ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 8 (2004) 539–564
Published: 17 March 2004

The metric space of geodesic laminations on a surface: I

Xiaodong Zhu
Francis Bonahon

NetScreen Technologies, 805 11-th Avenue
Building 3, Sunnyvale, CA 94089, USA
and
Department of Mathematics, University of Southern California
Los Angeles, CA 90089-1113, USA

Email: `xzhu@netscreen.com` and `fbonahon@math.usc.edu`

URL: `http://math.usc.edu/~fbonahon`

Abstract

We consider the space of geodesic laminations on a surface, endowed with the Hausdorff metric d_{H} and with a variation of this metric called the $d_{\log}$ metric. We compute and/or estimate the Hausdorff dimensions of these two metrics. We also relate these two metrics to another metric which is combinatorially defined in terms of train tracks.

AMS Classification numbers Primary: 57M99

Secondary: 37E35

Keywords: Geodesic lamination, simple closed curve

Proposed: Jean-Pierre Otal
Seconded: David Gabai, Martin Bridson

Received: 9 October 2003
Revised: 16 March 2004

Let S be a compact surface, possibly with boundary. We restrict attention to the case where the Euler characteristic of S is negative.

Geodesic laminations on S were introduced by Bill Thurston [20, 21] to provide a completion for the space of simple closed curves on S. They also occur in various problems in low-dimensional topology and geometry, and have been used as a very successful tool to attack these problems. So far, most of their use and analysis has been focused on *measured* geodesic laminations, namely geodesic laminations endowed with the additional structure of a transverse measure. However, there are several contexts where geodesic laminations occur without a preferred transverse measure. A fundamental example of this is the ending lamination of a geometrically infinite hyperbolic 3–manifold; see for instance [17] for a discussion, and compare [14, 15, 13]. The current article and its sequel [6] find their motivation in an attempt to systematically develop an 'unmeasured' theory of geodesic laminations. In particular, one could hope that the space $\mathcal{L}(S)$ of all geodesic laminations on S would share many features with the space $\mathcal{PML}(S)$ of projective measured geodesic laminations, which was analyzed in [20, 21, 9, 19] and in subsequent work. To some extent, our results show that this is not the case.

We endow the space $\mathcal{L}(S)$ of geodesic laminations on S with the Hausdorff metric d_{H}, whose definition can be found in Section 1. Our first result is the following.

Theorem 1 *The Hausdorff dimension of the metric space* $(\mathcal{L}(S), d_{\mathrm{H}})$ *is equal to* 0.

This should not be confused with the result of Joan Birman and Caroline Series [1] which says that the union of all geodesic laminations forms a subset of S of Hausdorff dimension 1, although the method of proof is closely related to the main ingredient of [1]. Thurston already proved in [22, Section 10] that the subset $\mathcal{L}^{\mathrm{cr}}(S)$ of $\mathcal{L}(S)$ that is the closure of the set of multicurves (= disjoint unions of finitely many simple closed geodesics) has Hausdorff dimension 0. Note that $\mathcal{L}(S)$ is significantly larger than $\mathcal{L}^{\mathrm{cr}}(S)$, as can for instance be seen in the examples of [6].

An immediate corollary of Theorem 1 is that the space $\mathcal{L}(S)$ also has topological dimension 0 (see for instance [10, Proposition 2.5])

Corollary 2 *The space* $\mathcal{L}(S)$ *is totally disconnected.*

Our analysis [6] of the simple examples where S is the once punctured torus or the four-times punctured sphere shows the kind of intricacies which one can expect for the topology of $\mathcal{L}(S)$. In these examples, $\mathcal{L}(S)$ consists of a standard Cantor set K in $\mathbb{R}$ with finitely many isolated points in each of the gaps of $\mathbb{R} - K$.

Theorem 1 is at first somewhat disappointing, as one would prefer a dimension which reflects the topological complexity of S. On second thoughts, its proof suggests a different metric on $\mathcal{L}(S)$ which has better invariance properties than the Hausdorff metric d_{H} but induces the same topology. Indeed, the definition of d_{H} requires that we choose a negatively curved metric on S, and a different choice for this metric changes d_{H} by a Hölder equivalence. In particular, the only way the Hausdorff dimension of d_{H} could be independent of this choice is if the dimension is 0 or ∞, which in retrospect makes Theorem 1 much more predictable.

For this reason, it is much better to consider the metric $d_{\log}$ introduced in Section 4. For small distances, $d_{\log}$ is just $1/\left|\log d_{\mathrm{H}}\right|$. The metric $d_{\log}$ depends only on the topology of S up to Lipschitz equivalence. In particular, its Hausdorff dimension depends only on the topology of S. Note that $d_{\log}$ and d_{H} induce the same topology on $\mathcal{L}(S)$.

A similar idea of modifying the metric, or at least of using a different gauge to compute Hausdorff dimensions, was already used by Geoff Mess in [16, Section 5], in the context of [1].

Theorem 3 *If the surface S is different from the three times punctured sphere, the twice punctured projective plane or the once punctured Klein bottle, the Hausdorff dimension $(\mathcal{L}(S), d_{\log})$ is finite, and at least equal to 2.*

In the three exceptional cases where S is the three times punctured sphere, the twice punctured projective plane or the once punctured Klein bottle, the space $\mathcal{L}(S)$ is finite, and consequently has Hausdorff dimension 0; see for instance [6, Proposition 16].

The lower estimate comes from [6], where we prove that the Hausdorff dimension of $(\mathcal{L}(S), d_{\log})$ is exactly equal to 2 when S is the once punctured torus or the 4-times punctured sphere. Proposition 19 provides an explicit upper bound for the Hausdorff dimension, but this upper bound is far from being sharp. By analogy with the arguments of [6], we conjecture that the Hausdorff dimension δ of $(\mathcal{L}(S), d_{\log})$ is always an integer (although the value of the conjectured integer tends to depend on which heuristic argument we are using). However, in the cases of the once punctured torus and the 4-times punctured sphere the δ–Hausdorff measure of $(\mathcal{L}(S), d_{\log})$ is the zero measure, and we conjecture that this is always the case. To some extent, this is another disappointment as it does not yield as interesting a dynamical systems as one could have hoped for.

In the last part of the paper, we show that the metric $d_{\log}$ has a strong combinatorial flavor, in terms of train tracks on the surface S. More precisely, the geodesic lamination λ is *weakly carried* by the train track Θ if, for every leaf l of λ, there is a bi-infinite curve c which is immersed in Θ and is homotopic to l by a homotopy which moves points by a bounded amount. In this case, we say that the leaf l is *tracked* by the curve c. When c crosses a succession of edges e_1, e_2, …, e_n of Θ, in this order, we say that λ *realizes* the edge path $\langle e_1, e_2, \dots, e_n\rangle$.

Given two geodesic laminations λ, λ' carried by the train track Θ, we define their *combinatorial distance* as follows.

$$d_\Theta(\lambda,\lambda')=\min\left\{\frac{1}{r+1}; \lambda \text{ and } \lambda' \text{ realize the same edge paths of length } r \text{ in } \Theta\right\}$$

The sets $\mathcal{L}^{\mathrm{w}}(\Theta)$, consisting of those geodesic laminations which are weakly carried by the train track Θ, can be used as local charts for the space $\mathcal{L}(S)$ of all geodesic laminations, in the sense that $\mathcal{L}(S)$ can be written as the union of finitely many $\mathcal{L}^{\mathrm{w}}(\Theta_i)$.

We prove the following estimate, which relates the Hausdorff metric d_{H} to the combinatorial metric d_Θ.

Theorem 4 *Let Θ be a train track on the negatively curved surface S. Then, there exists two constants c_1, $c_2 > 0$ such that, for every geodesic laminations λ, λ' which are weakly carried by Θ,*

$$\mathrm{e}^{-c_1/d_\Theta(\lambda,\lambda')} \leqslant d_{\mathrm{H}}\left(\lambda,\lambda'\right) \leqslant \mathrm{e}^{-c_2/d_\Theta(\lambda,\lambda')}$$

The first inequality is harder to prove, and deeper, than the second one. Indeed, this first inequality asserts in particular that two geodesic laminations which are close for the Hausdorff topology must be combinatorially close.

Since, for small distances, the metric $d_{\log}$ is just $1/\left|\log d_{\mathrm{H}}\right|$, the above estimate can be rephrased in the following way.

Corollary 5 *There exists two constants c, $c' > 0$ such that, for every pair of geodesic laminations λ, λ' which are weakly carried by Θ,*

$$c d_\Theta\left(\lambda,\lambda'\right) \leqslant d_{\log}\left(\lambda,\lambda'\right) \leqslant c' d_\Theta\left(\lambda,\lambda'\right).$$

In other words, the restrictions of the metrics $d_{\log}$ and d_Θ to the subspace of $\mathcal{L}(S)$ consisting of those geodesic laminations which are weakly carried by the train track are Lipschitz equivalent.

Theorem 4 is a key ingredient for the arguments of [6], which themselves provide the lower estimate of Theorem 3.

The results of this article were first proved in the dissertation [23]. The authors are grateful to the referee for several valuable comments on the first version of this article.

This work was partially supported by grants DMS-9504282, DMS-9803445 and DMS-0103511 from the National Science Foundation.

1 Geodesic laminations

Let S be a compact surface, possibly with boundary, whose Euler characteristic $\chi(S)$ is negative. The hypothesis on the Euler characteristic implies that S admits a riemannian metric m of negative curvature[1] for which the boundary ∂S is totally geodesic. Fix such a metric.

A *geodesic lamination* is a non-empty closed subset of S which can be decomposed as the union of a family of disjoint simple complete geodesics. Recall that a complete geodesic is one which admits a parametrization by arc length defined over all of $]-\infty, +\infty[$. In particular, a complete geodesic does not transversely hit the boundary. A closed geodesic is always complete, and a geodesic is allowed to be contained in the boundary ∂S. A geodesic is *simple* it it has no transverse self-intersection. It can be shown that the decomposition of a geodesic lamination into disjoint simple geodesics is unique. These geodesics are called the *leaves* of the geodesic lamination. We refer to [7, 8, 19, 5] for general background on geodesic laminations.

A fundamental example of geodesic lamination is provided by the union of finitely many disjoint simple closed geodesics. A geodesic lamination of this type is called a *multicurve*. However, typical geodesic laminations are much more complex from a topological and dynamical point of view.

This paper is devoted to studying the space $\mathcal{L}(S)$ of all geodesic laminations of S.

[1] The metric can even be chosen to have constant curvature -1. Those readers who are more comfortable with hyperbolic geometry can therefore assume that this is the case, although this does not really simplify the arguments.

To define a metric on $\mathcal{L}(S)$, we will use the classical *Hausdorff metric* on the set $C(X)$ of all non-empty bounded closed subsets of a metric space (X,d), defined by

$$d_{\mathrm{H}}(C,C') = \inf\left\{\varepsilon; C \subset N_\varepsilon(C') \text{ and } C' \subset N_\varepsilon(C)\right\} \quad \text{if} \quad C, C' \in C(X),$$

where $N_\varepsilon(C)$ denotes the ε–neighborhood $\{x \in X; \exists y \in C, d(x,y) \leqslant \varepsilon\}$ of C. It is well-known that d_{H} defines a metric on $C(X)$; see for instance [18].

To define a metric on $\mathcal{L}(S)$, we could consider it as a subset of the set $C(S)$ of all closed non-empty subsets of S, and use the restriction of the Hausdorff metric d_{H}. However, for reasons explained later in this section, it makes more sense to use the *projective tangent bundle* $PT(S)$ of S, defined as the set of pairs (x,l) where $x \in S$ and where l is a line (through the origin) in the tangent space T_xS. A geodesic lamination lifts to a closed subset $\widehat{\lambda}$ of $PT(S)$ by considering all $(x, T_x\lambda)$ where $x \in \lambda$ and $T_x\lambda \subset T_xS$ is the line tangent to the leaf of λ passing through x.

In this way, we have identified a geodesic lamination $\lambda \in \mathcal{L}(S)$ to a closed subset $\widehat{\lambda}$ of $PT(S)$. The Levi-Civita connection of the riemannian metric m on S enables us to lift m to a riemannian metric on $PT(S)$, thereby turning $PT(S)$ into a metric space. We then define the *Hausdorff distance* between the geodesic laminations λ and $\lambda' \in \mathcal{L}(S)$ as the Hausdorff distance $d_{\mathrm{H}}(\widehat{\lambda}, \widehat{\lambda}')$ between the closed subsets $\widehat{\lambda}$, $\widehat{\lambda}' \in C(PT(S))$. We will write $d_{\mathrm{H}}(\lambda,\lambda')$ for $d_{\mathrm{H}}(\widehat{\lambda}, \widehat{\lambda}')$.

This definition of $d_{\mathrm{H}}(\lambda,\lambda')$ can also be rephrased as

$$d_{\mathrm{H}}(\lambda,\lambda') = \min\left\{\varepsilon; \begin{array}{c} \forall x \in \lambda, \exists x' \in \lambda', d\left((x,T_x\lambda),(x',T_{x'}\lambda')\right) < \varepsilon \\ \forall x' \in \lambda', \exists x \in \lambda, d\left((x,T_x\lambda),(x',T_{x'}\lambda')\right) < \varepsilon \end{array}\right\}$$

where d here denotes the metric in the projective tangent bundle $PT(S)$.

Lemma 6 ([8, Section 3], [7, Section 4.1], [5, Section 1.2]) *The metric space* $(\mathcal{L}(S), d_{\mathrm{H}})$ *is compact.* □

We now investigate what happens when we replace the negatively curved metric m on S by another negatively curved metric m' with totally geodesic boundary. By compactness of S, the distortion between m and m' is bounded, and a complete m–geodesic g is quasi-geodesic for m'; a general property of quasi-geodesics in negatively curved manifolds then says that there exists a unique m'–geodesic g' which is homotopic to g by a homotopy moving points by a bounded amount. This correspondence between m–geodesics and m'–geodesics sends simple geodesics to simple geodesics, and disjoint geodesics to disjoint geodesics, and therefore induces a one-to-one correspondence between the set $\mathcal{L}(S,m)$ of m–geodesic laminations and the set $\mathcal{L}(S,m')$ of m'–geodesic laminations. In other words, the set $\mathcal{L}(S,m)$ is really independent of the metric m, which justifies our original notation $\mathcal{L}(S)=\mathcal{L}(S,m)$.

Let us see how the Hausdorff metric d_{H} depends on the metric m. Two metrics d and d' on a set X are said to be *Hölder equivalent* if there exists constants C_1, C_2, ν_1, $\nu_2>0$ such that

$$C_1 d(x,y)^{\nu_1} \leqslant d'(x,y) \leqslant C_2 d(x,y)^{\nu_2}$$

for every x, $y\in X$. They are *Lipschitz equivalent* if they are Hölder equivalent with exponents $\nu_1=\nu_2=1$.

Lemma 7 *As we change the negatively curved metric m on S, the Hausdorff metric d_{H} of $\mathcal{L}(S)$ varies only by a Hölder equivalence.*

Proof Let m and m' be two negatively curved metrics on S for which the boundary is totally geodesic.

In general, there is no homeomorphism $S\to S$ which sends each m–geodesic g to the corresponding m'–geodesic. However, there is always a homeomorphism $\varphi\colon PT(S)\to PT(S)$ which sends the lift $\widehat{g}\subset PT(S)$ of each m–geodesic g to the lift $\widehat{g}'\subset PT(S)$ of the corresponding m'–geodesic g'. See [11, Section 7.2] or [12, Section 19.1]. In addition, φ is *Hölder bicontinuous* in the sense that there exists constants C_1, C_2, ν_1, $\nu_2>0$ such that

$$C_1 d(x,y)^{\nu_1} \leqslant d'(\varphi(x),\varphi(y)) \leqslant C_2 d(x,y)^{\nu_2}$$

for every x, $y\in PT(S)$, where d and $'$ are the metrics on $PT(S)$ respectively defined by m and m'. It immediately follows that the Hausdorff metrics d_{H} and d'_{H} defined by m and m' on $\mathcal{L}(S)=\mathcal{L}(S,m)=\mathcal{L}(S,m')$ are Hölder equivalent. □

As indicated earlier, instead of d_{H}, we could have considered on $\mathcal{L}(S)$ the restriction δ_{H} of the Hausdorff metric of the spaces of closed subsets of S (as opposed to $PT(S)$). The identity map $(\mathcal{L}(S),d_{\mathrm{H}})\to(\mathcal{L}(S),\delta_{\mathrm{H}})$ is continuous by definition of the topologies, and is therefore a homeomorphism by the compactness property of Lemma 6. In other words, δ_{H} induces the same topology as d_{H} on $\mathcal{L}(S)$ and, consequently, is exactly as good as d_{H} from a topological point of view. However, from a metric point of view, Lemma 7 does not seem to hold for δ_{H} in any easy way. In addition to this technical problem, it may philosophically make more sense to consider a geodesic lamination as a subset of $PT(S)$, as indicated by the general framework of geodesic currents developed in [2, 3].

2 Fattened train tracks

A *fattened train track* Φ on the surface S is a family of finitely many 'long' rectangles $e_1, \ldots, e_n$ which are foliated by arcs parallel to the 'short' sides and which meet only along arcs (possibly reduced to a point) contained in their short sides. In addition, a fattened train track Φ must satisfy the following:

(i) each point of the 'short' side of a rectangle also belongs to another rectangle, and each component of the union of the short sides of all rectangles is an arc, as opposed to a closed curve;

(ii) each component of the boundary ∂S is either disjoint from Φ or contained in it;

(iii) note that the closure $\overline{S-\Phi}$ of the complement $S-\Phi$ has a certain number of 'spikes', corresponding to the points where at least 3 rectangles meet; we require that no component of $\overline{S-\Phi}$ is a disc with 0, 1 or 2 spikes or an annulus with no spike.

The rectangles are called the *edges* of Φ. The foliations of the edges of Φ induce a foliation of Φ, whose leaves are the *ties* of the fattened train track. The finitely many ties where several edges meet are the *switches* of the fattened train track Φ. A tie which is not a switch is *generic*. Finally, a point of $\partial\Phi$ which is contained in three edges (and consequently is the tip of a spike of $\overline{S-\Phi}$) is a *switch point*.

An m-geodesic lamination λ is *strongly carried* by the fattened train track Φ if it is contained in the interior of Φ in S and if its leaves are transverse to the ties of Φ. For a fattened train track Φ, let $\mathcal{L}^{\mathrm{s}}(\Phi) \subset \mathcal{L}(S)$ be the subset consisting of those geodesic laminations which are strongly carried by Φ.

Given a geodesic lamination λ, there are several constructions which provide a fattened train track Φ strongly carrying λ; see for instance [19, 7, 5]. In particular, $\mathcal{L}(S)$ is the union of all $\mathcal{L}^{\mathrm{s}}(\Phi)$ as Φ ranges over all fattened train tracks in S. Also, it immediately follows from definitions that $\mathcal{L}^{\mathrm{s}}(\Phi)$ is open in $\mathcal{L}(S)$. Since $\mathcal{L}^{\mathrm{s}}(\Phi)$ is compact by Lemma 6, this proves:

Lemma 8 *There exists finitely many fattened train tracks $\Phi_1, \Phi_2, \ldots, \Phi_n$ in S such that*

$$\mathcal{L}(S) = \mathcal{L}^{\mathrm{s}}(\Phi_1) \cup \mathcal{L}^{\mathrm{s}}(\Phi_2) \cup \ldots \mathcal{L}^{\mathrm{s}}(\Phi_n) \qquad \square$$

Let a be an oriented arc carried by the fattened train track Φ, namely such that a is immersed in the interior of Φ and is transverse to its ties. If a meets the edges e_1, e_2, $\ldots$, e_n of Φ in this order, we will say that $\langle e_1, e_2, \ldots, e_n\rangle$ is the *edge path followed by* the arc a.

If the geodesic lamination λ is strongly carried by Φ, it *realizes* the fattened train track $\langle e_1, e_2, \ldots, e_n\rangle$ if $\langle e_1, e_2, \ldots, e_n\rangle$ is followed an arc a immersed in a leaf of λ.

3 The Hausdorff dimension of $(\mathcal{L}(S), d_{\mathrm{H}})$

We want to show that the Hausdorff dimension $\dim_{\mathrm{H}}(\mathcal{L}(S), d_{\mathrm{H}})$ is equal to 0. By Lemma 8, it suffices to show that, for an arbitrary fattened train track Φ, the subset $\mathcal{L}^{\mathrm{s}}(\Phi) \subset \mathcal{L}(S)$ consisting of those geodesic laminations which are strongly carried by Φ has Hausdorff dimension 0.

Fix such a fattened train track Φ. For a geodesic lamination λ strongly carried by Φ and for each integer $r \geqslant 0$, we can consider the set $\{\gamma_1, \gamma_2, \ldots \gamma_n\}$ of all edge paths of length $2r+1$ that are realized by λ. Conversely, given a finite family of edge paths γ_1, γ_2, $\ldots$, γ_n of length $2r+1$, let $P_{\gamma_1\gamma_2\ldots\gamma_n} \subset \mathcal{L}^{\mathrm{s}}(\Phi)$ consist of those geodesic laminations such that the family of edge paths of length $2r+1$ realized by λ is exactly $\{\gamma_1, \gamma_2, \ldots \gamma_n\}$, namely such that λ realizes all the γ_i and no other edge path of length $2r+1$.

Lemma 9 *There exists constants a, $b > 0$, depending only on the fattened train track Φ and on the negatively curved metric m of S, such that, for every family of edge paths γ_1, γ_2, $\ldots$, γ_n of length $2r+1$, the diameter of $P_{\gamma_1\gamma_2\ldots\gamma_n} \subset \mathcal{L}^{\mathrm{s}}(\Phi)$ for the Hausdorff metric d_{H} is bounded by $a\mathrm{e}^{-br}$.*

Proof Consider two geodesic laminations λ, $\lambda' \in P_{\gamma_1\gamma_2\ldots\gamma_n}$, and their lifts $\widehat{\lambda}$, $\widehat{\lambda}'$ to the projective tangent bundle $PT(S)$.

Let x be a point of λ. In the leaf of λ containing x, let k be an immersed arc passing through x, and chosen so that each of the two halves of k delimited by x crosses exactly $r+1$ edges (counted with multiplicities) of Φ. Let γ be the edge path of length $2r+1$ thus followed by k.

Because λ and λ' realize the same edge paths of length $2r+1$, there is an arc k' immersed in a leaf of λ' which follows the same edge path γ as k.

Because they follow the same edge path γ, the two m–geodesic arcs k and k' are homotopic by a homotopy which moves their end points by a distance at most $c_2 > 0$, where c_2 depends only on the diameter of the edges of Φ.

The point x separates the arc k into two arcs whose length is at least $c_1 r$, for some constant $c_1 > 0$ depending on the lengths of the edges of Φ, suitably defined. Because the curvature of the metric m is negative, a Jacobi field argument then shows that there exists a point $x' \in k'$ such that the distance from $(x', T_{x'}k') = (x', T_{x'}\lambda')$ to $(x, T_x k) = (x, T_x \lambda)$ in $PT(S)$ is bounded by $a\mathrm{e}^{-br}$, where a and b depend only on c_1, c_2 and on a negative upper bound for the curvature of m.

If $\varepsilon = a\mathrm{e}^{-br}$, this shows that $\widehat{\lambda}$ is contained in the ε–neighborhood $N_\varepsilon\big(\widehat{\lambda}'\big)$ of $\widehat{\lambda}'$ in $PT(S)$.

Symmetrically, $\widehat{\lambda}'$ is contained in $N_\varepsilon\big(\widehat{\lambda}\big)$, and it follows that $d_{\mathrm{H}}(\lambda, \lambda') = d_{\mathrm{H}}(\widehat{\lambda}, \widehat{\lambda}') \leqslant \varepsilon = a\mathrm{e}^{-br}$. Since this holds for any λ, $\lambda' \in P_{\gamma_1\gamma_2\dots\gamma_n}$, this proves that the diameter of $P_{\gamma_1\gamma_2\dots\gamma_n}$ is bounded by $a\mathrm{e}^{-br}$. □

Let Γ_r denote the set of edge paths of length $2r+1$ in Φ. As r tends to ∞, the cardinal of Γ_r usually grows exponentially. However, most elements of Γ_r are not realized by any geodesic lamination:

Proposition 10 *The number of subsets $\{\gamma_1, \gamma_2, \dots, \gamma_n\}$ of Γ_r for which $P_{\gamma_1\gamma_2\dots\gamma_n}$ is non-empty is bounded by a polynomial function of r.*

Proof The argument is in spirit similar to the core ingredient of [1] (see also [22, Section 10]), which is that the number of elements of Γ_r that are realized by geodesic laminations grows at most polynomially. However, it is significantly more difficult because one has no *a priori* control on the number n of elements of the subsets considered.

Let $\mathcal{P}_r$ denote the set of those subsets $\{\gamma_1, \gamma_2, \dots, \gamma_n\} \subset \Gamma_r$ for which there exists a geodesic lamination $\lambda \in P_{\gamma_1,\gamma_2\dots\gamma_n}$ which crosses every edge of Φ. By consideration of the finitely many fattened train tracks which can be obtained from Φ by removing some edges, it suffices to show that the number of elements of $\mathcal{P}_r$ is bounded by a polynomial function of r.

Consider an edge path family $\{\gamma_1, \gamma_2, \dots, \gamma_n\} \in \mathcal{P}_r$ and a geodesic lamination $\lambda \in P_{\gamma_1,\gamma_2\dots\gamma_n}$ crossing every edge of Φ. Let s_1, s_2, …, s_p be the switch points of Φ.

For a given switch point s_i, consider the two leaves l and l' of λ that are closest to s_i, namely that can be joined to s_i by an arc contained in a tie and whose interior is disjoint from λ. Note that l and l' exist because λ crosses every edge of Φ. Then, one of the following happens:

(1) Starting from the switch (tie) containing s_i, the leaves l and l' follow the same edge path of length $2r+1$.

(2) The leaves l and l' follow a common edge path of length at most $2r$, and then diverge at some switch.

In the first case, draw an arc z_i which is carried by Φ, is disjoint from λ, begins at the switch point s_i, crosses exactly r edges (counted with multiplicities), and finally ends on the switch (tie) at the end of the r–th edge.

In the second case, note that there is a unique switch point s_j which separates l and l' at the switch where they diverge, and that l and l' are closest to s_j at that switch; this again follows from the fact that λ crosses each edge of Φ. Then, let $z_i = z_j$ be an arc which is carried by Φ, is disjoint from λ and joins s_i to s_j. Note that $z_i = z_j$ crosses at most $2r$ edges.

The arcs z_i are the *r–zippers* associated to λ. Those zippers arising in Case 2 are called *switch connections.*

Lemma 11 *The zippers z_1, z_2, $\dots$, z_p are embedded. The zippers z_i and z_j are disjoint unless they are equal (and correspond to a switch connection). Their union z^r_λ is unique up to isotopy of Φ respecting each tie.*

Proof We will prove that, if Q is a component of $e - \lambda$ where e is an edge of Φ, the intersection of z^r_λ with Q consists of at most one arc. The statements of the lemma automatically follow from this property.

For a given z_i, consider the components Q of the above type that meet z_i. By construction, these form a chain bounded on the sides by two leaves l and l' of λ and beginning at the switch point s_i. It immediately follows that z_i cannot visit the same Q twice. Similarly, if a Q meeting z_i is also visited by z_j, then the Q' meeting $z_i \cup z_j$ form a longer chain, which shows that the two leaves l and l' follow a common edge path of length at most $2r$ before diverging at s_j; it follows that $z_i = z_j$ is a switch connection. □

Let $\mathcal{Z}_r$ denote the set of all r–zipper families of this type, considered up to isotopy of Φ respecting its ties. Namely an element of $\mathcal{Z}_r$ is a family $\{z_1, z_2, \dots, z_p\}$ of embedded arcs carried by Φ and of the following two types:

(1) either z_i begins at the switch point s_i, crosses exactly r edges, and finally ends on the switch at the end of the r–th edge;

(2) or z_i joins the switch point s_i to another switch point s_j and crosses at most $2r$ edges of Φ, in which case $z_j = z_i$.

In addition, distinct z_i are disjoint. We identify two such zipper families when they differ only by an isotopy of Φ respecting its ties.

Lemma 11 associates an element $z^r_\lambda \in \mathcal{Z}_r$ to each geodesic lamination λ which is carried by Φ and crosses every edge of Φ.

It is convenient to consider the set $\mathcal{Q}_r$ of all pairs $(\lambda, \{\gamma_1, \gamma_2, \dots, \gamma_n\})$ where $\lambda \in P_{\gamma_1,\gamma_2,\dots,\gamma_n}$ crosses every edge of Φ. Lemma 11 defines a map $a\colon\ \mathcal{Q}_r \to \mathcal{Z}_r$.

If 2^{Γ_r} denotes the set of all subsets of the set Γ_r of all edge paths of length $2r+1$, there is also a forgetful map $b\colon\ \mathcal{Q}_r \to 2^{\Gamma_r}$ defined by $b\,(\lambda, \{\gamma_1, \gamma_2, \dots, \gamma_n\}) = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$. By construction, its image $b(\mathcal{Q}_r)$ is equal to $\mathcal{P}_r$.

We now construct a third map $c\colon\ \mathcal{Z}_r \to 2^{\Gamma_r}$.

Lemma 12 *Consider an r–zipper family $z \in \mathcal{Z}_r$. Let k_1 and k_2 be two arcs carried by Φ which cross $r+1$ edges of Φ, and which are disjoint from z. Suppose that k_1 and k_2 start from the same tie t of Φ and in the same direction, and that their starting point is in the same component of $t-z$. Then k_1 and k_2 follow the same edge path of length $r+1$.*

Proof Suppose that k_1 and k_2 diverge at some switch after following a common edge path of length $r' < r+1$. Let s_i be a switch point separating k_1 from k_2 at that switch, and follow the corresponding component z_i of z. If we backtrack from s_i we see that, for every tie t' located between t and s_i along k_1 and k_2, there is a point of $t' \cap z_i$ which separates the two points of k_1 and k_2 located on t'. But for $t' = t$ this would contradict our hypothesis that the starting points of k_1 and k_2 are in the same component of $t - z$. □

Lemma 12 defines a map $c\colon\ \mathcal{Z}_r \to 2^{\Gamma_r}$ as follows. Given $z \in \mathcal{Z}_r$, for each edge e of Φ and each component f of $e - z$, draw two arcs k and k' carried by Φ, starting from the same point in f but going in opposite directions, and each crossing $r+1$ edges before stopping. Lemma 12 shows that the edge path γ_f followed by $k \cup k'$ depends only on f. Then, $c(z) \in 2^{\Gamma_r}$ is the set of all edge paths γ_f as f ranges over all components of $e - z$ and e ranges over all edges of Φ.

If $\lambda \in P_{\gamma_1,\gamma_2,\dots,\gamma_n}$ crosses every edge of Φ, it follows from the construction of $z^r_\gamma = a(\lambda, \{\gamma_1, \gamma_2, \dots, \gamma_n\})$ that λ crosses every component f of $e - z^r_\gamma$ for every edge e of Φ. As a consequence, $c \circ a = b$, and the restriction $c_|\colon\ a\,(\mathcal{Q}_r) \to b(\mathcal{Q}_r) = \mathcal{P}_r$ is surjective.

The main conclusion of this is that the cardinal $\#\mathcal{P}_r$ of $\mathcal{P}_r$ is bounded by $\#a\,(\mathcal{Q}_r)$, and therefore by $\#\mathcal{Z}_r$ since $a\,(\mathcal{Q}_r) \subset \mathcal{Z}_r$.

Lemma 13 *If Φ has p switch points and q edges, then*

$$\#\mathcal{P}_r \leqslant \#\mathcal{Z}_r \leqslant 2^p p^{p+q} r^{p+q}.$$

Proof We already proved the first inequality.

Given an r–zipper family $z \in \mathcal{Z}_r$, we can consider the number n_e of times it crosses each edge e of Φ. For each edge e, the number n_e is bounded by the total number of edges crossed by z, counted with multiplicities, and this number is itself bounded by pr. Consequently, there are at most $(pr)^q$ possible assignments for these numbers n_e.

Now, for each edge e, $e \cap z$ consists of n_e disjoint parallel arcs transverse to the ties. Once we know the n_e, we can read some additional information from z, namely the location of its end points. More precisely, for each end point, we want to specify the component k of $e \cap z$, for some edge e of Φ, that contains it as well as which end point of k is the end point of z considered; note that there are $\sum_e n_e \leqslant pr$ such components of $e \cap z$, so that we have at most $2pr$ possible choices for each end point. Some of these end points are located at the switch points, and the arc of $e \cap z$ on which they sit is completely determined by the $n_{e'}$. There are at most p of the other end points, which leaves us with at most $(2pr)^p$ possible choices.

Once we know that each $e \cap k$ consists of n_e disjoint parallel arcs transverse to the ties, as well as the location of the end points of z, then z is completely determined up to isotopy of Φ respecting each tie. It follows that the number of elements of $\mathcal{Z}_r$ is bounded by $(pr)^q(2pr)^p = 2^p p^{p+q} r^{p+q}$ □

Lemma 13 concludes the proof of Proposition 10. □

Remark 14 The introduction of the set $\mathcal{Q}_r$ in the above proof may at first seem unnecessarily complicated. However, it is not hard to find examples of geodesic laminations carried by Φ which realize exactly the same edge paths of length $2r+1$ but whose associated r–zipper families are different. In other words, the map $a\colon \mathcal{Q}_r \to \mathcal{Z}_r$ does not necessarily factor through a right inverse $\mathcal{P}_r \to \mathcal{Z}_r$ for c.

Theorem 15 *The Hausdorff dimension of $(\mathcal{L}(S), d_{\mathrm{H}})$ is equal to* 0.

Proof By Lemma 8, we can write $\mathcal{L}(S)$ as the union of $\mathcal{L}^{\mathrm{s}}(\Phi_i)$ for finitely many fattened train tracks Φ_i. It therefore suffices to show that $(\mathcal{L}^{\mathrm{s}}(\Phi), d_{\mathrm{H}})$ has Hausdorff dimension 0 for an arbitrary fattened train track Φ.

Let Φ be such a fattened train track. By Lemma 9 and Proposition 10, there are constants a, b, c and $N > 0$ such that, for every integer $r \geqslant 1$, we can cover $\mathcal{L}^{\mathrm{s}}(\Phi)$ by at most cr^N sets of diameter bounded by $a\mathrm{e}^{-br}$. Therefore, for every $\varepsilon > 0$, we can cover $\mathcal{L}^{\mathrm{s}}(\Phi)$ by at most $c\left(1 + \frac{1}{b}\log\frac{a}{\varepsilon}\right)^N$ balls of radius ε (by taking r the smallest integer such that $a\mathrm{e}^{-br} < \varepsilon$). It immediately follows that the Hausdorff dimension of $(\mathcal{L}^{\mathrm{s}}(\Phi), d_{\mathrm{H}})$ is equal to 0. $\square$

For future reference, let us slightly improve the growth exponent in Lemma 13.

Lemma 16 *With the same data as in Lemma 13,*

$$\#\mathcal{P}_r \leqslant \#\mathcal{Z}_r \leqslant cr^{9|\chi(S)|-1}$$

for some constant $c > 0$, where $\chi(S)$ is the Euler characteristic of S.

Proof We first consider the subset $\mathcal{Z}'_r$ of $\mathcal{Z}_r$ consisting of those r–zipper families which have no switch connection. We will use the notation of the proof of Lemma 13.

For $z \in \mathcal{Z}'_r$, we again consider the number n_e of times z crosses the edge e of Φ. Since z has no switch connection, the sum $\sum_e n_e$ is equal to pr. The numbers n_e are far from being independent. Indeed, at each switch σ, the sum of the n_e corresponding to the edges e coming in on one side of σ is almost equal to the sum of the $n_{e'}$ corresponding to the edges going out on the other side of σ, with a small discrepancy coming from the end points of z that are located on σ; in particular, the discrepancy is bounded between $-2p$ and $+2p$.

If z and $z' \in \mathcal{Z}'_r$ have the same discrepancy type at the switches of Φ, then their associated families of crossing numbers n_e and n'_e are such that $n'_e - n_e$ now satisfy these switch relations with no discrepancy. The vector space $\mathcal{W}(\Phi)$ of all real edge weight systems satisfying the switch relations has dimension bounded by $3\,|\chi(S)|$; see for instance [19, Section 2.1] or [4, Theorem 15]. Since $\sum_e n_e = pr$ for $z \in \mathcal{Z}'_r$, we conclude that there exists a constant c_1 such that, within a given discrepancy type at the switches, the number of possible assignments of crossing numbers n_e for $z \in \mathcal{Z}'_r$ is bounded by $c_1 r^{3|\chi(S)|-1}$ for some constant c_1.

There are at most $(4p+1)^p$ discrepancy types at the switches of Φ. The same analysis of end point locations as in Lemma 13 then gives that $\#\mathcal{Z}'_r \leqslant c_1(4p+1)^p r^{3|\chi(S)|-1}(2pr)^p$. The number p of switch points is bounded by $6\,|\chi(S)|$ by a counting argument. It follows that $\#\mathcal{Z}'_r \leqslant c_2 r^{9|\chi(S)|-1}$.

Now, consider the subset $\mathcal{Z}_r''$ of $\mathcal{Z}_r$ consisting of those r–zipper families which have at least one switch connection. In this case, if $z \in \mathcal{Z}_r''$ crosses n_e times the edge e, we can only conclude that $\sum_e n_e \leqslant pr$, so that the number of possible assignments of crossing numbers n_e for $z \in \mathcal{Z}_r''$ is bounded by $c_3 r^{|3\chi(S)|}$ for some constant c_3. However, we now have to worry about 2 fewer end point locations, so that $\#\mathcal{Z}_r \leqslant c_3 r^{3|\chi(S)|}(2pr)^{p-2} \leqslant c_4 r^{9|\chi(S)|-2}$.

Since $\mathcal{Z}_r = \mathcal{Z}_r' \cup \mathcal{Z}_r''$, this concludes the proof. □

The estimate of Lemma 16 is stronger than that of Lemma 13 (where the growth exponent $p+q$ can be as large as $15\,|\chi(S)|$), but it is still quite crude.

4 The $d_{\log}$ metric

In a metric space (X,d), consider

$$d_{\log}(x,y) = \frac{1}{\left|\log\left(\min\left\{d(x,y), \frac{1}{4}\right\}\right)\right|}$$

for x, $y \in X$. In particular, $d_{\log}(x,y) = 1/\left|\log d(x,y)\right|$ when $d(x,y) \leqslant \frac{1}{4}$. As we will see in the proof of Proposition 17 below, the number 4 can be replaced by any number larger than 4.

Proposition 17 *The above formula defines a metric $d_{\log}$ on X, which induces on X the same topology as the original metric d.*

Proof The only non-trivial thing is to prove that $d_{\log}$ satisfies the triangle inequality. Because the function $f(u) = -1/\log u$ is increasing between 0 and 1, it suffices to show that f satisfies the inequality

$$f(u+v) \leqslant f(u) + f(v)$$

for every u, $v \in [0, \frac{1}{4}]$. For this, we will compute the minimum of the function

$$h(u,v) = f(u) + f(v) - f(u+v)$$

over the square $\left[0, \frac{1}{4}\right] \times \left[0, \frac{1}{4}\right]$, and show that it is equal to 0.

The critical points of h are the points (u,v) where $f'(u) = f'(v) = f'(u+v)$. However, the function $f'(u) = 1/\left(u \log^2 u\right)$ is decreasing on the interval $\left]0, \mathrm{e}^{-2}\right[$ and increasing on $\left]\mathrm{e}^{-2}, 1\right[$. In particular, f' is at most two-to-one on $]0,1[$. It follows that a critical point (u,v) of h must satisfy $u = v$. Solving the equation

$f'(u) = f'(2u)$ then shows that the only critical point of h in $\left]0, \frac{1}{4}\right[\times \left]0, \frac{1}{4}\right[$ is $\left(2^{-2-\sqrt{2}}, 2^{-2-\sqrt{2}}\right)$, where $h\left(2^{-2-\sqrt{2}}, 2^{-2-\sqrt{2}}\right) = (3 - 2\sqrt{2})/\log 2 > 0$.

On the boundary of the square, $h(u,0) = h(0,u) = 0$ while $h(u, \frac{1}{4}) = h(\frac{1}{4}, u) = g(u)$ with $g(u) = f(u) + f(\frac{1}{4}) - f(u + \frac{1}{4})$. From the graph of f', we see that g has a unique critical point, located between 0 and e^{-2}. Graphing g, we conclude that $g(u) \geqslant g(0) = g(\frac{1}{4}) = 0$ for every $u \in \left[0, \frac{1}{4}\right]$.

Therefore, the minimum of $h(u,v)$ over the square $\left[0, \frac{1}{4}\right] \times \left[0, \frac{1}{4}\right]$ is equal to 0. This proves that $f(u+v) \leqslant f(u) + f(v)$ for every u, $v \in [0, \frac{1}{4}]$, and concludes the proof that $d_{\log}$ satisfies the triangle inequality. □

Let us apply this to the case where (X, d) is the space $(\mathcal{L}(S), d_{\mathrm{H}})$ of geodesic laminations. This defines a new metric $d_{\log}$ on $\mathcal{L}(S)$.

This metric $d_{\log}$ is more intrinsic than d_{H}. Recall that two metrics d and d' on a space X are *Lipschitz equivalent* is there exists constants C_1, $C_2 > 0$ such that

$$C_1 d(x,y) \leqslant d'(x,y) \leqslant C_2 d(x,y)$$

for every x, $y \in X$.

Proposition 18 *The metric $d_{\log}$ associated to d_{H} on $\mathcal{L}(S)$ is, up to Lipschitz equivalence, independent of the choice of the metric m on S.*

Proof This immediately follows from the fact that d_{H} is well-defined up to Hölder equivalence (Lemma 7). □

In particular, the Hausdorff dimension of $(\mathcal{L}(S), d_{\log})$ is now well-defined, independently of the choice of a metric on S.

We have the following estimate for this Hausdorff dimension $\dim_{\mathrm{H}}(\mathcal{L}(S), d_{\log})$. This estimate is far from sharp, but at least proves that the dimension is strictly between 0 and ∞.

Proposition 19 *If the surface S is neither the three times punctured sphere, nor the twice punctured projective plane, nor the once punctured Klein bottle,*

$$2 \leqslant \dim_{\mathrm{H}}(\mathcal{L}(S), d_{\log}) \leqslant 9\,|\chi(S)| - 1$$

where $\chi(S) < 0$ is the Euler characteristic of S.

As indicated in the introduction, $\mathcal{L}(S)$ is finite when S is the three times punctured sphere, the twice punctured projective plane, or the once punctured Klein bottle; see for instance [6, Proposition 16] for a proof. In particular, $\mathcal{L}(S)$ has Hausdorff dimension equal to 0 for these exceptional surfaces.

Proof of Proposition 19 The second inequality is a by-product of the proof of Theorem 15. Indeed, for every integer $r \geqslant 1$, Lemmas 9 and 16 show that we can cover $\mathcal{L}(S)$ by at most $cr^{9|\chi(S)|-1}$ sets whose d_{H}–diameter is bounded by $a\mathrm{e}^{-br}$, for some constants a, b, $c > 0$. By definition of $d_{\log}$, the $d_{\log}$–diameter of these sets is bounded by $a'r^{-1}$ for some constant a' depending only on a and b. It immediately follows that $\dim_{\mathrm{H}}(\mathcal{L}(S), d_{\log}) \leqslant 9\,|\chi(S)| - 1$.

The first inequality is a consequence of [6], which itself uses Section 5 of the current paper. Indeed, if S is neither the three times punctured sphere, nor the twice punctured projective plane, nor the once punctured Klein bottle, then S contains a simple closed geodesic γ such that one component T of the surface obtained by splitting S open along γ is either a once punctured torus or a four times punctured sphere. Then, the set $\mathcal{L}_0(T)$, consisting of those geodesic laminations which are contained in the interior of T, is a natural subspace of $\mathcal{L}(S)$. Since we show in [6] that $\dim_{\mathrm{H}}(\mathcal{L}_0(T), d_{\log}) = 2$, this proves the first inequality. □

5 The combinatorial distance d_Θ

Having used fattened tracks so far, it is now convenient to switch to (unfattened) train tracks. A *train track* on the surface S is a graph Θ contained in the interior of S which consists of finitely many vertices, also called *switches*, and of finitely many edges joining them such that:

(1) the edges of Θ are differentiable arcs whose interiors are embedded and pairwise disjoint (the two end points of an edge may coincide);

(2) at each switch s of Θ, the edges of Θ that contain s are all tangent to the same line L_s in the tangent space T_sS and, for each of the two directions of L_s, there is at least one edge which is tangent to that direction;

(3) observe that the complement $S - \Theta$ has a certain number of spikes, each leading to a switch s and locally delimited by two edges that are tangent to the same direction at s; we require that no component of $S - \Theta$ is a disc with 0, 1 or 2 spikes or an open annulus with no spike .

A curve c *carried* by the train track Θ is a differentiable immersed curve in $c\colon I \to S$ whose image is contained in Θ, where I is an interval in $\mathbb{R}$. Such a curve is *bi-infinite* if its restriction to each component of $I - \{x\}$ has infinite length, where x is an arbitrary point in the interior of I.

In the definition of train tracks, the third condition is particularly crucial, as it has the following global corollaries (see for instance [5, Section 1.8] for proofs).

Lemma 20 *Let Θ be a train track in the surface S, and let $\widetilde{\Theta}$ be its preimage in the universal covering $\widetilde{S}$ of S. Then,*

1. *any curve carried by $\widetilde{\Theta}$ is embedded;*
2. *if two curves carried by $\widetilde{\Theta}$ coincide for a while and then diverge at some switch, they never meet again;*
3. *any bi-infinite curve which is carried by $\widetilde{\Theta}$ is quasi-geodesic in $\widetilde{S}$.* □

Consider a bi-infinite curve c carried by the train track Θ. By Lemma 20 and by the classical property of quasi-geodesics in negatively curved manifolds, there exists a unique geodesic g which, after a possible reparametrization of c, is homotopic to c by a homotopy which moves points by a uniformly bounded distance. In this situation, we will say that the geodesic g is *tracked by the curve* c carried by Θ.

The geodesic lamination λ is *weakly carried* by the train track Θ if every leaf of λ is tracked by a bi-infinite curve carried by Θ. Let $\mathcal{L}^{\mathrm{w}}(\Theta)$ denote the subset of $\mathcal{L}(S)$ consisting of those geodesic laminations which are weakly carried by Θ.

As the name indicates, a fattened train track Φ is a certain thickening of a train track Θ. Every geodesic lamination which is strongly carried by the fattened train track Φ is weakly carried by the train track Θ, so that $\mathcal{L}^{\mathrm{s}}(\Phi) \subset \mathcal{L}^{\mathrm{w}}(\Theta)$. Each point of view has its own technical advantages, and the framework of train tracks is somewhat better adapted to the context of the remainder of this paper. In particular, unlike $\mathcal{L}^{\mathrm{s}}(\Phi)$, $\mathcal{L}^{\mathrm{w}}(\Theta)$ does not depend on the choice of a negatively curved metric on S.

By Lemma 8, there exists finitely many train tracks $\Theta_1, \ldots, \Theta_n$ such that $\mathcal{L}(S)$ can be written as the union of the $\mathcal{L}^{\mathrm{w}}(\Theta_i)$. See [19] for an explicit such family of Φ_i. Our goal is to understand the metrics d_{H} and $d_{\log}$ on each of these $\mathcal{L}^{\mathrm{w}}(\Theta_i)$.

A bi-infinite curve carried by the train track Θ crosses a succession of oriented edges $\ldots, e_{-n}, \ldots, e_0, e_1, \ldots, e_n, \ldots$, in this order. The list

$\langle \dots, e_{-n}, \dots, e_0, e_1, \dots, e_n, \dots \rangle$ is the *bi-infinite edge path* followed by the curve c. If the geodesic g is tracked by a curve c on the train track Θ, the bi-infinite train track followed by c in Θ depends uniquely on g; see [5, Section 1.8]. A finite length edge path γ is *realized* by the geodesic lamination λ if it is contained in the bi-infinite edge path associated to a leaf of γ.

If λ is weakly carried by Θ and if $r \geqslant 1$ is an integer, we can consider the set of all the edge paths of length r of Θ which are realized by λ. Given two geodesic laminations λ, $\lambda' \in \mathcal{L}^{\mathrm{w}}(\Theta)$, we then define

$$d_\Theta(\lambda, \lambda') = \min\left\{\tfrac{1}{r+1}; \lambda \text{ and } \lambda' \text{ realize the same edge paths of length } r\right\}.$$

This $d_\Theta(\lambda, \lambda')$ is the *combinatorial distance* between λ and λ' in $\mathcal{L}^{\mathrm{w}}(\Theta)$.

Lemma 21 *The function d_Θ defines an ultrametric on the set $\mathcal{L}^{\mathrm{w}}(\Theta)$.*

Recall that an *ultrametric* on a set X is a metric d where the triangle inequality is replaced by the stronger condition that $d(x,z) \leqslant \max\{d(x,y), d(y,z)\}$ for every x, y, $z \in X$.

Proof of Lemma 21 The only non-trivial thing to prove is that $d_\Theta(\lambda, \lambda')$ can be 0 only if $\lambda = \lambda'$. (In particular, note that the ultrametric inequality is completely tautological.)

Suppose that $d_\Theta(\lambda, \lambda') = 0$, namely that λ and λ' realize exactly the same edge paths.

Let g be a leaf of λ. If γ is an edge path contained in the bi-infinite edge path realized by g, there consequently exists a leaf g'_γ of λ' whose associated bi-infinite edge path contains γ. Taking increasing larger edge paths γ and passing to a converging subsequence of leaves of λ, we conclude that there exists a leaf g' of λ' which follows exactly the same bi-infinite edge path as g. In this case, the two geodesics g and g' are homotopic by a homotopy which moves points by a uniformly bounded amount. Since the curvature of the metric m is negative, this implies that $g' = g$.

Since this property holds for every leaf g of λ, it follows that $\lambda \subset \lambda'$, and therefore that $\lambda = \lambda'$ by symmetry. □

6 Hausdorff distance and combinatorial distance

This section is devoted to proving Theorem 4 of the Introduction, which we restate here as:

Theorem 22 *Let Θ be a train track on the negatively curved surface S. Then, there exists two constants c_1, $c_2 > 0$ such that, for every geodesic laminations λ, λ' which are weakly carried by Θ,*

$$\mathrm{e}^{-c_1/d_\Theta(\lambda,\lambda')} \leqslant d_{\mathrm{H}}\left(\lambda, \lambda'\right) \leqslant \mathrm{e}^{-c_2/d_\Theta(\lambda,\lambda')}$$

We will split the proof of Theorem 4 into two lemmas, each devoted to one of the two inequalities. We begin with the easier one, namely the second inequality.

Lemma 23 *There exists constants c, $c' > 0$ such that*

$$d_{\mathrm{H}}(\lambda, \lambda') \leqslant c\,\mathrm{e}^{-c'/d_\Theta(\lambda,\lambda')}$$

for every λ, $\lambda' \in \mathcal{L}^{\mathrm{w}}(\Theta)$.

Proof The proof is in spirit very similar to that of Lemma 9, except that we now have to worry about quasi-geodesics. We can assume $\lambda \neq \lambda'$, since otherwise there is nothing to prove, and we can restrict attention to the case where $d_\Theta(\lambda, \lambda') \leqslant \frac{1}{2}$ since d_{H} is bounded. Consider the integer $r \geqslant 0$ such that $d_\Theta(\lambda, \lambda') = \frac{1}{2r+2}$ or $\frac{1}{2r+3}$. In particular, λ and λ' realize the same edge paths of length $2r+1$.

Consider a point $x \in \lambda$. Since $\lambda \in \mathcal{L}^{\mathrm{w}}(\Theta)$ is weakly carried by Θ, the leaf g of λ passing through x is tracked by a bi-infinite curve c carried by Θ. In particular, there is a homotopy from g to c which moves points by a distance bounded by a constant $c_1 > 0$ depending only on Θ and on the metric m. Let $y \in c$ be the image of $x \in g$ under this homotopy. Let a be an arc contained in c, containing y and crossing exactly $r+1$ edges in each direction when starting from y, and let γ be the edge path of length $2r+1$ that is followed by a. Let $b \subset g$ be the image of a by the homotopy from c to g.

Because $d_\Theta(\lambda, \lambda') \leqslant \frac{1}{2r+2}$, λ' also realizes the edge path γ. Namely, there is a leaf g' of λ' which is tracked by a bi-infinite curve c' carried by Θ, and an arc $a' \subset c'$ which follows γ.

The two arcs a and a' are homotopic by a homotopy which moves points by a distance $\leqslant c_2$, where c_2 depends only on the diameters of the edges of Θ, and

a' is also homotopic to an arc $b' \subset g'$ by a homotopy which moves points by a distance $\leqslant c_1$.

Now, the two geodesic arcs $b \subset g$ and $b' \subset g'$ are homotopic by a homotopy which moves their end points by at most $2c_1 + c_2$. In addition, by quasi-geodesicity of c, the point x is at distance at least $c_3 r$ from the two end points of b, for some constant $c_3 > 0$. Since the curvature of m is negative, it follows that there exists an $x' \in b'$ such that $(x, T_x b)$ is at distance at most $c_4 \mathrm{e}^{-c_5 r}$ from $(x', T_{x'} b')$ in the projective tangent bungle $PT(S)$, for some constants c_4, $c_5 > 0$ depending only on Θ and on the metric m.

Since this property holds for every $x \in \lambda$, we conclude that the lift $\widehat{\lambda}$ of λ in $PT(S)$ is contained in the $c_4 \mathrm{e}^{-c_5 r}$–neighborhood of the lift $\widehat{\lambda}'$ of λ'. Symmetrically, $\widehat{\lambda}$ of λ in $PT(S)$ is also contained in the $c_4 \mathrm{e}^{-c_5 r}$–neighborhood of $\widehat{\lambda}'$, and it follows that

$$d_{\mathrm{H}}(\lambda, \lambda') = d_{\mathrm{H}}(\widehat{\lambda}, \widehat{\lambda}') \leqslant c_4 \mathrm{e}^{-c_5 r} \leqslant c_6 \mathrm{e}^{-c_5 / 2 d_\Theta(\lambda, \lambda')}$$

for $c_6 = c_4 \mathrm{e}^{3c_5/2}$ This proves the inequality of Lemma 23. □

We now prove the first equality of Theorem 4. Its proof is somewhat more elaborate than that of the other inequality, and has a more topological flavor.

Lemma 24 *There exists constants c, $c' > 0$ such that*

$$d_{\mathrm{H}}(\lambda, \lambda') \geqslant c\, \mathrm{e}^{-c'/d_\Theta(\lambda, \lambda')}$$

for every λ, $\lambda' \in \mathcal{L}^{\mathrm{w}}(\Theta)$.

Proof It is convenient to arrange, by adding a few edges to Θ if necessary, that Θ is maximal for inclusion in the set of all train tracks. This is equivalent to the property that the complement $S - \Theta$ consists of finitely many triangles and finitely semi-open annuli (each containing a component of ∂S).

To simplify the notation, let the integer $r \geqslant 0$ be such that $d_\Theta(\lambda, \lambda') = \frac{1}{2r}$ or $\frac{1}{2r+1}$ (assuming $\lambda \neq \lambda'$ without loss of generality), and let $\delta = 2 d_{\mathrm{H}}(\lambda, \lambda')$.

Since $d_\Theta(\lambda, \lambda') > \frac{1}{2r+2}$ and exchanging the rôles of λ and λ' if necessary, there exists an edge path γ of length $2r+1$ which is realized by λ and not by λ'.

By definition, the fact that $\lambda \in \mathcal{L}^{\mathrm{w}}(\Theta)$ realizes γ means that there exists a leaf g which is tracked by a bi-infinite curve c carried by Θ, and an arc $a \subset c$ which follows the edge path γ. Pick a point $y \in c$ in the part of a corresponding to the $(r+1)$–th edge which it crosses, namely to the central edge of γ. By quasi-geodesicity, there is a homotopy from c to g which moves points by a distance bounded by c_1. Let $x \in g$ be the image of y under this homotopy.

Since $d_{\mathrm{H}}(\lambda, \lambda') < \delta$, there is an $x' \in \lambda'$, located in a leaf g' of λ', such that the distance from $(x, T_x g)$ to $(x', T_{x'} g')$ in $PT(S)$ is bounded by δ. The leaf g' is tracked by c' carried by Θ, and there is a homotopy from c' to g' which moves points by a distance bounded by c_1.

At this point, it is convenient to lift the situation to the universal covering $\widetilde{S}$ of S. Let $\widetilde{\Theta}$, $\widetilde{\lambda}$ and $\widetilde{\lambda}'$ be the respective preimages of Θ, λ and λ'. Lift x to a point $\widetilde{x}$, contained in a leaf $\widetilde{g}$ of $\widetilde{\lambda}$. Lift x' to a point $\widetilde{x}'$, contained in the leaf $\widetilde{g}'$ of $\widetilde{\lambda}'$, such that the distance from $(\widetilde{x}, T_{\widetilde{x}}\widetilde{g})$ to $(\widetilde{x}', T_{\widetilde{x}'}\widetilde{g}')$ in $PT\big(\widetilde{S}\big)$ is bounded by δ. Lift c and c' to bi-infinite curves $\widetilde{c}$ and $\widetilde{c}'$ carried by $\widetilde{\Theta}$, in such a way that they are respectively homotopic to $\widetilde{g}$ and $\widetilde{g}'$ by a homotopy which moves points by a distance bounded by c_1. Finally, the homotopy from $\widetilde{g}$ to $\widetilde{c}$ specifies preferred lifts $\widetilde{y}$, $\widetilde{a}$ of y, and a, as well as a lift of γ to an edge path $\widetilde{\gamma}$ of $\widetilde{\Theta}$.

Note that $\widetilde{g}$ and $\widetilde{c}$ have the same end points on the boundary at infinity $\partial_\infty \widetilde{S}$. Similarly, $\widetilde{g}'$ and $\widetilde{c}'$ have the same end points in $\partial_\infty \widetilde{S}$. Since the distance from $(\widetilde{x}, T_{\widetilde{x}}\widetilde{g})$ to $(\widetilde{x}', T_{\widetilde{x}'}\widetilde{g}')$ in $PT\big(\widetilde{S}\big)$ is bounded by δ, we conclude that each end point of $\widetilde{c}$ is seen from $\widetilde{x}$ within an angle of $\leqslant c_7\delta$ from an end point of $\widetilde{c}'$, where the constant c_7 depends only on c_1 and on the curvature of m.

We now split the argument into three cases.

Case 1 *The curve $\widetilde{c}'$ crosses the edge $\widetilde{e}_0$ of $\widetilde{\Theta}$ that contains the point $y \in \widetilde{c}$.*

By choice of γ, the curve $\widetilde{c}'$ does not realize the edge path $\widetilde{\gamma}$. Consequently, we can follow $\widetilde{c}$ and $\widetilde{c}'$ in one direction from $\widetilde{e}_0$ until they diverge at some switch $\widetilde{s}$ of $\widetilde{\Theta}$ after crossing at most $r+1$ edges of $\widetilde{\Theta}$. Orient $\widetilde{c}$ and $\widetilde{c}'$ in the direction from $\widetilde{e}_0$ to $\widetilde{s}$.

Since $\widetilde{c}$ and $\widetilde{c}'$ diverge at the switch $\widetilde{s}$, there is a spike of $\widetilde{S} - \widetilde{\Theta}$ which separates $\widetilde{c}$ from $\widetilde{c}'$ at $\widetilde{s}$. Let T be the component of $\widetilde{S} - \widetilde{\Theta}$ that contains that spike. By the maximality hypothesis on Θ, either T is a triangle or it is the preimage of a semi-open annulus component of $S - \Theta$.

We first analyze the case where T is a triangle. Consider the side of T that does not contain $\widetilde{s}$. Extend this side of T to a bi-infinite curve $\widetilde{c}''$ which is carried by $\widetilde{\Theta}$. By the second conclusion of Lemma 20, either $\widetilde{c}''$ remains disjoint from $\widetilde{c}$, or it meets $\widetilde{c}$ at one spike of T; in the second case, we can therefore assume that $\widetilde{c}''$ coincides with $\widetilde{c}$ after leaving the closure of T. Similarly, we can assume

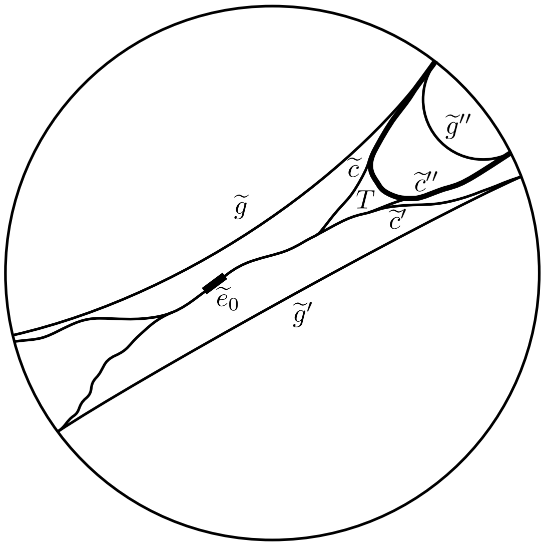

Figure 1: Diverging curves carried by the train track $\widetilde{\Theta}$

that either $\widetilde{c}''$ is disjoint from $\widetilde{c}'$ or it coincides with $\widetilde{c}'$ after leaving the closure T. See Figure 1. Again, $\widetilde{c}''$ is quasi-geodesic, and there is a geodesic $\widetilde{g}''$ which is homotopic to $\widetilde{c}''$ by a homotopy moving points by a distance bounded by c_1.

By construction, the distance from $\widetilde{y}$ to the geodesic $\widetilde{g}''$ is bounded by $c_8 r$ for a constant c_8 depending only on c_1, the length of the edges of Θ, and the diameter of the components of $S - \Theta$. It follows that the angle under which its end points are seen from $\widetilde{y}$ is at least $\mathrm{e}^{-c_9 r}$, where c_9 depends on c_8 and on a negative bound for the curvature of m. By construction, these end points separate the positive end points of $\widetilde{c}$ and $\widetilde{c}'$. We conclude that these positive end points are seen from $\widetilde{x}$ under an angle larger than $e^{-c_9 r}$.

On the other hand, splicing one half of $\widetilde{c}$ to one half of $\widetilde{c}'$ in $\widetilde{e}_0$, we can construct a curve $\widetilde{c}'''$ which is carried by $\widetilde{\Theta}$, which crosses $\widetilde{e}_0$, and which goes from the negative end point of $\widetilde{c}'$ to the positive end point of $\widetilde{c}$. By quasi-geodesicity of $\widetilde{c}'''$, it follows that the angle under which these two end points are seen from $\widetilde{x}$ is bounded from below by a constant $c_{10} > 0$.

However, we had concluded earlier that each end point of $\widetilde{c}$ is seen from $\widetilde{x}$ within an angle of $\leqslant c_7\delta$ from at least one end point of $\widetilde{c}'$. It follows that $e^{-c_9 r} \leqslant c_7\delta$ for r sufficiently large.

There remains to consider the case where T is the preimage of a semi-open annulus component of $S - \Theta$. Namely, T is an infinite strip containing a component $\widetilde{g}''$ of $\partial\widetilde{S}$. The geodesic $\widetilde{g}''$ separates the two positive end points of $\widetilde{c}$ and $\widetilde{c}'$, and its distance from $\widetilde{y}$ is again bounded by $c_8 r$ for a constant c_8 depending only on the length of the edges of Θ and on the diameter of the components of $S-\Theta$. The same argument as before then shows that $e^{-c_9 r} \leqslant c_7\delta$ for r sufficiently large in this case as well.

Case 2 *The curve $\widetilde{c}'$ does not cross the edge $\widetilde{e}_0$ containing $y \in \widetilde{c}$, but meets $\widetilde{c}$ at some edge $\widetilde{e}_1 \neq \widetilde{e}_0$.*

Orient $\widetilde{c}$ in the direction from $\widetilde{e}_1$ to $\widetilde{e}_0$, and orient $\widetilde{c}'$ so that its orientation coincides with that of $\widetilde{c}$ on $\widetilde{e}_1$.

The curve $\widetilde{c}'$ diverges from $\widetilde{c}$ at some point between $\widetilde{e}_1$ and $\widetilde{e}_0$. In particular, the direction in which it diverges specifies a side of $\widetilde{c}$. Among the two components of $\widetilde{S}-\widetilde{\Theta}$ which are adjacent to $\widetilde{e}_0$, let T be the one which is on this preferred side of $\widetilde{c}$. Again, either T is a triangle or it is the preimage of a semi-open annulus component of $S-\Theta$.

First consider the case where T is a triangle. Among the two sides of T which do not touch $\widetilde{e}_0$, the curve $\widetilde{c}'$ may meet one of them in $\widetilde{\Theta}$, but not both. Pick a side of T which is not followed by $\widetilde{c}'$ and extend it to a bi-infinite curve $\widetilde{c}''$ which is carried by $\widetilde{\Theta}$. As in Case 1, we can arrange that $\widetilde{c}''$ either is disjoint from $\widetilde{c}$ (resp. $\widetilde{c}'$), or coincides with this curve after it meets it.

This time the geodesic $\widetilde{g}''$ homotopic to $\widetilde{c}''$ passes at uniformly bounded distance from $\widetilde{x}$. Since it separates the positive end points of $\widetilde{c}$ and $\widetilde{c}'$, we conclude that the angle under which these two end points are seen from $\widetilde{x}$ is uniformly bounded from below by $c_{10} > 0$.

As in Case 1, the angle under which the positive end point of $\widetilde{c}$ and the negative end point $\widetilde{c}'$s are seen from $\widetilde{x}$ is bounded from below by $c_{10} > 0$.

We conclude in this case that $c_{10} < c_7\delta$, and therefore that $e^{-c_9 r} \leqslant c_7\delta$ for r sufficiently large.

The argument is similar when T is the preimage of a semi-open annulus component of $S-\Theta$, by considering the geodesic component $\widetilde{g}''$ of ∂S contained in T.

Case 3 *The curves $\widetilde{c}$ and $\widetilde{c}'$ never cross the same edge.*

The same argument as in Case 2, using the component T of $\widetilde{S}-\widetilde{\Theta}$ adjacent to $\widetilde{e}_0$ and on the same side as $\widetilde{c}'$, gives that $c_{10} < c_7\delta$, and therefore that $e^{-c_9 r} \leqslant c_7\delta$ for r sufficiently large.

Therefore, in all cases, $d_{\mathrm{H}}(\lambda, \lambda') = \frac{1}{2}\delta \geqslant \frac{1}{2}c_7^{-1}e^{-c_9 r} \geqslant \frac{1}{2}c_7^{-1}e^{-c_9/2 d_\Theta(\lambda,\lambda')}$ for r sufficiently large. This concludes the proof of Lemma 24. □

The combination of Lemmas 23 and 24 proves Theorem 22.

References

[1] **Joan S Birman**, **Caroline Series**, *Geodesics with bounded intersection number on surfaces are sparsely distributed*, Topology 24 (1985) 217–225

[2] **Francis Bonahon**, *Bouts des variétés hyperboliques de dimension* 3, Ann. of Math. 124 (1986) 71–158

[3] **Francis Bonahon**, *The geometry of Teichmüller space via geodesic currents*, Invent. Math. 92 (1988) 139–162

[4] **Francis Bonahon**, *Transverse Hölder distributions for geodesic laminations*, Topology 36 (1997) 103–122

[5] **Francis Bonahon**, *Closed curves on surfaces*, monograph in preparation, draft available at `http://math.usc.edu/~fbonahon`

[6] **Francis Bonahon**, **Xiaodong Zhu**, *The metric space of geodesic laminations on a surface II: small surfaces*, `arXiv:math.GT/0308268`

[7] **Richard D Canary**, **David B A Epstein**, **Paul Green**, *Notes on notes of Thurston*, from: "Analytical and geometric aspects of hyperbolic space (Coventry/Durham, 1984)", 3–92, London Math. Soc. Lecture Note Ser. vol. 111, Cambridge Univ. Press, Cambridge (1987)

[8] **Andrew J Casson**, **Steven A Bleiler**, *Automorphisms of surfaces after Nielsen and Thurston*, London Mathematical Society Student Texts 9, Cambridge University Press, Cambridge (1988)

[9] **Albert Fathi**, **François Laudenbach**, **Valentin Poénaru**, *Travaux de Thurston sur les surfaces*, Astérisque 66–67, Société Mathématique de France (1979)

[10] **Kenneth Falconer**, *Fractal geometry: mathematical foundations and applications*, John Wiley & Sons, Chichester (1990)

[11] **Mikhael L Gromov**, *Hyperbolic groups*, from: "Essays in group theory", 75–263, Math. Sci. Res. Inst. Publ. vol. 8, Springer, New York (1987)

[12] **Anatole Katok**, **Boris Hasselblatt**, *Introduction to the modern theory of dynamical systems*, Cambridge University Press, Cambridge (1995)

[13] **Erica Klarreich**, *Semiconjugacies between Kleinian group actions on the Riemann sphere*, Amer. J. Math. 121 (1999) 1031–1078

[14] **Howard A Masur**, **Yair N Minsky**, *Geometry of the complex of curves I. Hyperbolicity*, Invent. Math. 138 (1999) 103–149

[15] **Howard A Masur**, **Yair N Minsky**, *Geometry of the complex of curves II. Hierarchical structure*, Geom. Funct. Anal. 10 (2000) 902–974

[16] **Geoffrey Mess**, *Fuchsian groups with the same simple axes*, preprint IHES/M/90/16 (1990)

[17] **Yair N Minsky**, *Combinatorial and geometrical aspects of hyperbolic* 3*–manifolds*, `arXiv:math.GT/0205173`

[18] **James R Munkres**, *Topology: a first course*, Prentice-Hall, Englewood Cliffs (1975)

[19] **Robert C Penner**, **John L Harer**, *Combinatorics of train tracks*, Annals of Mathematics Studies vol. 125, Princeton University Press, Princeton (1992)

[20] **William P Thurston**, *The geometry and topology of three-manifolds*, lecture notes, Princeton University (1976-80)
`http://msri.org/publications/books/gt3m/`

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

[22] **William P Thurston**, *Minimal stretch maps between hyperbolic surfaces*, unpublished preprint (1986) `arXiv:math.GT/9801039`

[23] **Xiaodong Zhu**, *Fractal dimensions of the space of geodesic laminations*, doctoral dissertation, University of Southern California, Los Angeles (2000)