V(q), \cdot)^2$, which by the proposition is always contained in $\Pi^{-1}(q)$.

The minimality is essentially derived from the following well-known inequality.

Lemma 6.4 (Titu's lemma) *For any positive reals $a_1, \dots, a_n$ and $b_1, \dots, b_n$ we have*

$$\sum_j \frac{a_j^2}{b_j} \geq \frac{\left(\sum_j a_j\right)^2}{\sum_j b_j}.$$

Proof The inequality can be written as

$$\sum_i b_i \sum_j \frac{a_j^2}{b_j} \geq \left(\sum_j a_j\right)^2,$$

so the result follows after applying the Cauchy–Schwartz inequality. □

The implication this lemma has for our discussion is that $\mathcal{W}^q(\cdot)$ is convex, in the sense that for any $G = \sum_i G_i$ and any measured foliation F we have

$$\sum_i \mathcal{W}^q(G_i)(F) \geq \mathcal{W}^q(G)(F).$$

Proof of Proposition 6.3 If q is infusible then each G_j is indecomposable, so $M(q) = B(q)$, the fiber $\Pi^{-1}(q)$ has one point and the proposition is satisfied.

Consider then q fusible and $\xi \in \Pi^{-1}(q)$. Let $(x_n) = (R(q_n; t_n)) \subset \mathcal{T}$ converging to ξ. By Lemma 4.5 we have $\Xi^{-1}(h(x_n)) \geq \Xi^{-1}B(q_n)$. Hence, $\Xi^{-1}\xi \geq \liminf_{n\to\infty} \Xi^{-1}B(q_n)$.

Given a measured foliation F, take a subsequence so that

$$\liminf_{n\to\infty} \Xi^{-1} B(q_n)(F) = \lim_{n\to\infty} \Xi^{-1} B(q_n)(F).$$

The foliations $V(q_n)$ converge to $V(q)$, so by Proposition 5.5 for n big enough all boundary components P_i are contained within $V(q_n)$. Hence, for n big enough the foliations $V(q_n)$ can be split to the interior parts Z_j by cutting along the proper arcs. Denote by G_j^n the interior components of the foliation $V(q_n)$ restricted to Z_j. Let $G_{j,k}^n$ be the indecomposable components of G_j^n. The sequence G_j^n converges to G_j, so we can take a subsequence such that each $G_{j,k}^n$ converges to some foliation $G_{j,k}$ with $\sum_k G_{j,k} = G_j$. Applying Lemma 6.1 we have

$$\lim_{n\to\infty} \Xi^{-1} B(q_n)(F) = \lim_{n\to\infty} \sum_i^n \mathcal{W}^{q_n}(P_i) + \sum_j \sum_k \mathcal{W}^{q_n}(G_{j,k}^n) = \sum_i \mathcal{W}^q(P_i) + \sum_j \sum_k \mathcal{W}^q(G_{j,k}).$$

Hence, applying Lemma 6.4 to the second sum we get the first part of the proposition.

To observe that the limit is actually reached we can repeat the proof of Proposition 5.11 and observe that a proper arc for each interior part is enough to approach the foliation whenever each interior part of the foliation has at most two annuli parallel to boundaries with marked points. □

By Corollary 4.8 this lower bound is carried to the horofunction representation and by Proposition 3.9 we have an upper bound. Hence, we have the chain of inequalities

$$M(q) \le \xi \le B(q),$$

for any $\xi \in \Pi^{-1}(q)$. As we see in the next proposition, this chain can be translated as well to the Gardiner–Masur compactification.

Proposition 6.5 *Let $\xi \in \Pi^{-1}(q)$. Then,*

$$\Xi^{-1}\xi \le \Xi^{-1} B(q).$$

Proof We have a sequence of points $R(q_n; t_n)$ converging to ξ, with q_n converging to q. By Lemma 3.3 we have $\xi(R(q;t)) = -t$. Further, $R(q_n; t_n)$ converges in the Gardiner–Masur compactification to the function $f(G)^2 = \lim_{n\to\infty} e^{-2t_n} \operatorname{Ext}_{R(q_n;t_n)}(G)$, and we have $\Xi f(x) = \xi(x)$. Hence,

$$\tfrac{1}{2} \log \frac{f(F)}{\operatorname{Ext}_{R(q;t)}(F)} \le \tfrac{1}{2} \log \sup_{G\in P} \frac{f(G)}{\operatorname{Ext}_{R(q;t)}(G)} = -t.$$

Upon exponentiating and reordering the terms, we get

$$\lim_{n\to\infty} e^{-2t_n} \operatorname{Ext}_{R(q_n;t_n)}(F) = f^2(F) \le e^{-2t} \operatorname{Ext}_{R(q;t)}(F)$$

for all t. Letting $t \to \infty$, the right hand side converges to $(\Xi^{-1} B(q)(F))^2$, so we get the proposition. □

Using these bounds we can further refine the characterization of points in $\Xi^{-1}\Pi^{-1}(q)$.

Proposition 6.6 *Let q be a quadratic differential, let V_i, $i \in \{1, \dots, k\}$ be the indecomposable components of $V(q)$ and let $x_i(F) = i(V_i, F)/i(V_i, H(q))$. Given $f \in \Xi^{-1}\Pi^{-1}(q)$ and $c > 0$ we have, for all $F \in \mathcal{MF}$,*

$$f^2(F) = c^2 + 2c \sum_i i(V_i, H(q))(x_i(F) - c) + \sum_{i,j} O\big((x_i(F) - c)(x_j(F) - c)\big).$$

In particular, as a function of the values $x_i(F)$ at the point $x_i = c$ for all i, $f^2(x_1, \dots, x_k)$ takes value c^2, is differentiable and satisfies

$$\frac{\partial}{\partial x_i} f^2(x_1, \dots, x_k) = 2c\, i(V_i, H(q)).$$

Proof We have that $(\Xi^{-1}M(q))^2 \le f^2 \le (\Xi^{-1}B(q))^2$. Letting $a_i = i(V_i, H(q))$ and $x_i = x_i(\cdot)$, we have by Lemmas 6.4 and 6.3 that $\left(\sum a_i x_i\right)^2 \le (\Xi^{-1}M(q))^2$. Writing the bounds on f^2 in terms of the variables x_i, we obtain

$$\left(\sum a_i x_i\right)^2 \le f^2 \le \sum a_i x_i^2.$$

Adding that $\sum a_i = 1$, we have that f^2 is bounded below by the arithmetic mean, and above by the quadratic mean. Rewriting both sides as a polynomial in $x_i - c$, we get

$$c^2 + 2c \sum a_i(x_i - c) + \left(\sum a_i(x_i - c)\right)^2 \le f^2 \le c^2 + 2c \sum a_i(x_i - c) + \sum a_i(x_i - c)^2,$$

so the first part of the proposition is satisfied. Subbing in the value $x_i(F) = c$ we get the second part. □

By Propositions 3.4 and 3.13 all members of $\Pi^{-1}(q)$ share their values along $R(q; \cdot)$, as well as the directional derivatives at the points of the geodesic. For a given q we have $x_i(\lambda H(q)) = \lambda$ for all i and all $\lambda > 0$. Hence, Proposition 6.6 shows a similar relation for the representations of the elements of $\Pi^{-1}(q)$ in the Gardiner–Masur compactification, as they share their value, as well as some derivatives, at all foliations of the form $\lambda H(q)$.

As shown by Fortier Bourque [13], the Gardiner–Masur boundary contains extremal length functions, so we can use Proposition 6.6 to get some information on the differentials of these functions. Namely, we recover in a more restricted setting the following result, proven in [33, Theorem 1.1].

Theorem 6.7 (Miyachi) *Let G_t, $t \in [0, t_0]$ be a path in the space of measured foliations on X which admits a tangent vector $\dot{G}_0$ at $t = 0$ with respect to the canonical piecewise linear structure. Then, the extremal length $\mathrm{Ext}(G, X)$ is right-differentiable at $t = 0$ and satisfies*

$$\frac{d}{dt^+} \mathrm{Ext}(G_t, X)\bigg|_{t=0} = 2i(\dot{G}_0, F_{G_0,X}),$$

where $F_{G_0,X}$ is the horizontal foliation of the Hubbard–Masur differential associated to G_0 on X.

The concrete extremal length functions in the Gardiner–Masur boundary we are going to use are given by the following theorem.

Theorem 6.8 (Fortier Bourque) *Let $\{w_1, \dots, w_k\}$ be weights with $w_i > 0$, let $\phi_n = \tau_1^{\lfloor nw_1 \rfloor} \circ \dots \circ \tau_k^{\lfloor nw_k \rfloor}$ be a sequence of Dehn multitwist around a multicurve $\{\alpha_1, \dots, \alpha_k\}$ in a surface S and let $X \in \mathcal{T}(S)$. Then the sequence $\phi_n(X)$ converges to*

$$\left[\operatorname{Ext}^{1/2}\left(\sum_{i=1}^{k} w_i i(F, \alpha_i)\alpha_i, X\right)\right]_{F \in \mathcal{MF}(S)}$$

in the projective Gardiner–Masur compactification as $n \to \infty$.

The precise statement of this result is slightly weaker [13, Corollary 3.4], but the same proof yields this extension.

Fix a multicurve $\{\alpha_1, \dots, \alpha_k\}$, weights $\{w_1, \dots, w_k\}$ and let $\alpha = \sum w_i \alpha_i$. Furthermore, normalize the weights $\{w_1, \dots, w_k\}$ so that there is a unit area quadratic differential q such that $V(q) = \alpha$. Denote by V_i the vertical components of $V(q)$. That is, $V_i = w_i \alpha_i$. We are able to recover Miyachi's formula when $i(V_i, H(q)) = w_i$ for all i. The sequence $\phi_n(X)$ converges in the visual compactification based at X to $q \in T_X\mathcal{T}(S)$. By Theorem 6.8 the function $f(F) = \lambda^{1/2} \operatorname{Ext}^{1/2}\left(\sum_{i=1}^{k} w_i i(F, \alpha_i)\alpha_i, X\right)$ is in $\Xi^{-1}\Pi^{-1}(q)$ for some $\lambda > 0$. We have $i(F, \alpha_i) = x_i(F) i(V_i, H(q))/w_i$. So, assuming $i(V_i, H(q)) = w_i$ we can write

$$f^2(F) = \lambda \operatorname{Ext}\left(\sum_{i=1}^{k} x_i(F) V_i, X\right).$$

We have $x_i(H(q)) = 1$ for all i, so by Proposition 6.6 the value of λ satisfies

$$f^2(H(q)) = \lambda \operatorname{Ext}(V(q), X) = 1.$$

Since q has unit area, $\operatorname{Ext}(V(q), X) = 1$, so $\lambda = 1$. Let I be any foliation such that $H(q) + I$ is well defined, and let $F_t = H(q) + tI$. We have

$$f^2(F_t) = \operatorname{Ext}\left(\sum_i V_i + t \sum_i x_i(I) V_i, X\right).$$

Hence, letting $J = \sum x_i(I) V_i$ and $G_t = V(q) + tJ$, we can apply Proposition 6.6 to get

$$\frac{d}{dt^+} \operatorname{Ext}(G_t, X)\bigg|_{t=0} = \sum_i \frac{dx_i}{dt}\bigg|_{t=0} \frac{\partial f^2}{\partial x_i}\bigg|_{x_i=1} = \sum_i \frac{i(V_i, I)}{i(V_i, H(q))} \cdot 2i(V_i, H(q)) = 2i(V(q), I).$$

On the other hand, applying Miyachi's Theorem 6.7 directly, we get

$$\begin{aligned} \frac{d}{dt^+} \operatorname{Ext}(G_t, X)\bigg|_{t=0} &= 2i(H(q), J) = 2\sum_i i(H(q), V_i) x_i(I) \\ &= 2\sum_i i(H(q), V_i) \frac{i(V_i, I)}{i(H(q), V_i)} = 2i(V(q), I), \end{aligned}$$

so both expressions coincide, and we have recovered Theorem 6.7 in this rather restricted setting. We

would like to note that Proposition 6.6 also gives some information for finding the second derivatives around the point $H(q)$. Namely, the second derivatives cannot diverge to infinity as we approach $H(q)$.

Combining Proposition 6.6 with Proposition 6.2 we get fairly restrictive necessary conditions for the points in $\bar{\mathcal{B}}(q)$ for surfaces without boundary. We shall be using these conditions in Section 7 to prove that Busemann points are not dense in the horoboundary. Now we prove a more straightforward consequence. For a topological space U, denote by $\dim(U)$ its Lebesgue dimension. See the book by Munkres [35, Chapter 5.80] for some background on basic dimension theory. Given an embedding $U \hookrightarrow V$ we have $\dim(U) \leq \dim(V)$, so the conditions for the points on $\bar{\mathcal{B}}(q)$ gives us the following result.

Corollary 6.9 *Let S be a surface without boundary. Let q be a quadratic differential such that $V(q)$ has n indecomposable components. Then,*

$$\dim(\bar{\mathcal{B}}(q)) \leq \tfrac{1}{2}(n(n-1)).$$

Proof By Proposition 6.2 we have an embedding of $\bar{\mathcal{B}}(q)$ into the space of homogeneous polynomials of degree 2. For a given $\xi \in \bar{\mathcal{B}}(q)$, let $b^{\xi}_{i,j}$ be the coefficient of $x_i x_j$. Adding the restriction $b_{i,j} = b_{j,i}$ we have a coefficient for each possible pair, so the dimension of homogeneous polynomials of degree 2 is equal to the number of possible pairs, that is, $n(n+1)/2$. Furthermore, by Proposition 6.6 we know the value of the first derivatives at $x_i = c$ for all i. For each i this gives us the linear equation

$$\sum_{j \neq i} b^{\xi}_{i,j} + 2b^{\xi}_{i,i} = 2i\,(V_i, H(q)).$$

These n equations are linearly independent, as $b^{\xi}_{i,i}$ is only contained on the equation related to x_i. As such, the dimension of the coefficients is at most $n(n+1)/2 - n = n(n-1)/2$.

We note that the sum of the coefficients being 1 is the equation we get when summing the n equations given by the derivatives, so we cannot use that to restrict further the dimension. □

Recall that the number of indecomposable components n is bounded in terms of the topology of the surface. Hence, the previous corollary gives us a uniform upper bound on the dimension of $\bar{\mathcal{B}}(q)$. More interestingly, we can also get a lower bound for the dimension of $\bar{\mathcal{B}}(q)$. This allows us to get a lower bound on the dimension of $\Pi^{-1}(q)$. Furthermore, as this is a lower bound, we do not need to restrict ourselves to surfaces without boundary, as the set of Busemann points always contains the set of Busemann points of the form $B(q)$. The bound is obtained by finding a dimensionally big set of different ways to approach a certain q along the boundary and showing that each of these different approaches results in different limits for the associated Busemann points.

Theorem 6.10 *Let S be a surface of genus g with b_m and b_u boundaries with and without marked points respectively and p interior marked points. Then there is some unit quadratic differential q such that*

$$\dim(\bar{\mathcal{B}}(q)) \geq 2\left\lfloor \tfrac{1}{2}(g + b_m) + \tfrac{1}{4}(b_u + p) - \sigma(g, b_u + p) \right\rfloor,$$

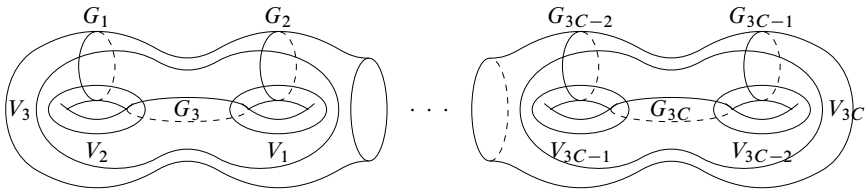

Figure 10: Labeling of the curves when the surface has no boundaries nor marked points. If g is odd then there is an unused handle.

where

$$\sigma = \begin{cases} 0 & \text{if } g \geq 2, \\ \frac{1}{4} & \text{if } g = 1 \text{ and } b_u + p \geq 1, \\ \frac{1}{2} & \text{if } g = 1 \text{ and } b_u + p = 0 \text{ or } g = 0 \text{ and } b_u + p \geq 2, \\ \frac{3}{4} & \text{if } g = 0 \text{ and } b_u + p = 1, \\ 1 & \text{if } g = 0 \text{ and } b_u + p = 0. \end{cases}$$

Proof For simplicity we shall first do the proof in the case where $b_m = b_u = p = 0$, and $g \geq 2$. Let q be the quadratic differential such that $V(q)$ is the union of the closed curves $V_1, \dots, V_{3C}$ shown in Figure 10, where $C = \lfloor g/2 \rfloor$. Let $U \subset \mathbb{R}^{3C}$ be the space of vectors $(\alpha_1, \alpha_2, \dots, \alpha_{3C})$ with positive coefficients and such that

$$\alpha_{3k+1} + \alpha_{3k+2} + \alpha_{3k+3} = \frac{1}{C}. \tag{3}$$

Each independent linear restriction reduces the dimension of the set U by 1, so $\dim U = 2C$. Hence, to prove the simplest case of the theorem it suffices to build an injective continuous map from U to $\overline{\mathcal{B}}(q)$.

Choose $\alpha \in U$ and consider the multicurve $\gamma^\alpha = \sum \alpha_i G_i$, where G_i are as in Figure 10. We will shortly show that by applying Dehn twists about the closed curves V_i to γ^α we can get a sequence of multicurves approaching $V(q)$. We can then take the sequences of associated Busemann points, which as we will see converge to distinct points in $\Pi^{-1}(q)$. We will define the injective continuous map from U to $\Pi^{-1}(q)$ by setting it as the limit of the associated sequence of Busemann points, giving us the theorem.

Let τ_i be the Dehn twist around V_i, and let w_i^α be such that

$$w_{3k+1}^\alpha(\alpha_{3k+2} + \alpha_{3k+3}) = w_{3k+2}^\alpha(\alpha_{3k+3} + \alpha_{3k+1}) = w_{3k+3}(\alpha_{3k+1} + \alpha_{3k+2}) = \frac{1}{3C}. \tag{4}$$

Define

$$\phi_n^\alpha = \tau_1^{\lfloor w_1^\alpha n \rfloor} \circ \tau_2^{\lfloor w_2^\alpha n \rfloor} \circ \cdots \circ \tau_{3C}^{\lfloor w_{3C}^\alpha n \rfloor}.$$

For $1 \leq k \leq C$ and $j \in \{1, 2, 3\}$, let

$$F_{k,j}^\alpha = \sum_{i \in \{1,2,3\} - j} w_{3k+i}^\alpha V_{3k+i}.$$

By counting the intersections between the curves V_i and G_i we have that there is some sequence λ_n such that $\lambda_n \phi_n^\alpha G_{3k+j}$ converges to $F_{k,j}^\alpha$ for all k, j as $n \to \infty$. By the conditions on the weights, $\lambda_n \phi_n^\alpha \gamma^\alpha$

converges to $V(q)$. Let q_n^α be the quadratic differential associated to $\lambda_n\phi_n^\alpha\gamma^\alpha$. Since $\lambda_n\phi_n^\alpha\gamma^\alpha$ converges to $V(q)$, we have that q_n converges to q, so all accumulation points of $(B(q_n))$ are in $\Pi^{-1}(q)$. We know that $(\Xi^{-1}B(q_n^\alpha))^2 = \sum_i \mathcal{W}^q(\alpha_i\lambda_n\phi_n^\alpha G_i)$, so by Lemma 6.1 we have

$$(\xi^\alpha)^2 = \lim_{n\to\infty}(\Xi^{-1}B(q_n^\alpha))^2 = \sum_{k=0}^{C-1}\sum_{j\in\{1,2,3\}}\alpha_{3k+j}\mathcal{W}^q(F_{k,j}^\alpha).$$

Define then the map from U to $\Pi^{-1}(q)$ sending $\alpha\in U$ to $\Xi\xi^\alpha\in\Pi^{-1}(q)$. As before, we shall let $x_i := i(V_i,\cdot)/i(V_i,H(q)) = 3Ci(V_i,\cdot)$. With this notation we have

$$\mathcal{W}^q(F_{k,j}^\alpha) = \frac{i(F_{k,j}^\alpha,\cdot)^2}{i(F_{k,j}^\alpha,H(q))} = \frac{\left(\sum_{i\notin\{1,2,3\}-j} w_{3k+i}^\alpha x_{3k+i}\right)^2}{3C\sum_{i\notin\{1,2,3\}-j} w_{3k+i}^\alpha}.$$

That is, given α we know precisely the shape of the polynomial ξ^α. Since α has positive coefficients, each of the w_i^α depends continuously on α, so ξ^α depends continuously on α.

It remains to show injectivity. Let $\beta\in U$ be such that $\xi^\alpha=\xi^\beta$. While we have equated two polynomials, we cannot conclude directly that the coefficients are equal, as these cannot be evaluated for arbitrary values. However, we can evaluate at elements of the form $b_1G_{3k+1}+b_2G_{3k+2}+b_3G_{3k+3}$ for $b_1,b_2,b_3\geq 0$, which is enough to prove that ξ^α and ξ^β have the same coefficients.

Equating then the coefficients for $x_{3k+1}x_{3k+2}$, $x_{3k+2}x_{3k+3}$ and $x_{3k+1}x_{3k+3}$ we get

$$\begin{aligned}
\frac{\alpha_{3k+1}w_{3k+2}^\alpha w_{3k+3}^\alpha}{w_{3k+2}^\alpha+w_{3k+3}^\alpha} &= \frac{\beta_{3k+1}w_{3k+2}^\beta w_{3k+3}^\beta}{w_{3k+2}^\beta+w_{3k+3}^\beta},\\
\frac{\alpha_{3k+2}w_{3k+1}^\alpha w_{3k+3}^\alpha}{w_{3k+1}^\alpha+w_{3k+3}^\alpha} &= \frac{\beta_{3k+2}w_{3k+1}^\beta w_{3k+3}^\beta}{w_{3k+1}^\beta+w_{3k+3}^\beta} \quad\text{and}\\
\frac{\alpha_{3k+3}w_{3k+1}^\alpha w_{3k+2}^\alpha}{w_{3k+1}^\alpha+w_{3k+2}^\alpha} &= \frac{\beta_{3k+3}w_{3k+1}^\beta w_{3k+2}^\beta}{w_{3k+1}^\beta+w_{3k+2}^\beta}.
\end{aligned}$$

Dividing these equalities and using equations (3) and (4) we get

$$\begin{aligned}
\frac{\alpha_{3k+1}}{\alpha_{3k+2}}\frac{(1/C+\alpha_{3k+2})}{(1/C+\alpha_{3k+1})} &= \frac{\beta_{3k+1}}{\beta_{3k+2}}\frac{(1/C+\beta_{3k+2})}{(1/C+\beta_{3k+1})},\\
\frac{\alpha_{3k+2}}{\alpha_{3k+3}}\frac{(1/C+\alpha_{3k+3})}{(1/C+\alpha_{3k+2})} &= \frac{\beta_{3k+2}}{\beta_{3k+3}}\frac{(1/C+\beta_{3k+3})}{(1/C+\beta_{3k+2})} \quad\text{and}\\
\frac{\alpha_{3k+3}}{\alpha_{3k+1}}\frac{(1/C+\alpha_{3k+1})}{(1/C+\alpha_{3k+3})} &= \frac{\beta_{3k+3}}{\beta_{3k+1}}\frac{(1/C+\beta_{3k+1})}{(1/C+\beta_{3k+3})}.
\end{aligned}$$

Rearranging the first equality we have

$$\frac{\alpha_{3k+1}}{\beta_{3k+1}}\frac{\beta_{3k+2}}{\alpha_{3k+2}} = \frac{(1/C+\alpha_{3k+1})}{(1/C+\beta_{3k+1})}\frac{(1/C+\beta_{3k+2})}{(1/C+\alpha_{3k+1})}. \tag{5}$$

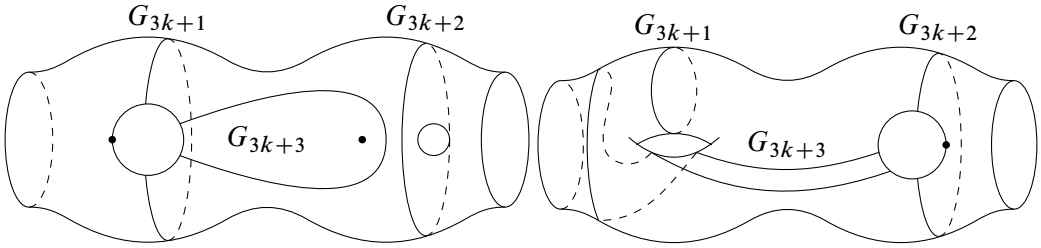

Figure 11: Each pair of marked points and boundary components without marked points can replace a genus, as well as each boundary with marked points.

If

$$\frac{\alpha_{3k+1}}{\beta_{3k+1}} < 1 \quad \text{then} \quad \frac{(1/C + \alpha_{3k+1})}{(1/C + \beta_{3k+1})} > \frac{\alpha_{3k+1}}{\beta_{3k+1}},$$

and if

$$\frac{\alpha_{3k+2}}{\beta_{3k+2}} > 1 \quad \text{then} \quad \frac{(1/C + \alpha_{3k+1})}{(1/C + \beta_{3k+1})} < \frac{\alpha_{3k+1}}{\beta_{3k+1}}.$$

Assume then that $\alpha_{3k+1} < \beta_{3k+1}$. One of the factors of the left hand side of the product in (5) is replaced in the right hand side by a larger value. Hence, the other factor has to be replaced by a smaller value. That is, the inequality $\alpha_{3k+2} < \beta_{3k+2}$ has to be satisfied. Similarly, if $\alpha_{3k+2} < \beta_{3k+2}$ we have $\alpha_{3k+3} < \beta_{3k+3}$. Equation (3) leads to

$$\frac{1}{C} = \alpha_{3k+1} + \alpha_{3k+2} + \alpha_{3k+3} < \beta_{3k+1} + \beta_{3k+2} + \beta_{3k+3} = \frac{1}{C},$$

which is a contradiction. Similarly, $\alpha_{3k+1} > \beta_{3k+1}$ leads to another contradiction, so $\alpha_{3k+1} = \beta_{3k+1}$, which leads to $\alpha = \beta$. Therefore, $\dim(\overline{\mathcal{B}}(q)) \geq \dim(U) = 2\lfloor g/2 \rfloor$.

Assume now that $g \geq 2$ and there are some marked points or boundaries. For each pair of marked points or unmarked boundaries, or for each marked boundary we can repeat the proof with an extra genus, by replacing the curves G_i by the curves shown in Figure 11, and halving the associated weights for w_i, as the curves intersect now twice the vertical components instead of once.

If $g = 1$ we need to place at least one feature at one of the ends to prevent the curve G_1 from being contractible or parallel to a unmarked boundary, so if we have marked points or boundaries without marked points we place these, as boundaries with marked points are more effective at increasing the dimension. In this way we get that if $b_u + p \geq 1$ then

$$\dim(\overline{\mathcal{B}}(q)) \geq 2\left\lfloor \tfrac{1}{2}(g + b_m) + \tfrac{1}{4}(b_u + p - 1) \right\rfloor$$

and if $b_u + p = 0$ then

$$\dim(\overline{\mathcal{B}}(q)) \geq 2\left\lfloor \tfrac{1}{2}(g + b_m - 1) \right\rfloor.$$

Lastly, if $g = 0$ we need to place two elements, one at each end. Using the same choice as we took for $g = 1$ we get

$$\begin{aligned}
\dim(\bar{\mathcal{B}}(q)) &\geq 2\left\lfloor \tfrac{1}{2}(b_m) + \tfrac{1}{4}(b_u + p - 2)\right\rfloor && \text{for } b_u + p \geq 2,\\
\dim(\bar{\mathcal{B}}(q)) &\geq 2\left\lfloor \tfrac{1}{2}(b_m - 1)\right\rfloor && \text{for } b_u + p = 1,\\
\dim(\bar{\mathcal{B}}(q)) &\geq 2\left\lfloor \tfrac{1}{2}(b_m - 2)\right\rfloor && \text{for } b_u + p = 0.
\end{aligned}$$

□

We would like to note that this lower bound is does not look optimal to us. Furthermore, the method used is restricted to getting to the dimension of the closure of Busemann points, so the dimension of the whole fiber may be significantly larger than what could be achieved by refining the strategy from the proof.

7 Nondensity of the Busemann points

7.1 Busemann points are not dense in the horoboundary

By Proposition 6.2 we know that points in the closure of Busemann points are smooth in the Gardiner–Masur representation with respect to certain variables. By showing that at least one point in the horoboundary is not smooth with respect to the corresponding variables we will prove that Busemann points are not dense. The points we use for this analysis are once again the ones found by Fortier Bourque in Theorem 6.8.

Following Fortier Bourque's reasoning, we shall first prove the nondensity for the sphere with five marked points, and then lift to general closed surfaces by using the branched coverings given by the following lemma, found in [15, Lemma 7.1].

Lemma 7.1 (Gekhtman–Markovic) *Let S be a closed surface of genus g with p marked points, such that $3g + p \geq 5$. Then there is a branched cover $\overline{S_{g,p}} \to \overline{S_{0,5}}$ that branches at all preimages of marked points that are not marked and induces an isometric embedding $\mathcal{T}(S_{0,5}) \hookrightarrow \mathcal{T}(S_{g,p})$.*

The particular conformal structure given to $S_{0,5}$ is obtained as follows. Let $S^1 = \mathbb{R}/\mathbb{Z}$ and let $C = S^1 \times [-1, 1]$. We obtain a sphere Σ by sealing the top and bottom of C via the relation $(x, y) \sim (-x, y)$ for all $(x, y) \in S^1 \times \{-1, 1\}$. Let P be set consisting of the five points $(0, \pm 1)$, $\left(\frac{1}{2}, \pm 1\right)$ and $(0, 0)$. The pair $S = (\Sigma, P)$, where we view Σ as a topological space, is the sphere with five marked points. We get a point X in $\mathcal{T}(S)$ by considering the complex structure on Σ obtained by the construction, using the identity map as our marking.

Let $\alpha(t) = \left(t, \frac{1}{2}\right)$ and $\beta(t) = \left(t, -\frac{1}{2}\right)$ for $t \in S^1$. Denote by τ_α and τ_β the Dehn twists along α and β. By Fortier Bourque's theorem, the sequence $(X_n) = ((\tau_\alpha \circ \tau_\beta)^n X)$ converges to a multiple of $\mathrm{Ext}^{1/2}(i(\alpha, \cdot)\alpha + i(\beta, \cdot)\beta, X))$ in the Gardiner–Masur compactification. Furthermore, the sequence (X_n) converges in the visual compactification based at X to the geodesic spawned by the quadratic differential

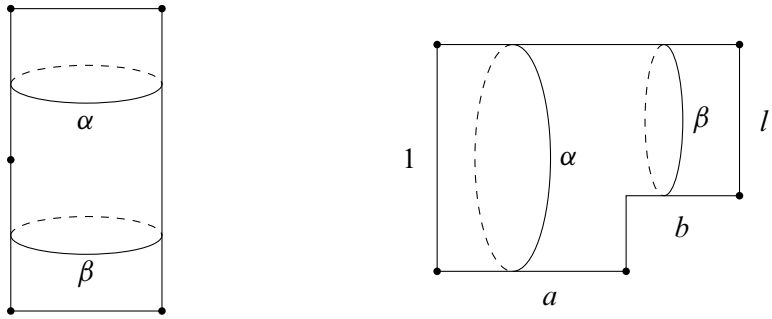

Figure 12: (Lemma 7.2) Left: sphere with five marked points, with curves α and β. We show that the extremal length is not C^2 along the path $\alpha + t\beta,\ t \in [0, t_0]$. Right: doubling of the L–shaped polygon together with the curves α and β.

$q_{\alpha+\beta,X}$. Indeed, as detailed in [13, Section 4], the elements (X_n) diverge to infinity along the horocycle defined by the quadratic differential $q_{\alpha+\beta,X}$. Hence, inside embedded hyperbolic plane associated to $q_{\alpha+\beta,X}$, the sequence (X_n) converges in the visual boundary to the geodesic spawned by $q_{\alpha+\beta,X}$, and so the same occurs in the ambient space. That is, $\Xi\,\mathrm{Ext}^{1/2}(i(\alpha,\cdot)\alpha + i(\beta,\cdot)\beta, X) \in \Pi^{-1}(q_{\alpha+\beta,X})$, so by Proposition 6.2 if we show that $\mathrm{Ext}(i(\alpha,\cdot)\alpha + i(\beta,\cdot)\beta, X)$ is not smooth with respect to the values of $i(\alpha,\cdot)$ and $i(\beta,\cdot)$, then $\Xi\,\mathrm{Ext}^{1/2}(i(\alpha,\cdot)\alpha + i(\beta,\cdot)\beta, X) \notin \bar{\mathcal{B}}(q_{\alpha+\beta,X})$, and hence it is also not in $\bar{\mathcal{B}}$.

Lemma 7.2 *Let $X \in \mathcal{T}(S_{0,5})$ and G_t, $t \in [0, t_0]$ be the foliation $\alpha + t\beta$ on $S_{0,5}$. The map $f(t) := \mathrm{Ext}(G_t, X)$ is not C^2.*

Proof By Miyachi's Theorem 6.7 we have

$$\frac{d}{dt}\,\mathrm{Ext}(G_t, X) = 2i(\beta, F_{G_t,X}),$$

where we remind that $F_{G_t,X}$ is the horizontal foliation of the unique Hubbard–Masur differential associated to G_t on X. Hence, the Lemma is equivalent to proving that $g(t) = i(\beta, F_{G_t,X})$ is not C^1.

For a general surface finding a precise expression of $F_{G,X}$ is a complicated problem, as the relation established by Hubbard and Masur is not explicit. However, in our case the surface is topologically simple, and one can use Schwartz–Christoffel maps to get a map from G to $F_{G,X}$. In particular, it is possible to show that the sphere with 5 marked points is conformally equivalent to the Riemannian surface obtained by doubling an L–shaped polygon, marking the inner angles as shown in Figure 12, right, and setting certain values for a, b and l. Furthermore, the quadratic differential obtained by dz^2 has α and β as vertical foliations, with weights a and b. Hence $q_{G_t,X}$ is dz^2 on the L–shaped pillowcase where $a = 1$ and $b = t$, so $i(\beta, F_{G_t,X}) = 2l$. Markovic estimated in [30, Section 9] the values of a, b and l around $b = 0$ depending on a common parameter r. Up to rescaling, these values are given by

$$a(r) = a(0) + D_1 r + O(r^2), \quad b(r) = D_2 r + O(r^2) \quad \text{and} \quad l(r) = l(0) + D_3 r \log\frac{1}{r} + o\Big(r \log\frac{1}{r}\Big),$$

where $A(r) = B(r) + O(f(r))$ means $|A(r) - B(r)|/f(r)$ is bounded around $r = 0$, and $A(r) = B(r) + o(f(r))$ means $|A(r) - B(r)|/f(r)$ converges to 0 as r converges to 0.

Rescaling the pillowcase by $1/a(r)$ we see that the parameter t can be expressed as $t(r) = b(r)/a(r)$, and $g(t(r)) = i(\beta, F_{G_t,X}) = 2l(r)/a(r)$. Observing that $t(0) = 0$, we can evaluate the first derivative of $g(t)$ at 0 by evaluating the limit

$$\begin{aligned}\lim_{h\to 0}\frac{g(h)-g(0)}{h} = \lim_{r\to 0}\frac{g(t(r))-g(0)}{t(r)} &= \lim_{r\to 0}\frac{2l(r)/a(r)-2l(0)/a(0)}{b(r)/a(r)}\\ &= 2\lim_{r\to 0}\frac{l(r)-l(0)a(r)/a(0)}{b(r)}\\ &= 2\lim_{r\to 0}\frac{D_3 r\log(1/r)+o(r\log(1/r))-(l(0)D_1/a(0))r}{D_2 r+O(r^2)}\\ &= \infty.\end{aligned}$$

And so, $g(t)$ is not differentiable at $t = 0$, and hence $f(t)$ is not C^2. □

Repeating Fortier Bourque's reasoning we can lift this example to any surface of genus g with p marked points as long as $3g + p \geq 5$. Besides the Gekhtman–Markovic lemma (Lemma 7.1), the other key ingredient for the lifting is the following result.

Lemma 7.3 (Fortier Bourque) *Let $\pi : S_{g,p} \to S_{0,5}$ be a branched cover of degree d and let*

$$\iota : \mathcal{T}(S_{0,5}) \hookrightarrow \mathcal{T}(S_{g,p})$$

be the induced isometric embedding. For any measured foliation F on $S_{0,5}$ and any $X \in \mathcal{T}(S_{0,5})$, we have the identity

$$\operatorname{Ext}(\pi^{-1}(F), \iota(X)) = d\operatorname{Ext}(F, X).$$

Proof Recall that $q_{F,X}$ is the Hubbard–Masur differential associated to γ. We have that $\pi^* q_{F,X} = q_{\pi^{-1}(F),\iota(X)}$, so

$$\operatorname{Ext}(\pi^{-1}(F), \iota(X)) = \int_{\iota(X)} |q_{\pi^{-1}(F),\iota(X)}| = d\int_X |q_{F,X}| = d\operatorname{Ext}(F, X).$$ □

Lifting the foliation G_t from Lemma 7.2 we get an upper bound for the smoothness of the extremal length.

Theorem 7.4 *Let S be a closed surface of genus g with p marked points, such that $3g+p \geq 5$. Then there exist two nonintersecting multicurves $\hat{\alpha}$, $\hat{\beta}$ and some $X \in \mathcal{T}(S)$ such that the map $f(t) := \operatorname{Ext}(\hat{\alpha} + t\hat{\beta}, X)$, $t \in [0, t_0]$ is not C^2.*

Proof Since $3g + p \geq 5$ we have a map $\pi\colon S_{g,p} \to S_{0,5}$, with an induced isometric embedding $\iota\colon \mathcal{T}(S_{0,5}) \hookrightarrow \mathcal{T}(S_{g,p})$. By Lemma 7.2 we have two curves $\alpha, \beta \in S_{0,5}$ such that, for any $X \in \mathcal{T}(S_{0,5})$ the map $t \to \operatorname{Ext}(\alpha + t\beta, X)$ is not C^2. Let $\hat{\alpha} = \pi^{-1}(\alpha)$ and $\hat{\beta} = \pi^{-1}(\beta)$. We have $\hat{\alpha} + t\hat{\beta} = \pi^{-1}(\alpha + t\beta)$, so applying Lemma 7.3 we get $\operatorname{Ext}(\hat{\alpha} + t\hat{\beta}, i(X)) = d \operatorname{Ext}(\alpha + t\beta, X)$. By Lemma 7.2 the function $\operatorname{Ext}(\alpha + t\beta, X)$ is not C^2, so we get the theorem. □

Theorem 1.9 is essentially a rephrasing of the previous theorem. Finally, we are able to prove that Busemann points are not dense.

Proof of Theorem 1.8 Let α and β be as in Lemma 7.2. Furthermore, let $\pi\colon S_{g,p} \to S_{0,5}$ and $\iota\colon \mathcal{T}(S_{0,5}) \hookrightarrow \mathcal{T}(S_{g,p})$ be as in Lemma 7.1. For the $X \in \mathcal{T}(S_{0,5})$ described before Lemma 7.2 the sequence $(X_n) = (\tau_\beta \circ \tau_\alpha)^n X$ is contained in the horocycle generated by $q_{\alpha+\beta,X}$ and the distance $d(X_n, X)$ goes to infinity. Therefore (X_n) converges in $\overline{\mathcal{T}(S_{0,5})}^v_X$ to the geodesic spawned by $q_{\alpha+\beta,X}$. Following Fortier Bourque's reasoning in the proof of [13, Theorem 1.1], using half translation structures, applying the Dehn twist $\tau_\alpha \circ \tau_\beta$ to X is equivalent to applying the shearing transformation

$$h_m = \begin{pmatrix} 1 & m \\ 0 & 1 \end{pmatrix}$$

to the half translation structure defined by $q_{\alpha+\beta,X}$. This action commutes with the pull-back coming from the branched cover, so the elements (X_n) are associated with the half translation structure defined by $h_n \pi^*(q_{\alpha+\beta,X})$. These points diverge to infinity along the horocycle defined by $\pi^*(q_{\alpha+\beta,X})$, and so converge in $\overline{\mathcal{T}(S_{g,p})}^v_{\iota(X)}$ to the geodesic spawned by $q_{\pi^{-1}(\alpha)+\pi^{-1}(\beta),\iota(X)}$.

Let c_i, $1 \leq i \leq k$, be the components of the half translation structure associated to $\pi^{-1}(\alpha + \beta, X)$. Each c_i covers either α or β with some degree $d_i \in \mathbb{N}$. Hence, each component c_i corresponds to a curve and is a cylindrical with height 1 and circumference d_i. Therefore, if m is the common multiple between all d_i, and γ_i is the curve associated to the component c_i, shifting the flat metric via the matrix h_m is equivalent to performing m/d_i Dehn twists around each curve γ_i. Letting ϕ be the composition of such Dehn twists, we have $\iota(X_{mn}) = \phi^n \iota(X)$. Hence, by Fortier Bourque's Theorem 6.8, in the Gardiner–Masur compactification the sequence $(\iota(X_{mn}))_n$ converges, as $n \to \infty$, to

$$\xi = \left[\operatorname{Ext}^{1/2}\left(\sum_{i=1}^{k} \frac{1}{d_i} i(F, \gamma_i)\gamma_i, \iota(X) \right) \right]_{F \in \mathcal{MF}(S_{g,n})}.$$

Therefore, $\Xi\xi \in \Pi^{-1}(q_{\pi^{-1}(\alpha)+\pi^{-1}(\beta),\iota(X)})$. To see that $\Xi\xi$ is not in $\overline{\mathcal{B}}$ it remains to see that it is not in $\overline{\mathcal{B}}(q_{\pi^{-1}(\alpha)+\pi^{-1}(\beta),\iota(X)})$. We have, $i(c_i, H(q_{\pi^{-1}(\alpha)+\pi^{-1}(\beta),\iota(X)}) = d_i$, so by Proposition 6.2 it remains to prove that there is some path of foliations G_t such that the functions $x_i = i(\gamma_i, G_t)/d_i$ vary smoothly, while the function $f(x_1, \dots, x_k) = \operatorname{Ext}\bigl(\sum_{i=1}^k (1/d_i) x_i \gamma_i, \iota(X)\bigr)$ does not. Reorder the curves so there is some $p \geq 1$ such that $\pi^{-1}\alpha = \gamma_1 + \cdots + \gamma_p$ and $\pi^{-1}\beta = \gamma_{p+1} + \cdots + \gamma_k$. It follows from Dehn–Thurston coordinates that for any natural numbers n_j, $1 \leq j \leq k$ there is a multicurve $G_{(n_j)}$ such that $i(G_{(n_j)}, \gamma_i) = n_i$. See, for example, the book by Penner and Harer [38, Theorem 1.2.1]. Allowing

renormalizations of the multicurves we get that n_j can be any nonnegative rationals. Finally, doing a limit argument in the space of projective measured foliations we can take n_j to be any nonnegative real numbers. That is, for any $t \geq 0$ there exists a measured foliation G_t such that $i(G_t, \gamma_i) = d_i$ for $i \leq p$, and $i(G_t, \gamma_i) = t d_i$ otherwise. Hence, along such foliations we have $x_i = 1$ for $i \leq p$ and $x_i = t$ otherwise. Therefore, along this path,

$$f(1, \dots, 1, t, \dots, t) = \mathrm{Ext}\big(\pi^{-1}(\alpha) + t\pi^{-1}(\beta), \iota(X)\big),$$

which by Theorem 7.4 is not smooth, as $\pi^{-1}(\alpha)$ and $\pi^{-1}(\beta)$ are the curves used in the proof of the theorem. □

7.2 Busemann points with one indecomposable component are nowhere dense

The Thurston compactification can be build in a similar way as the Gardiner–Masur compactification, by using the hyperbolic length of the curves instead of the extremal length. Let ϕ be the map between $\mathcal{T}(S)$ and $P\mathbb{R}_+^{\mathcal{S}}$ defined by sending $X \in \mathcal{T}(S)$ to the projective vector $[\ell(\alpha, X)]_{\alpha \in \mathcal{S}}$. The pair $\big(\phi, \overline{\phi(\mathcal{T}(S))}\big)$ defines a compactification, and the boundary is given by the space of projective measured foliations, denoted $\mathcal{PMF}$.

As explained by Miyachi [34], neither the Thurston nor the horofunction compactification is finer than the other one. However, it is possible to get some relation. Let $\mathcal{PMF}^{\mathrm{UE}} \subset \mathcal{PMF}$ be the set of uniquely ergodic foliations. Following the work of Masur [31], $\mathcal{PMF}^{\mathrm{UE}}$ has full Lebesgue measure within $\mathcal{PMF}$. Miyachi [34, Corollary 1] shows that the mapping ϕ on $\mathcal{T}(S)$ can be extended to an homeomorphism f between $\phi(\mathcal{T}(S)) \cup \mathcal{PMF}^{\mathrm{UE}}$ and $h(\mathcal{T}(S)) \cup B_{\mathrm{UE}}$ such that for $x \in \mathcal{T}(S)$ we have $f(\phi(x)) = h$, where B_{UE} are the Busemann points associated to quadratic differentials whose vertical foliation is uniquely ergodic. One might understand this result as stating that the two compactifications are the same almost everywhere with respect to the Lebesgue measure on $\mathcal{PMF}$. As we shall see, the same does not follow with respect to any strictly positive measure on the horoboundary.

The homeomorphism f described by Miyachi is obtained by first defining a map between the boundaries. For a given $x \in \mathcal{T}(S)$, the map on the boundary is denoted $\mathcal{G}_x$, and by its definition we have $\mathcal{G}_x(F) = B(q_{F,x})$, where we recall that $q_{F,x}$ is the quadratic differential on x with $V(q_{F,x}) = F$. Denote by $\mathcal{B}_1$ the set of Busemann points associated to foliations with one indecomposable component. We have $\mathcal{G}_x(\mathcal{PMF}^{\mathrm{UE}}) = \mathcal{B}_{\mathrm{UE}} \subset \mathcal{B}_1$. However, the following is also satisfied.

Theorem 7.5 *Let S be a closed surface of genus g with p marked points, such that $3g + p \geq 5$. Then the set $\mathcal{B}_1$ is nowhere dense in the horoboundary.*

Proof The action of $\mathrm{MCG}(S)$ on $\mathcal{T}(S)$ is extended to the projectivized version of the Gardiner–Masur compactification by $\psi[f(\alpha)]_{\alpha \in \mathcal{S}} = [f(\psi\alpha)]_{\alpha \in \mathcal{S}}$. For any q such that $V(q)$ is an indecomposable measured foliation, $\mathcal{E}_q = \Xi^{-1} B(q) = [i(V(q), \alpha)]_{\alpha \in \mathcal{S}}$, so $\psi \mathcal{E}_q = [i(V(q), \psi(\alpha))]_{\alpha \in \mathcal{S}} = [i(\psi^{-1}(V(q)), \alpha)]_{\alpha \in \mathcal{S}}$.

Hence, $\psi\mathcal{E}_q$ is equal to the representation of the Busemann point in the Gardiner–Masur compactification associated to the quadratic differential with vertical foliation $\psi^{-1}V(q)$, which also is an indecomposable measured foliation. Therefore, $\mathcal{B}_1$ is invariant under the action of $\mathrm{MCG}(S)$, and since $\mathrm{MCG}(S)$ acts by homeomorphisms, the complement of the closure is also invariant.

Let q_0 be a quadratic differential such that there is some $f \in \Xi^{-1}\Pi^{-1}(q_0)$ not in $\Xi^{-1}\bar{\mathcal{B}}$. Such a quadratic differential exists, by Theorem 1.8. By the proof of the theorem, we can assume that $V(q_0)$ is a multicurve. Furthermore, let q be a quadratic differential such that $V(q)$ and $H(q)$ are the stable and unstable foliations respectively of some pseudo-Anosov element $\phi \in \mathrm{MCG}(S)$. It is well known [12, exposé 12] that for any closed curve α we have that $\lambda^{-n}\phi^n(\alpha)$ converges to $\big(i(\alpha, V(q))/i(H(q), V(q))\big)H(q)$, where λ is the stretch factor of ϕ. For any foliation F we have that $\Xi^{-1}M(q_0)(F) = 0$ if and only if $i(V(q_0), F) = 0$, where $M(q_0)$ is the minimal point defined in Section 6. Hence, since $H(q)$ is the unstable foliation of a pseudo-Anosov element and $V(q_0)$ is a multicurve, we have $i(V(q_0), H(q)) \neq 0$, and so $f(H(q)) \geq \Xi^{-1}M(q_0)(H(q)) > 0$. We have $\phi^n[f(\alpha)]_{\alpha\in\mathcal{S}} = [f(\phi^n(\alpha))]_{\alpha\in\mathcal{S}}$. Taking limits and using that the functions in the Gardiner–Masur compactification are homogeneous of degree 1, we get that

$$\lim_{n\to\infty}[\phi^n f(\alpha)]_{\alpha\in\mathcal{S}} = \left[i(\alpha, V(q))f\left(\frac{H(q)}{i(V(q), H(q))}\right)\right]_{\alpha\in\mathcal{S}} = [i(\alpha, V(q))]_{\alpha\in\mathcal{S}},$$

Hence, in the normalized version, $\phi^n f$ converges to $i(\cdot, V(q)) = \Xi^{-1}B(q)$, as $V(q)$ is uniquely ergodic and therefore indecomposable. That is, $B(q)$ can be approached through a sequence of elements contained in the complement of the closure of $\mathcal{B}_1$.

Let $B(q')$ be any element in $\mathcal{B}_1$, where q' is any quadratic differential such that $V(q')$ has one indecomposable component. The set of pseudo-Anosov foliations is dense in $\mathcal{MF}(S)$, so we have a sequence of quadratic differentials (q_n) converging to q' with $V(q_n)$ being a pseudo-Anosov foliation. Since q' has one indecomposable component, the convergence is strong, and so $B(q_n)$ converges to $B(q')$. Each $B(q_n)$ can be approached through a sequence of elements contained in the complement of the closure of $\mathcal{B}_1$, so taking a diagonal sequence the same can be said for $B(q')$. □

Corollary 7.6 *Let S be a closed surface of genus g with p marked points, such that $3g + p \geq 5$. Then, for any finite strictly positive measure ν on the horoboundary, the set $\bar{\mathcal{B}}_1$ does not have full ν–measure.*

Proof By Theorem 7.5, the complement of $\bar{\mathcal{B}}_1$ is open and nonempty, so it must have positive ν–measure. □

This last result tells us that the image of Miyachi's homeomorphism does not have full ν–measure within the horoboundary for any strictly positive measure ν. However, as announced in the introduction, any attempt to extend the identity from the Thurston compactification to the horoboundary compactification to a set of full measure within the Thurston compactification results in the same problem. We restate here the result as we shall use the notation for the proof.

Corollary 1.10 *Let ν be any finite strictly positive measure on the horoboundary and let μ be the Lebesgue measure on the Thurston boundary. Furthermore, let ϕ be a map from the Thurston compactification to the horofunction compactification satisfying $\phi|_{\mathcal{T}(S)} = h$, where h is as in* (1)*. Then there is no subset U of the Thurston boundary with full μ–measure such that ϕ is continuous at every point in U and $\phi(U)$ has full ν–measure.*

Proof Assume such a U exists. Choose then a basepoint $x \in \mathcal{T}(S)$ and let $U' = U \cap \mathcal{PMF}^{\mathrm{UE}}$. For each element of $F \in U'$ the associated Hubbard–Masur quadratic differential $q_{F,x}$ satisfies $R(q_{F,x}; t) \to F$ as $t \to \infty$. Hence, since ϕ is continuous at F we have $\phi(F) = B(q_{F,x})$. That is, $\phi(U') \subset \mathcal{B}_1$.

Let $G \in U$. The set $\mathcal{PMF}^{\mathrm{UE}}$ has full μ measure, so $U' = \mathcal{PMF}^{\mathrm{UE}} \cap U$ also has full measure. Hence, since the Lebesgue measure is strictly positive, U' is dense within $\mathcal{PMF}$. Therefore G can be accessed through a sequence $(F_n) \subset U'$. Hence, since ϕ is continuous in G we have $\phi(G) = \lim \phi(F_n)$, so $\phi(U) \subset \overline{\mathcal{B}}_1$ and $\phi(U)$ cannot have full ν–measure. □

Another natural family of measures on the boundary is obtained by considering harmonic measures. Given a nonelementary measure μ on $\mathrm{MCG}(S)$ it is possible to define a random walk (w_n) as the sequence of random variables defined by

$$w_n = g_0 g_1 g_2 \cdots g_n,$$

where g_i are independent, identically distributed random variables on $\mathrm{MCG}(S)$ sampled according to the distribution μ. As proven by Kaimanovich and Masur [22, Theorem 2.2.4], random walks generated by a nonelementary probability measure converges almost surely in Thurston's compactification, so we can define the hitting measure ν in $\mathcal{PMF}$. Furthermore, the walk converges almost surely to uniquely ergodic projective foliations, so we can translate this result to the horofunction compactification in the following way.

Corollary 7.7 *Let μ be a nonelementary measure on* $\mathrm{MCG}(S)$*. Then the associated harmonic measure on the horoboundary is supported in a nowhere dense set.*

Proof For any $x \in \mathcal{T}(S)$ the sequence $(w_n x)$ converges almost surely in Thurston compactification to some $F \in \mathcal{PMF}^{\mathrm{UE}}$. Hence, by [34, Corollary 1], the sequence $(w_n x)$ converges almost surely to the Busemann point generated by a quadratic differential q with $V(q)$ being a multiple of F. Hence, the support of the harmonic measure is contained in $\mathcal{B}_1$, which is nowhere dense by Theorem 7.5. □

8 Topology of the horoboundary

In this section we make some progress towards determining the global topology of the horoboundary. We begin by showing that the minimal point $M(q)$ introduced in Proposition 6.3 serves as a section for the map Π whenever S does not have a boundary. Our main goal for this section is proving the following theorem.

Theorem 8.1 *Let S be a surface of genus g with b_m and b_u boundaries with and without marked points respectively and p interior marked points. Then, the map Π restricted to the boundary has a global continuous section $\partial\overline{\mathcal{T}}^v \to \partial\overline{\mathcal{T}}^h$ if and only if at least one of the two following conditions is satisfied:*

- $b_m = b_u = 0$, *or*
- $2g + 2b_m + b_u + p - \max(1 - b_u, 0) \leq 4$.

The section is given by sending the ray in the direction of q to the point $M(q)$ defined before Proposition 6.3.

Furthermore, if the map does not admit a global section, then it does not admit any local section around some points.

We begin by proving the theorem for surfaces without boundary, as it is significantly easier to prove.

Proposition 8.2 *Let S be a surface without boundary. Then the projection map Π restricted to the boundary admits a global section, given by the map $M : \partial\overline{\mathcal{T}}^v \to \partial\overline{\mathcal{T}}^h$.*

Proof By Proposition 6.3 every preimage $\Pi^{-1}(q)$ contains $M(q)$. We have $M(q) = \Xi\big(i(V(q), \cdot)\big)$, which is continuous, as the map Ξ is continuous. □

The rest of the cases of Theorem 8.1 require a more careful analysis.

Proposition 8.3 *Let S be either*

- *a torus with up at most two unmarked boundaries or interior marked points,*
- *a torus with one marked boundary and one interior marked point,*
- *a sphere with one marked boundary and up to three interior marked points, or*
- *a sphere with two marked boundaries and interior marked point.*

Then the projection map Π restricted to the boundary admits a global section, given by the map $M : \partial\overline{\mathcal{T}}^v \to \partial\overline{\mathcal{T}}^h$.

Proof We shall build the section in the same way we built it in Proposition 8.2, that is, sending q to $M(q)$.

Our first step in the proof is seeing that if $V(q)$ contains a separating proper arc then only one of the two parts separated by the proper arc admit interior components. We shall do this by inspecting each possible case. Assume then that $V(q)$ has a separating proper arc.

If S is a torus with up to two unmarked boundaries or marked points or a torus with one marked boundary and one marked point, then the separating proper arc splits the surface into a torus with a marked boundary

and a sphere with a marked boundary and a marked point or unmarked boundary. The latter does not admit an interior component.

If S is a sphere with one marked boundary and up to three boundaries then the separating proper arc splits the surface into two spheres, both with one marked boundary, one of them with two marked points and the other one with one marked point. Again, the latter does not admit an interior component.

Finally, if S is a sphere with two marked boundaries and one marked point or unmarked boundary, the proper arc splits the surface into one sphere with two marked boundaries and a sphere with one marked boundary and one marked point, which again does not admit an interior component.

Take then a sequence of unit quadratic differentials (q_n) converging to q. Let P_i, $i \in \{1, \dots, c\}$ be the boundary components of $V(q)$. Furthermore, denote by G the union of the interior components. By the first part of the proof, all the interior components are contained in the same interior part. We thus have

$$\Xi^{-1} M(q) = \left(\sum_i \mathcal{W}^q(P_i) + \mathcal{W}^q(G) \right)^{1/2}.$$

By Proposition 5.5 all boundary components of $V(q)$ are contained in $V(q_n)$ for n big enough, and all other boundary components of $V(q_n)$, denoted P^n, vanish in the limit. Denote by G^n the union of the interior components of $V(q_n)$. As before, each indecomposable component of G^n is contained in the same interior part, so we have

$$\Xi^{-1} M(q_n) = \left(\sum_i \mathcal{W}^{q_n}(\alpha_i^n P_i) + \mathcal{W}^{q_n}(P^n) + \mathcal{W}^{q_n}(G^n) \right)^{1/2},$$

which converges to $\Xi^{-1} M(q)$. □

Proposition 8.4 *Let S be either*

- *a surface of genus at least two and at least one boundary,*
- *a torus with at least one boundary and two more boundaries or interior marked points,*
- *a torus with at least two boundaries, one being marked, and possibly interior marked points,*
- *a sphere with at least one boundary, and four more boundaries or interior marked points,*
- *a sphere with at least two boundaries, one being marked, and two interior marked points, or*
- *a sphere with at least three boundaries, two being marked, and possibly interior marked points.*

Then the projection map Π restricted to the boundary does not admit a local section around some points.

Proof We shall prove this by finding a quadratic differential q and sequences (q_n^1) and (q_n^2) converging to q such that their preimages by Π are singletons, but such that $\Pi^{-1}(q_n^1)$ and $\Pi^{-1}(q_n^2)$ converge to

different points in $\Pi^{-1}(q)$. If we had a section around q, then its value at q_n^1 and q_n^2 would be $\Pi^{-1}(q_n^1)$ and $\Pi^{-1}(q_n^2)$ respectively, giving us a contradiction.

In all cases the construction will be similar. For q_n^1 we build a foliation with a separating proper arc P such that each of the parts has precisely one interior component consisting of a closed curve, which we denote by G_1 and G_2. Letting the weight of the proper arc diminish to 0 we can get a sequence of quadratic differentials (q_n^1) converging to a quadratic differential q such that $V(q) = G_1 + G_2$. Let $F_n^1 = P + nG_1 + nG_2$, A_n^1 and A the area of the Hubbard–Masur differentials $q_{F_n^1,X}$ and $q_{G_1+G_2,X}$, respectively. Denote $(1/\sqrt{A_n^1})q_{F_n^1,X}$ as q_n^1. These quadratic differentials have unit area, and converge to $(1/\sqrt{A})q_{G_1+G_2,X}$, which we denote by q. By construction, $V(q_n^1)$ is internally indecomposable, so $\Pi^{-1}(q_n^1)$ is a singleton, and

$$\Xi^{-1}\Pi^{-1}(q_n^1) = \left\{\left(\frac{\mathcal{W}^{q_n^1}(P) + n\mathcal{W}^{q_n^1}(G_1) + n\mathcal{W}^{q_n^1}(G_2)}{\sqrt{A_n^1}}\right)^{1/2}\right\}.$$

The sequences $P/\sqrt{A_n^1}$, $nG_1/\sqrt{A_n^1}$ and $nG_2/\sqrt{A_n^1}$ converge, respectively, to 0, $G_1/\sqrt{A}$ and $G_2/\sqrt{A}$. Hence, by Lemma 6.1

$$\text{the sequence } \Pi^{-1}(q_n^1) \quad \text{converges to } \left\{\left(\frac{\mathcal{W}^{q}(G_1) + \mathcal{W}^{q}(G_2)}{\sqrt{A}}\right)^{1/2}\right\}.$$

For building q_n^2 we take a curve γ intersecting G_1 and G_2 at b_1 and b_2 times, where $b_1, b_2 \in \{1, 2\}$. Denote by τ_1 and τ_2 the Dehn twists around G_1 and G_2. Let $F_n^2 = \tau_1^{2n/b_1}\tau_2^{2n/b_2}\gamma$ and A_n^2 the area of the Hubbard–Masur differential $q_{F_n^2,X}$. As before, denote by $(1/\sqrt{A_n^2})q_n^2$ the quadratic differentials $(1/\sqrt{A_n^2})q_{F_n^2,X}$. These quadratic differentials have unit area, and converge to q. Furthermore, each $V(q_n^2)$ is a singleton and

$$\Xi^{-1}\Pi^{-1}(q_n^2) = \left\{\left(\frac{\mathcal{W}^{q_n^2}((\tau_1\tau_2)^n\gamma)}{\sqrt{A_n^2}}\right)^{1/2}\right\}.$$

The sequence $(\tau_1\tau_2)^n\gamma/\sqrt{A_n^2}$ converges to $(G_1 + G_2)/A$, so by Lemma 6.1

$$\text{the sequence } \Xi^{-1}\Pi^{-1}(q_n^2) \quad \text{converges to } \left\{\left(\frac{\mathcal{W}^{q}(G_1 + G_2)}{\sqrt{A}}\right)^{1/2}\right\},$$

which is different than the limit of $\Xi^{-1}\Pi^{-1}(q_n^1)$.

It remains then to find such a multicurve. For genus at least two we take P to be a separating proper arc such that each of the parts is of genus at least one, and G_1 and G_2 to be noncontractible curves, not parallel to unmarked boundaries on each part, as shown in Figure 13, left.

For the torus we take P to be a separating proper arc with both endpoints in the unmarked boundary, or a marked boundary if there are no unmarked boundaries. Further, we choose the proper arc such that, after cutting along the arc, one part is a torus with one boundary. That is, every other feature of the surface lies in the other part. Then we let G_1 and G_2 be noncontractible curves on each part, as shown in Figure 13, middle.

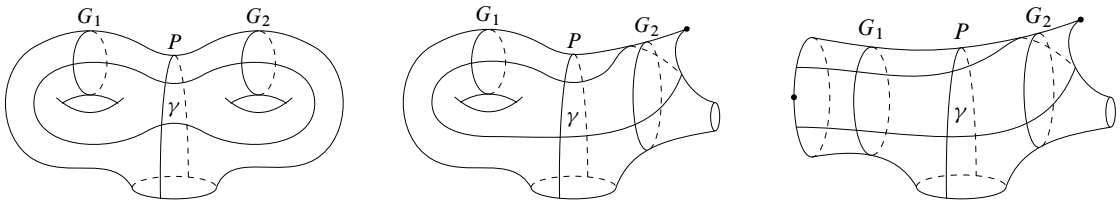

Figure 13: Curves chosen in the proof of Proposition 8.4.

Finally, for the sphere we let P be a separating proper arc with both endpoints on an unmarked boundary, or a marked boundary if there are no boundaries without marked points. Further, we choose the arc such that each interior part has at least either a combination of two marked points or boundaries without marked points, or a boundary with marked points. Hence, each interior part supports an interior component formed by a curve, as shown in Figure 13, right. □

Proof of Theorem 8.1 This is a combination of the results from Propositions 8.2, 8.3 and 8.4. □

By Proposition 1.2 we know that the horoboundary is connected whenever the real dimension of Teichmüller space is at least 2. In the following result we go a bit further, by showing that it is actually path connected.

Proof of Theorem 1.13 Let $x, y \in \partial\overline{\mathcal{T}(S)}^h$. If S does not have boundary then Π has a global section, so we can lift any path between $\Pi(x)$ and $\Pi(y)$ to a path between $M(\Pi(x))$ and $M(\Pi(y))$. Then, since $\Pi^{-1}(x)$ and $\Pi^{-1}(y)$ are path connected, we can connect x to $M(\Pi(x))$ and y to $M(\Pi(y))$ via paths.

If S has boundary we might have to be a bit more careful, as we might not have a global section. However, as we shall see, we can take a path q_t between $\Pi(x)$ and $\Pi(y)$ such that $B(q_t)$ has finitely many discontinuities. Then, since each of the preimages is path connected these discontinuities can be fixed by using paths in the fibers, so we will have a path between x and y.

Choose a boundary component of S, denote by b a curve parallel to that boundary and let $F_x = V(\Pi(x))$. If F_x contains b then all the expressions of the form $(1-t)F_x + tb$ with $t \in [0, 1]$ correspond to foliations on S, which we denote by F_t. Denote by q_t the unit area quadratic differential such that $V(q_t)$ is a multiple of F_t. This defines a continuous path joining $\Pi(x)$ and the unit area quadratic differential associated to a multiple of b. Let V_i be the vertical components of F_x that are not b, and let w_0 be the weight of b in F_x. Then,

$$B(q_t)^2 = \frac{1}{\sqrt{\mathrm{Area}(q_{F_t,X})}}\left((1-t)\sum \mathcal{W}^{q_t}(V_i) + (t + (1-t)w_0)\mathcal{W}^{q_t}(b)\right),$$

which gives a continuous path from $B(q_0) \in \Pi^{-1}\Pi(x)$ to $B(q_1) \in \Pi^{-1}(q_1)$. If F_x does not contain b, but b can be added to the foliation then we proceed just as before. Hence, if both x and y result in foliations where b can be added, we create a path by concatenating the paths between x, the Busemann point in $\Pi^{-1}\Pi(x)$, the Busemann point associated to b, the Busemann point in $\Pi^{-1}\Pi(y)$ and y.

If b cannot be added to the foliation F_x then there must be some set P of proper arcs in F_x incident to the boundary component associated to b. Let F'_x be the foliation F_x without the proper arcs P and assume F'_x is nonempty. Denote by F_t the foliations $(1-t)P+(1+t)F'_x$, $t\in[0,1]$, and q_t the unit area quadratic differentials such that $V(q_t)$ is a multiple of F_t. Denoting V_i the vertical components of F'_x, and P_j the proper arcs incident to the boundary component associated to b, we have

$$B(q_t)^2=\frac{1}{\sqrt{\mathrm{Area}(q_{F_t,X})}}\Big((1-t)\sum_j \mathcal{W}^{q_t}(P_j)+(1+t)\sum \mathcal{W}^{q_t}(V_i)\Big)$$

for $t<1$, which is continuous. Furthermore, $\lim_{t\to 1} B(q_t)\in\Pi^{-1}(q_1)$. Hence, we can concatenate a paths between x, the Busemann point in $\Pi^{-1}\Pi(x)$, the limit $\lim_{t\to 1} B(q_t)$, the Busemann point $B(q_1)$ and Busemann point associated to b.

If F'_x is empty we want to add some other components to F_x. If it admits some other component k then we repeat the previous reasoning with $F_t=\big(1-\frac{t}{2}\big)F_x+\frac{t}{2}k$, which does not result in any discontinuity. If F_x does not admit any other component then there must be at least 2 proper arcs incident to the boundary component associated to b, so we choose one of them, denoted p, and repeat the previous reasoning with $F_t=(1-t)F_x+tp$, which does not result in any discontinuity. Finally, we concatenate this last path with the previous paths. □

9 Formulas for limits of extremal lengths

We finish by reframing the bounds we got for the elements of $\Xi^{-1}\Pi^{-1}(q)$ as results regarding limits of extremal lengths, getting in this way some extensions of [44, Theorem 1].

Proposition 9.1 *Let F be a measured foliation, (q_n) be a sequence of unit area quadratic differentials converging to a quadratic differential q and (t_n) be a sequence of real numbers converging to infinity. Then,*

$$(\Xi^{-1}M(q))^2\le\liminf_{n\to\infty} e^{-2t_n}\,\mathrm{Ext}_{R(q_n;t_n)}(F)\le\limsup_{n\to\infty} e^{-2t_n}\,\mathrm{Ext}_{R(q_n;t_n)}(F)\le(\Xi^{-1}B(q))^2.$$

Proof Take a subsequence such that $e^{-2t_n}\,\mathrm{Ext}_{R(q_n;t_n)}(F)$ converges to the liminf. Furthermore, take a subsequence such that $R(q_n;t_n)$ converge to a point $\xi\in\Pi^{-1}(q)$. By Proposition 6.3 we have $(\Xi^{-1}M(q))^2\le\xi^2$. Since $e^{-2t_n}\,\mathrm{Ext}_{R(q_n;t_n)}(F)$ converges to $\xi^2(F)$ we have the lower bound. For the upper bound we repeat the process taking the limsup and using Proposition 6.5. □

By noting that $\Xi^{-1}M(q)(F)$ and $\Xi^{-1}B(q)(F)$ evaluate to 0 if and only if $i(V(q),F)=0$, we get the following corollary, which has also been proven for surfaces without boundary by Liu and Shi [27, Corollary 3.11].

Corollary 9.2 *Let (q_n) be a sequence of unit area quadratic differentials converging to a quadratic differential q, and (t_n) be a sequence of real numbers converging to infinity. Then,*

$$\liminf_{n\to\infty} e^{-2t_n}\operatorname{Ext}_{R(q_n;t_n)}(F)=0 \iff i(V(q),F)=0.$$

Proposition 9.1 can be strengthened slightly in the following manner.

Proposition 9.3 *Let (q_n) be a sequence of unit area quadratic differentials converging to a quadratic differential q. Furthermore, denote by V_i^n the indecomposable components of q_n. If the vertical components can be reordered so that for each i we have that V_i^n converges to a foliation V_i, then*

$$\liminf_{n\to\infty} e^{-2t_n}\operatorname{Ext}_{R(q_n;t_n)}(F)\geq \sum_i \mathcal{W}^q(V_i).$$

Proof Take a sequence such that the limit is equal to the liminf, and such that we have convergence in the Gardiner–Masur compactification. Let ξ be the limit in the horofunction compactification. By Lemma 4.5 we have $e^{-2t_n}\operatorname{Ext}_{R(q_n;t_n)}(F)\geq(\Xi^{-1}B(q_n))^2$, and by Corollary 4.6 we have $(\Xi^{-1}B(q_n))^2=\sum_i \mathcal{W}^{q_n}(V_i^n)$. Hence, by Lemma 6.1, taking limits on both sides we get the proposition. □

If we have strong convergence the upper bound from Proposition 9.1 and the lower bound from Proposition 9.3 coincide, so adding Walsh's formula for the Busemann points [44, Theorem 1] we have a proof of Theorem 1.14.

Finally, the path connectedness of the fibers can be translated to the following result.

Proposition 9.4 *Let (q_n) be a sequence of unit quadratic differentials converging to q, and (t_n) be a sequence of times converging to infinity. Further, for any $F\in\mathcal{MF}$ let $L(F):=\liminf_{n\to\infty}\operatorname{Ext}_{R(q_n;t_n)}(F)$. Then, for any $s\in[L(F),\mathcal{E}_q^2(F)]$ there is a subsequence of $q_{n_k^s}$ and a sequence (t_k^s) of times such that, for any $G\in\mathcal{MF}$ the limit*

$$\lim_{k\to\infty} e^{-2t_k^s}\operatorname{Ext}_{R(q_{n_k^s};t_k^s)}(G)$$

is defined, and if $G=F$ it has value s.

Proof We can take a subsequence such that $\lim_{n\to\infty}\operatorname{Ext}_{R(q_n;t_n)}(F)$ converges to the liminf, and a further subsequence such that we have convergence in the Gardiner–Masur compactification to a point $\Xi^{-1}\xi\in\Xi^{-1}\Pi^{-1}(q)$. By Proposition 3.11 we have a path between ξ and $B(q)$ contained in $\Pi^{-1}(q)$, and hence a path γ between $\Xi^{-1}\xi$ and $\Xi^{-1}B(q)$ contained in $\Xi^{-1}\Pi^{-1}(q)$. By continuity there is a point in that path such that $\gamma_t(F)=\sqrt{s}$, and by the way we constructed γ_t, it is reached by taking a subsequence of $(q_{n_k^s})$ and a sequence (t_k^s) of times converging to infinity. Finally, since γ_t is a point in the Gardiner–Masur compactification approached by $R(q_{n_k^s}^s;t_k^s)$, the value of $\gamma_t(G)^2$ is equal to the limit from the proposition. □

References

[1] **W Abikoff**, *The real analytic theory of Teichmüller space*, Lecture Notes in Math. 820, Springer (1980) MR Zbl

[2] **V Alberge**, *Convergence of some horocyclic deformations to the Gardiner–Masur boundary*, Ann. Acad. Sci. Fenn. Math. 41 (2016) 439–455 MR Zbl

[3] **F Bonahon**, *Shearing hyperbolic surfaces, bending pleated surfaces and Thurston's symplectic form*, Ann. Fac. Sci. Toulouse Math. 5 (1996) 233–297 MR Zbl

[4] **F Bonahon**, *Geodesic laminations with transverse Hölder distributions*, Ann. Sci. École Norm. Sup. 30 (1997) 205–240 MR Zbl

[5] **F Bonahon**, *Transverse Hölder distributions for geodesic laminations*, Topology 36 (1997) 103–122 MR Zbl

[6] **D Burago**, **Y Burago**, **S Ivanov**, *A course in metric geometry*, Grad. Stud. Math. 33, Amer. Math. Soc., Providence, RI (2001) MR Zbl

[7] **Y Chen**, **R Chernov**, **M Flores**, **M F Bourque**, **S Lee**, **B Yang**, *Toy Teichmüller spaces of real dimension 2: the pentagon and the punctured triangle*, Geom. Dedicata 197 (2018) 193–227 MR Zbl

[8] **F Dal'bo**, **M Peigné**, **A Sambusetti**, *On the horoboundary and the geometry of rays of negatively curved manifolds*, Pacific J. Math. 259 (2012) 55–100 MR Zbl

[9] **C J Earle**, *The Teichmüller distance is differentiable*, Duke Math. J. 44 (1977) 389–397 MR Zbl

[10] **P Eberlein**, **B O'Neill**, *Visibility manifolds*, Pacific J. Math. 46 (1973) 45–109 MR

[11] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Math. Ser. 49, Princeton Univ. Press (2012) MR Zbl

[12] **A Fathi**, **F Laudenbach**, **V Poénaru**, *Thurston's work on surfaces*, Math. Notes 48, Princeton Univ. Press (2012) MR Zbl

[13] **M Fortier Bourque**, *A divergent horocycle in the horofunction compactification of the Teichmüller metric*, Ann. Inst. Fourier (Grenoble) 73 (2023) 1885–1902 MR Zbl

[14] **F P Gardiner**, **H Masur**, *Extremal length geometry of Teichmüller space*, Complex Variables Theory Appl. 16 (1991) 209–237 MR Zbl

[15] **D Gekhtman**, **V Markovic**, *Classifying complex geodesics for the Carathéodory metric on low-dimensional Teichmüller spaces*, J. Anal. Math. 140 (2020) 669–694 MR Zbl

[16] **M Gromov**, *Hyperbolic manifolds, groups and actions*, from "Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference", Ann. of Math. Stud. 97, Princeton Univ. Press (1981) 183–213 MR Zbl

[17] **P de la Harpe**, *On Hilbert's metric for simplices*, from "Geometric group theory, I", Lond. Math. Soc. Lect. Note Ser. 181, Cambridge Univ. Press (1993) 97–119 MR Zbl

[18] **F Herrlich**, **G Schmithüsen**, *On the boundary of Teichmüller disks in Teichmüller and in Schottky space*, from "Handbook of Teichmüller theory, I", IRMA Lect. Math. Theor. Phys. 11, Eur. Math. Soc., Zürich (2007) 293–349 MR Zbl

[19] **J Hubbard**, **H Masur**, *Quadratic differentials and foliations*, Acta Math. 142 (1979) 221–274 MR Zbl

[20] **M Jiang**, **W Su**, *Convergence of earthquake and horocycle paths to the boundary of Teichmüller space*, Sci. China Math. 59 (2016) 1937–1948 MR Zbl

[21] **J Kahn**, **K M Pilgrim**, **D P Thurston**, *Conformal surface embeddings and extremal length*, Groups Geom. Dyn. 16 (2022) 403–435 MR Zbl

[22] **V A Kaimanovich**, **H Masur**, *The Poisson boundary of the mapping class group*, Invent. Math. 125 (1996) 221–264 MR Zbl

[23] **S P Kerckhoff**, *The asymptotic geometry of Teichmüller space*, Topology 19 (1980) 23–41 MR Zbl

[24] **T Klein**, **A Nicas**, *The horofunction boundary of the Heisenberg group: the Carnot–Carathéodory metric*, Conform. Geom. Dyn. 14 (2010) 269–295 MR Zbl

[25] **B Lemmens**, **C Walsh**, *Isometries of polyhedral Hilbert geometries*, J. Topol. Anal. 3 (2011) 213–241 MR Zbl

[26] **A Lenzhen**, **B Modami**, **K Rafi**, *Teichmüller geodesics with d–dimensional limit sets*, J. Mod. Dyn. 12 (2018) 261–283 MR Zbl

[27] **L Liu**, **Y Shi**, *On the properties of various compactifications of Teichmüller space*, Monatsh. Math. 198 (2022) 371–391 MR Zbl

[28] **L Liu**, **W Su**, *The horofunction compactification of the Teichmüller metric*, from "Handbook of Teichmüller theory, IV", IRMA Lect. Math. Theor. Phys. 19, Eur. Math. Soc., Zürich (2014) 355–374 MR Zbl

[29] **J Maher**, **G Tiozzo**, *Random walks on weakly hyperbolic groups*, J. Reine Angew. Math. 742 (2018) 187–239 MR Zbl

[30] **V Markovic**, *Carathéodory's metrics on Teichmüller spaces and L–shaped pillowcases*, Duke Math. J. 167 (2018) 497–535 MR Zbl

[31] **H Masur**, *Interval exchange transformations and measured foliations*, Ann. of Math. 115 (1982) 169–200 MR Zbl

[32] **H Miyachi**, *Teichmüller rays and the Gardiner–Masur boundary of Teichmüller space*, Geom. Dedicata 137 (2008) 113–141 MR Zbl

[33] **H Miyachi**, *A differential formula for extremal length*, from "In the tradition of Ahlfors–Bers, VI", Contemp. Math. 590, Amer. Math. Soc., Providence, RI (2013) 137–152 MR Zbl

[34] **H Miyachi**, *Teichmüller rays and the Gardiner–Masur boundary of Teichmüller space, II*, Geom. Dedicata 162 (2013) 283–304 MR Zbl

[35] **J R Munkres**, *Topology: a first course*, Prentice-Hall, Englewood Cliffs, NJ (1975) MR Zbl

[36] **K Ohshika**, *Compactifications of Teichmüller spaces*, from "Handbook of Teichmüller theory, IV", IRMA Lect. Math. Theor. Phys. 19, Eur. Math. Soc., Zürich (2014) 235–254 MR Zbl

[37] **A Papadopoulos**, *Deux remarques sur la géométrie symplectique de l'espace des feuilletages mesurés sur une surface*, Ann. Inst. Fourier (Grenoble) 36 (1986) 127–141 MR Zbl

[38] **R C Penner**, **J L Harer**, *Combinatorics of train tracks*, Ann. of Math. Stud. 125, Princeton Univ. Press (1992) MR Zbl

[39] **M A Rieffel**, *Group C^*–algebras as compact quantum metric spaces*, Doc. Math. 7 (2002) 605–651 MR Zbl

[40] **K Strebel**, *Quadratic differentials: a survey*, from "On the work of Leonhard Euler", Birkhäuser, Basel (1984) 219–238 MR Zbl

[41] **W P Thurston**, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. 19 (1988) 417–431 MR Zbl

[42] **C Walsh**, *The horoboundary and isometry group of Thurston's Lipschitz metric*, from "Handbook of Teichmüller theory, IV", IRMA Lect. Math. Theor. Phys. 19, Eur. Math. Soc., Zürich (2014) 327–353 MR Zbl

[43] **C Walsh**, *The horofunction boundary and isometry group of the Hilbert geometry*, from "Handbook of Hilbert geometry", IRMA Lect. Math. Theor. Phys. 22, Eur. Math. Soc., Zürich (2014) 127–146 MR Zbl

[44] **C Walsh**, *The asymptotic geometry of the Teichmüller metric*, Geom. Dedicata 200 (2019) 115–152 MR Zbl

[45] **C Webster**, **A Winchester**, *Boundaries of hyperbolic metric spaces*, Pacific J. Math. 221 (2005) 147–158 MR Zbl

School of Mathematics and Statistics, University of Glasgow
Glasgow, United Kingdom

aitor@azemar.xyz

https://aitor.azemar.xyz

Received: 28 November 2022 Revised: 3 November 2023

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2024 is US $705/year for the electronic version, and $1040/year (+$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

https://msp.org/

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 24 Issue 7 (pages 3571–4137) 2024