Closed Curves on Surfaces

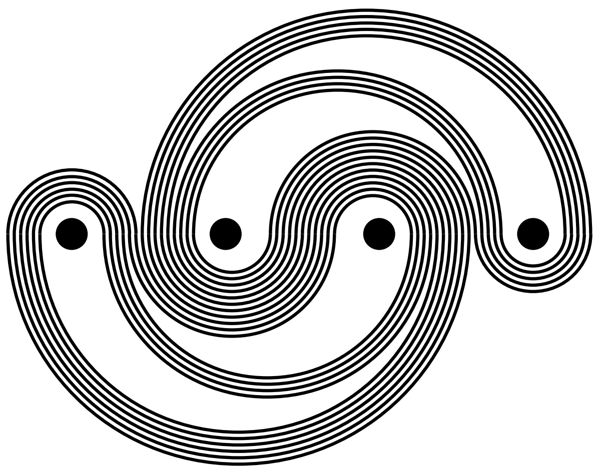

Francis Bonahon

Department of Mathematics
University of Southern California
Los Angeles, CA 90089-1113, U.S.A.
e-mail: fbonahon@math.usc.edu

16 November 2000

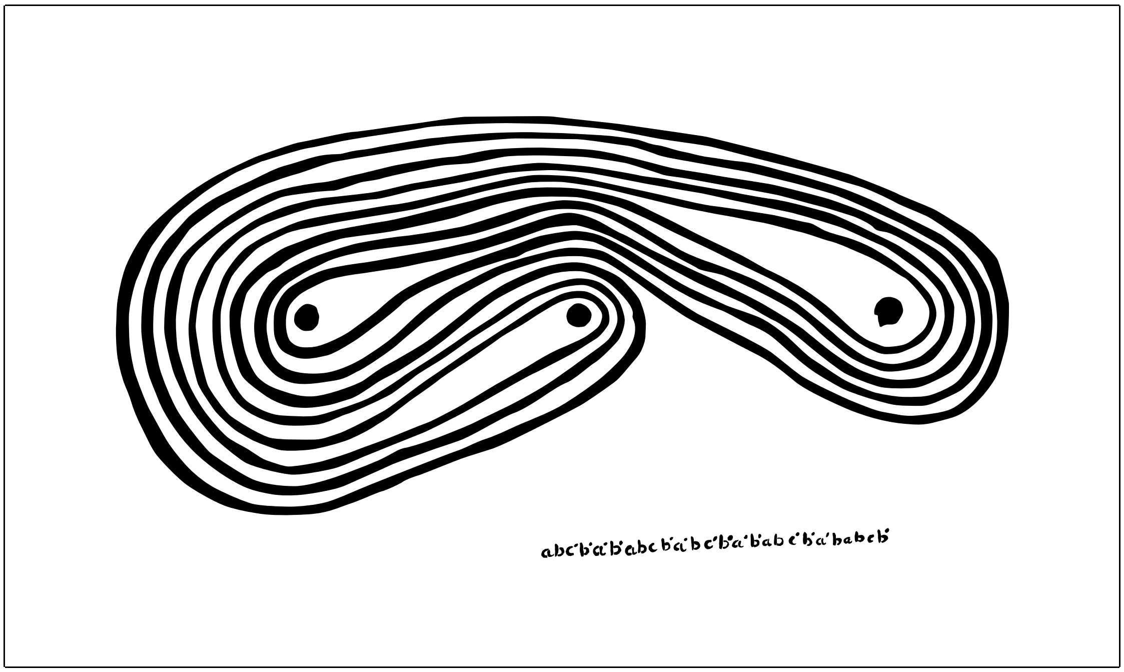

Mural in the Department of Mathematics of the University of California at Berkeley. The artwork is unsigned and undated, but a nearby mural on the same wall bears the date 12/9/71, and is signed DS BT.

INTRODUCTION

There will be an introduction here, at some point.

TABLE OF CONTENTS

CHAPTER I

GEODESIC LAMINATIONS

1.1. An example

To motivate the notion of geodesic laminations, let us consider a piece of experimental mathematics. On a surface S, take a simple closed curve γ and a diffeomorphism $\varphi : S \to S$. Then, draw the iterates $\varphi(\gamma)$, $\varphi^2(\gamma)$, ... , $\varphi^n(\gamma)$, ... , considered up to isotopy, and see what happens. Here is a specific example.

The surfaces that are easiest to visualize are subsets of the plane $\mathbb{R}^2$. Choose four points x_1, x_2, x_3, x_4 in this order in $0\times\mathbb{R} \subset \mathbb{R}^2$, and let S be the complement $\mathbb{R}^2 - \{x_1, x_2, x_3, x_4\}$, as in Figure 1.1(a). We will consider the diffeomorphism $\varphi : S \to S$ described as follows.

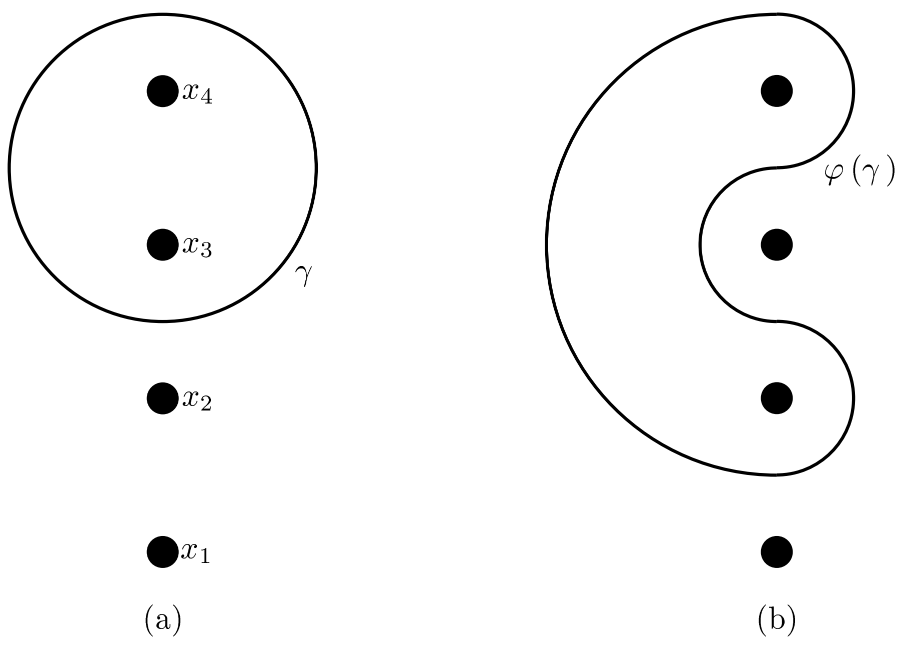

FIGURE 1.1. The curve γ and its image under φ

For $i = 1$, 2 or 3, let D_i be a disk in $\mathbb{R}^2$ which contains x_i and x_{i+1} in its interior, and which is disjoint from the other two x_j. Let the diffeomorphism $\sigma_i : S \to S$ fix the complement of D_i and perform a right-handed half-twist in D_i, as indicated in Figure 1.2. Up to isotopy, such a σ_i depends only on i, and

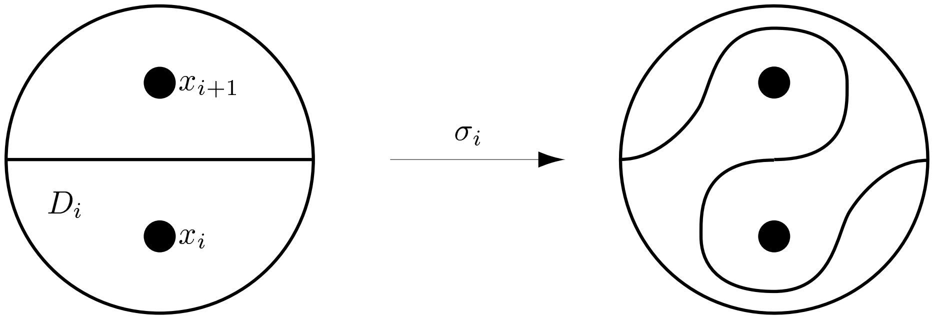

FIGURE 1.2

the reader will recognize here one of the standard generators of the braid group on 4 strings $B_4 = \pi_0 \mathrm{Diff}^+(S)$; compare [**Bir**].

Consider the diffeomorphism $\varphi = \sigma_2^{-1}\sigma_1\sigma_3$, well defined up to isotopy. If γ is the curve represented on Figure 1.1(a) then, after rearrangement by an isotopy, $\varphi(\gamma)$, $\varphi^2(\gamma)$ and $\varphi^3(\gamma)$ are represented on Figures 1.1 and 1.3. Looking at these pictures, two things become clear. First of all, the two curves $\varphi^2(\gamma)$ and $\varphi^3(\gamma)$ look very much alike, except that $\varphi^3(\gamma)$ is between 3 and 4 times as long as $\varphi^2(\gamma)$. The other thing is that, if this increase in complexity continues, it will not be very easy to draw $\varphi^4(\gamma)$, $\varphi^5(\gamma)$, ... if we only use brute force.

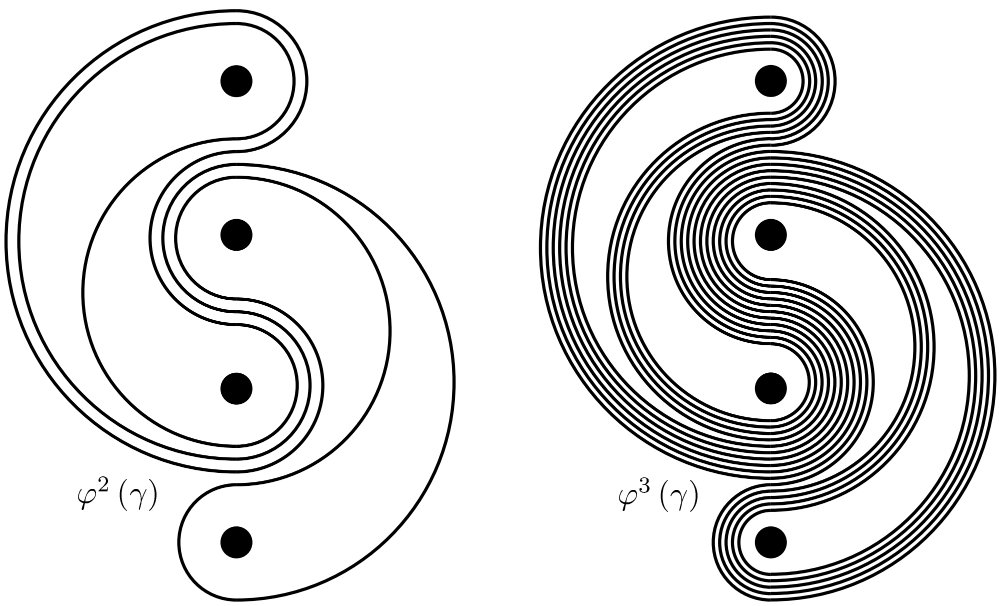

FIGURE 1.3. Two more iterates of the curve γ

To determine $\varphi^4(\gamma)$, it is convenient to use a short hand notation for $\varphi^3(\gamma)$. By squeezing together parallel strands of $\varphi^3(\gamma)$, we can make this curve look like the graph Θ of Figure 1.4, where the edge of Θ labelled by a corresponds to 6 parallel strands of γ, the edge labelled by b to 4 strands, the edge labelled by c to 11 strands, the edge labelled by d to 4 strands and the edge labelled by e to 5 strands.

Up to isotopy, we can recover the curve $\varphi^3(\gamma)$ from Θ by replacing each edge

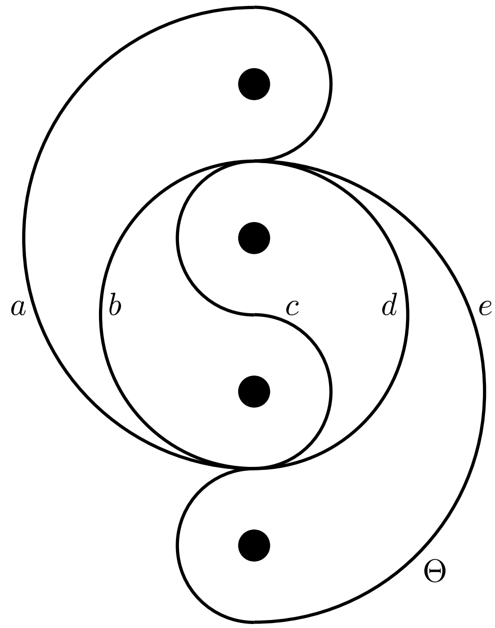

FIGURE 1.4

of Θ by the appropriate number of parallel strands, and suitably connecting them near the two vertices of Θ. Therefore, $\varphi^4(\gamma)$ can be obtained by replacing each edge of $\varphi(\Theta)$ by a certain number of parallel strands, and by suitably connecting them. This image $\varphi(\Theta)$ is represented on Figure 1.5(a). From Figures 1.5(a) and (b), we see that $\varphi(\Theta)$ can also be collapsed to Θ. It then becomes apparent that $\varphi^4(\gamma)$ can also be represented by replacing by $6+4+11 = 21$ parallel strands the edge of Θ labelled in Figure 1.4 by a, by $6+4+5 = 15$ parallel strands the edge labelled by b, by $2\times 6+4+11+4+2\times 5 = 41$ parallel strands the edge labelled by c, by $6+4+5 = 15$ parallel strands the edge labelled by d and by $11+4+5 = 20$ parallel strands the edge labelled by e.

We can clearly iterate the process to determine $\varphi^5(\gamma)$, $\varphi^6(\gamma)$, ... More generally, given non-negative integers a, b, c, d, $e \in \mathbb{N}$, we can replace each edge of Θ by the corresponding number of parallel strands. The graph Θ has an additional structure at its vertices, namely all edges are tangent to a given line at each vertex. If, in addition, $b+c = a+d+e$ and $a+b+e = c+d$, this determines a unique way to connect these strands at the two vertices without introducing any crossing point or any corner. We obtain in this way a ***simple multicurve***, namely a 1–dimensional submanifold of S, not necessarily connected. This multicurve $C_{(a,b,c,d,e)}$ is uniquely determined by a, b, c, d, e up to isotopy. For instance, $\varphi^3(\gamma) = C_{(6,4,11,4,5)}$. Also, $\varphi(\gamma) = C_{(1,0,1,0,0)}$ and $\varphi^2(\gamma) = C_{(2,1,3,1,1)}$. It can be shown that the original curve γ is not of the form $C_{(a,b,c,d,e)}$ for any integers a, b, c, d, e.

What Figure 1.5 shows is that

$$\varphi\left(C_{(a,b,c,d,e)}\right) = C_{\Phi(a,b,c,d,e)},$$

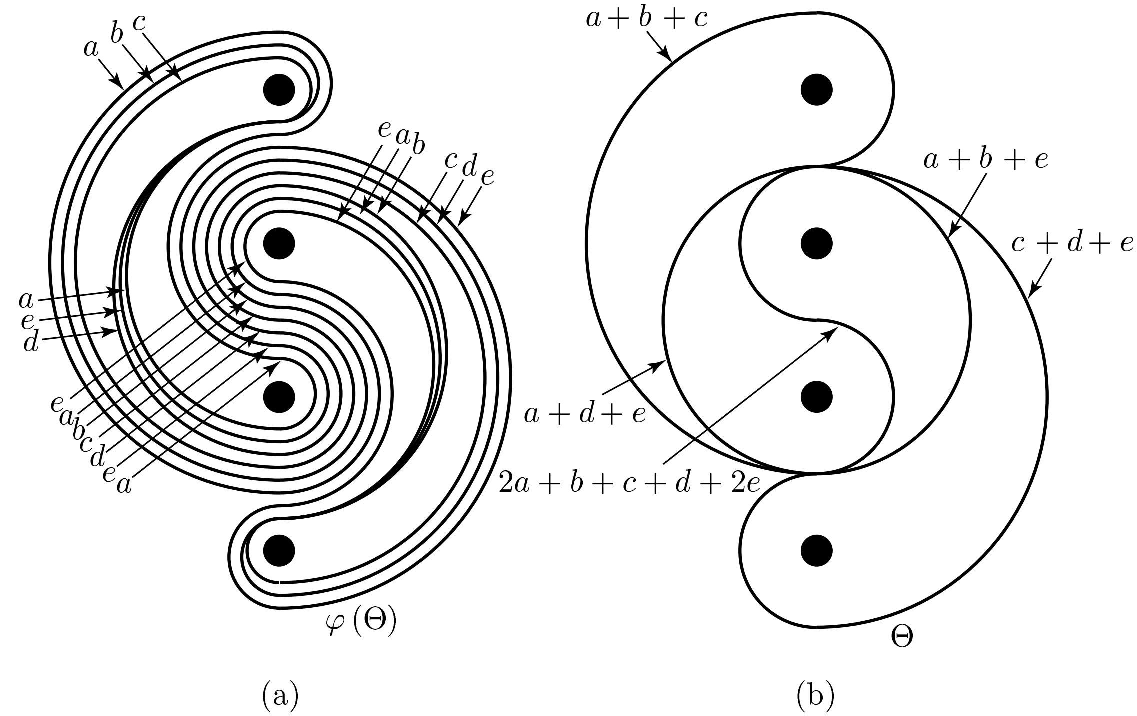

FIGURE 1.5

where Φ is the linear map defined by

$$\Phi\left(a, b, c, d, e\right) = \left(a + b + c, a + d + e, 2a + b + c + d + 2e, a + b + e, c + d + e\right).$$

In particular,

$$\varphi^n\left(\gamma\right) = C_{\Phi^{n-1}(1,0,1,0,0)}$$

for every $n \geqslant 1$.

We now give a heuristic argument for why the shape of $\varphi^n\left(\gamma\right)$ converges to a limit configuration. For any integer $n > 0$, the multicurve $C_{(na,nb,nc,nd,ne)}$ just consists of n parallel copies of $C_{(a,b,c,d,e)}$. Since these n parallel copies can be drawn so close to each other that they are undistinguishable, the two multicurves $C_{(na,nb,nc,nd,ne)}$ and $C_{(a,b,c,d,e)}$ have the same shape, in an intuitive sense since this notion is so far mathematically undefined. More generally, the two multicurves $C_{(a,b,c,d,e)}$ and $C_{(a',b',c',d',e')}$ have the same shape when there exists a number $\lambda \in \mathbb{Q}$ such that $(a, b, c, d, e) = \lambda\left(a', b', c', d', e'\right)$, namely when (a, b, c, d, e) and (a', b', c', d', e') represent the same class in the projective space $\mathbb{RP}^4$. Therefore, the shape of $C_{(a,b,c,d,e)}$ depends only on the class of (a, b, c, d, e) in $\mathbb{RP}^4$.

We can now apply to Φ the Perron-Frobenius Theorem (see for instance [**BerP**]):

LEMMA 1.1. *Let A be a $p \times p$–matrix with real non-negative entries, such that some power A^n, $n \in \mathbb{N}$, has all its entries positive. Then, the largest modulus eigenvalue λ of A is real positive, and its eigenspace is generated by a vector*

X_λ with non-negative coordinates. In addition, for every vector $X \neq 0$ with non-negative coefficients, the class of $A^n X$ in $\mathbb{RP}^{p-1}$ converges to the class of X_λ as n tends to ∞. □

The matrix of Φ has non-negative entries, and all the entries of the matrix of Φ^2 are positive. An elementary computation gives $\lambda = 2+\sqrt{3}$ as the largest modulus eigenvalue of Φ, with eigenvector $X_\lambda = \big(1, \sqrt{3}-1, 2, \sqrt{3}-1, 1\big)$. From Lemma 1.1, we conclude that the class of $\Phi^{n-1}(1,0,1,0,0)$ in $\mathbb{RP}^4$ converges to the class of $\big(1, \sqrt{3}-1, 2, \sqrt{3}-1, 1\big)$ as n tends to ∞.

Therefore, as n tends to ∞, the shape of the curve $\varphi^n(\gamma) = C_{\Phi^{n-1}(1,0,1,0,0)}$ should converge to the shape of $C_{\left(1,\sqrt{3}-1,2,\sqrt{3}-1,1\right)}$.

Of course, there are two problems with this statement. First of all, we did not give any mathematical definition of these notions of shape and shape convergence. In addition, because the coordinates $\big(1, \sqrt{3}-1, 2, \sqrt{3}-1, 1\big)$ are not integers, the type of geometric object represented by $C_{\left(1,\sqrt{3}-1,2,\sqrt{3}-1,1\right)}$ is unclear. The notion of geodesic lamination was introduced in part to deal with these problems.

We would like to add a few comments before closing this section.

First, there is nothing special about the curve γ. Indeed, for every simple closed curve γ' which is not homotopic to 0 in S, the curve $\varphi^n(\gamma')$ is isotopic to some curve $C_{(a,b,c,d,e)}$ for n sufficient large, with the (a,b,c,d,e) in $\mathbb{RP}^4$ close to the class of $\big(1, \sqrt{3}-1, 2, \sqrt{3}-1, 1\big)$. This is a consequence of Thurston's theory of surface diffeomorphisms; see [**Thu**][**FLP**].

Also, the same phenomenon can be observed for 'most' surface diffeomorphisms φ. This reflects the property that 'most' surface diffeomorphisms can be isotoped to pseudo-Anosov homeomorphisms, a fact which, although difficult to formulate precisely, is definitely in agreement with all experiments.

Exercise. Perform some of these experiments. First, for the above diffeomorphism $\varphi = \sigma_2^{-1}\sigma_1\sigma_3$ of the 4 times punctured plane, choose several examples of curves γ' and check that $\varphi^n(\gamma')$ is isotopic to some curve $C_{(a,b,c,d,e)}$ for n sufficient large. Then, take a 'random' diffeomorphism φ of an n–punctured plane (or a more general surface S) and a curve γ in S; draw several iterates $\varphi^n(\gamma)$ and try to guess a graph Θ as in Figure 1.4 such that, for n large enough, $\varphi^n(\gamma)$ is described by a set of integer weights associated to the edges of Θ as above and so that $\varphi(\Theta)$ collapses to Θ as in Figure 1.5.

1.2. Geodesic laminations

In §1.1, we observed that the simple closed curves $\varphi^2(\gamma)$ and $\varphi^3(\gamma)$ looked very much alike in Figure 1.3. However, these two curves were only defined up to isotopy, and we could equally well have drawn them in such a way that they do not look alike at all. It turns out that we used an implicit normalization in Figure 1.3, by drawing the two curves as a union of semi-circles centered on the axis $0 \times \mathbb{R}$. The fact that the two curves look similar for this normalization does reflect an intrinsic property of the corresponding isotopy classes. Such a normalization is very convenient to study simple curves on punctured planes, but it does not easily generalize to other surfaces. We will therefore use another type of normalization to achieve the same goals.

First of all we will restrict attention to compact surfaces S, possibly non-orientable and/or with non-empty boundary. This is not a great loss of generality. For instance, in the example of §1.1, this amounts to replacing our original surface by the complement of 4 small disjoint open disks in a large closed disk. Among such compact surfaces S, those with non-negative Euler characteristic $\chi(S)$ have only finitely many isotopy classes of simple closed curves, with the exception of the torus. For the torus, isotopy classes of simple closed curves are classified by their (rational) slope, and are therefore well understood; see for instance [**Rol**]. Therefore, we can restrict attention to compact surfaces S with $\chi(S) < 0$.

The hypothesis that $\chi(S) < 0$ is equivalent to the existence of a Riemannian metric of negative curvature on S for which the boundary ∂S is geodesic. Fix such a metric m_0.

Then, each closed curve γ is homotopic to a closed m_0–geodesic γ^*, unless it is homotopic to 0, and this geodesic is unique because m_0 has negative curvature; see for instance [??]. If, in addition, γ is a simple closed curve, then the closed geodesic γ^* is also simple and is isotopic to γ, unless we are in the following exceptional case: the surface S is non-orientable and the curve γ bounds a Möbius strip in S; in this case, γ^* wraps twice arounds the closed geodesic homotopic to the core of the Möbius strip (see for instance [**Eps**], etc...). Let us call a simple closed curve γ in S ***indivisible*** if it does not bound a disk or a Möbius strip in S. We just observed that an indivisible simple closed curve in S is isotopic to a unique simple closed m_0–geodesic.

In this way, we have established a one-to-one correspondence between the set of isotopy classes of indivisible simple closed curves on one hand, and the set $\mathcal{S}(S)$ of simple closed m_0–geodesics on the other hand. What this accomplishes is to get rid of the isotopy undetermination.

The set $\mathcal{S}(S)$ is contained in the set of all closed subsets of S, which carries a natural metric called the ***Hausdorff metric***. The Hausdorff distance $d_H(C, C')$ between two closed subsets C and C' is defined as the smallest ε such that C' is contained in the ε–neighborhood of C and C is contained in the ε–neighborhood of C'. It is relatively easy to check that this indeed defines a metric; see for instance [**Mun**, §??].

We have now endowed the set of isotopy classes of indivisible simple closed curves, identified to $\mathcal{S}(S)$, with a metric d_H. In the example of §1.1, it can be shown that the $\varphi^n(\gamma)$ form a Cauchy sequence for d_H; this is proved as Exercise ?? in Chapter ??. This gives a sense to our empiric observation that the $\varphi^n(\gamma)$ 'look alike' for n large enough and seem to converge to some limit shape. To understand this limit configuration, we have to consider the closure of $\mathcal{S}(S)$ in the space of all closed subsets of S.

Let an m_0–***geodesic lamination*** be a closed subset λ of S which can be decomposed as the union of a family of disjoint simple m_0–geodesics. Here, a simple geodesic is one that does not intersect itself, but may be closed or infinite. Among these geodesics, we allow components of the boundary ∂S, but not geodesics which transversely hit the boundary.

Given such a decomposition of λ as a union of disjoint simple geodesics, these geodesics are the ***leaves*** of this decomposition. We will see in §1.4 that this decomposition of λ is unique, and these geodesics will eventually be called the leaves of the geodesic lamination λ.

Let $\mathcal{L}(S)$ denote the set of geodesic laminations of S, endowed with the Hausdorff metric d_H.

PROPOSITION ??. *The space $\mathcal{L}(S)$ of all m_0–geodesic laminations is compact.*

PROOF. We will take a longer route to prove this, by a side trip through the ***projective tangent bundle*** $PT(S)$ of S, consisting of all pairs (x, l) where $x \in S$ and l is a line through the origin in the tangent space T_xS. Note that $PT(S)$ is a fiber bundle over S with fiber the projective line $\mathbb{RP}^1 \cong S^1$. In particular, $PT(S)$ is compact.

An immersed curve α in S has a well defined lift to $PT(S)$, by considering the tangent line l of α at each $x \in S$. In particular, a geodesic lamination λ determines a subset $\widehat{\lambda}$ of $PT(S)$, defined as the union of the lifts of all the leaves of the decomposition of λ.

The metric m_0 induces a metric on $PT(S)$. A geometric estimate provides a constant C, such that, if (x, l), $(x', l') \in PT(S)$ and if the m_0–geodesic passing through x and tangent to l is disjoint from the m_0–geodesic passing through x' and tangent to l', then $d((x, l), (x', l')) \leqslant C d(x, x')$; see ??. This proves that the subset $\widehat{\lambda}$ associated to the geodesic lamination λ is closed in $PT(S)$.

In the space of closed subsets of $PT(S)$, let $\mathcal{L}_{PT}(S)$ be the subset of those $\widehat{\lambda}$ associated to geodesic laminations $\lambda \in \mathcal{L}(S)$. The map $\lambda \mapsto \widehat{\lambda}$ defines a natural bijection between $\mathcal{L}(S)$ and $\mathcal{L}_{PT}(S)$, but we can now endow $\mathcal{L}_{PT}(S)$ with the Hausdorff metric d_H^{PT} coming from the metric of $PT(S)$.

LEMMA. *The space $\mathcal{L}_{PT}(S)$ is compact.*

PROOF. A classical property of the Hausdorff metric is that the space of closed subsets of a copmpact space is compact; see for instance [**Mun**, §??]. Therefore, it suffices to show that $\mathcal{L}_{PT}(S)$ is closed in the space of closed subsets of $PT(S)$.

Consider a sequence $\widehat{\lambda}_n \in \mathcal{L}_{PT}(S)$, $n \in \mathbb{N}$, corresponding to a sequence of geodesic laminations$\widehat{\lambda}_n \in \mathcal{L}_{PT}(S)$, which converges to some closed subset Λ of $PT(S)$ for the Hausdorff metric d_H^{PT}. We want to show that Λ is the lift $\widehat{\lambda}$ of a geodesic lamination λ.

By definition, every $(x, l) \in \Lambda$ if the limit of a sequence $(x_n, l_n) \in PT(S)$, $n \in \mathbb{N}$, such that each (x_n, l_n) is in $\widehat{\lambda}_n$. For each (x_n, l_n), there is a unique geodesic g_n of S which passes through x_n and is tangent to the line l_n at this point. By construction, g_n is a leaf of λ_n, and its lift $\widehat{g}_n$ is completely contained in $\widehat{\lambda}_n$. We conclude that, if g is the geodesic tangent to l at x, its lift $\widehat{g}$ is completely contained in Λ.

This proves that Λ is a union of lifts $\widehat{g}$ as g ranges over a family Γ of geodesics of S. Note that no two geodesics g and $g' \in \Gamma$ can intersect transversely at some point $x \in S$. Otherwise, by approximating the two lifts $\widehat{g}$, $\widehat{g}' \subset \Lambda$ of g, g' by the lifts $\widehat{g}_n$, $\widehat{g}'_n$ of two leaves g_n, g'_n of λ_n, we would have two leaves g_n, g'_n of λ_n which intersect transversely near x, contradicting the fact that λ_n is a geodesic lamination. Therefore, the image λ of Λ by the projection $PT(S) \to S$ is a union of disjoint simple geodesics. It is also closed since Λ is compact. It follows that λ is a geodesic lamination, whose lift $\widehat{\lambda}$ is equal to Λ by construction.

This concludes the proof that $\mathcal{L}_{PT}(S)$ is closed in the space of closed subsets of $PT(S)$, which implies that $\mathcal{L}_{PT}(S)$ is compact. □

The map $PT(S) \to S$ is distance non-increasing. It follows that the map $\mathcal{L}_{PT}(S) \to \mathcal{L}(S)$ is also distance non-increasing for the metrics d_H^{PT} and d_H, and in particular continuous. It follows that $\mathcal{L}(S)$ is compact. □

Note that the continuous bijection $\mathcal{L}_{PT}(S) \to \mathcal{L}(S)$ must be a homeomorphism since $\mathcal{L}_{PT}(S)$ is compact. As a side benefit of Lemma ??, we therefore obtain the following result, which is of interest by itself and will often be useful.

Proposition. *The metrics d_H and d_H^{PT} induce the same topology on $\mathcal{L}(S) = \mathcal{L}_{PT}(S)$. In particular, if $\lambda_n \in \mathcal{L}(S)$ converges to λ in $\mathcal{L}(S)$ and if $x_n \in \lambda_n$ converges to $x \in \lambda$ in S, then the direction l_n tangent to the leaf of λ_n at x_n converges to the direction l tangent to the leaf of λ at x.*

Proof. We already discussed the first statement. For the second statement note that, if a subsequence of the sequence (x_n, l_n) converges to some (x', l') in $PT(S)$, necessarily (x', l') is in the lift $\widehat{\lambda}$ since $\widehat{\lambda}_n$ converges to $\widehat{\lambda}$, and $x' = x$ by continuity, from which we conclude that l' must be the direction l tangent to the leaf of λ at x. Since this holds for every converging subsequence, the result follows. □

A corollary of Proposition ?? is that, if a sequence of simple closed geodesic has a Hausdorff limit, this limit must be a geodesic lamination. Therefore, the completion of the space $\mathcal{S}(S)$ of simple closed geodesics for the Huadorff metric d_H is contained in the space $\mathcal{L}(S)$ of geodesic laminations.

Not every geodesic lamination is a limit of simple closed geodesics for the Hausdorff metric. We will characterize limits of simple closed geodesics in §??.

1.3. Examples of geodesic laminations

We want to give a few examples of geodesic laminations, of varying complexity. At this point, we do not yet have the technology to completely justify some of these constructions. As a consequence, the reader will be asked to accept by an act of faith that some of the phenomena described in this section do really occur. The rigorous justifications will be given later, mostly in §1.??. However, we thought that the reader would find it useful to have several examples to keep in mind when reading our general analysis of geodesic laminations in later sections.

An example with finitely many leaves. The first example of an m_0–geodesic lamination is the one which we have already encountered, namely one

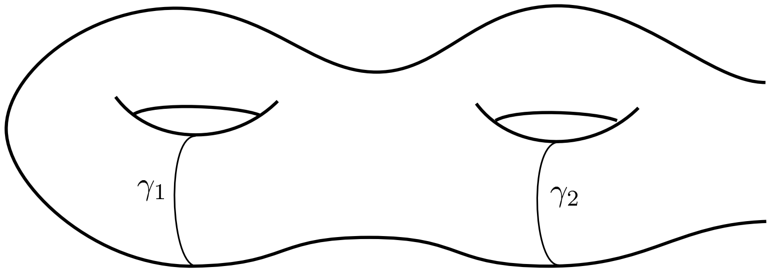

FIGURE 1.6. A geodesic lamination consisting of two closed leaves.

simple closed m_0–geodesics or, more generally, a finite union of disjoint simple closed m_0–geodesics.

Another example with finitely many leaves. We can enlarge the first example as follows. Let γ_1 and γ_2 be two disjoint simple closed geodesics in S, and let c be a simple bi-infinite curve, disjoint from γ_1 and γ_2, such that one end of c spirals around γ_1 and the other end spirals around γ_2; see Figure 1.6. In §1.8, we will show that we can pull c tight to get a simple geodesic γ which is disjoint from γ_1 and γ_2 and whose ends spiral around γ_1 and γ_2, respectively. The union $\lambda = \gamma \cup \gamma_1 \cup \gamma_2$ forms a geodesic lamination with an infinite leaf.

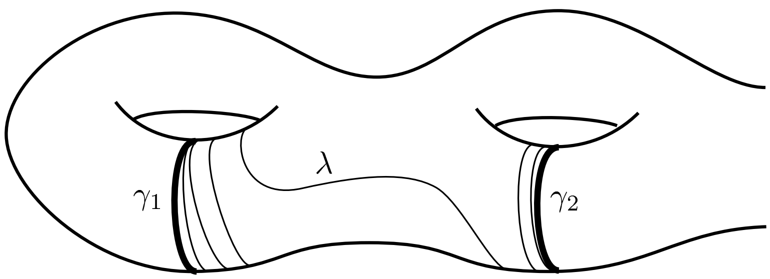

FIGURE 1.6. A geodesic lamination consisting of two closed leaves and one infinite leaf.

This example can be extended by adding several of these spiralling leaves. For instance, if S is a 'pair of pants', namely a disk minus two disjoint open disks, there are several geodesic laminations of S consisting of the boundary ∂S together with 3 infinite simple geodesics spiralling along this boundary. We have illustrated two of these in Figure 1.7. Note that the complement of these geodesic laminations consist of three infinite triangles, each bounded by the three infinite geodesics.

EXERCISE. Show that the pair of pants admits exactly 159 non-empty geodesic laminations, 32 of which consist of the boundary together with 3 spiralling geodesics as in Figure 1.7. (Hint: A simple geodesic must, either be a boundary component, or spiral around one or two boundary components.)

An example with uncountably many leaves. We will now construct an example with uncountably many leaves. Given two positive real numbers a, $a' > 0$, we glue the vertical sides of two rectangles $R \cong [0,1] \times [0,a]$ and $R' \cong [0,1] \times$

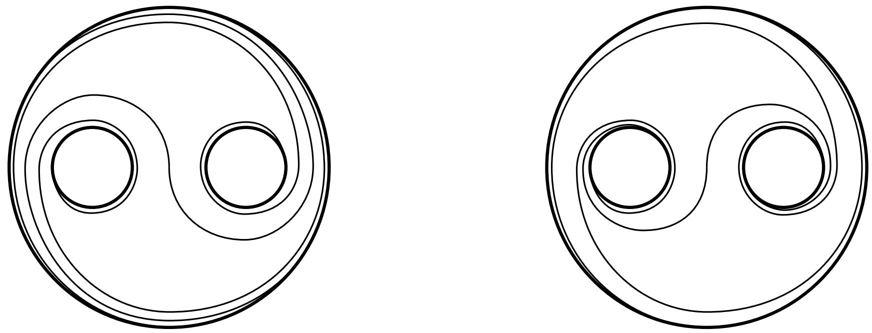

FIGURE 1.7. Two geodesic laminations on a pair of pants.

$[0, a']$ to the interval $I \cong [0, a + a']$ as follows: Using the obvious translation maps, we identify the side $\{0\} \times [0, a]$ of R to the interval $[0, a] \subset I$, $\{1\} \times [0, a] \subset R$ to $[a', a + a'] \subset I$, $\{0\} \times [0, a'] \subset R'$ to $[a, a + a'] \subset I$, and $\{1\} \times [0, a'] \subset R'$ to $[0, a'] \subset I$. The corresponding identification space Φ is homeomorphic to a torus minus an open disk, and is represented on Figure 1.??(a) and (b) in the case where $a > a'$. We are using in Figure 1.??(a) the usual presentation of the torus as a square with opposite sides identified, and a more global picture in Figure 1.??(b). Let S be the differentiable surface with boundary obtained by removing a smaller disk from the same torus, so that Φ is contained in the interior of S and so that the complement $S - \Phi$ topologically is a semi-open annulus. In Figure 1.??(a), S is obtained from the torus by removing a small disk around the point corresponding to the corners of the square.

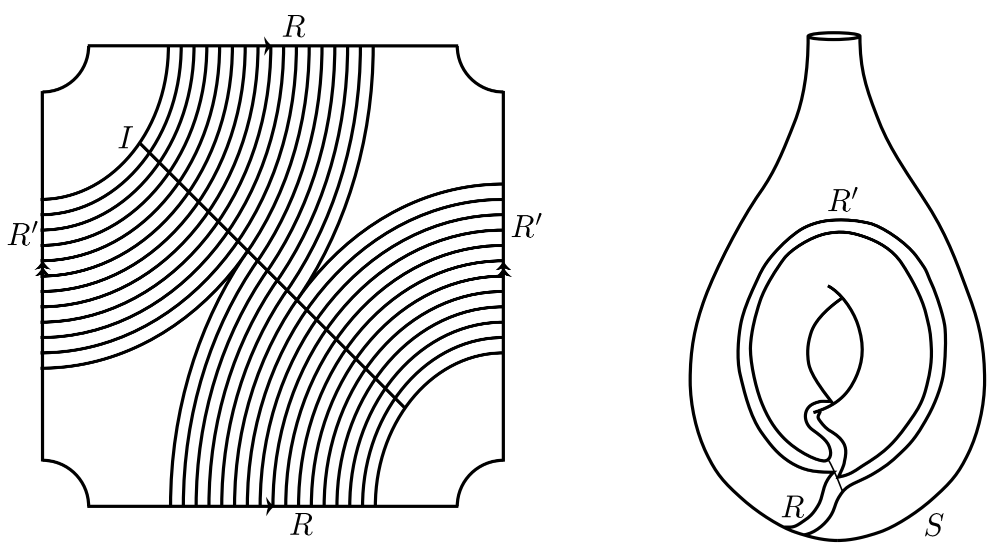

(a) (b)

FIGURE 1.??

We can even choose the differentiable structure of S so that, whenever a point $(0, x)$ of R or R' is glued to another point $(1, y)$ of R or R', the union of the arcs $[0, 1] \times \{x\}$ and $[0, 1] \times \{y\}$ forms a differentiable curve in S. Let a

horizontal simple curve in Φ be a closed or bi-infinite differentiable curve which is made up of arcs $[0,1] \times \{x\}$ of R or R' and does not have any self-intersection. In general, horizontal curves are pairwise disjoint, except for those which pass through the points a, $a' \in [0, a+a'] = I$. A horizontal curve is ***singular*** if it passes through one of these points a, $a' \in I$, and ***regular*** otherwise.

If the ratio a/a' is rational, we can split the intervals $[0,a]$, $[0,a']$ and $[0,a+a']$ into finitely many intervals of the same length, and the gluings respect this decomposition. It follows that all regular horizontal curves are closed, and pairwise parallel. Just a little more thought shows that there are only two singular horizontal simple curves, and that they are parallel to the regular horizontal curves. Since they always intersect the arc $I \subset \Phi$ in the same direction, these closed curves are not homologous to 0, and in particular not homotopic to 0. Therefore, if we endow S with an arbitrary metric m of negative curvature for which ∂S is geodesic, there is a unique simple closed m–geodesic which is homotopic to the horizontal simple curves. This associates to our configuration the m–geodesic lamination consisting of this simple closed geodesic.

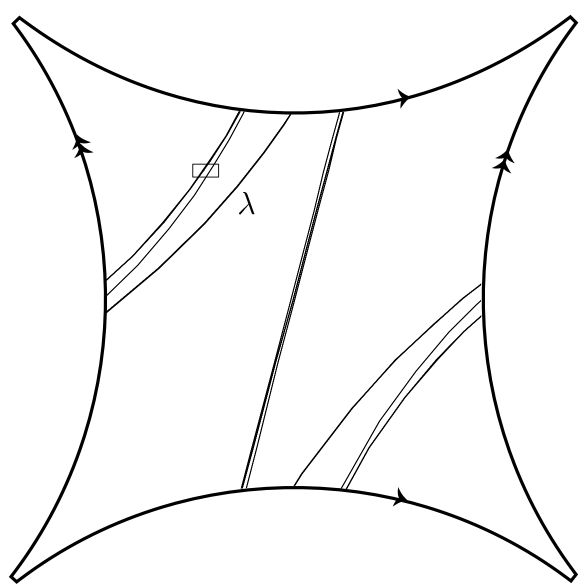

FIGURE 1.??. A geodesic lamination with uncountably many leaves.

The situation is much more interesting when a/a' is irrational. Note that the successive intersection points of a regular horizontal curve with $I = [0, a+a']$

are of the form $\dots$, $x - a'$, x, $x + a'$, $\dots$, $x + na'$, $\dots$ modulo $a + a'$. It follows that all regular horizontal curves are bi-infinite. Similarly, the singular horizontal simple curves are bi-infinite and there are exactly two of them.

In §1.?? (see also §3.??), we will see that we can pull these horizontal simple curves to make them m–geodesic. More precisely, there is a constant C such that every horizontal simple curve is homotopic to a (unique) m–geodesic by a homotopy which moves points by a distance bounded by C. In particular, the boundedness of the homotopy, combined with the Exponential Divergence Principle of Appendix A, shows that the corresponding geodesics are simple and pairwise disjoint. A little argument (see Proposition 1.?? in §1.??) also shows that the union of these geodesics is closed in S, namely that they form an m–geodesic lamination λ.

Figure 1.?? describes λ in a case where $a/a' \approx 1.548$. Here, we have presented the surface S in a manner reminiscent of Figure 1.??(a), by identifying two pairs of opposite sides of an octogon. Here, the octogon is contained in the (Poincaré model) for the hyperbolic plane, and the identifications are chosen so that the metric of the octogon induces on S a hyperbolic metric with geodesic boundary. In particular, for every point of intersection of λ with a side of the octogon, the identification sends x to a point y of the opposite side which is also in λ, in such a way that λ makes the same angle with the corresponding side at x and at y.

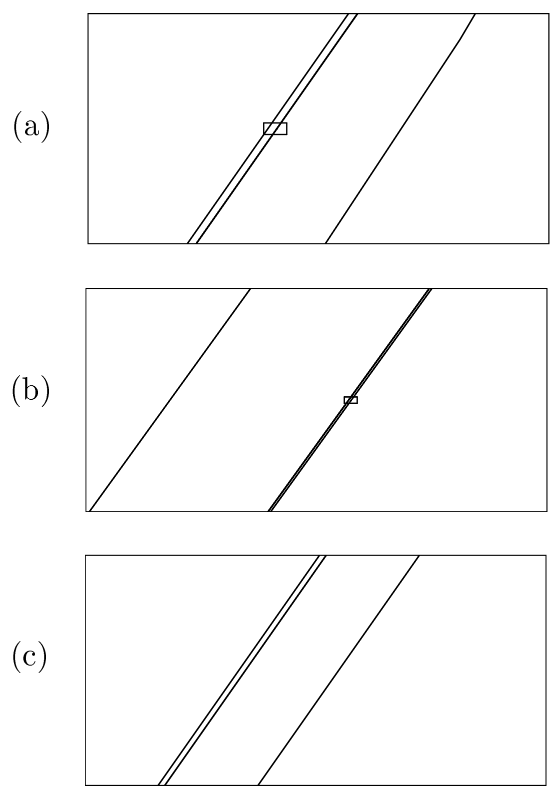

FIGURE 1.??. Details of the geodesic lamination of Figure 1.??.

When inspecting Figure 1.?? and comparing it to Figure 1.??, one may be struck by the fact that λ occupies remarkably little space on the surface S. This property tends to obscure the fact that the topology of λ is much more complex than immediately apparent on the picture. Indeed, if we zoom on the little rectangle shown on Figure 1.??, the enlarged rectangle is the one represented on Figure 1.??(a). We then see that λ crosses this rectangle in more lines than was apparent on Figure 1.??. Similarly, Figure 1.??(b) is an elargement of the little rectangle shown on Figure 1.??(a), and Figure 1.??(c) represents the small rectangle of Figure 1.??(b). These pictures suggest that the uncountably many leaves of λ are grouped in patterns which are reminiscent of a Cantor set, and more precisely that λ seems to be locally homeorphic to the product of the Cantor with an interval. This fact will be rigorously proved in §1.?? (Proposition 1.??).

EXAMPLE 4. We now illustrate the Hausdorff topology on $\mathcal{L}(S)$, by exhibiting a sequence of simple closed geodesics which converges to a geodesic lamination.

For this, go back to Example 3 and, in this construction, let λ_n correspond to $a = 1$ and $a' = 1/n$. Since a and a' are both rational, these λ_n are simple closed geodesics. More precisely, λ_n is isotopic to the simple closed curve which crosses the rectangle R' once and turns n times around R. For instance, Figure 1.??(a) represents the curve λ_4.

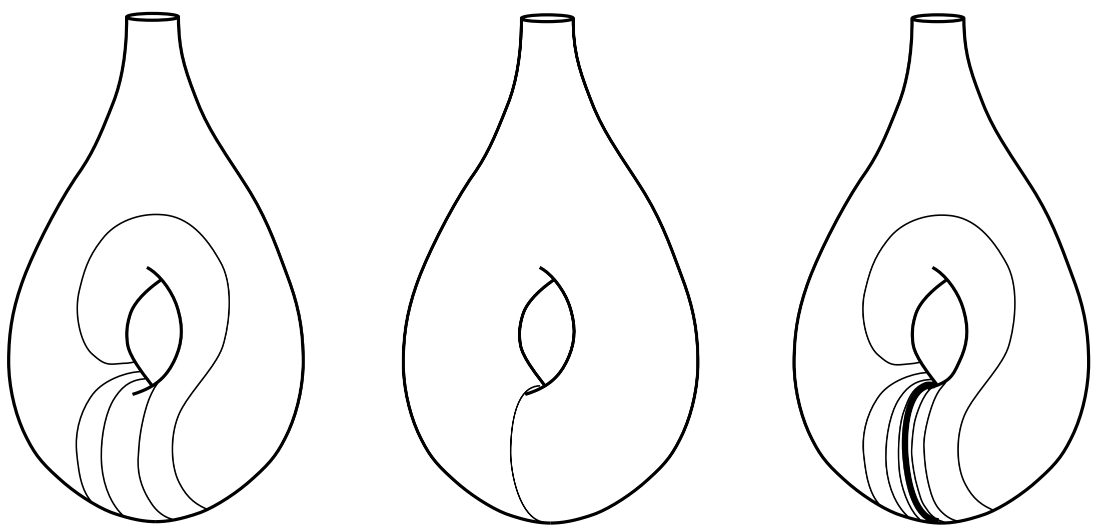

(a) (b) (c)

FIGURE 1.??. A converging sequence of geodesic laminations.

One could think that, as n tends to ∞, the closed geodesic converges to the simple closed geodesic λ_∞ associated to $a = 1$ and $a' = 0$, namely to the closed geodesic λ_∞ which is isotopic to the simple closed curve which turns once around R and 0 time around R'. However, the Hausdorff limit of the λ_n is actually larger. Indeed, while λ_n turns many times around R and consequently has a large part which is very close to λ_∞, there is always a small part of λ_n which 'goes once in the direction of' R'. As a consequence, for n very large, λ_n is very close to the geodesic lamination λ_∞^+ that consists of the closed geodesic

λ_∞ and of one infinite leaf which goes from one side of λ_∞ to the other and spirals in opposite directions around λ_∞. See Figure 1.??(c).

The fact that λ_n converges to λ_∞^+ for the Hausdorff metric will be rigorously proved as Exercise ?? in §1.8.

1.4. Topological properties of geodesic laminations

Consider a decomposition of the geodesic lamination λ into leaves, namely into disjoint simple geodesics. Associate to each $x \in \lambda$ the direction of the leaf passing through λ. By Lemma A??, this defines a Lipschitz direction field on λ, which we can extend to a Lipschitz direction field v with isolated singularities over all of S; in addition, we can require v to be tangent to the boundary along ∂S.

By integration along v, we get the following local picture for $\lambda \subset S$.

LEMMA. *Every $x \in S$ has a neighborhood U that is $C^{1,1}$ diffeomorphic to the product $I \times J$ of two intervals I, J of $\mathbb{R}$, in which $U \cap \lambda$ corresponds to $K \times J$ for some closed subset K of I.* □

Endow the complement $S - \lambda$ with the path metric induced by the metric of S. Namely, the distance between two points x, $y \in S - \lambda$ for this path metric is the infimum of the length of all paths going from x to y in $S - \lambda$. Let $\widehat{S - \lambda}$ be the completion of this path metric.

PROPOSITION. *The completion $\widehat{S - \lambda}$ is a complete finite area surface with negative curvature and geodesic boundary. Its topological type its finite, its boundary consists of finitely many isometric copies of S^1 or $\mathbb{R}$, and it is the union of a compact set and of finitely many 'spikes' homeomorphic to $[0, 1] \times [0, \infty[$, on which the metric induced on the boundary pieces $0 \times [0, \infty[$ and $1 \times [0, \infty[$ corresponds to the standard metric and where the distance from $(0, t)$ to $(1, t)$ decreases exponentially with t as t tends to ∞. In addition, the number s of spikes is even and is equal to the number of non-compact components of the boundary, and*

$$\chi\left(\widehat{S - \lambda}\right) = \chi(S) + \tfrac{1}{2}s.$$

PROOF. Because a Cauchy sequence for the path metric of $S - \lambda$ gives a Cauchy sequence in S, Lemma ?? shows that $\widehat{S - \lambda}$ is a surface with geodesic boundary. Each point $\widehat{x} \in \widehat{S - \lambda}$ projects to a unique point $x \in S$, and the curvature of $\widehat{S - \lambda}$ at $\widehat{x}$ is equal to the curvature of S at x. It also follows from Lemma ?? that the area of $\widehat{S - \lambda}$ is equal to the area of $S - \lambda$, and is therefore finite.

The fact that each boundary component of $\widehat{S - \lambda}$ is isometric to S^1 or $\mathbb{R}$ just follows from the completeness property.

Because S is compact, there is a lower bound $\varepsilon > 0$ for the injectivity radius at all points of S. Namely, two distinct geodesic arcs of length less than ε and issued from the same point cannot intersect again. Let K_ε consist of those $\widehat{x} \in \widehat{S - \lambda}$ for which the set of geodesic arcs of length less than ε joining $\widehat{x}$ to the boundary is either empty or path connected. Because the curvature has a negative upper bound, the areas of the balls of radius ε centered at the points of K_ε is bounded away from 0. It follows that K_ε is compact surface with boundary.

In particular, each end of the boundary must be separated from another end by a stripe homeomorphic to $[0,1] \times [0,\infty[$, on which the metric induced on the boundary pieces $0 \times [0,\infty[$ and $1 \times [0,\infty[$ corresponds to the standard metric and where the distance from each $(0,t)$ to $(1,t)$ is bounded by ε. Because of the negative curvature, this actually implies that the distance from $(0,t)$ to $(1,t)$ decreases exponentially; see ??.

The surface $\widehat{S-\lambda}$ is obtained from K_ε by adding spikes as above, rectangles of width at most ε going from K_ε to the boundary, and rectangles of width at most ε going from one component of the boundary to another one. By compactness of K_ε, it follows that $\widehat{S-\lambda}$ has finite type, and has finitely many boundary components and spikes. The number s of spikes is equal to half the number of ends of boundary components of $\widehat{S-\lambda}$, and is therefore equal to the number of non-compact boundary components.

It remains to prove the Euler characteristic formula. For this, we will use the Lipshitz direction field v on S which has only isolated singularities and which is tangent to the bounday along ∂S. By the Poincaré formula, the Euler characteristic $\chi(S)$ is equal to the sum of the indices of the singularities of v. By the local picture given by Lemma ??, v induces a direction field $\widehat{v}$ on $\widehat{S-\lambda}$ which is tangent to the boundary. If we add a point at the end of each spike, we get a compact surface with the same Euler characteristic as $\widehat{S-\lambda}$, and $\widehat{v}$ extends to a direction field on this surface with boundary singularities of index $+\frac{1}{2}$ at each of these spike ends. Applying the Poincaré formula to this compact surface, we get that its Euler characteristic $\chi\left(\widehat{S-\lambda}\right)$ is equal to the sum of the indices of its interior singularities, namely $\chi(S)$, and of the sum of the indices of its boundary singularities, namely $\frac{1}{2}s$. □

Lemma. *For a geodesic lamination λ, the number of components of λ and the number of components of $\partial\widehat{S-\lambda}$ are bounded by constants depending only on S.*

Proof. If a component A_i of $\widehat{S-\lambda}$ has s_i spikes, the Gauss-Bonnet formula shows that $\chi(A_i) - \frac{1}{2}s_i$ is negative, and therefore is less than or equal to $-\frac{1}{2}$. Since the $\chi(A_i) - \frac{1}{2}s_i$ add up to $\chi(A) - \frac{1}{2}s = \chi(S)$, it follows that $\widehat{S-\lambda}$ has at most $2|\chi(S)|$ components. Also, this shows that $\chi(A_i) - \frac{1}{2}s_i$ is greater than or equal to $\chi(S)$, and therefore that s_i and the number of ends of A_i are uniformly bounded. It follows that the number of components of $\partial\widehat{S-\lambda}$ is uniformly bounded. Finally, since each component of λ contains the image of at least one boundary component of $\partial\widehat{S-\lambda}$, this proves that the number of components of λ is uniformly bounded. □

Exercise. Show that the maximum number of connected components of a geodesic lamination of S is $\frac{3}{2}|\chi(S)| + \frac{1}{2}b$ if S is orientable and has b boundary components, and is $2\left|\chi\left(S\right)\right|+1$ if S is non-orientable. (Hint: First show that this maximum number of components is realized by a geodesic lamination consisting only of closed geodesics).

Exercise. Show that the maximum number of components of $\partial\widehat{S-\lambda}$ for a geodesic lamination λ is equal to $6\left|\chi\left(S\right)\right|$.

From these topological considerations, we obtain a geometric property of geodesic laminations.

Proposition. *An m_0–geodesic lamination λ has m_0–area 0.*

Proof. Let K_0 denote the Gauss curvature of m_0. The Gauss-Bonnet formula gives

$$\iint_{S-\lambda} K_0 = 2\pi\chi\left(\widehat{S-\lambda}\right) - s\pi$$

because the boundary of $\widehat{S-\lambda}$ is geodesic and each spike basically corresponds to a turn of π in the boundary (this formula can be obtained by truncating the spikes of $\widehat{S-\lambda}$, applying Gauss-Bonnet to the compact surface so obtained, and passing to the limit as we truncate the spikes arbitrarily far). From Proposition ?? and the Gauss-Bonnet formula applied to S, we conclude that

$$\iint_{S-\lambda} K_0 = 2\pi\chi\left(S\right) = \iint_S K_0.$$

Since K_0 is bounded away from 0, the result follows. □

Corrollary. *The decomposition of a geodesic lamination as a union of disjoint simple geodesics is unique.*

Proof. In Lemma ??, the set K must have 1–dimensional Lebesgue measure 0 in the interval I, and is therefore totally discontinuous. □

A ***sublamination*** of a geodesic lamination λ is a geodesic lamination contained in λ. The geodesic lamination λ is ***minimal*** if it does not contain any proper sublamination.

By Lemma ??, each point x of the geodesic lamination λ has a neighborhod U homeomorphic to a square $I\times J$ such that $U\cap\lambda$ corresponds to $K\times J$ for some closed subset K of I. We say that the leaf containing x is a ***boundary leaf*** if x is in the closure of some component of $U-\lambda$. Equivalently, a boundary leaf is the image of a boundary component of $\widehat{S-\lambda}$ under the canonical projection $\widehat{S-\lambda}\to S$. If, in addition, we can choose U so that K consists of a single point, this leaf is an ***isolated leaf***.

▮define spiraling

Proposition. *A geodesic lamination λ is the disjoint union of finitely many minimal sublaminations and of finitely many infinite isolated leaves. In addition, each end of an infinite isolated leaf spirals along a unique minimal sublamination.*

Proof. By definition, minimal sublaminations are pairwise disjoint. The union of finitely many such minimal sublaminations forms a geodesic lamination.

Lemma ?? then shows that their number is uniformly bounded. Therefore, λ admits only finitely many minimal sublaminations. Let λ' denote the union of these minimal sublaminations.

Let g be a leaf of λ which is not in λ'. This leaf must be infinite and have no cluster points in $S-\lambda'$ (for the path metric), since this would otherwise provide additional minimal sublaminations for λ. Therefore, each end of g gets arbitrarily close to a boundary component or a spike of $\widehat{S-\lambda'}$. If g enters a spike, it cannot escape and must be asymptotic to the two boundary leaves bounding this spike. Then, the corresponding end of g spirals along the component of λ' containing these two boundary leaves. If an end of g gets arbitrarily close to a non-compact component of $\partial\widehat{S-\lambda'}$, then it must be asymptotic to an end of that leaf; otherwise, it would hit the two boundary leaves that are adjacent to that leaf. Again, this end of g spirals along a component of λ'. Finally, if an end of g gets arbitrarily close to a closed component of $\partial\widehat{S-\lambda'}$, then it must spirals along that component; otherwise, g would intersect itself. We conclude that each end of g spirals along a component of λ'.

Let $g_1, \dots, g_n$ be distinct leaves of λ which are not in λ'. Since each end of each g_i spirals along a component of λ', the union λ'' of λ' and of the g_i is closed, and therefore is a geodesic lamination. By induction, $\partial\widehat{S-\lambda''}$ has $2n$ more non-compact components than $\widehat{S-\lambda'}$. By Lemma ??, it follows that n is uniformly bounded.

Therefore, $\lambda-\lambda'$ consists of finitely many infinite leaves. Since there are only finitely many of them, these leaves must be isolated. □

1.5. Geodesic laminations from a metric independent viewpoint

We introduced geodesic laminations as a way to complete the space of isotopy classes of simple closed curves on S. So far, our construction depends on the choice of a negatively curved metric m_0 on S for which the boundary ∂S is geodesic. We now show that it is possible to define geodesic laminations in a metric independent way. One way to do so is by using of the ***boundary at infinity*** $\partial_\infty\widetilde{S}$ of the universal covering $\widetilde{S}$, as defined in the Appendix.

If we endow S (and consequently $\widetilde{S}$) with a negatively curved metric m_0 making the boundary geodesic, every bi-infinite m_0–geodesic in $\widetilde{S}$ is asymptotic to two distinct points of $\partial_\infty\widetilde{S}$. Conversely, every pair of distinct points of $\partial_\infty\widetilde{S}$ is asymptotic to a unique m_0–geodesic. This establishes a natural corespondence between the space of unoriented m_0–geodesics of $\widetilde{S}$ and the space $G(\widetilde{S})$ of pairs of distinct points of $\partial_\infty\widetilde{S}$. Note that the space $G(\widetilde{S})$ is homeomorphic to $\big(\partial_\infty\widetilde{S}\times\partial_\infty\widetilde{S}-\Delta\big)/\mathbb{Z}_2$, where Δ denotes the diagonal and where the group $\mathbb{Z}_2=\{+1,-1\}$ acts by permutation of the two factors. When the surface S has no boundary, the boundary at infinity $\partial_\infty\widetilde{S}$ is a circle, and $G(\widetilde{S})$ is an open Möbius strip. When S has non-empty boundary, the spaces $\partial_\infty\widetilde{S}$ and $G(\widetilde{S})$ are both Cantor sets.

Given an m_0–geodesic lamination λ, its pre-image $\widetilde{\lambda}$ forms an m_0–geodesic lamination of $\widetilde{S}$. The leaves of $\widetilde{\lambda}$ define a subset $b(\lambda)$ of $G(\widetilde{S})$. If, for a sequence of geodesics, the end points of these geodesics in $\partial_\infty \widetilde{S}$ converge to a pair of distinct points, then the geodesics converge to the corresponding geodesic, uniformly on compact sets. Since λ and $\widetilde{\lambda}$ are closed, it follows that the set $b(\lambda)$ is closed in $G(\widetilde{S})$. Also, $b(\lambda)$ is clearly invariant under the action of the fundamental group $\pi_1(S)$ on $G(\widetilde{S})$. Finally, any two elements $\{x_1, y_1\}$ and $\{x_2, y_2\}$ of $b(\lambda)$ are ***unlinked*** in the space $\widetilde{S} \cup \partial_\infty \widetilde{S}$, in the sense that it is possible to connect x_1 to y_1 and x_2 to y_2 in $\widetilde{S} \cup \partial_\infty \widetilde{S}$ by two arcs with disjoint interiors; indeed, the corresponding leaves of $\widetilde{\lambda}$ provide such arcs. We will say that a subset B of $G(\widetilde{S})$ is ***totally unlinked*** if any two pairs in B are unlinked.

Conversely, let B be a closed totally unlinked $\pi_1(S)$–invariant subset of $G(\widetilde{S})$. Let $\widetilde{\lambda}$ be the union of the m_0–geodesics corresponding to the elements of B. Because B is totally unlinked, any two distinct geodesics of $\widetilde{\lambda}$ are disjoint. Since B is $\pi_1(S)$–invariant, so is $\widetilde{\lambda}$. Finally, the fact that B is closed in $G(\widetilde{S})$ easily implies that $\widetilde{\lambda}$ is closed in $\widetilde{S}$. Therefore, $\widetilde{\lambda}$ is a $\pi_1(S)$–invariant m_0–geodesic lamination of $\widetilde{S}$. This geodesic lamination $\widetilde{\lambda}$ projects to an m_0–geodesic lamination λ in the quotient $S = \widetilde{S}/\pi_1(S)$. By construction, $b(\lambda) = B$.

This shows:

PROPOSITION. *There is a one-to-one correspondence between geodesic laminations in S and closed totally unlinked $\pi_1(S)$–invariant subsets of $G(\widetilde{S})$.* □

The boundary at infinity $\partial_\infty \widetilde{S}$ can be defined without reference to a negatively curved metric m_0; see the Appendix. Therefore, Proposition ?? provides a metric independent description of geodesic laminations.

This shows that, as a set, $\mathcal{L}(S)$ is independent of the choice of the negatively curved metric m_0. Let us show that this is also true for the topology of $\mathcal{L}(S)$.

PROPOSITION. *A sequence of m_0–geodesic laminations λ_n converges to a geodesic lamination λ_0 for the Hausdorff metric in S if and only if the corresponding subsets $b(\lambda_n)$ converge to $b(\lambda_0)$ for the Hausdorff distance in $G(\widetilde{S})$.*

PROOF. □

Changing the negatively curved metric m_0 on S does not change the topology of $\mathcal{L}(S)$, but it does change the Hausdorff metric on this space. On the other hand, the proof of Proposition ?? shows that the Hölder class of the Hausdorff metric stays unchanged. In other words, $\mathcal{L}(S)$ admits a canonical Hölder structure, as defined in §??.

1.6. Train tracks

A ***train track*** Θ in S is an embedded graph of the type we already encountered in §1.1. Namely, Θ is contained in the interior of S and consists of finitely many vertices, also called ***switches***, and of finitely many edges joining them such that:

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

(ii) 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;

(iii) 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 .

The origin of the terminology should be obvious[1]. Of the above conditions, Condition (iii) is probably the less natural. Perhaps it would make sense to omit this property from the definition of a train track, and to call a train track essential if it satisfies (iii) in addition. However, because we will need (iii) for all the train tracks which we are going to consider (and also because of the weight of tradition), it is more convenient to include it in the definition.

Note that Condition (iii) allows a component of $S - \Theta$ to be a semi-open annulus with no spike, containing a boundary component of S.

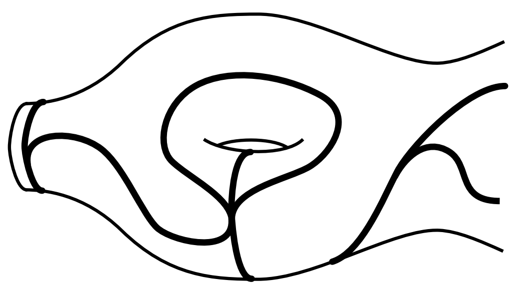

FIGURE 1.8. A train track Θ on the surface S

A ***fattened train track*** Φ on the surface S consists of a finite family of 'long' rectangles e_i in S, each foliated by arcs parallel to the 'short' sides, and meeting as illustrated in Figure 1.8. Namely, two rectangles meet only along their short sides, and every point of the short side of a rectangle is contained in another short side of rectangle; note that this allows the two short sides of a same rectangle to meet along an arc. In addition, we require the following conditions:

(a) each component of the union of the short sides of the e_i is an arc, as opposed to a closed curve;

(b) if a component of the boundary ∂S meets Φ, then this whole component is contained in Φ;

(c) observe that the closure $\overline{S - \Phi}$ of the complement $S - \Phi$ has a certain number of spikes, corresponding to points belonging to three rectangles; we require that no component of $\overline{S - \Phi}$ is a disc with 0, 1 or 2 spikes or an annulus with no spike.

Note that Condition (c) does not allow any component of $\overline{S - \Phi}$ to be an annulus with 0 spike, even for annuli containing a component of ∂S.

The rectangles e_i are the ***edges*** of the fattened train track Φ. The leaves of the foliation of Φ induced by the foliation of the e_i by arcs parallel to the short

[1] Regarding this terminology, the non-american english speaker may benefit from the footnote of [**Bon**, p.??]

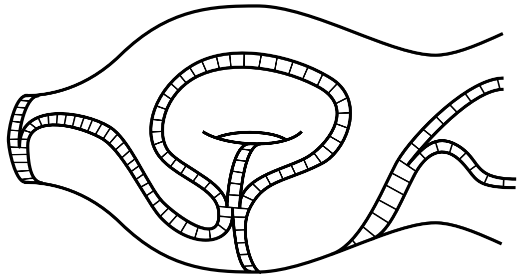

FIGURE 1.9. A fattened train track Φ on the surface S

sides are the ***ties*** of Φ. The (finitely many) ties where several edges meet are the ***switches*** of Φ. A tie which is not a switch is said to be ***generic***.

Given a fattened train track Φ in S, we can associate to Φ a train track Θ in S by collapsing each of the ties of Φ to a point; this train track Θ is well defined up to isotopy of S. In this collapsing process, edges of Φ correspond to edges of Θ, switches of Φ corrrespond to switches of Θ, and ties of Φ correspond to points of Θ. The terminology of train tracks and fattened train tracks is therefore consistent.

Conversely, given a train track Θ in S, we can thicken the edges of Θ, adjust this thickening near the switches, and possibly add a few collar neighborhoods of boundary components of S to obtain a fattened train track Φ whose collapse is isotopic to Θ. The reader should however beware that a train track can have several distinct fattenings if one of its switches has more than one edge on each of its two sides, as indicated by Figure 1.9.

In this context, it is useful to introduce the following definition: A switch of a train track or of a fattened train track is ***simple*** if there is only one edge coming in on one of the two sides of the switch. Then, if all the switches of the train track Θ are simple, Θ admits a unique fattening up to isotopy.

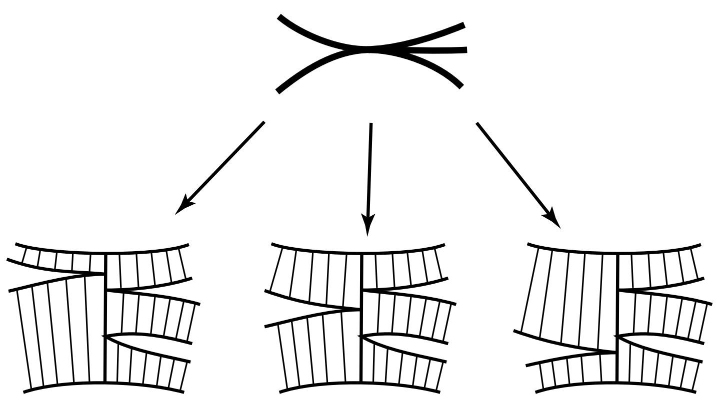

FIGURE 1.10. Three fattenings of a train track near a given switch

An m_0–geodesic lamination λ is said to be ***(strongly) carried*** by the fattened train track Φ if it is contained in the interior of Φ and if each leaf of λ is transverse to the ties of Φ. We will denote by $\mathcal{L}(\Phi, m_0)$, or $\mathcal{L}(\Phi)$ for short if there is no ambiguity about the underlying metric m_0, the subset of $\mathcal{L}(S)$ consisting of those m_0–geodesic laminations which are carried by Φ.

PROPOSITION. *Each m_0–geodesic lamination λ is carried by a fattened train track Φ.*

PROOF. In §1.4, we already used a Lipschitz direction field v defined on a neighborhood of λ which, at each $x \in \lambda$, is orthogonal to the leaf passing by x. By compactness, there is an $\varepsilon > 0$ such that no integral curve of v stays within a distance of ε from λ.

Consider $S - \lambda$ and its completion $\widehat{S - \lambda}$, as in §1.4. By Proposition ??, $\widehat{S - \lambda}$ is the union of a compact piece and of finitely many spikes parametrized by $[0, 1] \times [0, \infty[$. We can choose the parametrization of each spike so that each arc $[0, 1] \times x$ is a piece of integral curve of v of length at most ε. Let Φ' be the union of these spikes. Add to Φ' a collar neighborhood of $\partial\widehat{S - \lambda} - \Phi'$, homeomorphic to $\left(\partial\widehat{S - \lambda} - \Phi'\right) \times [0, 1]$ in such a way that each $x \times [0, 1]$ is apiece of integral curve of v issued from x and of length at most ε. We can choose these collars so that the frontier $[0, 1] \times 0$ of each spike $[0, 1] \times [0, \infty[$ is covered by the closures of the two adjacent collars. In this way, we get a neighborhood Φ'' of $\partial\widehat{S - \lambda}$ in $\widehat{S - \lambda}$, whose frontier has the same local type as the frontier of a train track.

Let $\Phi \subset S$ be the union of λ and of the projection of Φ''. By our choice of ε, no leaf of the foliation of Φ by integral curves of v is a closed curve. Considering the leaves through the spikes of $\overline{S - \Phi}$, we get a decomposition of Φ into rectangles meeting along switches as in a fattened train track. Conditions (i) and (ii) are clearly satisfied. As for Condition (iii), it follows from the fact that each component of $\widehat{S - \lambda}$ is a disk with at most 2 spikes or an annulus with no spike, which follows from the Gauss-Bonnet formula applied to the negatively curved metric induced by the metric of S.

This proves that Φ is a fattened train track, whose ties are contained in integral curves of v. By construction, the leaves of λ are transverse to the ties of Φ, and λ is therefore carried by Φ. ☐

1.7. Hausdorff dimension of simple geodesics

Recall that a geodesic is ***simple*** if it does not intersect itself in a transverse way. This section is devoted to the proof of the following result, due to Joan S. Birman and Caroline Series [**BiS**], which shows that simple geodesics are very scarce.

THEOREM. *If the surface S is endowed with a negatively curved metric m_0 with geodesic boundary, the union of all simple m_0–geodesics of S has Hausdorff dimension 1.*

This gives the following immediate corollary, which strengthen Proposition 1.??.

COROLLARY. *Every m_0–geodesic lamination has Hausdorff dimension 1.* ☐

PROOF OF THEOREM ??. We follow the lines of the argument given in [**Thu**, §??].

The closure of a simple geodesic is clearly a geodesic lamination. Therefore, we have to show that the union $\bigcup_{\lambda\in\mathcal{L}(S)}\lambda$ of all m_0–geodesic laminations has Hausdorff dimension 1.

For a fattened train track Φ, consider the set $\mathcal{L}(\Phi)=\mathcal{L}(\Phi,m_0)\subset\mathcal{L}(S)$ consisting of those geodesic laminations which are (strongly) carried by Φ. Clearly, $\mathcal{L}(\Phi)$ is open in $\mathcal{L}(S)$ for the Hausdorff topology ■need lemma for transversality. Since $\mathcal{L}(S)$ is compact (Proposition ??), we can cover it by finitely many $\mathcal{L}(\Phi)$. Therefore, it suffices to show that, for each fattened train track Φ, the subset $\bigcup_{\lambda\in\mathcal{L}(\Phi)}\lambda$ of S has Hausdorff dimension 1. And for this it suffices to show that $\bigcup_{\lambda\in\mathcal{L}(\Phi)}\lambda$ meets each generic tie k of Φ in a set of Hausdorff dimension 0.

First, let us introduce some terminology.

An immersed arc c is ***carried*** by the fattened train track Φ if it is contained in the interior of Φ and if it is transverse to its ties. Such an arc c traverses some oriented edges e_1, e_2, … , e_n of Φ, in this order, where an orientation of an edge amounts to a coherent transverse orientation of its ties. Let an ***oriented edge path*** be any ordered finite sequence $\gamma=\langle e_1,e_2,\dots,e_n\rangle$ associated in this way to a compact oriented curve carried by Φ. In this situation, we will say that the curve c ***follows*** the edge path γ. An (unoriented) ***edge path*** is obtained by identifying the edge path $\langle e_1,e_2,\dots,e_n\rangle$ to $\langle \bar e_n,\bar e_{n-1},\dots,\bar e_1\rangle$, where $\bar e_i$ is the edge e_i with its orientation reversed. The ***length*** of the edge path $\langle e_1,e_2,\dots,e_n\rangle$ is the number n of edges in the sequence.

Fix a generic tie k_0 of the fattened train track Φ. Let P be the intersection of k_0 with $\bigcup_{\lambda\in\mathcal{L}(\Phi)}\lambda$.

For $r\geqslant 0$, let Γ_r be the set of edge paths $\langle e_{-r},e_{-r+1},\dots,e_0,\dots,e_{r-1},e_r\rangle$ of length $2r+1$, where e_0 is the edge containing k_0. If c is an arc contained in a leaf of $\lambda\in\mathcal{L}(\Phi)$ and follows the edge path $\gamma\in\Gamma_r$, consider the point x_c where it hits k_0 when crossing the central edge e_0. For a fixed $\gamma\in\Gamma_r$, let $P_\gamma\subset k_0$ consist of all points x_c obtained in this way.

Clearly, $P=\bigcup_{\gamma\in\Gamma_r}P_\gamma$ for every $r\geqslant 0$.

Let c and c' be geodesic arcs which are carried by Φ and follow the same edge path $\gamma\in\Gamma_r$. There is a homotopy sending c to c' which moves points by a distance bounded by a contant depending only on the width of the edges of Φ. Also, for $r\geqslant 1$, the lengths of c and c' are at least equal to r times a constant depending only on the length of the edges of Φ. By Lemma A.??, it follows that the corresponding intersection points x_c, $x_{c'}\in k_0$ are at most $O\left(e^{-Ar}\right)$ apart, for some constant $A>0$.

This proves that the diameter of the set $P_\gamma\subset k_0$ is an $O\left(e^{-Ar}\right)$ for every $\gamma\in\Gamma_r$.

In general, the cardinal of Γ_r grows exponentially with r. However, we are interested only in those edge paths which are followed by arcs c immersed in simple geodesics. In particular, these edge paths are all followed by embedded arcs (this requires a small perturbation if c is contained in a simple closed geodesic). This restriction drastically reduces the number of such edge paths, and is the key point of the proof of Theorem 1.??.

Lemma. *For any fattened train track Φ, the number of edge paths of Φ of length r that are followed by embedded arcs grows polynomially with r.*

PROOF. Up to isotopy respecting the edges, an embedded arc c carried by Φ is uniquely determined by the way it crosses each edge of Φ. Indeed, there is a unique way to glue these pieces together at the switches without introducing any intersection point. The intersection of c with an edge e is determined by the number p_e of times c crosses e plus, if any of the two end points of c lies in e, some information about where the component of $c \cap e$ containing the end point sits with respect to the other components and in which direction this component is going. If n denotes the number of edges of Φ, it follows that there are at most $4n^2 r^2 (r+1)^n$ edge paths of length r that are followed by embedded arcs, since each p_e is between 0 and r. $\square$

The growth rate estimate given by the above proof of Lemma ?? is very crude, and can be much improved by the use of Lemma ?? in Chapter ??.

Lemma ?? shows that the number of $\gamma \in \Gamma_r$ for which P_γ is non-empty is an $O(r^n)$ for some $n \geqslant 1$. Therefore, for every $r \geqslant 1$, we can cover the set P by $O(r^n)$ subsets of diameter $O\left(e^{-Ar}\right)$. It follows that, for every $d > 0$, the d–dimensional Hausdorff measure of P is 0.

This proves that, for every train track Φ, the set $\bigcup_{\lambda \in \mathcal{L}(\Phi)} \lambda$ meets each tie of Φ in a set of Hausdorff dimension 0, and therefore has Hausdorff dimension 1.

This completes the proof of Theorem ??. $\square$

1.8. Geodesic laminations weakly carried by train tracks

In §1.7, we made a strong use of the subsets $\mathcal{L}(\Phi) = \mathcal{L}(\Phi, m_0)$ of $\mathcal{L}(S)$, consisting of those m_0–geodesic laminations which are carried by the fattened train track Φ. However, these $\mathcal{L}(\Phi, m_0)$ strongly depend on the base metric m_0. It is therefore convenient to relax the definition of a geodesic lamination carried by a train track, in order to make it metric independent. This is the main purpose of this section.

A curve $c : I \to S$, defined on an interval I of $\mathbb{R}$, is ***strongly carried*** by the train track Θ if its image is contained in Θ and if it is an immersion, namely if the derivative $c'(t)$ is never 0.

The main point is that it is possible to associate a geodesic to each bi-infinite curve that is carried by a train track. Here is a warm up.

LEMMA. *Let Θ be a train track in S, and let $\widetilde{\Theta}$ be its pre-image in the universal covering $\widetilde{S}$ of S. Then, any curve c that is strongly carried by $\widetilde{\Theta}$ crosses each point of $\widetilde{\Theta}$ at most once; in particular, c is embedded.*

PROOF. Without loss of generality, we can assume that each every switch s of Θ is simple, namely that there is a single edge coming in on one of the two sides of s. Indeed, we can otherwise replace s by a small edge, as indicated on Figure 1.11, and it clearly suffices to prove the property for the modified train track so obtained. Then, there is a thickening Φ of Θ and a curve c_0 carried by the pre-image $\widetilde{\Phi} \subset \widetilde{S}$ of Φ such that the collapse $\Phi \to \Theta$ sends c_0 to c. We need to prove that c_0 meets each tie of Φ at most once.

Suppose that c_0 meets the tie k at least twice, and let c_1 be an arc in c joining two distinct points x, $y \in k \cap c$. Shortening c_1 if necessary, we can assume that

FIGURE 1.11

FIGURE 1.12

k is the only tie which c_1 meets more than once. In particular, c_1 is embedded.

Let k_0 be the subarc of k that joins x to y. The simple closed curve $c_1 \cup k_0$ can be perturbed to a curve c_2 which is transverse to the ties except possibly at one point; see Figure 1.??. Let D be the disk bounded by c_2 in $\widetilde{S}$. Because no component of $\widetilde{S} - \widetilde{\Theta}$ is a disk with at most 1 spike, the ties of $\widetilde{\Theta}$ can be extended to a foliation of D whose only interior singularities are of saddle type with 3 prongs. See Figure ??; such saddle type singularities have index $-\frac{1}{2}$. On the boundary c', this foliation has 0 or 1 singularity, of index $\pm\frac{1}{2}$. This contradicts the Poincaré-Hopf Theorem (see for instance ??) which says that the indices of these singularities should add up to $\chi(D) = 1$. $\square$

FIGURE 1.13. Foliations with saddle-type singularities

PROPOSITION. *Let Θ be a train track in S, and let $\widetilde{\Theta}$ be its pre-image in the universal covering $\widetilde{S}$ of S. Then, for any metric m_0 on S, there is a constant K such that every curve that is strongly carried by $\widetilde{\Theta}$ is K–geodesic in $\widetilde{S}$ for m_0.*

PROOF. The train track Θ can always be enlarged to a train track which is

maximal, namely cannot be extended to a larger train track. This is easily seen to be equivalent to the property that Θ contains ∂S and that each component of its complement is a disk with 3 spikes. Therefore, we can assume that Θ is maximal.

In addition, as in the proof of Lemma ??, we can arrange that all switches of Θ are simple. Let Φ be a fattening of Θ. Then, for every curve c that is strongly carried by $\widetilde{\Theta}$, there is a curve c_0 carried by the pre-image $\widetilde{\Phi} \subset \widetilde{S}$ of Φ such that the collapse $\Phi \to \Theta$ sends c_0 to c. Since this collapse moves points by a bounded amount, it suffices to show that every curved carried by $\widetilde{\Phi}$ is K'–quasi geodesic for some K' depending only on S, Φ and m_0.

Let $\widetilde{X}$ be the complement in $\widetilde{S}$ of the union of the boundary of $\widetilde{\Phi}$ and of its switches. Namely, $\widetilde{X}$ consists of the interior of the edges of $\widetilde{\Phi}$ and of the components of $\widetilde{S} - \widetilde{\Phi}$. For x, $y \in \widetilde{X}$, consider all arcs γ in $\widetilde{S}$ that join x to y and are transverse to the boundary of $\widetilde{\Phi}$ and to its switches. Define $\delta(x, y)$ as -1 plus the minimum, over all such arcs γ, of the number of components of $\gamma \cap \widetilde{X}$. This defines a pseudo-metric δ on $\widetilde{X}$.

By the Milnor-Svarc Lemma (see ??), the pseudo-metric δ is quasi-isometric to the metric m_0. In particular, an arc is quasi-isometric for m_0 if and only if it is quasi-isometric for δ. Therefore, we only have to show the existence of a constant K_0 such that, if c is an arc carried by $\widetilde{\Phi}$ with end points x, $y \in \widetilde{X}$, then $l_\delta(c) \leqslant K_0 \delta(x, y)$, where $l_\delta(c)$ is the δ–length of c, namely $l_\delta(c)$ is equal to -1 plus the number of components of $c \cap \widetilde{X}$. We have now converted our original problem into a purely combinatorial one.

So, consider an arc c which is carried by $\widetilde{\Phi}$ and with end points x, $y \in \widetilde{X}$. Let c' be an arc joining x to y and minimizing the number of components of $c' \cap \widetilde{X}$, namely such that $l_\delta(c') = \delta(x, y)$.

Consider a component of $c' - \widetilde{\Phi}$. Because of our assumption that Θ is maximal, every component of $\widetilde{S} - \widetilde{\Phi}$ is a relatively compact disk and we can deform this arc to an arc with the same end points and contained in $\widetilde{\Phi}$. Performing this operation on all the components of $c' - \widetilde{\Phi}$, we eventually get an arc c'' joining x to y in $\widetilde{\Phi}$ and such that $l_\delta(c'') \leqslant K_1 l_\delta(c') = K_1 \delta(x, y)$, where K_1 is an upper bound for the number of edges of Φ that are adjacent to each of the components of $S - \Phi$.

Pulling c'' tight in $\widetilde{\Phi}$, we can arrange that c'' is transverse to the ties of $\widetilde{\Phi}$, except along finitely many arcs (possibly reduced to a point) contained in switches of $\widetilde{\Phi}$ and delimited by spikes of $\widetilde{S} - \widetilde{\Phi}$, as in Figure ??. Note that c'' makes a U-turn at each of these non-transverse arcs k, in the sense that it branches out on the same side at the end points of k. We can also assume that c'' is embedded and transverse to c. Let $x = x_0, x_1, \ldots, x_n = y$ be the intersection points of c'' and c, occurring in this order on c''.

Let c_i and c''_i be the subarcs of c and c'', respectively, that join x_{i-1} to x_i. Although c_i may contain additional intersection points of $c \cap c''$, the curve $c_i \cup c''_i$ is simple and bounds a disk D in $\widetilde{S}$. The partial foliation of D by the ties of $\widetilde{\Phi}$ extends to a foliation of D whose interior singularities are of saddle type, each with at least 3 prongs. On the boundary of D, the foliation has no singularity

Figure 1.14

along the interior of c_i, two singularities of type half-center or half-saddle at x_{i-1} and x_i, and possibly singularities of type half-center or half-saddle on the interior of c_i''. Counting indices, we see that the number of interior saddle points of the foliation is bounded by the number of singularities occurring in the interior of c_i'', and therefore by $l_\delta\left(c_i''\right)$ (since each singularity in the interior of c_i'' corresponds to a point where c_i'' makes a U-turn at a switch of $\widetilde{\Phi}$). As a consequence, D contains at most $l_\delta\left(c_i''\right)$ components of $\widetilde{S}-\widetilde{\Phi}$.

Consider a tie of $\widetilde{\Phi}$ that meets c_i, and follow this tie in D. Because the above foliation of D has no singularity in the interior of c_i, and has only saddle type singularities in the interior of D, this tie cannot return to c_i; therefore, it must hit either c_i'' or a component of $\widetilde{S}-\widetilde{\Phi}$. As a consequence, each edge of $\widetilde{\Phi}$ that is traversed by c_i meets, either c_i'', or a component of $\widetilde{S}-\widetilde{\Phi}$ that is contained in D. Since there are at most $l_\delta\left(c_i''\right)$ such components of $\widetilde{S}-\widetilde{\Phi}$, and since each such component meets at most K_1 edges of $\widetilde{\Phi}$, we conclude that $l_\delta\left(c_i\right) \leqslant \left(K_1+1\right) l_\delta\left(c_i''\right)$.

By construction, $l_\delta\left(c''\right)=\sum_{i=1}^n l_\delta\left(c_i''\right)$ and $l_\delta(c) \leqslant \sum_{i=1}^n l_\delta\left(c_i\right)$, where the equality holds only when the intersection points $x=x_0, x_1, \ldots, x_n=y$ occur in this order on c''. We conclude that $l_\delta(c) \leqslant \left(K_1+1\right) l_\delta\left(c''\right) \leqslant K_1\left(K_1+1\right) \delta(x, y)$.

This concludes the proof that every curve that is carried by $\widetilde{\Phi}$ is quasi-geodesic for the pseudo-metric δ, and therefore for the negatively curved metric m_0. □

Consider a bi-infinite curve c which is strongly carried by the pre-image $\widetilde{\Theta}$ of the train track Θ, namely a curve $c:]-\infty,+\infty[\rightarrow \widetilde{S}$ carried by $\widetilde{\Theta}$ such that both $c([0,+\infty[)$ and $c(]-\infty, 0])$ have infinite length. Let m_0 be a negatively curved metric on S for which ∂S is geodesic. By Proposition ??, c is quasi-geodesic for m_0, and therefore is (■weakly?) asymptotic to a unique m_0–geodesic g_0 of $\widetilde{S}$. This associates to c a unique element $G(c)$ of the space of geodesics $G(\widetilde{S})=\left(\partial_{\infty} \widetilde{S} \times \partial_{\infty} \widetilde{S}-\Delta\right) / \mathbb{Z}_2$. This $G(c)$ does not depend on the choice of the metric m_0; indeed, for another choice of metric m_0', the m_0'–geodesic g' that is asymptotic to c is asymptotic to the m_0–geodesic $G(c)$, and therefore defines the same element of $G(\widetilde{S})$. We will say that a geodesic $g \in G(\widetilde{S})$ is ***weakly carried*** by $\widetilde{\Theta}$ if there is a bi-infinite curve c which is strongly carried by $\widetilde{\Theta}$ and such that $g=G(c)$.

We will say that the geodesic lamination λ is ***(weakly) carried*** by the train track Θ if every leaf of the pre-image $\widetilde{\lambda} \subset \widetilde{S}$ of λ is weakly carried by the pre-image $\widetilde{\Theta}$ of Θ. Let $\mathcal{L}(\Theta)$ denote the subset of $\mathcal{L}(S)$ that consists of those geodesic laminations which are carried by Θ. Note that $\mathcal{L}(\Theta)$ is independent of the choice of the negatively curved metric m_0.

We can emphasize the metric independence of $\mathcal{L}(\Theta)$ by giving a purely combinatorial description of its elements. Indeed, an element of $\mathcal{L}(\Theta)$ is a $\pi_1(S)$–invariant closed subset $\widetilde{\lambda}$ of $G(\widetilde{S})$ such that any $g \in \widetilde{\lambda}$ is weakly carried by $\widetilde{\Theta}$ and any two distinct g, $g' \in \widetilde{\lambda}$ correspond to disjoint geodesics. We can translate these notions in terms of bi-infinite edge paths in $\widetilde{\Theta}$.

Indeed, the geodesic $G(c) \in G(\widetilde{S})$ that is asymptotic to the bi-infinite curve c carried by $\widetilde{\Theta}$ clearly depends only on the bi-infinite edge path $\gamma = \langle \dots, e_{-n}, \dots, e_{-1}, e_0, e_1, \dots, e_n, \dots \rangle$ that is followed by c (where we identify two edge paths when they differ only by orientation reversal and shift in the indexing of the edges). Therefore, we can write $G(c) = G(\gamma)$.

Conversely, if the geodesic g is carried by $\widetilde{\Theta}$, we are going to prove that the edge path γ such that $g = G(\gamma)$ is unique. We begin with two combinatorial lemmas.

Lemma. *Consider two curves c, $c' : [0,1] \to \widetilde{\Theta}$ which are strongly carried by $\widetilde{\Theta}$ and with the same end points $c(0) = c'(0)$ and $c(1) = c'(1)$. Then, c and c' follow the same edge path in $\widetilde{\Theta}$.*

Proof. This result is easily proved by an index argument analogous to the one used in the proof of Lemma ??, and we leave its proof as an exercise. □

Lemma. *Consider two closed curves c, c' in S which are strongly carried by Θ. If c and c' are homotopic, then they follow the same cyclic edge path in Θ (modulo reindexing of the edges).*

Proof. We can lift c to a bi-infinite curve $\widetilde{c}$ which is carried by $\widetilde{\Theta}$ and is invariant under an element $\alpha \in \pi_1(S)$ realized by c. If c' is homotopic to c, we can lift it to a similar bi-infinite curve $\widetilde{c}'$ carried by $\widetilde{\Theta}$ which is invariant under α and is weakly asymptotic to $\widetilde{c}$ in both directions. As usual, we consider a fattening Φ of Θ and two bi-infinite curves $\widetilde{c}_0$, $\widetilde{c}'_0$ which are carried by the pre-image $\widetilde{\Phi} \subset \widetilde{S}$ of Φ and which the collapse $\widetilde{\Phi} \to \widetilde{\Theta}$ projects to $\widetilde{c}$, $\widetilde{c}'$ (assuming all switches of Θ simple without loss of generality).

If $\widetilde{c}$ and $\widetilde{c}'$ cross the same edge e of $\widetilde{\Theta}$, they also cross αe. By Lemma ??, they must follow the same edge path $\widetilde{\gamma}$ between e and αe. It follows that c and c' follow the same cyclic edge path γ in Θ, namely the projection of $\widetilde{\gamma}$.

Otherwise, $\widetilde{c}_0$ and $\widetilde{c}'_0$ are disjoint and are separated by an α–invariant strip $\widetilde{\Sigma}$ in $\widetilde{S}$. The ties of $\widetilde{\Phi}$ give a partial foliation of $\widetilde{\Sigma}$, which projects to a partial foliation of the quotient $\Sigma = \widetilde{\Sigma}/\alpha$. This foliation extends to a foliation of Σ whose only singularities are interior saddle type singularities of index $-\frac{1}{2}$. But Σ is an annulus, with Euler characteristic 0. The Poincaré-Hopf Theorem therefore implies that this foliation of Σ cannot have any singularity. It follows that $\widetilde{c}_0$ and $\widetilde{c}'_0$ cross the same edges of $\widetilde{\Phi}$, and in the same order. As a consequence, in this case again, c and c' follow the same edge path in Θ. □

We can rephrase Lemma ?? by saying that two distinct edge paths in $\widetilde{\Theta}$ cannot have the same two end edges. We now show that this holds even for bi-infinite edge paths and their associated points at infinity.

If $c : [0, \infty[\to \widetilde{S}$ is an infinite curve carried by $\widetilde{\Theta}$, it follows from Proposition ?? that it converges to a unique point of the boundary at infinity $\partial_\infty \widetilde{S}$.

LEMMA. *If two infinite curves c, $c' : [0, \infty[\to \widetilde{S}$ are carried by $\widetilde{\Theta}$ and converge to the same point in $\partial_\infty \widetilde{S}$, then the edge paths $\gamma = \langle e_0, e_1, \ldots, e_n, \ldots \rangle$ and $\gamma' = \langle e'_0, e'_1, \ldots, e'_n, \ldots \rangle$ respectively followed by c and c' eventually coincide, in the sense that there is a number m such that $e'_n = e_{n+m}$ for every n sufficiently large.*

PROOF. If c and c' converge to the same point at infinity, they are asymptotic and we can find sequences $n_0 < n_1 < \ldots < n_k < \ldots$ and $n'_0 < n'_1 < \ldots < n'_k < \ldots$ such that each oriented edge e_{n_k} is at uniformly bounded distance from $e'_{n'_k}$ in $\widetilde{S}$. Modulo the action of $\pi_1(S)$, there are only finitely many pairs of oriented edges of $\widetilde{\Theta}$ whose distance in $\widetilde{S}$ is bounded by a given constant, since Θ is finite. Therefore, we can assume that there are elements $\alpha_k \in \pi_1(S)$ such that $e_{n_k} = \alpha_k e_{n_0}$ and $e'_{n'_k} = \alpha_k e'_{n'_0}$ for every k.

The oriented edge paths $\gamma_k = \langle e_{n_0}, e_{n_0+1}, \ldots, e_{n_k-1}, e_{n_k} \rangle$ and $\gamma'_k = \Big\langle e'_{n'_0}, e'_{n'_0+1}, \ldots, e'_{n'_k-1}, e$ project to edge paths in Θ which are both realized by closed curves representing α_k. By Lemma ??, the projections of these two edge paths are therefore equal modulo cyclic permutation. Actually, the proof of Lemma ?? gives more: If $\widetilde{\gamma}_k$ is the bi-infinite oriented edge path in $\widetilde{\Theta}$ obtained by chaining together all the $\alpha_k^n \gamma_k$, $n \in \mathbb{Z}$, and if $\widetilde{\gamma}'_k$ is similarly obtained by chaining the $\alpha_k^n \gamma'_k$ together, then $\widetilde{\gamma}_k = \widetilde{\gamma}'_k$.

In particular, exchanging the rôles of c and c' if necessary, there is an oriented edge path in $\widetilde{\Theta}$ beginning with the oriented edge e_{n_0} and ending with $e'_{n'_0}$. This path is unique by Lemma ??. If n_k is larger than the length N of this edge path, we conclude from the equality $\widetilde{\gamma}_k = \widetilde{\gamma}'_k$ that $e'_{n+n'_0} = e_{n+n_0+N}$ for every n between 0 and $n_k - N$.

Letting k (and therefore n_k) tend to ∞, we conclude that the edge paths γ and γ' eventually coincide. □

An immediate corollary of Lemmas ?? and ?? is that, if two bi-infinite edge paths γ and γ' in $\widetilde{\Theta}$ are asymptotic to the same geodesic $G(\gamma) = G(\gamma') \in G(\widetilde{S})$, then γ and γ' must be equal (modulo orientation reversal and shift of indices). This proves:

PROPOSITION. *The map $\gamma \mapsto G(\gamma)$ defines a one-to-one correspondence between bi-infinite edge paths in $\widetilde{\Theta}$ and geodesics weakly carried by $\widetilde{\Theta}$.* □

Let $G(\widetilde{\Theta})$ denote the subset of $G(\widetilde{S})$ consisting of those geodesics which are weakly carried by $\widetilde{\Theta}$. Namely, $G(\widetilde{\Theta})$ is the image of the map $\gamma \mapsto G(\gamma)$. As seen in §1.4, the data of a geodesic lamination carried by Θ is equivalent to the data of a closed subset $\widetilde{\lambda}$ of $G(\widetilde{S})$ which is contained in $G(\widetilde{\Theta})$, which is

$\pi_1(S)$–invariant, and such that any two distinct g, $g' \in \widetilde{\lambda}$ are disjoint in $\widetilde{S}$. We now translate these properties in terms of edge paths.

The map $\gamma \mapsto G(\gamma)$ is clearly equivariant with respect to the actions of $\pi_1(S)$. Consequently, a subset of $G(\widetilde{\Theta})$ is $\pi_1(S)$–invariant if and only if the corresponding set of edge paths in $\pi_1(S)$–invariant.

Let us say that two bi-infinite train tracks $\gamma = \langle \dots, e_{-n}, \dots, e_{-1}, e_0, e_1, \dots, e_n, \dots \rangle$ and $\gamma' = \langle \dots, e'_{-n}, \dots, e'_{-1}, e'_0, e'_1, \dots, e'_n, \dots \rangle$ ***cross each other*** when, after possible orientation reversal and shift of indices, there exists indices p, q such that $e_r = e'_r$ for every r with $p < r < q$, $e_p \neq e'_p$, $e_q \neq e'_q$, and e'_p and e'_q sit on opposite sides γ. It easily follows from Lemmas ?? and ?? that the geodesics $G(\gamma)$ and $G(\gamma')$ cross each other if and only if γ and γ' intersect.

Finally, we have to understand when a subset of $G(\widetilde{\Theta})$ is closed in $G(\widetilde{S})$.

LEMMA. *The set $G(\widetilde{\Theta})$ is closed in $G(\widetilde{S})$. A sequence of geodesics $G(\gamma_n)$ in $G(\widetilde{\Theta})$ converges to $G(\gamma)$ as n tends to ∞ if and only if, for every finite edge path γ' contained in the bi-infinite edge path γ, γ_n contains γ' for n sufficiently large.*

PROOF. Suppose that the sequence of geodesics $G(\gamma_n)$ converges to a geodesic $g \in G(\widetilde{S})$ as n tends to ∞. For every n, choose a bi-infinite curve c_n that is strongly carried by $\widetilde{\Theta}$ and follows γ_n. By Proposition ??, there is a constant M such that c_n stays within a distance of M from the geodesic $G(\gamma_n)$.

Since $G(\gamma_n)$ converges to g as n tends to ∞, for every compact subset K of $\widetilde{S}$, there is an n_0 such that c_p stays within a distance of $3M$ from c_q over K for every p, q large enough. The argument used in the proof of Lemma ?? then provides an increasing sequence of finite edge paths $\gamma'_n = \langle e_{-n}, \dots, e_{-1}, e_0, e_1, \dots, e_n \rangle$ such that γ'_n is contained in γ_p for every p large enough (depending on n). Let γ be the union of the γ'_n.

If c is a bi-infinite curve following γ, it stays within a distance of $3M$ from g by construction. It follows that $g = G(\gamma)$. In particular, g is an element of $G(\widetilde{\Theta})$.

This proves that $G(\widetilde{\Theta})$ is closed. The second statement of Lemma ?? is an immediate consequence of the construction of the γ'_n. □

In practice, it is often more convenient to consider edge paths in the train track Θ rather in its pre-image $\widetilde{\Theta} \subset \widetilde{S}$. We can paraphrase the above discussion as follows.

PROPOSITION ??. *Given a train track Θ in S, the rule $\gamma \mapsto G(\gamma)$ establishes a one-to-one correspondence between geodesic laminations that are weakly carried by Θ and families Γ of bi-infinite edge paths in Θ such that:*

(i) no two edge paths γ, $\gamma' \in \Gamma$ cross each other;

(ii) if γ is an edge path in Θ and if every finite edge path δ contained in γ is also contained in some $\gamma' \in \Gamma$, then γ is also in Γ. □

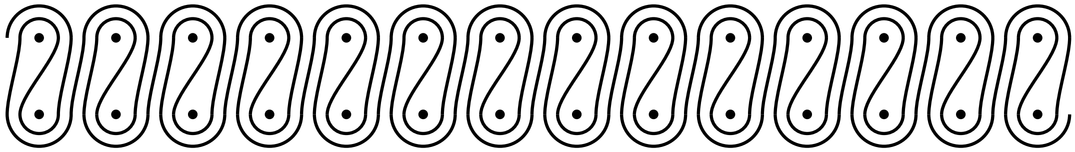

CHAPTER II

MEASURED GEODESIC LAMINATIONS

2.1. Transverse measures for geodesic laminations

Let us go back to the example of §1.1. With the notation used there, consider the two multicurves $C_{(3,2,6,2,3)}$ and $C_{(5,4,12,4,7)}$. If we draw these two multicurves, we see that $C_{(3,2,6,2,3)}$ consists of two disjoint simple closed curves γ_1 and γ_2, and that $C_{(5,4,12,4,7)}$ is isotopic to the union of γ_1 and of three parallel copies of γ_2 (it is fun to draw the pictures, but it may be easier to notice that $C_{(3,2,6,2,3)} = \varphi\left(C_{(1,0,2,0,1)}\right)$ and $C_{(5,4,12,4,7)} = \varphi\left(C_{(1,0,4,0,3)}\right)$). Therefore, the two multicurves $C_{(3,2,6,2,3)}$ and $C_{(5,4,12,4,7)}$ define the same geodesic lamination λ, consisting of the disjoint simple closed geodesics λ_1 and λ_2 homotopic to γ_1 and γ_2, respectively. On the other hand, there is a definite difference between $C_{(3,2,6,2,3)}$ and $C_{(5,4,12,4,7)}$, since γ_2 occurs 3 times as much in $C_{(5,4,12,4,7)}$ as in $C_{(3,2,6,2,3)}$. A natural way to keep track of this type of information is to associate to $C_{(5,4,12,4,7)}$ the geodesic lamination λ with multiplicities 1 and 3 respectively attached to its closed leaves γ_1 and γ_2. A transverse measure for a geodesic lamination λ can be thought of as a way to generalize this idea of leaf multiplicity to the case where λ has infinite leaves.

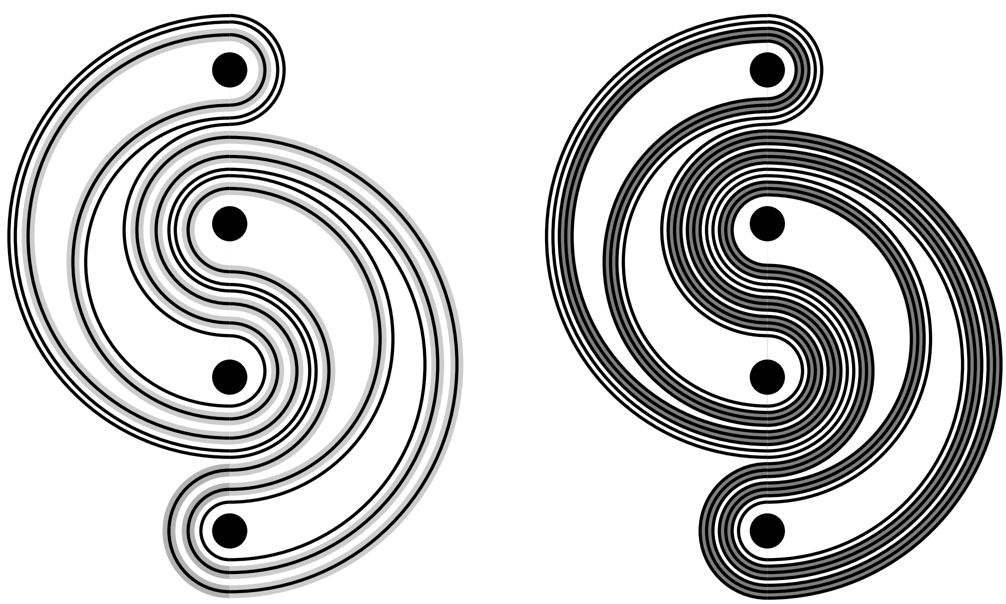

FIGURE 2.1

A ***transverse (invariant) measure*** for a geodesic lamination λ is a measure defined on each arc k transverse to λ, such that every homotopy sending k to another arc k' while respecting λ sends the measure defined on k to the measure defined on k'. Since this notion plays a fundamental rôle in this monograph, it is probably worth it to spend some time clarifying the motions involved in this definition.

By convention, throughout this monograph, the word "measure" will mean ***Radon measure***, namely a measure on a locally compact space X which is defined on the σ–algebra of Borel subsets of X and which assigns non-negative finite mass to each compact set. For such a Radon measure α, a continuous φ function has a finite integral $\int \varphi\, d\alpha$, and this associates to α a positive linear form on the space of all continuous function with compact support on X. Conversely, the Riesz Representation Theorem (see for instance [**Rudin**]) asserts that a positive linear form on the space of continuous functions with compact support on X is associated in this way to a (unique) Radon measure α.

A ***transverse arc*** for the measured geodesic lamination λ is a compact differentiable arc k embedded in the surface S such that, at each point x of intersection of k with λ, the tangent to k is distinct from the tangent to the leaf of λ passing through x. We also require the additional condition that the end points of k are not in λ. From the fact that the the union of all lifts of geodesic laminations is a closed subset of Hausdorff dimension 1 in $PT(S)$ (Theorem 1.??), it easily follows that every differentiable arc in S can be differentiably approximated by an arc which is transverse to all geodesic laminations.

Given two arcs k, k' transverse to the geodesic lamination λ, a ***homotopy respecting*** λ from k to k' is a continuous map $H : [0,1] \times [0,1] \to S$ such that the restrictions $s \mapsto H(s,0)$ and $s \mapsto H(s,1)$ are homeomorphisms from $[0,1]$ to k and k', respectively, and such that the pre-image $H^{-1}(\lambda)$ is of the form $K \times [0,1]$ for some compact subset K of $[0,1]$. For instance, if $t \mapsto k_t$, $t \in [0,1]$, is a differentiable family of arcs which are all transverse to λ, an easy application of ?? provides a homotopy H respecting λ such that each map $s \mapsto H(s,t)$ is a homeomorphism from $[0,1]$ to k_t. We should emphasize the such the homotopy H is only required to be continuous, and not necessarily differentiable. Indeed, we will see in ?? that, if the two transverse arcs k and k' are homotopic respecting λ, there may not necessarily exist a differentiable homotopy respecting λ from k to k'. This fact is only a minor technical point right now, but will play a much more crucial rôle in later chapters.

If the two transverse arcs k and k' are homotopic by a homotopy $H : [0,1] \times [0,1] \to S$ respecting λ, the two maps $s \to H(s,0)$ and $s \mapsto H(s,1)$ define homeomorphisms from $[0,1]$ to k and k', and therefore define a homeomorphism $k \to k'$ which, because H respects λ, sends $k \cap \lambda$ to $k' \cap \lambda$.

By definition, a transverse measure α for λ is the data of a (Radon) measure defined on each arc k transverse to λ, in such a way that, for every homotopy respecting λ between two transverse arcs k and k', the homeomorphism $k \to k'$ induced by this homotopy sends the measure defined on k to the measure defined on k'.

From the property of invariance under homotopy respecting λ, it easily follows that the support of the measure defined by α on a transverse arc k is necessarily contained in $k \cap \lambda$.

There is another description of transverse measures which is perhaps more natural. Indeed, a transverse measure for λ can also be locally interpreted as a measure on the space of leavesof λ. We now explain this.

In general, the space of leaves of the geodesic lamination λ is not a Hausdorff topological space. However, by Lemma 1.??, every point of S admits an open neighborhood U such that the space $G(\lambda \cap U)$ of leaves of $\lambda \cap U$ is Hausdorff. For such a neighborhood U and for $l \in G(\lambda \cap U)$, let $k \subset U$ be an arc transverse to λ which meets the leaf l in one point and which is small enough that it meets each leaf of $\lambda \cap U$ in at most one point; then, the set $V_k \subset G(\lambda \cap U)$ that consists of those leaves which meet k is a neighborhood of l and, as k ranges over all such transverse arcs k, the V_k form a neighborhood basis for l in $G(\lambda \cap U)$. In addition, by hypothesis on k, each leaf of V_k meets k in exactly one point, and this defines a homeomorphism between V_k and $k \cap \lambda$. If α is a transverse measure for λ, the measure induced by α on k therefore defines a measure on $V_k \subset G(\lambda \cap U)$. Also, if $V_k = V_{k'}$, one easily sees that the two transverse arcs k and k' are homotopic respecting λ. Therefore, the measure so defined on the neighborhood V_k of l depends only α, and not on the transverse arc k. This also shows that, for two such neighborhoods V_k and $V_{k''}$, the measures defined by α coincide on the overlap $V_k \cap V_{k''}$. In particular, these measures fit together to define a measure on the space $G(\lambda \cap U)$ of leaves of $\lambda \cap U$.

Therefore, for every open subset U of S such that the space $G(\lambda \cap U)$ of leaves of $\lambda \cap U$ is Hausdorff, the transverse measure α defines a measure on $G(\lambda \cap U)$. In addition, if V is another open subset with $G(\lambda \cap V)$, the corresponding measures fit together on the overlap of U and V, in the following sense: The canonical maps $G(\lambda \cap U \cap V) \to G(\lambda \cap U)$ and $G(\lambda \cap U \cap V) \to G(\lambda \cap U)$ locally are measure preserving homeomorphisms (use small transverse arcs to check this).

In this sense, the transverse measure α can be locally interpreted as a transverse measure on the space of leaves of λ.

Conversely, suppose that we are given, for every open subset U such that $G(\lambda \cap U)$ is Hausdorff, a measure α in such a way that, for any two such open subsets U, V, the canonical maps $G(\lambda \cap U \cap V) \to G(\lambda \cap U)$ and $G(\lambda \cap U \cap V) \to G(\lambda \cap U)$ locally are measure preserving. If k is an arc transverse to λ, it can be decomposed into a union of small transverse arcs k' which are each contained in some open U with $G(\lambda \cap U)$ Hausdorff. Pulling back the measure defined on $G(\lambda \cap U)$ by the local homeomorphism $k' \cap \lambda \to G(\lambda \cap U)$, we obtain a measure on $k' \cap \lambda$, and therefore on k'. Doing this for all the small arcs k', this provides a measure on k. Using the invariance property of the measures defined on the $G(\lambda \cap U)$ with respect to overlaps, one easily sees that this measure on-the transverse arc k is independent of all the choices involved, and is invariant under homotopy of k respecting λ.

Therefore, a transverse measure for the geodesic lamination λ is exactly equivalent to a locally defined measure on the space of leaves of λ, as described above.

It is now time to see some examples of transverse measures.

The simplest example of a transverse measure for the geodesic laminations λ is associated to any collection of closed leaves $g_1, \dots, g_n$ of λ (if any) and of positive multiplicities $a_1, \dots, a_n \in \mathbb{R}$. Indeed, on each arc k transverse to λ, we can consider the measure which to any Borelian set $A \subset k$ associates $\sum_{i=1}^{n} a_i \# (A \cap g_i)$ (where $\# X$ denotes the cardinal of X), namely the sum of the Dirac measures of weight a_i based at the finitely many points $k \cap g_i$; since every homotopy respecting λ must also respect each g_i, the invariance under homotopy respecting λ is automatic.

This example of course only applies to geodesic laminations withe closed leaves. Here is a more general construction.

LEMMA. *Every non-empty geodesic lamination λ admits a non-trivial transverse measure.*

PROOF. Let $I_1, \dots, I_n, \dots$ be a sequence of geodesic arcs immersed in λ, such that the length of I_n is equal to n. For each arc k transverse to λ, consider the Dirac measure α_n of mass $1/n$ based at $k \cap I_n$. In particular, for every continuous function $\varphi : k \to \mathbb{R}$, $\alpha_n(\varphi) = \frac{1}{n} \sum_{x \in k \cap I_n} \varphi(x)$ (considering intersection points with multiplicity when I_n is not embedded). This does not define a transverse measure for λ, because it fails to satisfy the property of invariance under homotopy respecting λ (when the homotopy sweeps through one of the end points of I_n).

By transversality, the lengths of the leaves of $\lambda - k$ admit a lower bound $L > 0$. Because I_n has length n, it follows that the cardinal of $k \cap I_n$ is at most $1 + n/L$, and therefore that $|\alpha_n(\varphi)| \leqslant (1 + 1/L) \max |\varphi|$.

This proves that, for every continuous function $\varphi : k \to \mathbb{R}$ defined on an arc k transverse to λ, the $|\alpha_n(\varphi)|$ are all bounded independently of n, and consequetly are in some compact interval I_φ in $\mathbb{R}$. Consider the product $\prod_\varphi I_\varphi$ of all such I_φ where φ ranges over all continuous functions defined on arcs transverse to λ, and note that each α_n can be interpreted as the element $\alpha_n \in \prod_\varphi I_\varphi$ whose φ–coordinate is $\alpha_n(\varphi)$ for every φ. . This product is compact by the Tychonov Compactness Theorem (see [**Munkres**]) and the sequence $(\alpha_n)_{n \in \mathbb{N}}$ therefore has a limit point $\alpha \in \prod_\varphi I_\varphi$. This means that α associates a number $\alpha(\varphi) \in \mathbb{R}$ to each φ in such a way that, for every finite family of functions $\varphi_1, \dots, \varphi_k$ defined on arcs transverse to λ, we can extract a subsequence $\left(\alpha_{n_p}\right)_{p \in \mathbb{N}}$ such that $\alpha_{n_p}(\varphi_i)$ converges to $\alpha(\varphi_i)$ as p tends to ∞ for every $i = 1, \dots, k$. (One cannot right away find a subsequence which converges in $\prod_\varphi I_\varphi$ becasue the product does not have a countable basis of neighborhoods). Note that $\alpha(\varphi) \geqslant 0$ for every φ by continuity. On each transverse arc k, this construction therefore gives a positive linear form α on the vector space of all continuous functions on k, namely a measure α by the Riesz Representation Theorem (see [**Rudin**]).

Let us check that α is invariant is we homotop k to another transverse arc k' by a homotopy respecting λ. Consider a continuous $\varphi : k \to \mathbb{R}$, and let $\varphi' : k' \to \mathbb{R}$ be its image under the homotopy. The homotopy associates to each

$x \in k \cap \lambda$ a point $x' \in k' \cap \lambda$ and conversely, and $\varphi'(x') = \varphi(x)$ by definition of φ'. There is an upper bound K to the length of the paths followed by points of $k \cap \lambda$ under the homotopy. Therefore, if $x \in k \cap I_n$ is at least K away from the end points of I_n (as measured in I_n), then the corresponding point $x' \in k' \cap \lambda$ is also in I_n. Also, as above, if $L > 0$ is a lower bound for the length of the leaves of $\lambda - k$ and $\lambda - k'$, there are at most $2(1 + K/L)$ points of $I_n \cap k$ which are within K of the end points of I_n. Since the same property holds fo k', it follows that

$$|\alpha_n(\varphi) - \alpha_n(\varphi')| \leqslant \frac{4}{n}(1 + K/L) \max|\varphi|.$$

Letting n go to infinity in the appropriate subsequence, we conclude that $\alpha(\varphi) = \alpha(\varphi')$. This proves that α is invariant under homotopy respecting λ, and therefore defines a transverse measure for λ.

It remains to show that this transverse measure is non-trivial. By compactness, we can find finitely many transverse arcs $k_1, k_2, \dots, k_N$ such that every arc of length 1 immersed in a leaf of λ must hit at least one of the k_i. Then, I_n has at least $n - 1$ intersection points with the k_i, and

$$\sum_{i=1}^{N} \alpha_n(k_i) \geqslant \frac{n-1}{n}$$

where $\alpha_n(k_i)$ is the mass of the measure defined by α_n, namely the integral of the constant function 1 for the measure defined by α_n on k_i. Passing to the limit in a suitable subsequence, we conclude that

$$\sum_{i=1}^{N} \alpha(k_i) \geqslant 1,$$

so that the transverse measure α is really non-trivial. □

Lemma ?? is a special case of a more general construction of transverse measures for foliations whose leaves have subexponential growth [**Plante**].

Given a transverse measure α for a geodesic lamination λ, its ***support*** is the union of the supports of the measures induced by α on all arcs k transverse to λ. This support is clearly a sublamination of λ. The transverse measure is said to have ***full support*** if its support is exactly equal to λ.

Although every geodesic lamination admits a non-trivial transverse measure, not every geodesic lamination admits a full support transverse measure.

LEMMA. *The support of a transverse measure for a geodesic lamination λ contains no infinite isolated leaf of λ.*

PROOF. Suppose that the infinite isolated leaf g is in the support of a transverse measure α for λ. Because S is compact, the geodesic g nust accumulate somewhere, and we can therefore find an arc k transverse to λ that g hits in infinitely many points $x_1, x_2, \dots, x_n, \dots$. Because g is an isolated leaf, we can find for every n an arc $k_n \subset k$ such that $k_n \cap \lambda = \{x_n\}$; in addition, we can choose the k_n pairwise disjoint. By invariance of α under homotopy respecting λ, the mass $\alpha(k_n)$ of k_n for the measure deposited by α must be equal to $\alpha(k_1)$. Since we assumed that g is in the support of α, $\alpha(k_1)$ is strictly positive. But this implies that $\alpha(k) \geqslant \sum_{n=1}^{\infty} \alpha(k_n) = \sum_{k=1}^{\infty} \alpha(k_1) = \infty$ since the k_n are disjoint and contained in k, contradicting the finiteness of $\alpha(k)$. □

Proposition. *A geodesic lamination admits a full support transverse measure if and only if each of its connected components is minimal.*

Proof. By Proposition 1.??, a geodesic lamination λ is the disjoint union of finitely many minimal sublaminations and of finitely many infinite isolated leaves. In particular, it has finitely many connected components.

If λ admits a full support transverse measure, Lemma 2.?? shows that it contains no infinite isolated leaf. Therefore it is the disjoint union of finitely many minimal sublaminations, each of which forms a connected component of λ.

Conversely, assume that each connected component λ_i of λ is minimal. By Lemma 2.??, each λ_i admits a non-trivial transverse measure α_i. The support of α_i must be equal to λ_i since λ_i is minimal. Then, the sum of the α_i provides a full support transverse measure for λ. □

A ***measured geodesic lamination*** α consists of a geodesic lamination λ_α together with a full support transverse measure for λ_α.

We will denote by $\mathcal{ML}(S)$ the set of all measured geodesic laminations of S.

2.2. Measured geodesic laminations from a metric independent viewpoint

In §1.??, we saw that the notion of geodesic lamination is independent of the choice of a negatively curved metric m_0 on S. In §2.1, the definition of a transverse measure for an m_0–geodesic lamination λ was based on arcs transverse to this m_0–geodesic lamination corresponding to λ, so that transverse measures to λ could conceivably depend on m_0. We now show that this is not the case, by giving a metric independent description of transverse measures for a geodesic lamination.

Again, the main idea is to lift the situation to the universal covering $\widetilde{S}$ of S. Given a transverse measure α for the m_0–geodesic lamination λ, the preimage $\widetilde{\lambda} \subset \widetilde{S}$ of λ is an m_0–geodesic lamination of $\widetilde{S}$. If $G(\widetilde{S})$ is the space of biinfinite geodesics of $\widetilde{S}$, the leaves of the m_0–geodesic lamination $\widetilde{\lambda}$ form a closed $\pi_1(S)$–invariant subset of $G(\widetilde{S})$, which we also denote by $\widetilde{\lambda}$. Remember that it is possible to define $G(\widetilde{S}) \cong \big(\partial_\infty \widetilde{S} \times \partial_\infty \widetilde{S} - \Delta\big)/\mathbb{Z}_2$, as well as the subset $\widetilde{\lambda} \subset G(\widetilde{S})$ in a metric independent way. We can now translate to this context the data of the transverse measure α for λ.

Given a geodesic $g \in \widetilde{\lambda} \subset G(\widetilde{S})$, choose an arc $\widetilde{k} \subset \widetilde{S}$ which is transverse to $\widetilde{\lambda}$, meets g and is small enough that it meets each leaf of $\widetilde{\lambda}$ in at most one point. Those leaves of $\widetilde{\lambda}$ which meet $\widetilde{k}$ form a neighborhood U of g in $\widetilde{\lambda} \subset G(\widetilde{S})$ and, since they are pairwise disjoint, they are determined by their intersection points with $\widetilde{k}$. The corresponding bijection $\gamma : U \to \widetilde{k} \cap \widetilde{\lambda}$ is clearly continuous, and its inverse is continuous by Lemma A.??; it consequently is a homeomorphism. The measure defined by α on the projection $k \subset S$ of $\widetilde{k}$ determines a measure α_U on $\widetilde{k} \cap \widetilde{\lambda}$, and we can transport this measure to a measure α_U on U by the homeomorphism γ: The measure α_U associates to a continuous $\varphi : U \to \mathbb{R}$ the integral of $\varphi \circ \gamma^{-1}$ with respect to the measure defined by α on $\widetilde{k} \cap \widetilde{\lambda}$. Given two such transverse arcs $\widetilde{k}_1$ and $\widetilde{k}_2$, which respectively define neighborhoods U_1' and U_2' of g, one easily sees that there exists sub-arcs $\widetilde{k}_1' \subset \widetilde{k}_1$ and $\widetilde{k}_2' \subset \widetilde{k}_2$ which meet g and are homotopic respecting $\widetilde{\lambda}$. (Hint: Use Lemma 1.??). From the invariance of α under homotopy respecting λ, it follows that the measures α_{U_1} and α_{U_2} coincide on the neighborhood of g consisting of those geodesics that meet $\widetilde{k}_1'$ (and therefore $\widetilde{k}_2'$). It follows that these measures can be glued together to define a measure on $\widetilde{\lambda} \subset G(\widetilde{S})$. If we consider its push forward under the inclusion map $\widetilde{\lambda} \to G(\widetilde{S})$, this gives a measure on $G(\widetilde{S})$, which we will denote by α.

Note that the measure α so defined on $G(\widetilde{S})$ is invariant under the action of $\pi_1(S)$. Also, by construction, its support is contained in $\widetilde{\lambda}$.

Conversely, suppose that we have a $\pi_1(S)$–invariant measure α on $G(\widetilde{S})$, whose support is contained in $\widetilde{\lambda}$. If k is a sufficiently small arc transverse to λ in S, it lifts to an arc $\widetilde{k}$ transverse to $\widetilde{\lambda}$ which cuts each leaf of $\widetilde{\lambda}$ at most once. Then the leaves of $\widetilde{\lambda}$ cutting $\widetilde{k}$ form a subset U which is both open and closed in $\widetilde{\lambda}$, and we can again consider the homeomorphism $\gamma : U \to \widetilde{k} \cap \widetilde{\lambda}$ which associates to a geodesic $g \in \widetilde{\lambda}$ its intersection point with U. Through this homeomorphism, we can now transport the measure defined by α on U to a measure on $\widetilde{k} \cap \widetilde{\lambda}$, and therefore to a measure on $k \cap \lambda$ by projection, and finally to a measure defined on k by extension. By invariance under the action of $\pi_1(S)$, this measure on k is independent of the choice of the lift $\widetilde{k}$.

Therefore, the $\pi_1(S)$–invariant measure α on $G(\widetilde{S})$ uniquely determines a measure on each arc k transverse to λ, which is small enough that it lifts to an arc $\widetilde{k}$ cutting each leaf of $\widetilde{\lambda}$ in at most one point. If the arc k is homotopic respecting λ to another arc k', this homotopy lifts to a homotopy respecting $\widetilde{\lambda}$ from the lift $\widetilde{k}$ to an arc $\widetilde{k}'$ lifting k'. By construction, those leaves of $\widetilde{\lambda}$ which meet $\widetilde{k}'$ are the same as those which meet $\widetilde{k}$, and form the same subset $U \subset \widetilde{\lambda} \subset G(\widetilde{S})$. Similarly, the map $\widetilde{k} \cap \widetilde{\lambda} \to \widetilde{k}' \cap \widetilde{\lambda}$ induced by the homotopy commutes with the maps $\gamma : U \to \widetilde{k} \cap \widetilde{\lambda}$ and $\gamma' : U \to \widetilde{k}' \cap \widetilde{\lambda}$. It follows that this map sends the measure on $\widetilde{k}$ to the measure on $\widetilde{k}'$, and therefore that the homotopy from k to k' sends the measure constructed on k to that constructed on k'. In other words, we constructed a transverse measure for the geodesic lamination λ.

This proves:

Proposition. *Let λ be a geodesic lamination and consider its pre-image $\widetilde{\lambda}$ in the universal covering $\widetilde{S}$, which we also interpret as a subset $\widetilde{\lambda} \subset G(\widetilde{S})$ as in §??. Then, the above construction provides a one-to-one correspondence* *between transverse measures for λ and $\pi_1(S)$–invariant measures on the space of geodesics $G(\widetilde{S})$ whose support is contained in $\widetilde{\lambda}$.* □

Recall from ?? that the space $G(\widetilde{S})$ endowed with the action of $\pi_1(S)$ admits a presentation which does not depends on the choice of a negatively curved metric m_0. We already used this property to give a metric independent description of the geodesic lamination λ. Proposition 2.?? now provides a metric independent definition of transverse measures for the geodesic lamination λ.

Another benefit of Proposition 2.?? is that it provides an embedding of the set $\mathcal{ML}(S)$ of all measured geodesic laminations in a large topological space.

PROPOSITION. *There is a natural one-to-one correspondence between measured geodesic laminations on S and $\pi_1(S)$–invariant measures on $G(\widetilde{S})$ whose support is a geodesic lamination, namely is such that any two distinct elements of this support correspond to disjoint geodesics of $\widetilde{S}$.* □

In addition to providing a metric independent description of measured geodesic laminations, this point of view will also enable us to endow with a topology the space $\mathcal{ML}(S)$ of all measured geodesic laminations.

2.3. The topology of the space $\mathcal{ML}(S)$ of measured geodesic laminations

Proposition 2.?? provides an embedding of the space $\mathcal{ML}(S)$ of all measured geodesic laminations into the set of all measures on $G(\widetilde{S})$.

One of the classical topologies on the space of (Radon) measures on a locally compact space is the ***weak* topology***. In our setting, this topology is defined by the family of semi-distances $d_\varphi(\alpha,\beta) = |\alpha(\varphi) - \beta(\varphi)|$, where $\varphi : G(\widetilde{S}) \to \mathbb{R}$ ranges over all continuous functions with compact support on $G(\widetilde{S})$, and where $\alpha(\varphi) = \int \varphi\, d\alpha$ denotes the integral of the function φ with respect to the measure α. In other words, let $\mathcal{M}(G(\widetilde{S}))$ denote the space of measures on $G(\widetilde{S})$, and consider the embedding of $\mathcal{M}(G(\widetilde{S}))$ in the product $\prod_\varphi \mathbb{R}_\varphi$ whose coordinates maps are the maps $\alpha \mapsto \alpha(\varphi)$, where φ ranges over all continuous functions $\varphi : G(\widetilde{S}) \to \mathbb{R}$ with compact support and where $\mathbb{R}_\varphi$ just denotes a copy of $\mathbb{R}$ associated to φ. Then the weak* topology is just the pull back of the product topology of $\prod_\varphi \mathbb{R}_\varphi$ under this embedding.

Every continuous function $\varphi : G(\widetilde{S}) \to \mathbb{R}$ can be written as the difference $\varphi = \varphi^+ - \varphi^-$ of the two non-negative continuous functions $\varphi^+ = \max\{\varphi(x), 0\}$ and $\varphi^- = \max\{-\varphi(x), 0\}$. It follows that we can restrict attention to non-negative functions in the definition of the weak* topology. Namely, the topology of $\mathcal{M}(G(\widetilde{S}))$ is the pullback of the product topology by the embedding $\mathcal{M}(G(\widetilde{S})) \to \prod_{\varphi \geqslant 0} \mathbb{R}_\varphi$, where φ ranges over all *non-negative* continuous functions $\varphi : G(\widetilde{S}) \to \mathbb{R}$ with compact support.

In §2.2, we interpreted the elements of $\mathcal{ML}(S)$ as measures on $G(\widetilde{S})$, namely we considered $\mathcal{ML}(S)$ as a subset of $\mathcal{M}(G(\widetilde{S}))$. We can therefore endow

$\mathcal{ML}(S)$ with the subspace topology endowed by the weak* topology of $\mathcal{M}\big(G(\widetilde{S})\big)$.

LEMMA. *The topologies of* $\mathcal{M}\big(G(\widetilde{S})\big)$ *and* $\mathcal{ML}(S)$ *are metrizable.*

PROOF. It clearly suffices to show that the topology of $\mathcal{M}\big(G(\widetilde{S})\big)$ is metrizable.

There exists a countable family of continuous functions $\varphi_n : G(\widetilde{S}) \to \mathbb{R}$, $n \in \mathbb{N}$, with compact support which is dense in the space of all continuous functions with compact support, in the following sense: given a continuous $\varphi : G(\widetilde{S}) \to \mathbb{R}$ with compact support, a neighborhood U of the support of φ, and a number $\varepsilon > 0$, there exists a φ_n whose support is contained in U and such that $|\varphi(g) - \varphi_n(g)| < \varepsilon$ for every $g \in G(\widetilde{S})$. The existence of such a countable family follows from some general principles (compare ??), but is easy to exhibit in our case: If S has no boundary, then $G(\widetilde{S})$ is a 2–manifold, more specifically an open Möbius strip, and we can for instance use the approximation of functions of $\mathbb{R}^2$ by polynomials whose coefficients are rational; if ∂S is non-empty, then $G(\widetilde{S})$ is a Cantor set, and we can take for the φ_n all locally constant functions with compact support and with rational values.

Adding some φ_n if necessary, we can also require that, for every compact subset K of $G(\widetilde{S})$, there is a φ_n in the family which is non-negative and such that $\varphi_n(g) \geqslant 1$ for every $g \in K$.

Consider the map $\mathcal{M}\big(G(\widetilde{S})\big) \to \prod_{n\in\mathbb{N}} \mathbb{R}_{\varphi_n}$ defined by composing the embedding $\mathcal{M}\big(G(\widetilde{S})\big) \to \prod_{\varphi} \mathbb{R}_{\varphi}$ with the natural projection. Pulling back the product topology of $\prod_{n\in\mathbb{N}} \mathbb{R}_{\varphi_n}$ under this map, one obtains a topology on $\mathcal{M}\big(G(\widetilde{S})\big)$ which we will temporarily call the 'new' topology, as opposed to the 'original' one. By construction, the new topology is weaker than or equal to the original topology, in the sense that every subset of $\mathcal{M}\big(G(\widetilde{S})\big)$ which is open for the new topology is also open for the original topology.

Given a continuous $\varphi : G(\widetilde{S}) \to \mathbb{R}$ with compact suport and a number $\varepsilon > 0$, the choice of the countable family of the φ_n, $n \in \mathbb{N}$, provides such a φ_n with $|\varphi(g) - \varphi_n(g)| < \varepsilon$ for every $g \in G(\widetilde{S})$, as well as a non-negative φ_p which takes only values greater than or equal to 1 on the support of φ and on the support of φ_n. As a consequence, $|\alpha(\varphi) - \alpha(\varphi_n)| < \varepsilon\alpha(\varphi_p)$ for every measure $\alpha \in \mathcal{M}\big(G(\widetilde{S})\big)$. It easily follows that every subset of $\mathcal{M}\big(G(\widetilde{S})\big)$ which is open for the original topology is also open for the new topology.

Therefore, the new and original topology coincide on $\mathcal{M}\big(G(\widetilde{S})\big)$. As a countable product of metric spaces, the product $\prod_{n\in\mathbb{N}} \mathbb{R}_{\varphi_n}$ is metrizable. It follows that the topology of $\mathcal{M}\big(G(\widetilde{S})\big)$ is metrizable. In view of the classical proof that a countable product of metric spaces is metrizable (see ??), an explicit metric defining the topology of $\mathcal{M}\big(G(\widetilde{S})\big)$ is the one defined by

$$d(\alpha, \beta) = \sum_{n=1}^{\infty} \frac{1}{2^n} \min\{1, |\alpha(\varphi_n) - \beta(\varphi_n)|\}$$

for every α, $\beta \in \mathcal{M}\big(G(\widetilde{S})\big)$. □

The following technical lemma will be useful on many occasions.

LEMMA. *Let $\varphi_0 : G(\widetilde{S}) \to \mathbb{R}$ be any non-negative continuous function with compact support and with the following property: For every $g \in G(\widetilde{S})$, there exists a $\gamma \in \pi_1(S)$ such that $\varphi_0(\gamma g) > 0$. Then, for every constant C, the set of those $\alpha \in \mathcal{ML}(S)$ with $\alpha(\varphi_0) \leqslant C$ is compact.*

PROOF. Let X consist of those $\alpha \in \mathcal{ML}(S)$ with $\alpha(\varphi_0) \leqslant C$. We want to show that X is compact.

Let $U \subset G(\widetilde{S})$ consist of those g with $\varphi_0(g) > 0$. Then, U is a non-empty open subset of $G(\widetilde{S})$. By hypothesis on φ_0, $G(\widetilde{S})$ is the union of all γU as γ ranges over all elements of $\pi_1(S)$. In particular, if $\varphi : G(\widetilde{S}) \to \mathbb{R}$ is continuous with compact support, we can cover the support of φ with finitely many $\gamma_1 U$, $\gamma_2 U, \ldots, \gamma_n U$, with each $\gamma_i \in \pi_1(S)$. As a consequence, the function $\sum_{i=1}^n \varphi \circ \gamma_i^{-1}$ is positive on the support of φ, and there exists a constant A_φ such that $|\varphi| \leqslant A_\varphi \sum_{i=1}^n \varphi \circ \gamma_i^{-1}$. It follows that

$$|\alpha(\varphi)| \leqslant A_\varphi \sum_{i=1}^{n} \alpha\left(\varphi \circ \gamma_i^{-1}\right) = A_\varphi n \alpha(\varphi)$$

by invariance of the measure α under the action of $\pi_1(S)$.

This proves that, for every continuous function $\varphi : G(\widetilde{S}) \to \mathbb{R}$ with compact support, there exists a constant B_φ such that $\alpha(\varphi) \in [-B_\varphi, B_\varphi]$ for every $\alpha \in X$. In particular, if we identify $\mathcal{ML}(S)$ to its image in the product $\prod \mathbb{R}_\varphi$, the subset X is contained in $\prod [-B_\varphi, B_\varphi]$. The space $\prod [-B_\varphi, B_\varphi]$ is compact by Tykhonov's theorem (see for instance [**Mun**, §??]). Also, Lemma ?? shows that $\mathcal{ML}(S)$ is closed in $\prod \mathbb{R}_\varphi$, and X is clearly closed in $\mathcal{ML}(S)$. It immediately follows that X is compact. □

Of course, for Lemma ?? to be of any use, we need to show that there exists such a φ_0.

LEMMA. *There exists a non-negative function $\varphi_0 : G(\widetilde{S}) \to \mathbb{R}$ with compact support which satisfies the hypothesis of Lemma ??, namely such that for every $g \in G(\widetilde{S})$ there exists a $\gamma \in \pi_1(S)$ such that $\varphi_0(\gamma g) > 0$.*

PROOF. Since S is compact, there is a compact subset K of $\widetilde{S}$ whose projection to S is equal to K. As a consequence, every bi-infinite geodesic of S has a lift to $\widetilde{S}$ which meets K. In other words, for every $g \in G(\widetilde{S})$, there exists a $\gamma \in \pi_1(S)$ such that γg belongs to the subset $X \subset G(\widetilde{S})$ which consists of those geodesics which meets K. Since K is compact, X is also compact. and there exists a non-negative continuous function $\varphi_0 : G(\widetilde{S}) \to \mathbb{R}$ which is strictly positive on X. By construction, such a function satisfies the required property. □

The topology of $\mathcal{ML}(S)$ is very natural, and metric independent. However, having to lift everything to the universal cover $\widetilde{S}$ is not always very convenient. Also, it is fairly easy to loose track of the geometry of the geodesic lamination in this framework. For this reason, we will often use a different presentation of the topology of $\mathcal{ML}(S)$, which is metric dependent and tends to be a little *ad hoc*, but is more geometric and can be directly seen in the surface S.

Let us say that an arc k in S is ***in general position with respect to simple*** m_0–***geodesics*** if it is transverse to every simple m_0–geodesic. For instance, Proposition 1.?? shows that almost every m_0–geodesic arc is in general position with respect to simple m_0–geodesics. Proposition ?? also shows that every differentiable arc can be C^1–approximated by an arc which is in general position with respect to simple m_0–geodesics.

If the arc k is in general position with respect to simple m_0–geodesics, then each measured geodesic lamination α defines a measure on k. For every continuous function $\psi : k \to \mathbb{R}$ and every measured geodesic lamination $\alpha \in \mathcal{ML}(S)$, we can then consider the integral $\alpha(\psi) = \int \psi \, d\alpha$ of ψ with respect to the measure defined by α on k.

Since the topology of $\mathcal{ML}(S)$ is metrizable, we only need to understand the convergence of sequences in $\mathcal{ML}(S)$ to understand this topology.

LEMMA. *The sequence $\alpha_n \in \mathcal{ML}(S)$, $n \in \mathbb{N}$, converges to $\alpha_\infty \in \mathcal{ML}(S)$ if and only if, for every continuous function $\psi : k \to \mathbb{R}$ defined on an arc k which is generic with respect to simple m_0–geodesics, the integrals $\alpha_n(\psi) \in \mathbb{R}$ converge to $\alpha_\infty(\psi)$.*

PROOF. Let us rephrase the property we want to prove. Consider the natural map $\Phi : \mathcal{ML}(S) \to \prod \mathbb{R}_\psi$, where ψ ranges over all continuous functions $\psi : k \to \mathbb{R}$ defined on arcs k which are generic with respect to simple m_0–geodesics, whose coordinates functions are the maps $\alpha \mapsto \alpha(\psi)$. We need to prove that this map $\mathcal{ML}(S) \to \prod \mathbb{R}_\psi$ induces a hoemomorphism from $\mathcal{ML}(S)$ to its image. The map Φ is easily seen to be injective.

To prove that the map Φ is continuous, suppose that $\alpha_n \in \mathcal{ML}(S)$ converges to α_∞. We want to show that $\alpha_n(\psi)$ converges to $\alpha_\infty(\psi)$ for every continuous $\psi : k \to \mathbb{R}$ defined on an arc k which is genenric with respect to simple m_0–geodesics. Lift k to an arc $\widetilde{k}$ in $\widetilde{S}$, and $\psi : k \to \mathbb{R}$ to a function $\widetilde{\psi} : \widetilde{k} \to \mathbb{R}$. Since k is generic with respect to simple m_0–geodesics, there is an $\varepsilon > 0$ such that every m_0–geodesic of $\widetilde{S}$ which meets $\widetilde{k}$ and which projects to a simple geodesic of S makes an angle of at least ε with k (note that the set of those m_0–geodesics which meet $\widetilde{k}$ and project to simple geodesics is compact). Also, subdividing k if necessary, we can assume that every m_0–geodesic of $\widetilde{S}$ which projects to a simple geodesics meets $\widetilde{k}$ at most once (the genericity hypothesis provides a lower bound for the distance between two disting intersection points). Then, we can construct a continuous function $\varphi : G(\widetilde{S}) \to \mathbb{R}$ such that every geodesic $g \in G(\widetilde{S})$ of the support of φ meets $\widetilde{k}$, and such that $\varphi(g) = \widetilde{\psi}\left(g \cap \widetilde{k}\right)$ for every $g \in G(\widetilde{S})$ which meets $\widetilde{k}$ and projects to a simple geodesic of S. By construction, $\alpha(\psi) = \alpha(\varphi)$ for every $\alpha \in \mathcal{ML}(S)$, if we also denote by α the measures defined by α on k and on $G(\widetilde{S})$. Since α_n converges to α for the topology of $\mathcal{ML}(S)$, then $\alpha_n(\varphi)$ converges to $\alpha_\infty(\varphi)$ by definition of this topology, and it follows that $\alpha_n(\psi)$ converges to $\alpha_\infty(\psi)$. This concludes the proof that $\Phi : \mathcal{ML}(S) \to \prod \mathbb{R}_\psi$ is continuous.

Let $\varphi_0 : G(\widetilde{S}) \to \mathbb{R}$ be a function satisfying the hypotheses of Lemma ??, and whose existence is provided by Lemma ??. For a given $\alpha \in \mathcal{ML}(S)$ and a constant $C > \alpha(\varphi_0)$, the set X of those $\beta \in \mathcal{ML}(S)$ with $\beta(\varphi) \leqslant C$ is a closed neighborhood of α in $\mathcal{ML}(S)$. By Lemma ??, X is compact and the restriction of Φ to $X \to \Phi(X)$ of Φ is a homeomorphism, since it is continuous and bijective. As a consequence, the restriction of Φ to $\mathcal{ML}(S) \to \Phi(\mathcal{ML}(S))$ is a local homeomorphism, and is therefore a homeomorphism since it is bijective. □

Given a measured geodesic lamination $\alpha \in \mathcal{ML}(S)$, we can always multiply its transverse measure by a positive number $a > 0$ to obtain a new measured geodesic lamination $a\alpha \in \mathcal{ML}(S)$. This defines a natural action of the multiplicative group $\mathbb{R}^+$ on $\mathcal{ML}(S)$. We will sometimes be interested in the space of orbits of this action, namely on the quotient $\mathcal{PML}(S)$ of $\mathcal{ML}(S) - \{0\}$ under the equivalence relation which identifies α and $\beta \in \mathcal{ML}(S) - \{0\}$ when there exists a positive number $a > 0$ such that $\beta = a\alpha$. This quotient space $\mathcal{PML}(S)$ is the space of ***projective measured geodesic laminations*** on the surface S. The terminology "projective" comes of course from the analogy with the construction of the projective space from $\mathbb{R}^n - \{0\}$.

PROPOSITION. *The space $\mathcal{PML}(S)$ of projective measured geodesic laminations is compact.*

PROOF. Let $\varphi_0 : G(\widetilde{S}) \to \mathbb{R}$ be a non-negative function with compact support satisfying the confditions of Lemmas ?? and ??, namely such that: for every $g \in G(\widetilde{S})$, there is a $\gamma \in \pi_1(S)$ such that $\varphi_0(\gamma g) > 0$. In particular, for every $\alpha \in \mathcal{ML}(S) - \{0\}$, the set of those g with $\varphi_0(g) > 0$ must meet the support of α, since this support is invariant under the action of $\pi_1(S)$. It follows that $\alpha(\varphi_0) > 0$ for every $\alpha \in \mathcal{ML}(S) - \{0\}$. As a consequence, we can consider the map $\alpha \mapsto \alpha/\alpha(\varphi_0)$, which induces a homeomorphism between $\mathcal{PML}(S)$ and the space W of those $\beta \in \mathcal{ML}(S)$ with $\beta(\varphi_0) = 1$. This space W is closed in $\mathcal{ML}(S)$, and is contained in a compact subset by Lemma ??. It is therefore compact. □

2.4. Measured geodesic laminations disjoint from the boundary

A ***measured geodesic lamination disjoint from the boundary*** ∂S is a measured geodesic lamination $\alpha \in \mathcal{ML}(S)$ whose underlying geodesic lamination λ_α is disjoint from ∂S. Let $\mathcal{ML}_0(S)$ denote the subspace of $\mathcal{ML}(S)$ consisting of those geodesic laminations which are disjoint from the boundary.

LEMMA. *Let m_0 be a negatively curved metric with geodesic boundary on S. There is a neighborhood U of the boundary ∂S such that every simple m_0–geodesic that meets U either is a component of ∂S, or spirals around a component of ∂S.*

PROOF. Let γ be a component of ∂S, and let k be a small m_0–geodesic arc orthogonal to γ. Suppose that $g \neq \gamma$ is a simple m_0–geodesic which cuts k in a point $y \in k \cap g$ which is at distance at most $\varepsilon > 0$ from the point $x = k \cap \gamma$. By Lemma A.??, the tangents to the two disjoint geodesics g and γ are very close if ε is small enough. As a consequence, g closely follows γ for a long time. In particular, it stays parallel to γ until it hits k a second time at some point $z \in k$, after going once around γ. See Figure ??. Note that $z \neq y$ since otherwise g would provide a closed geodesic which is parallel to γ, which is forbidden by the fact that m_0 has negative curvature. If z lies between y and x, then it is 'trapped' between γ and the part of g that lies between y and z; by induction, g closely follows γ for ever and we conclude that one end of g spirals around γ. Finally, it z is not between y and x, we can just reverse the orientation of g and the same argument shows that the other end of g spirals around γ.

Therefore, for every component γ of ∂S, we found an $\varepsilon > 0$, depending on the Lipschitz constant of Lemma A.?? (and therefore on the curvature of m_0) and on the length of γ, such that every simple m_0–geodesic which meets the ε–neighborhood of γ and is different from γ must spiral around γ. Since ∂S has only finitely many components, this completes the proof. $\square$

When m_0 is a hyperbolic metric, namely has constant curvature -1, the proof of Lemma 2.?? can be made more explicit and gives the value $\varepsilon < \blacksquare$. See Figure 2.??.

Let $\alpha \in \mathcal{ML}(S)$ be a measured m_0–geodesic lamination, with underlying m_0–geodesic lamination λ_α. Since every component of λ_α is minimal by Proposition 2.??, a leaf of λ_α cannot be asymptotic to a component of ∂S. Therefore, if U is the neighborhood of ∂S provided by Lemma ??, $U \cap \lambda_\alpha$ consists only on components of ∂S. This provides a canonical decomposition of the transverse measure α as $\alpha = \alpha_1 + \alpha_2$, where the support λ_{α_1} of α_1 is equal to $\lambda_\alpha - U$ and where the support of α_2 is contained in ∂S. The geodesic lamination λ_{α_1} together with the transverse measure α_1 defines an element of $\mathcal{ML}_0(S)$. The transverse measure α_2 for ∂S restricts to a Dirac transverse measure on each boundary component, and is therefore determined by the corresponding weights. Therefore the splitting $\alpha = \alpha_1 + \alpha_2$ induces a map $\mathcal{ML}(S) \to \mathcal{ML}_0(S) \times [0, \infty[^d$, where d is the number of components of ∂S. Since U is independent of α, this map is clearly a homeomorphism.

We summarize this discussion in the following statement.

PROPOSITION. *There is a natural homeomorphism between $\mathcal{ML}(S)$ and $\mathcal{ML}_0(S) \times [0, \infty[^d$, where d is the number of components of ∂S.* $\square$

This reduces the analysis of $\mathcal{ML}(S)$ to that of $\mathcal{ML}_0(S)$.

As a variation of the space $\mathcal{PML}(S)$, we can define the space $\mathcal{PML}_0(S)$ of ***projective measured geodesic laminations disjoint from the boundary***, quotient of $\mathcal{ML}_0(S) - \{0\}$ by the natural action of $\mathbb{R}^+$. This is also the image of $\mathcal{ML}_0(S) - \{0\}$ in $\mathcal{PML}(S)$.

EXERCISE. Recall that the ***joint*** of the two topological spaces X and Y is the space $J(X,Y)$ obtained by, abstractly, joining each point of X to each point of Y by an arc. More formally, $J(X,Y)$ is the quotient of the product $X \times Y \times [0,1]$ by the equivalence relation which collapses each $\{x\} \times Y \times \{0\}$ and each $X \times \{y\} \times \{1\}$ to a point. Show that, if ∂S has $d \geqslant 1$ components, the space $\mathcal{PML}(S)$ is homeomorphic to the joint of $\mathcal{PML}_0(S)$ and of the sphere S^{d-1}.

2.5. Interval exchanges

Interval exchanges were originally introduced, apparently by V.I. Arno'ld ?, as a simple example of measure-preserving maps with complex dynamical properties. But they also turn out to provide examples of measured geodesic lamination.

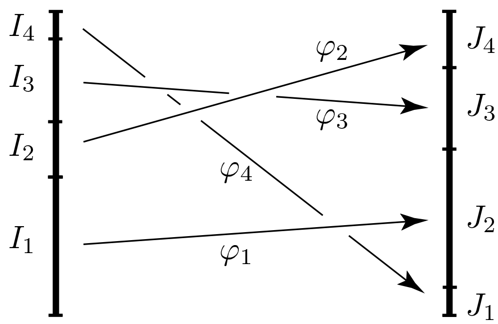

FIGURE ??. An interval exchange

An ***interval exchange*** of the interval $I = [a,b] \subset \mathbb{R}$ consists of: two subdivisions of I into n intervals $I_i = [x_i, x_{i+1}]$ and $J_j = [y_j, y_{j+1}]$, with $a = x_1 < x_2 \ldots < x_n < x_{n+1} = b$ and $a = y_1 < y_2 \ldots < y_n < y_{n+1} = b$; a permutation σ of the set $\{1, 2, \ldots, n\}$; for each i, an isometry $\varphi_i : I_i \to J_{\sigma(i)}$. In particular, I_i and $J_{\sigma(i)}$ must have the same length.

The union of the φ_i defines a 'map' $\varphi : I \to I$. Note that φ is not well defined at the cut point x_i for $i = 2, \ldots, n$, as the above definition of φ on the intervals I_{i-1} and I_i usually gives two distinct values for $\varphi(x_i)$ at $x_i = I_{i-1} \cap I_i$. We therefore have to think of φ as a multivalued function, which associates one or two values to each such x_i and associates a single value to any other $x \in I - \{x_2, \ldots, x_n\}$. In particular, φ is uniquely defined almost everywhere, and $\varphi(A)$ and $\varphi^{-1}(A)$ have the same Lebesgue measure as A for every measurable set $A \subset I$.

A multivalued map $\varphi : I \to I$ obtained in this way is an ***interval exchange map***. Note that an interval exchange involves more data than just the corresponding interval exchange map, as it also includes the two subdivisions of I into intervals I_i and J_j.

It may happen that φ is single valued at $x_i = I_{i-1} \cap I_i$; this occurs when, either $\sigma(i) = \sigma(i-1) + 1$ and both φ_{i-1} and φ_i are orientation-preserving, or

$\sigma(i) = \sigma(i-1) - 1$ and both φ_{i-1} and φ_i are orientation-reversing. In this case, φ induces an isometry from the interval $I_{i-1} \cup I_i$ to the interval $J_{\sigma(i-1)} \cup J_{\sigma(i)}$, and φ is now associated to an exchange of fewer intervals. We will say that the interval exchange is ***reduced*** if it is not of this type, namely if φ is really 2–valued at each cut point $x_i = I_{i-1} \cap I_i$ of the interval decomposition. One easily checks that each interval exchange map is associated to a unique reduced interval exchange.

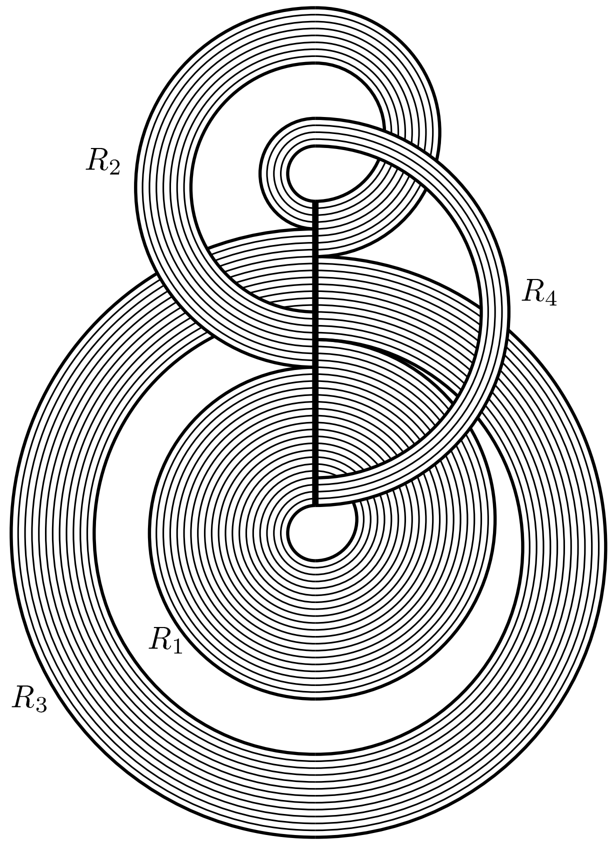

FIGURE ??. The suspension of the interval exchange of Figure ??

Consider the interval exchange, with interval exchange map $\varphi : I \to I$, associated to two decompositions of the interval I into n intervals I_i and J_j, respectively, and isometries $\varphi_i : I_i \to J_{\sigma(i)}$. The ***suspension*** Σ_φ of this interval exchange is the topological space defined as follows: For each $i = 1, \ldots, n$, consider the rectangle $R_i = I_i \times [0,1]$. Then, Σ_φ is the identification space obtained from the disjoint union of I and of the R_i by gluing each $I_i \times \{0\} \subset R_i$ to $I_i \subset I$ by the identity map, and each $I_i \times \{1\} \subset R_i$ to $J_{\sigma(i)} \subset I$ by the isometry $\varphi_i : I_i \to J_{\sigma(i)}$. In other words, Σ_φ is the natural extension to the multivalued context of the construction of the mapping torus of a continuous function, which is so omnipresent in topology. Note that the notation Σ_φ is somewhat abusive, in the sense that this suspension depends on the interval exchange, including the decomposition of I into intervals I_i, and not just on the interval exchange map $\varphi : I \to I$.

If, in the construction of the suspension Σ_φ, we think of the $I_i \times \{0\} \subset R_i$

as being glued 'on one side' of I, and of the $I_i \times \{1\} \subset R_i$ as being glued 'on the other side', then Σ_φ acquires a very familiar structure. Indeed, it is then presented as a special case of a fattened train track. The boundary of Σ_φ is the union of all arcs $\{x_i\} \times [0,1] \subset R_i$ and $\{x_{i+1}\} \times [0,1] \subset R_i$, where $I_i = [x_i, x_{i+1}]$ (note that there are two arcs $\{x_i\} \times [0,1]$, one in R_i and the other one in R_{i-1}, so that we need to indicate the rectangle to specify which arc we are considering). What would be the spikes of the complement of the fattened train track correspond to the points $x_i = I_{i-1} \cap I_i \in I$ where $\{x_i\} \times [0,1] \subset R_{i-1}$ meets $\{x_i\} \times [0,1] \subset R_i$, and to the points $y_j = J_{j-1} \cap J_j \in I$ where $\left\{\varphi^{-1}_{\sigma^{-1}(j-1)}(y_j)\right\} \times [0,1] \subset R_{\sigma^{-1}(j-1)}$ meets $\left\{\varphi^{-1}_{\sigma^{-1}(j)}(y_j)\right\} \subset R_{\sigma^{-1}(j)}$, for $2 \leqslant i \leqslant n$ and $2 \leqslant j \leqslant n$. The point $x_1 = y_1 \in I$ where $\{x_1\} \times [0,1] \subset R_1$ meets $\left\{\varphi^{-1}_{\sigma^{-1}(1)}(y_1)\right\} \subset R_{\sigma^{-1}(1)}$, and the point $x_{n+1} = y_{n+1} \in I$ where $\{x_{n+1}\} \times [0,1] \subset R_n$ meets $\left\{\varphi^{-1}_{\sigma^{-1}(n)}(y_{n+1})\right\} \subset R_{\sigma^{-1}(n)}$, are both smooth points of the boundary of this fattened train track.

Note that the boundary of Σ_φ is a union of finitely many cycles of arcs $\{x_i\} \times [0,1]$, where the orientation is reversed at each spike x_i, $y_j \in I$ with $2 \leqslant i, j \leqslant n$ and preserved at the two smooth points $x_1 = y_1$ and $x_{n+1} = y_{n+1}$. These cycles are embedded and disjoint in the 'generic' case where no x_i with $2 \leqslant i \leqslant n$ coincides with some y_j with $2 \leqslant j \leqslant n$. Also, note that the number of spikes of each spike is even: Indeed, the natural orientation of each arc $\{x_i\} \times [0,1]$ coming from the natural orientation of $[0,1]$ defines an orientation on the complement of the spikes, which is reversed at each spike.

To really turn Σ_φ into a thickened train track in a surface S, just pick a compact surface S' (not necessarily connected) and identifications between the cycles of Σ_φ and some components of ∂S, and then construct the surface S by gluing together Σ_φ and S' using these identifications. There are some conditions which need to be satisfied for Σ_φ to really be a thickened train track in S, namely: a cycle of $\partial\Sigma_\varphi$ with 0 spike is allowed to be glued to no component of $\partial S'$, and a cycle with $\geqslant 2$ spikes must be glued to some component of $\partial S'$; for a cycle of $\partial\Sigma_\varphi$ with 0 or 2 spikes, the corresponding component of $\partial S'$ is not allowed to bound a disk component of S'; for any annulus or Möbius strip component of S', at least one of its boundary components must be glued to a cycle with $\geqslant 2$ spikes. In addition, at least one boundary component of each component of S' must be glued to some cycle of Σ_φ if we want S to be connected. Any surface S constructed in this way will be called an ***enlargement*** of the suspension Σ_φ. The pair consisting of S and Σ_φ is an ***enlarged suspension*** of the interval exchange.

EXERCISE. Show that the Euler characteristic of the enlargement S is always negative.

Among all possible enlargements, there is a simplest one where each component of S' is an annulus. This enlargement is the ***minimum enlargement*** of Σ_φ. The Euler characteristic of S is then equal to $1 - n$, where n in the number of intervals of the exchange.

Another interesting case occurs when each component of S' is a disk. Such an elargement exists only when each cycle of Σ_φ has at least 4 spikes. If such an enlargement exists, the pair consisting of S and Σ_φ is the ***closed suspension*** of the interval exchange.

EXERCISE. Consider a reduced interval exchange of $n \geqslant 3$ intervals where the interval exchange map $\varphi : I \to I$ is orientation-preserving, wherever this makes sense namely outside of the cut points. Assume in addition that φ fixes no end point of I (otherwise, φ would split as the 'union' of the identity map on the interval touching this en point and of an exchange of fewer intervals). Show that every cycle of $\partial\Sigma_\varphi$ has at least 4 spikes, so that φ admits a closed suspension.

Let a ***horizontal curve*** in Σ_φ be a closed or bi-infinite differentiable curve in $\Sigma_\varphi \subset S$ which is made up of arcs $\{x\} \times [0,1]$ contained in the rectangles R_i. Such a horizontal curve is ***regular*** if does not pass through any of the cut points $x_i = I_{i-1} \cap I_i$, $i = 2, \dots, n$, and it is ***singular*** otherwise. Note that there are uncountably many horizontal curves, that all but finitely of them are regular, and that regular horizontal curves are simple and pairwise disjoint.

Let $\widetilde{\Sigma}_\varphi$ be the pre-image of Σ_φ in the universal covering $\widetilde{S}$, and define a horizontal curve of $\widetilde{\Sigma}_\varphi$ as above. If we consider $\widetilde{\Sigma}_\varphi$ as a thickened train track, every horizontal curve of $\widetilde{\Sigma}_\varphi$ is strongly carried by this thickened train track and, by Proposition 1.??, is quasi-geodesic in S. In particular, a horizontal curve of $\widetilde{\Sigma}_\varphi$ cannot be closed, and is therefore bi-infinite. By Proposition B.??, for each horizontal curve l of $\widetilde{\Sigma}_\varphi$, there consequently exists a unique geodesic l^* of S which is homotopic to this curve by a homotopy which moves points by a uniformly bounded distance.

By obvious extension of the case of Σ_φ, a horizontal curve l of $\widetilde{\Sigma}_\varphi$ it regular if it does not meet any spike of $\widetilde{\Sigma}_\varphi$, namely any point in the pre-image of the cut points $x_1, \dots, x_n$; otherwise, it is singular. Note that, when a horizontal curve l is singular if and only if it meets the boundary $\partial\widetilde{\Sigma}_\varphi$. More generally, a singular horizontal curve l is ***almost regular*** if there is a side of l such that, whenever l meets $\partial\widetilde{\Sigma}_\varphi$ at a point x which is not a spike, Σ_φ is locally on that side of l near x.

In $\widetilde{\Sigma}_\varphi$, consider all horizontal curves which are either regular or almost regular. In the space $G(\widetilde{S})$, let $\widetilde{\lambda} \subset G(\widetilde{S})$ consist of those geodesics of $\widetilde{S}$ which are associated to these regular and almost regular horizontal curves of $\widetilde{\Sigma}_\varphi$.

LEMMA. *The subset $\widetilde{\lambda} \subset G(\widetilde{S})$ defines a geodesic lamination of $\widetilde{S}$, which is the pre-image of a geodesic lamination $\lambda \subset S$.*

PROOF. We will use the coding of geodesics by bi-infinite edge paths in $\widetilde{\Sigma}_\varphi$ developed in §1.8. By construction, each geodesic l^* of $\widetilde{\lambda}$ is described by the bi-infinite edge path associated to corresponding horizontal curve l of $\widetilde{\Sigma}_\varphi$.

To show that $\widetilde{\lambda} \subset G(\widetilde{S})$ is a geodesic lamination of $\widetilde{S}$, we need to prove that any two geodesics of $\widetilde{\lambda}$ which are distinct are disjoint, and that $\widetilde{\lambda}$ is a closed subset of $G(\widetilde{S})$.

Let l_1 and l_2 be two distinct horizontal curves of $\widetilde{\Sigma}_\varphi$, and let γ_1 and γ_2 the bi-infinite edge paths which they respectively follow in this thickened train track. Suppose that γ_1 and γ_2 cross each other, in the sense of §1.8. Then l_1 and l_2 necessarily coincide on any edge and switch which they both cross; indeed, we could otherwise find a regular horizontal leaf l_3 which separates l_1 from l_2, and consequently prevents γ_1 from crossing γ_2. If we look at the largest common edge path of γ_1 and γ_2 and at the way how l_1 and l_2 can diverge at the end points of this edge path, we conclude from the fact that γ_1 and γ_2 cross each other that l_1 and l_2 are neither regular nor almost regular.

Consequently, if l_1 and l_2 are two leaves of $\widetilde{\Sigma}_\varphi$ which are regular or almost regular, the associated bi-infinite edge paths γ_1 and γ_2 cannot cross each other. By Lemma 1.??, it follows that the geodesics l_1^* and l_2^* associated to l_1 and l_2 (and γ_1 and γ_2) are disjoint or equal.

As a consequence, any two distinct geodesics of $\widetilde{\lambda}$ are disjoint.

To prove that $\widetilde{\lambda}$ is closed in $G(\widetilde{S})$, let $l^* \in G(\widetilde{S})$ be the limit of a sequence of geodesics $l_n^* \in \widetilde{\lambda}$. We want to show that l^* is in $\widetilde{\lambda}$. By definition, for every n, there is a horizontal curve l_n of $\widetilde{\Sigma}_\varphi$ which is homotopic to l_n^* by a bounded homotopy. By Lemma 1.??, the set of geodesics which are weakly carried by $\widetilde{\Sigma}_\varphi$ is closed, and l^* is therefore weakly carried by $\widetilde{\Sigma}_\varphi$. Let γ be the bi-infinite edge path of $\widetilde{\Sigma}_\varphi$ associated to l^*. By Lemma 1.??, the edge path γ_n associated to l_n converges to γ on every compact subset of $\widetilde{\Sigma}_\varphi$. In other words, for every finite union K of edges of $\widetilde{\Sigma}_\varphi$, the edge path γ_n followed by l_n coincides with γ on K. Passing to a subsequence through a diagonal argument, we can consequently assume that l_n converges, uniformly on every compact subset of $\widetilde{\Sigma}_\varphi$, to a horizontal curve l.

By construction, this horizontal curve l follows the edge path γ in $\widetilde{\Sigma}_\varphi$. As a consequence, it stays within bounded distance of the geodesic l^*, since this geodesic is associated to the edge path γ. Also, by construction, l is a limit of almost regular horizontal curves, and it easily follows that l is almost regular. Therefore, l^* is the geodesic associated to the almost regular horizontal curve l, the l^* consequently belongs to $\widetilde{\lambda}$. This concludes the proof that $\widetilde{\lambda}$ closed in $G(\widetilde{S})$.

As a consequence, $\widetilde{\lambda}$ is a geodesic lamination of $\widetilde{S}$. □

■use the terminology that a geodesic is tracked by an edge path and/or a horizontal curve?

■horizontal curves → leaves?

By construction, $\widetilde{\lambda}$ is invariant under the action of the fundamental group $\pi_1(S)$ on $\widetilde{S}$ and $G(\widetilde{S})$. By ??, it consequently is the pre-image of a geodesic lamination λ in S.

Let R be the space of regular horizontal curves of $\widetilde{\Sigma}_\varphi$, endowed for the topology of uniform comvergence on compact subsets. This set carries a natural measure. Indeed, for an element l of R, a small neighborhood U_k of l consists of those $l' \in R$ which meet a small arc k contained in the pre-image of $I \subset \Sigma_\varphi$, namely in the union of the switches of $\widetilde{\Sigma}_\varphi$, and meeting l; note that this holds only if l is regular. The Lebesgue measure of I defines a measure on k and, since each $l' \in U_k$ crosses k in exactly one point, this also defines a measure on U_k. If U_k and $U_{k'}$ are two such neighborhoods of l, it easily follows from the fact that the interval exchange map $\varphi : I \to I$ respects the Lebesgue measure that the measures so defined on U_k and $U_{k'}$ coincide on $U_k \cap U_{k'}$. In other words, this measure is locally independent of the choice of the neighborhood U_k. Therefore, this gives a well defined measure on R.

Consider the map $R \to \widetilde{\lambda}$ which associates to the regular horizontal leaf $l \in R$ its corresponding geodesic $l^* \in \widetilde{\lambda}$. This map is continuous by Lemma 1.??. We can then endow $\widetilde{\lambda}$ with the push forward of the measure of R, for which the measure of a subset of $\widetilde{\lambda}$ is the measure of its pre-image in R. By construction, this measure is invariant under the action of $\pi_1(S)$, and gives a transverse measure α for λ by ??. Using the point of view of §2.??, we can also extend α to a $\pi_1(S)$–invariant measure on $G(\widetilde{S})$, whose support is contained in $\widetilde{\lambda}$.

From the definition of "almost regular", every almost regular horizontal curve of $\widetilde{\Sigma}_\varphi$ is a limit of regular horizontal curves, for the topology of convergence on compact sets. It follows that the image of R is dense in $\widetilde{\lambda}$. Since the measure of R has full support, we conclude that the support of the measure α is exactly equal to $\widetilde{\lambda}$. In other words, α is a full support transverse measure for λ.

In this way, we have associated to the interval exchange $\varphi : I \to I$ a measured geodesic lamination $\alpha \in \mathcal{ML}(S)$, where S is an enlargement of the suspension Σ_φ of φ.

■finish

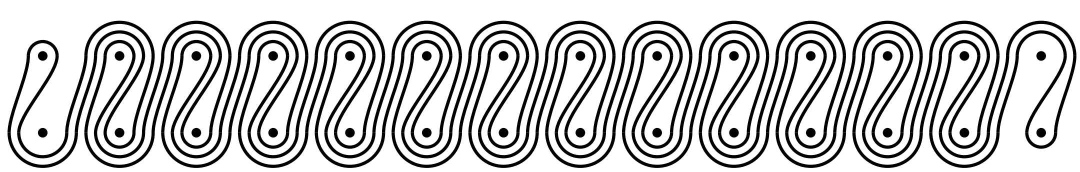

CHAPTER III

THE SPACE $\mathcal{ML}(S)$ OF MEASURED GEODESIC LAMINATIONS

3.1. Measured geodesic laminations carried by train tracks

A measured m_0–geodesic lamination α is ***(strongly) carried*** by the fattened train track Φ if its underlying geodesic lamination λ_α is strongly carried by Φ. It is ***(weakly) carried*** by the train track Θ if λ_α is weakly carried by Θ. We often omit the adverbs "weakly" and "strongly", since the first one only applies to a train track and the second one only to a fattened train track. We denote by $\mathcal{ML}(\Phi, m_0)$ the subset of $\mathcal{ML}(S)$ consisting of those m_0–geodesic laminations which are strongly carried by Φ, and by $\mathcal{ML}(\Theta)$ the subset of those geodesic laminations which are weakly carried by Θ.

We remind the reader that $\mathcal{ML}(\Phi, m_0)$ depends on the metric m_0 whereas $\mathcal{ML}(\Theta)$ is metric independent. We will use these subsets to locally parametrize the space $\mathcal{ML}(S)$.

Proposition. *Let α be a measured geodesic lamination whose underlying geodesic lamination is not contained in the boundary ∂S. Then, there is a negatively curved metric m and a fattened train track Φ such that every measured m–geodesic lamination which is sufficiently close to α in $\mathcal{ML}(S)$ is carried by Φ.*

In other words, the set $\mathcal{ML}(\Phi, m)$ is a neighborhood of α in $\mathcal{ML}(S)$. If Θ is the train track associated to Φ, well defined up to isotopy, the set $\mathcal{ML}(\Phi, m)$ is contained in $\mathcal{ML}(\Theta)$. We immediately conclude:

Corollary. *Let α be a measured geodesic lamination whose underlying geodesic lamination is not contained in the boundary ∂S. Then, there is a train track Θ such that $\mathcal{ML}(\Theta)$ is a neighborhood of α in $\mathcal{ML}(S)$.* □

We will use the neighborhoods $\mathcal{ML}(\Theta)$ and $\mathcal{ML}(\Phi, m)$ to parametrize $\mathcal{ML}(S)$ locally and globally. They each have their own advantages. The fact that $\mathcal{ML}(\Theta)$ is metric independent makes it conceptually simpler, and more convenient for global constructions. However when we will need explicit estimates for local computations, as in Chapters ?? and ??, it will be useful to assume that the geodesic laminations considered are carried by the same fattened train track, so that $\mathcal{ML}(\Phi, m)$ will be better adapted to that context.

The proof of Proposition ?? will take a while. It is based on the following lemma.

Lemma. *Let λ be a geodesic lamination which is not contained in the boundary of S. Then, there is a negatively curved metric m and a fattened train track Φ such that every m–geodesic lamination containing the m–geodesic lamination corresponding to λ is strongly carried by Φ.*

Proof of Proposition ?? from Lemma ??. Let m and Φ be provided by Lemma ?? applied to the geodesic lamination $\lambda = \lambda_\alpha$ underlying α. We want to show that the corresponding subset $\mathcal{ML}(\Phi, m)$ is a neighborhood of α.

Suppose that this is not the case. Then, since $\mathcal{ML}(S)$ has a countable neighborhood basis by Lemma ??, we can find a sequence $\alpha_n \in \mathcal{ML}(S)$ which converges to α and such that the corresponding underlying m–geodesic laminations λ_{α_n} are not carried by Φ. Passing to a subsequence if necessary, we can assume that the λ_{α_n} converge to some m–geodesic lamination λ' for the Hausdorff topology of $\mathcal{L}(S)$.

For every small arc k which is transverse to λ and is in generic position with respect to simple m–geodesics, the measure deposited by α on k is non trivial. Since the measure deposited by α_n on k converges to that deposited by α, it follows that the support λ_{α_n} of α_n crosses k for n large enough. As a consequence, λ' crosses k. Since this is true for every small arc k which is transverse to λ and in generic position with respect to simple m–geodesics, we conclude that λ is contained in λ'. By construction of Φ, the m–geodesic lamination λ' is therefore carried by Φ.

But then λ_{α_n}, which is close to λ' for the Hausdorff topology, would be carried by Φ for n sufficiently large, contradicting our choice of the α_n.

This concludes the proof of Proposition ??, modulo Lemma ??. □

It now remains to prove Lemma ??.

Proof of Lemma ??. Fix a negatively curved metric m_0. Because λ is not contained in the boundary of S, there is a 2–sided simple closed geodesic γ_1 which has non-trivial transverse intersection with λ. We can extend γ_1 to a family γ of disjoint 2–sided simple closed geodesics $\gamma_1, \gamma_2, \dots, \gamma_n$ splitting S into a family of pairs of pants or "Möbius pants" (see [**BoasMath**.**Mag**.**95**] for the terminology), namely of spheres minus 3 disks or projective planes minus 2 disks. ■γ_1 may not exist if S is a pair of pants or Möbius pants.

Given a family of pairwise disjoint arcs embedded in a pair of pants or Möbius pants, which go from the boundary to the boundary and cannot be deformed to the boundary, these arcs fall in at most 3 distinct isotopy classes. For each surface, there are two possibilities for such a triple of isotopy classes, and these are indicated on Figure ??. We can choose the γ_i so that, for each component P of $S - \gamma$, the leaves of $\lambda \cap P$ realize the maximum number of such isotopy classes of arcs. Namely, if P is a pair of pants containing 0, 1, 2 components of ∂S, then the leaves of $\lambda \cap P$ fall into exactly 3, 2 or 1 isotopy classes, respectively; if P is a pair of Möbius pants containing 0 or 1 component of ∂S, then the leaves of $\lambda \cap P$ fall into exactly 3 or 1 isotopy classes, respectively. The fact that we can choose the γ_i in this way is an easy exercise on the classification of surfaces, focusing attention on one leaf of $\lambda - \gamma_1$ going from γ_1 to itself in $S - \gamma_1$.

Choose an orientation of S near γ, and let $D : S \to S$ denote the left handed Dehn twist around γ defined by this orientation.

For each leaf k of $\lambda - \gamma$, consider the geodesic which is asymptotic to the curve that follows k and spirals to the right along the two components of γ adjacent to k. Since, up to isotopy in $S - \gamma$, there are only finitely many such k, this defines finitely many infinite geodesics in S. Let γ^+ be the union of γ and of these geodesics. By construction, the complement of γ^+ consists of finitely many infinite triangles, and the geodesic lamination γ^+ is therefore maximal.

Lemma. *If λ' is a geodesic lamination containing λ, then the geodesic lamination $D^n(\lambda')$ converges to γ^+ in $\mathcal{L}(S)$ as n tends to ∞.*

Let Φ be a fattened train track carrying the m_0–geodesic lamination corresponding to γ^+. Lemma ?? says that the set $\mathcal{L}(\Phi, m_0)$, consisting of all m_0–geodesic laminations carried by Φ, is open in $\mathcal{L}(S)$ for the Hausdorff topology. Since the set X of all geodesic laminations λ' containing λ is compact, Lemma ?? implies that there is an integer $n > 0$ such that X is contained in $D^{-n}(\mathcal{L}(\Phi, m_0)) = \mathcal{L}(D^{-n}(\Phi), D^{-n}(m_0))$. In other words, for every geodesic lamination λ' containing λ, the corresponding $D^{-n}(m_0)$–geodesic lamination is strongly carried by the fattened train track $D^{-n}(\Phi)$. This concludes the proof of Proposition ??. □

The train track provided by the proof of Proposition ?? satisfies some additional topological properties. Since we will need those later on, we summarize them here in the following statement.

A train track Θ_0 in S is said to be ***boundary-free*** if no component of ∂S is weakly carried by Θ_0. This is easily seen to be equivalent to the property that no component of $S - \Theta_0$ that meets ∂S is a semi-open annulus with 0 spike (Hint: Apply Lemma ?? to a train track obtained by adding to Θ_0 parallel copies of a few components of ∂S). The train track Θ_0 is ***maximal boundary-free*** if it is boundary-free and if it cannot be enlarged to another boundary-free train track. Simple combinatorics show that this is equivalent to the property that each component of $S - \Theta$ is, either a disk with 3 spikes, or a semi-open annulus with one spike and containing a component of ∂S. Beware that a maximal boundary-free train track is not necessarily maximal as a train track (if S has non-empty boundary, a boundary-free train track is always contained in a larger train track), but is maximal only among boundary-free train tracks.

The following complement then comes for free from the proof of Proposition ??.

Complement. *In the conclusions of Proposition ??, we can also add the following requirements for the train track Θ: For every component γ of ∂S, there is a component of Θ which is a closed curve parallel to γ; the train track Θ_0 obtained from Θ by removing these parallel copies of boundary components of S is maximal boundary-free; for every edge e of Θ_0, the geodesic lamination λ_α underlying α crosses e, in the sense that there is a bi-infinite edge path associated to a leaf of λ_α which contains e.* □

3.2. Characterizing measured geodesic laminations by edge weight systems

In §3.1, we saw that every measured geodesic lamination $\alpha \in \mathcal{ML}(S)$ which is not contained in the boundary admits a neighborhood in $\mathcal{ML}(S)$ of the form $\mathcal{ML}(\Theta)$, where Θ is a train track in S. In this section and the next one, we determine the topology of the space $\mathcal{ML}(\Theta)$ for a fixed train track Θ.

As usual, let $\widetilde{\Theta}$ denote the preimage of Θ in the universal covering $\widetilde{S}$. In §1.8, we saw that a bi-infinite edge path γ in $\widetilde{\Theta}$ determines a unique geodesic $G(\gamma)$ of $\widetilde{S}$. This establishes a one-to-one correspondence between the set $\Gamma\big(\widetilde{\Theta}\big)$ of bi-infinite edge paths in $\widetilde{\Theta}$ and the set $G(\widetilde{\Theta})$ of those geodesics of $\widetilde{S}$ which are weakly carried by $\widetilde{\Theta}$ (Proposition 1.??).

Given an edge path γ (finite or infinite) in $\widetilde{\Theta}$, let $G(\gamma) \subset G(\widetilde{\Theta})$ consist of those geodesics which are asymptotic to a bi-infinite edge path γ' containing γ. In particular, if γ is bi-infinite, the set $G(\gamma)$ consists of exactly the geodesic which we have denoted $G(\gamma)$ so far, so that our notation is consistent modulo the abuse of identifying a 1-element set to its element.

LEMMA. *For every finite edge path γ in $\widetilde{\Theta}$, the subset $G(\gamma)$ is compact and open in $G(\widetilde{\Theta})$.*

PROOF. Proposition 1.?? expresses the topology of $G(\widetilde{\Theta})$ in terms of edge paths. It is then immediate that $G(\gamma)$ is open and closed in $G(\widetilde{\Theta})$. By definition, every geodesic of $G(\gamma)$ is at uniformly bounded distance from a curve carried by $\widetilde{\Theta}$ and following γ. In particular, every geodesic of $G(\gamma)$ meets a compact subset K of $\widetilde{S}$. Since the set of those geodesics which meet K is compact in $G(\widetilde{S})$, and since $G(\gamma)$ is closed in $G(\widetilde{\Theta})$ which is closed in $G(\widetilde{S})$, this proves the compactness of $G(\widetilde{\Theta})$. □

Let α be a measured geodesic lamination which is weakly carried by Θ. For every edge path γ in $\widetilde{\Theta}$, let $\alpha(\gamma)$ denote the mass $\alpha(G(\gamma)) \geqslant 0$ of the compact subset $G(\gamma)$ defined above, for the measure induced by α on $G(\widetilde{S})$.

If γ is a finite edge path in Θ, we define $\alpha(\gamma) = \alpha\big(\widetilde{\gamma}\big)$ where $\widetilde{\gamma}$ is an arbitrary lift of γ to a finite edge path in $\widetilde{\Theta}$. By invariance under the action of $\pi_1(S)$, this does not depend on the choice of the lift $\widetilde{\gamma}$. Finally, if e is an edge of Θ or $\widetilde{\Theta}$, we can consider the edge path $\langle e\rangle$, and define $\alpha(e) = \alpha(\langle e\rangle)$.

The train track Θ has finitely many edges e. We now show that the finitely many numbers $\alpha(e)$ completely determine the measured geodesic lamination $\alpha \in \mathcal{ML}(\Theta)$.

LEMMA. *Fix a finite edge path γ of the train track Θ. Then, for every $\alpha \in \mathcal{ML}(\Theta)$, the number $\alpha(\gamma)$ is a piecewise linear function (depending only on the combinatorics of γ) of the weights $\alpha(e)$ assigned by α to the edges e of Θ.*

PROOF. It is convenient to lift γ to the preimage $\widetilde{\Theta}$ of Θ in the universal covering $\widetilde{\Theta}$. We will argue by induction on the length n of the edge path

$\gamma_n = \langle e_1, e_2, \ldots, e_n \rangle$ of $\widetilde{\Theta}$, assuming that this edge path starts at a fixed edge e_1.

Select a 'right' and a 'left' side for e_1. Then, this distinguishes a right and a left side for γ_n.

For $g \in G(e_n)$, the bi-infinite edge path corresponding to g contains, either γ_n, or some $\langle e_i', e_{i+1}, \ldots, e_n \rangle$ where $1 \leqslant i < n$ and where $e_i' \neq e_i$. Let $G_n^{\mathrm{l}} \subset G(\widetilde{\Theta})$ be the union of those $G(\langle e_i', e_{i+1}, \ldots, e_n \rangle)$ where e_i' branches in on the left of γ. Similarly, let G_n^{r} be the union of those $G(\langle e_i', e_{i+1}, \ldots, e_n \rangle)$ there e_i' branches in on the right of γ_n. Then $G(e_n)$ is the disjoint union of $G(\gamma_n)$, G_n^{l} and G_n^{r}.

We will prove by induction on n that $\alpha(\gamma_n)$, $\alpha\left(G_n^{\mathrm{l}}\right)$ and $\alpha(G_n^{\mathrm{r}})$ are piecewise linear functions of the weights $\alpha(e)$.

The property is trivially true when $n = 1$. Assume as induction hypothesis that it holds for $n-1$. We want to prove it for n.

Consider the switch s where e_{n-1} meets e_n. Let a_-^{l} (resp. a_-^{r}, a_+^{l}, a_+^{r}) denote the sum of the weights of the edges entering s on the same side as e_{n-1} (resp. e_{n-1}, e_n, e_n) and on the left (resp. right, left, right) side of $\langle e_1, e_2, \ldots, e_n \rangle$.

Then, analyzing what can happen to the geodesics which are carried by ■$\widetilde{\tau}$ and pass through the switch s, we find that

$$\begin{aligned}
\alpha\left(G_n^{\mathrm{l}}\right) &= \min\left\{\alpha(e_n), \max\left\{\alpha\left(G_{n-1}^{\mathrm{l}}\right) + a_-^{\mathrm{l}} - a_+^{\mathrm{l}}, 0\right\}\right\} \\
\alpha\left(G_n^{\mathrm{r}}\right) &= \min\left\{\alpha(e_n), \max\left\{\alpha\left(G_{n-1}^{\mathrm{r}}\right) + a_-^{\mathrm{r}} - a_+^{\mathrm{r}}, 0\right\}\right\} \\
\alpha(\gamma_n) &= \alpha(e_n) - \alpha\left(G_n^{\mathrm{l}}\right) - \alpha\left(G_n^{\mathrm{r}}\right),
\end{aligned}$$

■explain better. Restrict to generic case and draw 4 pictures? which clearly concludes the proof by induction. □

A piecewise linear map $F : \mathbb{R}^p \to \mathbb{R}^q$ is in general not differentiable. However, F has a well-defined ***tangent map*** $T_x F : \mathbb{R}^p \to \mathbb{R}^q$ at every $x \in \mathbb{R}^p$, defined as follows: For $v \in \mathbb{R}^p$, choose a curve $\sigma : [0, \varepsilon[\to \mathbb{R}^p$ which passes through x and is tangent to v at $t = 0$. Then, $T_x F(v)$ is defined to be the tangent vector of the curve $t \mapsto F(\sigma(t))$ at $t = 0$. It is immediate that this tangent vector exists and is independent of the choice of the curve σ.

The tangent map $T_x F$ is in general not linear. However, it is positive homogeneous, in the sense that $T_x F(\lambda v) = \lambda T_x F(v)$ for every $\lambda \geqslant 0$ (but $T_x F(-v)$ may be different from $-T_x F(v)$). If we have chosen norms $\| \ \|$ on $\mathbb{R}^p$ and $\mathbb{R}^q$, this enables us to define the ***norm*** of $T_x F$ as

$$\|T_x F\| = \max_{v \neq 0} \frac{\|T_x F(v)\|}{\|v\|}.$$

The following complement is an immediate by-product of the proof of Lemma ??.

Complement. *Let γ be a finite edge path of $\widetilde{\Theta}$. Let F_γ be the piecewise linear function provided by Lemma ?? which, for every $\alpha \in \mathcal{ML}(\Theta)$, expresses the number $\alpha(\gamma)$ in terms of the weights $\alpha(e)$ assigned by α to the (finitely many) edges e of Θ. Then, the norms of the tangent maps of F_γ are uniformly bounded by a constant times the length of γ.* □

PROPOSITION. *For a train track Θ, the measured geodesic lamination $\alpha \in \mathcal{ML}(\Theta)$ is completely determined by the weights $\alpha(e)$ it assigns to the edges e of Θ.*

PROOF. We will show that, for every continuous function $\varphi : G(\widetilde{S}) \to \mathbb{R}$ with compact support $\mathrm{Supp}(\varphi)$, the integral $\alpha(\varphi)$ is completely determined by φ and the weights $\alpha(e)$.

For every edge e_0 of $\widetilde{\Theta}$, the subset $G(e_0)$ is open in $G(\widetilde{\Theta})$ by Lemma ??. Since the support of α is contained in $G(\widetilde{\Theta})$, we can therefore, using a suitable partition of unity, restrict attention to the case where the intersection of $G(\widetilde{\Theta})$ with the support of φ is contained in some $G(e_0)$. Namely, every geodesic of $G(\widetilde{\Theta})$ is associated to a bi-infinite edge path containing e_0.

For a given integer $r \geqslant 0$, let Γ_r be the set of all edge paths $\langle e_{-r}, \dots, e_{-1}, e_0, e_1, \dots, e_r\rangle$ of length $2r+1$ centered at e_0. Then $G(e_0)$ is the union of the $G(\gamma)$ with $\gamma \in \Gamma_r$, and any two such $G(\gamma)$ are disjoint. Also, there is are two constants M_1 and M_2, depending only on the quasi-geodicity constant of Proposition 1.?? and on the diameters of the edges of $\widetilde{\Theta}$, such that any two geodesics of g, $g' \in G(\gamma)$ stay within a distance of M_1 from each other over the rM_2–neighborhood of e_0; by a geometric estimate, it follows that the distance from g to g' is an $O\left(e^{-Ar}\right)$ for some constant $A > 0$ independent of r. Therefore, the diameter of each $G(\gamma)$ with $\gamma \in \Gamma_r$ is an $O\left(e^{-Ar}\right)$; in particular, this diameter uniformly tends to 0 as r tends to ∞.

Approximating the restriction of φ to $G(\widetilde{\Theta})$ by suitable step functions, and remembering that $\alpha(\gamma)$ is by definition $\alpha(G(\gamma))$ for every edge path, we conclude that

$$\alpha(\varphi) = \lim_{r\to\infty} \sum_{\gamma\in\Gamma_r} \alpha(\gamma)\,\varphi(g_\gamma)$$

where, for an edge path γ, g_γ denotes an arbitrary element of $G(\gamma)$.

By Lemma ??, the number $\alpha(\gamma)$ associated to an edge path γ is completely determined by the edge weights $\alpha(e)$. The above formula then shows that the $\alpha(e)$ also determine the measure defined by α on $G(\widetilde{S})$. Since, by §1.??, the measured geodesic lamination α can be identified to the measure it defines on $G(\widetilde{S})$, this completes the proof. □

We can rephrase Proposition ?? in the following way.

For a measured geodesic lamination $\alpha \in \mathcal{ML}(\Theta)$ the edge weights $\alpha(e) \geqslant 0$ satisfy certain obvious ***switch relations***. Namely, for every switch s of Θ, if e_1, … , e_p are the edges adjacent to one side of s and if e_{p+1}, … , e_{p+q} are the edges adjacent to the other side (counting twice any edge whose two short sides are in s), then

$$\sum_{i=1}^{p} \alpha(e_i) = \sum_{j=p+1}^{p+q} \alpha(e_j).$$

Indeed, if we lift s to a switch $\widetilde{s}$ of $\widetilde{\Theta}$ lifting s, each side of this equation expresses the α–mass of the subset of $G(\widetilde{\Theta})$ consisting of those geodesics which correspond bi-infinite edge path in $\widetilde{\Theta}$ that cross the switch $\widetilde{s}$.

Let $\mathcal{W}(\Theta, [0, \infty[)$ be the space of all such ***edge weight systems*** for Θ. Namely, if $E(\Theta)$ is the set of edges of Θ, $\mathcal{W}(\Theta, [0, \infty[)$ consists of all maps $a : E(\Theta) \to [0, \infty[$ such that, for every switch s of Θ,

$$\sum_{i=1}^{p} a(e_i) = \sum_{j=p+1}^{p+q} a(e_j)$$

where $e_1, \ldots, e_p$ are the edges adjacent to one side of s and if $e_{p+1}, \ldots, e_{p+q}$ are the edges adjacent to the other side.

Then, we can rephrase Proposition ?? as saying that the natural map $\mathcal{ML}(\Theta) \to \mathcal{W}(\Theta, [0, \infty[)$ is injective.

LEMMA. *The map $\mathcal{ML}(\Theta) \to \mathcal{W}(\Theta, [0, \infty[)$ is continuous.*

PROOF. We need to show that, for every edge e of $\widetilde{\Theta}$, the map $\alpha \mapsto \alpha(e) = \alpha(G(e))$ is continuous.

For this, we have to remember the following property of the weak* convergence of measures. Let μ_0 be a (Radon) measure on a locally compact space X, and suppose that the measure μ tends to μ_0 for the weak* topology. Then, if A is a relatively compact subset of X whose frontier δA is such that $\alpha(\delta A) = 0$, the mass $\mu(A)$ converges to $\mu_0(A)$ as μ tends to μ_0. (See for instance [**Bourbaki**][**Rudin**?]). Note that the condition $\alpha(\delta A) = 0$ is absolutely necessary. For instance, in $X = \mathbb{R}$, the Dirac measure μ_x of weight 1 based at x converges to μ_0 as x tends to 0, since $\mu_x(\varphi) = \varphi(x)$ converges to $\mu_0(\varphi) = \varphi(0)$ for every continuous function; however, $\mu_x([0, 1]) = 0$ does not converge to $\mu_0([0, 1]) = 1$ as x tends to 0 on the left.

In our case, we saw that $G(e)$ is open and closed in $G(\widetilde{\Theta})$. Therefore, the frontier $\delta G(e)$ of $G(e)$ in $G(\widetilde{S})$ is disjoint from $G(\widetilde{\Theta})$. Given $\alpha_0 \in \mathcal{ML}(\Theta)$, this implies that $\alpha_0(\delta G(e))$ is equal to 0. Therefore, as $\alpha \in \mathcal{ML}(\Theta)$ tends to α_0, $\alpha(e) = \alpha(G(e))$ converges to $\alpha_0(e) = \alpha(G(e))$ since the measure induces by α on $G(\widetilde{S})$ converges to the measure induced by α_0 for the weak* topology. This concludes the proof of Lemma ??. □

LEMMA. *The map $\mathcal{ML}(\Theta) \to \mathcal{W}(\Theta, [0, \infty[)$ is proper.*

PROOF. □

3.3. Realizing edge weight systems by mesured geodesic laminations

This section is devoted to proving that, for a train track Θ, the map $\mathcal{ML}(\Theta) \to \mathcal{W}(\Theta, [0, \infty[)$ introduced in §3.2 is surjective. Note that, if we replace each non-simple switch of Θ by a small edge as in Figure 1.??, the train track Θ' so obtained is such that $\mathcal{ML}(\Theta') = \mathcal{ML}(\Theta)$ and $\mathcal{W}(\Theta', [0, \infty[)$ is naturally isomorphic to $\mathcal{W}(\Theta, [0, \infty[)$. Therefore, we can assume without loss of generality that all switches of Θ are simple. Let Φ be a fattening of Θ.

Consider an edge weight system $a \in \mathcal{W}(\Theta, [0, \infty[)$ namely, if $E(\Theta)$ denotes the set of edges of Θ, a map $a : E(\Theta) \to [0, \infty[$ satisfying the switch relations.

For each edge e of Θ, we will also denote by e the corresponding edge of Φ. Choose an identification of the edge e of Φ with a euclidean rectangle $I_e \times J_e$ where I_e and J_e are closed intervals in $\mathbb{R}$, where the length of I_e is equal to the weight $a(e)$, and where the ties of e correspond to the arcs $I_e \times *$. Because a satisfies the switch relations and all switches of Φ are simple, we can arrange that, on each switch s of Φ, the euclidean metrics of the adjacent edges $e \cong I_e \times J_e$ fit together to give a well-defined euclidean metric on s. (If a switch of Φ was not simple, we would need the weights $a(e)$ of the adjacent edges to satisfy certain linear inequalities in addition to the switch relation to be able to achieve this.)

Foliate each edge $e \cong I_e \times J_e$ by the arcs $* \times J_e$. This defines a foliation L_a of Φ, whose only singularities correspond to the spikes of $S - \Phi$. The euclidean structure of the intervals I_e defines a tranverse measure μ_a for L_a. Indeed, if k is an arc transverse to L_a and contained in an edge e of Φ (we allow the end points of k to be on L_a), we define $\mu_a(k)$ as the length of the image of k under the projection $e \cong I_e \times J_e \to I_e$. The compatibility condition on the identifications $e \cong I_e \times J_e$ at the switches of Φ guarantees the invariance of μ_a under homotopy respecting λ.

Lift L_a to a foliation $\widetilde{L}_a$ of the preimage $\widetilde{\Phi}$ of Φ in the universal covering $\widetilde{S}$. Every non-singular leaf of $\widetilde{L}_a$ is a bi-infinite curve carried by $\widetilde{\Phi}$, and is therefore asymptotic to a unique geodesic of $\widetilde{S}$. Two distinct non-singular leaves of $\widetilde{L}_a$ are disjoint; it follows that their associated geodesics are disjoint or equal. Therefore, the union $\widetilde{\lambda}$ of all such geodesics of $\widetilde{S}$ associated to non-singular leaves of $\widetilde{L}_a$ is a $\pi_1(S)$–invariant family of pairwise disjoint geodesics of $\widetilde{S}$. Its closure $\widetilde{\lambda}_\alpha$ is a $\pi_1(S)$–invariant geodesic lamination of $\widetilde{S}$, and is therefore the preimage of a geodesic lamination λ_α of S. By construction, λ_α is weakly carried by Θ.

The transverse measure μ_a defines a measure on the space $L(\widetilde{L}_a)$ of non-singular leaves of $\widetilde{L}_a$. Indeed, by Lemma 1.??, such non-singular leaves are locally parametrized by their intersection points with a tie k of $\widetilde{\Phi}$; therefore, the measure deposited by μ_a on k locally defines a measure on $L(\widetilde{L}_a)$. Since μ_a is invariant under homotopy respecting $\widetilde{L}_a$ this local measure does not depend on the choice of the tie k, and therefore extends to a global full support measure on $L(\widetilde{L}_a)$. Two non-singular leaves of $\widetilde{L}_a$ which are close to each other follow the same edge path for a long time. Therefore, the map $L(\widetilde{L}_a) \to \widetilde{\lambda}$ is continuous, and the above measure on $L(\widetilde{L}_a)$ pushes forward to a full support measure on $\widetilde{\lambda}$. This measure extends to a measure α on $G(\widetilde{S})$ whose support is the closure of $\widetilde{\lambda}$, namely $\widetilde{\lambda}_\alpha$. Since the measure α is invariant under the action of $\pi_1(S)$, it defines a full support transverse measure for λ_α. In this way, we have constructed a measured geodesic lamination α which is weakly carried by Θ.

By construction, for every edge e of $\widetilde{\Theta}$, $\alpha(e) = \alpha(G(e)) = \mu_a(L_e(\widetilde{L}_a)) = \mu_a(k_e) = a(e)$, where $L_e(\widetilde{L}_a) \subset L(\widetilde{L}_a)$ consists of those leaves which cross the corresponding edge e of Φ and where k_e is any tie of e.

Therefore, for every edge weight system $a \in \mathcal{W}(\Theta, [0,\infty[)$, we found a measure geodesic lamination $\alpha \in \mathcal{ML}(\Theta)$ such that $\alpha(e) = a(e)$ for every edge e of Θ. This shows that the map $\mathcal{ML}(\Theta) \to \mathcal{W}(\Theta, [0,\infty[)$ is surjective.

Combining this with Proposition ?? and Lemmas ?? and ??, we obtain:

PROPOSITION. *For every train track Θ, the map $\mathcal{ML}(\Theta) \to \mathcal{W}(\Theta, [0,\infty[)$ is a homeomorphism.* □

3.4. Edge weight systems for train tracks

Let Θ be a train track in S. If $\mathbb{M}$ is a module over a ring R, let $\mathcal{W}(\Theta, \mathbb{M})$ be the space of edge weight systems for Θ that are valued in $\mathbb{M}$ and satisfy the switch relations. Namely, if $E(\Theta)$ denotes the set of all edges of Θ, the set $\mathcal{W}(\Theta, \mathbb{M})$ consists of all maps $a : E(\Theta) \to \mathbb{M}$ such that, for every switch s of Θ, the sum of the weights $a(e)$ corresponding to all edges e coming in on one side of s is equal to the sum of the weights $a(e')$ corresponding to all edges e' going out on the other side of s. The structure of $\mathbb{M}$ makes $\mathcal{W}(\Theta, \mathbb{M})$ an R–module in a natural way.

PROPOSITION. *For a train track Θ, the module $\mathcal{W}(\Theta, \mathbb{M})$ is isomorphic to*

$$\mathbb{M}^{-\chi(\Theta)} \oplus \mathbb{M}^{N_{\mathrm{o}}} \oplus \{m \in \mathbb{M}; 2m = 0\}^{N_{\mathrm{no}}}$$

where $\chi(\Theta)$ is the Euler characteristic, N_{o} is the number of orientable components of Θ and N_{no} is the number of its non-orientable components.

Note that, because there are at least two edges meeting at each switch, the Euler characteristic of Θ is always non-positive, so that $-\chi(\Theta) = |\chi(\Theta)|$.

PROOF OF PROPOSITION ??. Each of the numbers $\chi(\Theta)$, N_{o}, N_{no} is the sum of the numbers similarly associated to the components of Θ. We can therefore restrict attention to the case where Θ is connected.

It is useful to leave the category of train tracks on S. If we forget the global embedding of the train track Θ in S, we can consider Θ only as an abstract graph. However, the tangential information at each switch (=vertex) s of Θ separates the edge ends adjacent to s into two sets, corresponding to each side of s. This induces an additional structure on the abstract graph Θ, namely a partition of the link of each vertex of Θ into two non-empty subsets. We will call such a graph Θ, endowed with a partition of the link of each vertex into two non-empty subsets, a ***train track graph***. There is an obvious way to translate to the train track graph setting the notions of edge weight systems and of orientability. We then want to show that, for connected train track graph Θ, the set $\mathcal{W}(\Theta, \mathbb{M})$ of all $\mathbb{M}$-valued edge weight systems (satisfing the appropriate switch relations) is isomorphic to $\mathbb{M}^{-\chi(\Theta)+1}$ if Θ is orientable, and to $\mathbb{M}^{-\chi(\Theta)} \oplus \{m \in \mathbb{M}; 2m = 0\}$ if Θ is non-orientable. This will clearly prove Proposition ??.

Let e be an edge of a train track graph Θ, whose end vertices s_1 and s_2 are distinct. Let Θ' be obtained from Θ by collapsing e to a point s'. The graph

Θ' can be turned into a train track graph by keeping the same link partitions as Θ outside of s_1 and s_2, and by partitioning the link of s' as follows. Let the partition of the link of s_i in Θ consist of A_i and B_i, where A_i contains the end of the edge e. In the link of s' in Θ', let A'_i and B'_i consist of the edge ends corresponding to edge ends in A_i and B_i, respectively. Then, we partition the link of s' as the union of $A'_1 \cup B'_2$ and $A'_2 \cup B'_1$. It is immediate that this new train track graph Θ' is orientable if and only if Θ is orientable, and that $\chi(\Theta') = \chi(\Theta)$. In addition, an easy algebraic manipulation shows that the restriction operator induces an isomorphism $\mathcal{W}(\Theta, \mathbb{M}) \cong \mathcal{W}(\Theta', \mathbb{M})$. In general, Θ' will not correspond to any train track in S even if Θ does.

After performing a succession of such edge collapses, we eventually reach a train track graph Θ_n with only one vertex, such that Θ_n is orientable if and only if Θ is orientable, such that $\chi(\Theta_n) = \chi(\Theta)$ and such that $\mathcal{W}(\Theta_n, \mathbb{M})$ is isomorphic to $\mathcal{W}(\Theta, \mathbb{M})$. Note that Θ_n has $-\chi(\Theta) + 1$ edges.

If every edge of Θ_n has its ends in opposite subsets of the link partition of the vertex of Θ_n, then the switch relation defining $\mathcal{W}(\Theta_n, \mathbb{M})$ is trivial. It follows that $\mathcal{W}(\Theta_n, \mathbb{M}) \cong \mathcal{W}(\Theta, \mathbb{M})$ is isomorphic to $\mathbb{M}^{-\chi(\Theta)+1}$. Note that, in this case, Θ_n is orientable and therefore that Θ is orientable.

If there are edges $e_1, e_2, \ldots, e_p$ of Θ_n whose ends are in the same subset of the link partition of the vertex of Θ_n, then the switch relation defining $\mathcal{W}(\Theta_n, \mathbb{M})$ is of the form $\pm a(e_1) \pm a(e_2) \pm \ldots \pm a(e_p) = 0$. It follows that $\mathcal{W}(\Theta_n, \mathbb{M})$ is isomorphic to $\mathbb{M}^{-\chi(\Theta)} \oplus \{m \in \mathbb{M}; 2m = 0\}$. In this case, Θ_n is non-orientable and therefore Θ is non-orientable.

This proves that, for a connected train track graph Θ, the group $\mathcal{W}(\Theta_n, \mathbb{M})$ is isomorphic to $\mathbb{M}^{-\chi(\Theta)} \oplus \mathbb{M}$ or $\mathbb{M}^{-\chi(\Theta)} \oplus \{m \in \mathbb{M}; 2m = 0\}$, according to whether Θ is orientable or not. This completes the proof of Lemma ??. □

In the case we are currently interested in, the module $\mathbb{M}$ is equal to $\mathbb{R}$, considered as a vector space over itself. Proposition ?? then restricts to the following statement.

COROLLARY. *For a train track Θ, the space $\mathcal{W}(\Theta, \mathbb{R})$ of $\mathbb{R}$–valued edge weight systems for Θ is a real vector space of dimension $-\chi(\Theta) + N_o$, where N_o is the number of orientable components of Θ.* □

3.5. The piecewise linear structure of $\mathcal{ML}_0(S)$

Consider a non-trivial measured geodesic lamination $\alpha_0 \in \mathcal{ML}_0(S) - \{0\}$ disjoint from the boundary. By Proposition ??, α_0 admits a neighborhood in $\mathcal{ML}(S)$ of the form $\mathcal{ML}(\Theta)$ for some train track Θ. By Complement ??, each component of ∂S is parallel to a closed curve component of Θ, the train track Θ_0 obtained from Θ by removing this parallel copy of ∂S is maximal boundary-free, and α_0 passes through each edge of Θ_0. Since geodesics of $\widetilde{S}$ that are weakly carried by Θ are characterized by their associated bi-infinite edge path (Lemma ??), the space $\mathcal{ML}(\Theta_0)$ is the intersection of $\mathcal{ML}(\Theta)$ with $\mathcal{ML}_0(S)$. As a consequence, $\mathcal{ML}(\Theta_0)$ is a neighborhood of α_0 in $\mathcal{ML}_0(S)$.

By Proposition ??, $\mathcal{ML}(\Theta_0)$ is homeomorphic to the space $\mathcal{W}(\Theta_0, [0, \infty[)$ of non-negative edge weight systems for Θ_0. This space $\mathcal{W}(\Theta_0, [0, \infty[)$ is a subset of $\mathcal{W}(\Theta_0, \mathbb{R}) \cong \mathbb{R}^{-\chi(\Theta_0)+N_o}$, where N_o is the number of orientable components of Θ_0, by Corollary ??.

Let Φ_0 be an arbitrary fattening of Θ_0. Because Θ_0 is maximal boundary-free, each component of the boundary of Φ_0 contains an odd number of spikes. This last property implies that Φ_0 and Θ_0 are non-orientable. Therefore, N_o is equal to 0. The fact that Θ_0 is maximal in $S - \partial S$ also implies that it is connected.

The foliation of Φ_0 by its ties has no interior singularities, and one half-saddle type singularity of index $-\frac{1}{2}$ for each spike. If s is the number of spikes of $S - \Phi_0$, the Poincaré-Hopf formula then shows that $\chi(\Theta_0) = \chi(\Phi_0)$ is equal to $-\frac{1}{2}s$. This foliation of Φ_0 by its ties also extends to a foliation of S which, on each triangle (= disk with 3 spikes) component of $S - \Phi_0$, has one saddle type singularity of index $-\frac{1}{2}$ and, on each annulus component (with one spike and one component of ∂S) of $S - \Phi_0$, has one half-saddle type boundary singularity of index $-\frac{1}{2}$. If S has b boundary components and if $S - \Phi_0$ has t triangle components, we conclude that $\chi(S) = -\frac{1}{2}t - \frac{1}{2}b$ and $s = 3t + b$. Combining these equalities and eliminating s and t, we obtain that $\chi(\Theta_0) = 3\chi(S) + b$.

Therefore, $\mathcal{W}(\Theta_0; \mathbb{R})$ is isomorphic to $\mathbb{R}^{-3\chi(S)-b}$.

The space $\mathcal{W}(\Theta_0, [0, \infty[)$ is the closed convex subset of $\mathcal{W}(\Theta_0; \mathbb{R})$ consisting of those $a \in \mathcal{W}(\Theta_0; \mathbb{R})$ such that the weight $a(e)$ is non-negative for every edge e of Θ_0. The measured geodesic lamination α_0 corresponds to a system $a_0 \in \mathcal{W}(\Theta_0; \mathbb{R})$ of positive edge weights since α_0 passes through each edge of Θ_0. Therefore, a_0 is in the interior $\mathcal{W}(\Theta_0,]0, \infty[)$ of $\mathcal{W}(\Theta_0, [0, \infty[)$, and this interior is open in $\mathcal{W}(\Theta_0; \mathbb{R}) \cong \mathbb{R}^{-3\chi(S)-b}$.

The subset of $\mathcal{ML}(\Theta_0) \cong \mathcal{W}(\Theta_0, [0, \infty[)$ corresponding to $\mathcal{W}(\Theta_0,]0, \infty[)$ contains an open neighborhood of α_0. Therefore, for every $\alpha_0 \in \mathcal{ML}_0(S) - \{0\}$, we found an open neighborhood of α_0 in $\mathcal{ML}_0(S) - \{0\}$ which is homeomorphic to an open subset of $\mathbb{R}^{-3\chi(S)-b}$. This proves:

LEMMA. *The space $\mathcal{ML}_0(S) - \{0\}$ is a topological manifold of dimension $-3\chi(S) - b$, where b is the number of boundary components of S.* □

We can be a little more precise in our choice of an atlas for $\mathcal{ML}_0(S) - \{0\}$. Say that $\alpha \in \mathcal{ML}_0(S)$ is ***carried by Θ with positive weights*** if α is weakly carried by the train track Θ and if the weight $\alpha(e)$ it assigns to each edge e of Θ is positive. Let $\mathcal{ML}^+(\Theta)$ denote the subset of $\mathcal{ML}(\Theta)$ consisting of those measured geodesic laminations which are carried with positive weights by Θ. By Proposition ??, this space is homeomorphic to $\mathcal{W}(\Theta,]0, \infty[)$.

Note that $\mathcal{ML}^+(\Theta)$ and $\mathcal{W}(\Theta,]0, \infty[)$ may very well be empty. Figure ?? gives a simple example of a train track for which this occurs. A train track Θ is said to be ***recurrent*** if $\mathcal{W}(\Theta,]0, \infty[)$ is non-empty. The terminology is explained by the following exercise.

EXERCISE. Show that a train track is recurrent if and only it admits an edge path which passes infinitely many times through each of its edges. (Hint: If Θ is recurrent, then it admits a system of positive integer weights, and it carries a closed curve which passes through each of its edges).

Assume now that the train track Θ is recurrent and maximal boundary-free. Then, the map $\mathcal{W}(\Theta,]0,\infty[) \to \mathcal{ML}^+(\Theta) \subset \mathcal{ML}_0(S) - \{0\}$ induces a homeomorphism from $\mathcal{W}(\Theta,]0,\infty[)$ to $\mathcal{ML}^+(\Theta)$. Since Θ is recurrent maximal boundary-free, we already determined that $\mathcal{W}(\Theta,]0,\infty[)$ is a non-empty open subset of $\mathcal{W}(\Theta;\mathbb{R}) \cong \mathbb{R}^{-3\chi(S)-b}$, and $\mathcal{ML}_0(S) - \{0\}$ is a manifold of dimension $-3\chi(S) - b$ by Proposition ??. It follows from the Theorem of Invariance of the Domain (see for instance ??) that $\mathcal{ML}^+(\Theta)$ is open in $\mathcal{ML}_0(S) - \{0\}$.

Therefore, for every recurrent maximal boundary-free train track Θ, $\mathcal{ML}^+(\Theta)$ is open in $\mathcal{ML}_0(S) - \{0\}$ and the natural map $\mathcal{ML}^+(\Theta) \to \mathcal{W}(\Theta,]0,\infty[)$ gives a homeomorphism from $\mathcal{ML}^+(\Theta)$ to the open convex subset $\mathcal{W}(\Theta,]0,\infty[)$ of $\mathcal{W}(\Theta;\mathbb{R}) \cong \mathbb{R}^{-3\chi(S)-b}$. In addition, Proposition ?? shows that the union of the $\mathcal{ML}^+(\Theta)$, for all such train tracks Θ, is equal to $\mathcal{ML}_0(S) - \{0\}$.

We will now investigate the regularity of the changes of charts of this atlas for $\mathcal{ML}_0(S) - \{0\}$.

We begin with a technical lemma.

LEMMA. *Let γ' and γ'' be two finite edge paths in the preimage $\widetilde{\Theta}$ of the train track Θ. Then, there are finitely many finite edge paths $\gamma_1, \dots, \gamma_n$ such that the subsets $G(\gamma_i) \subset G(\widetilde{S})$ are pairwise disjoint, and such that $G(\gamma') \cap G(\gamma'')$, $G(\gamma') - G(\gamma'')$ and $G(\gamma'') - G(\gamma')$ are all a union of such sets $G(\gamma_i)$.*

PROOF. If $G(\gamma') \cap G(\gamma'')$ is empty, the property is trivial. Otherwise, γ' and γ'' are contained in a finite edge path γ of $\widetilde{\Theta}$. Let e_0 be an edge of γ' and, for r larger than the length of γ, let $\gamma_1, \dots, \gamma_n$ be the list of all edge paths $\langle e_{-r}, \dots, e_{-1}, e_0, e_1, \dots, e_r \rangle$.

The $G(\gamma_i)$ are pairwise disjoint. Their union is equal to $G(e_0)$, and therefore contains $G(\gamma')$. Also, by construction, $G(\gamma_i)$ meets γ' (resp, γ'') exactly when γ' (resp. γ'') is contained in γ_i, in which case $G(\gamma_i)$ is contained in $G(\gamma')$ (resp. $G(\gamma'')$). These observations complete the proof. □

Now, consider two train tracks Θ, Θ' which are recurrent and maximal in $S - \partial S$, and let $\Phi_\Theta : \mathcal{ML}^+(\Theta) \to \mathcal{W}(\Theta,]0,\infty[) \subset \mathcal{W}(\Theta;\mathbb{R}) \cong \mathbb{R}^{-3\chi(S)-b}$ and $\Phi_{\Theta'} : \mathcal{ML}^+(\Theta') \to \mathcal{W}(\Theta',]0,\infty[) \subset \mathcal{W}(\Theta';\mathbb{R}) \cong \mathbb{R}^{-3\chi(S)-b}$ be the associated charts.

If α is in the intersection of $\mathcal{ML}^+(\Theta)$ and $\mathcal{ML}^+(\Theta')$, the support of the measure defined by α on $G(\widetilde{S})$ is contained in $G(\widetilde{\Theta}) \cap G(\widetilde{\Theta}')$, where $\widetilde{\Theta}$ and $\widetilde{\Theta}'$ are the preimages of Θ and Θ' in $\widetilde{S}$, respectively.

Let e' be an edge of $\widetilde{\Theta}'$. By Lemma ??, $G(\widetilde{\Theta}) \cap G(e')$ is compact and open in $G(\widetilde{\Theta}) \cap G(\widetilde{\Theta}')$. By Lemma ??, the $G(\gamma)$ associated to all finite edge path of $\widetilde{\Theta}$ form a neighborhood basis for $G(\widetilde{\Theta})$. Therefore, we can write $G(\widetilde{\Theta}) \cap G(e')$ as the union of finitely many $G(\gamma_i) \cap G(\widetilde{\Theta}')$ where the γ_i are edge paths in $\widetilde{\Theta}$. Using Lemma ??, we can arrange that the $G(\gamma_i)$ are pairwise disjoint.

Then, if α is in $\mathcal{ML}^+(\Theta) \cap \mathcal{ML}^+(\Theta')$, $\alpha(e')$ is equal to the sum of the $\alpha(\gamma_i)$. By Lemma ??, each $\alpha(\gamma_i)$ is a piecewise linear function of the $\alpha(e)$ where e ranges over all edges of Θ. This implies that $\alpha(e')$ is a piecewise linear function of the $\alpha(e)$. Another way to say this is that $\alpha(e')$ is a piecewise linear function of $\Phi_\Theta(\alpha)$.

Since this is true for every edge e' of $\widetilde{\Theta}'$, this shows that $\Phi_{\Theta'}(\alpha)$ is a piecewise linear function of $\Phi_\Theta(\alpha)$. In other words, the change of charts $\Phi_{\Theta'} \circ \Phi_\Theta^{-1}$ is piecewise linear. This proves:

PROPOSITION. *The family of charts $\Phi_\Theta : \mathcal{ML}^+(\Theta) \to \mathcal{W}(\Theta,]0,\infty[) \subset \mathcal{W}(\Theta;\mathbb{R}) \cong \mathbb{R}^{-3\chi(S)-b}$, associated to all train tracks Θ which are recurrent and maximal in $S - \partial S$, forms a piecewise linear atlas for $\mathcal{ML}_0(S) - \{0\}$.* □

A structure on a space associated to the surface S is ***natural*** if it is invariant under the action of the homeomorphism group of S. As a consequence of Proposition ??:

THEOREM. *The space $\mathcal{ML}_0(S) - \{0\}$ admits a natural structure of piecewise linear manifold of dimension $-3\chi(S) - b$.* □

Note that the piecewise linear changes of charts $\Phi_{\Theta'} \circ \Phi_\Theta^{-1}$ have some additional properties. First of all, they are ***positive homogeneous***, namely $\Phi_{\Theta'} \circ \Phi_\Theta^{-1}(\lambda a) = \lambda \Phi_{\Theta'} \circ \Phi_\Theta^{-1}(a)$ for every $a \in \mathcal{W}(\Theta,]0,\infty[)$ and every $\lambda > 0$. In addition, they are ***piecewise integral linear***, in the sense that they locally coincide with linear maps whose matrices have integer entries, and that the pieces where the changes of maps are linear are defined by linear (in)-equalities with integer coefficients. Therefore, the piecewise linear structure of $\mathcal{ML}_0(S) - \{0\}$ is actually a ***piecewise homogenous integral linear structure***, for the obvious definition.

In §§3.6 and 3.7, we will see that $\mathcal{ML}_0(S)$ is a piecewise linear manifold even at 0.

3.6. The global structure of $\mathcal{ML}_0(S)$ and $\mathcal{ML}(S)$ when S has non-empty boundary

The global topology of $\mathcal{ML}(S)$ and $\mathcal{ML}_0(S)$ is relatively easy to analyze when the boundary ∂S is non empty. Indeed, in this case, we can find disjoint simple arcs $k_1, \ldots, k_n$ going from the boundary to the boundary in S, and decomposing S into hexagons where the sides alternately correspond to some k_i and to arcs in ∂S. Figure ?? gives an example of such a decomposition when S is a disk minus 4 holes.

A counting argument gives that the number n of the arcs k_i in the decomposition is equal to $-3\chi(S)$.

Consider a train track Θ which, on each hexagon of the decomposition, is as indicated on Figure ??. Namely, the switches of Θ are located at the midpoints of the k_i and each hexagon contains exactly 3 edges of Θ. In particular, the complement consists of one semiopen annulus with no spike for each boundary component of S and of one disk with 3 spikes for each hexagon.

After isotopy, we can assume that the k_i are geodesic for the base metric m_0. This guarantees that every m_0–geodesic cannot hit the same side of a hexagon two times in a row (of course, it can hit the same k_i twice in a row if this k_i corresponds to two distinct sides of the hexagon). In each hexagon, an arc contained in the interior of S that joins two distinct sides of the hexagon can be deformed fixing end points to an arc that is strongly carried by θ. It follows that every m_0–geodesic of S is weakly carried by Θ. Therefore, the space $\mathcal{ML}(S)$ is equal to the space $\mathcal{ML}(\Theta) \cong \mathcal{W}(\Theta; [0,\infty[)$.

We now determine the space $\mathcal{W}(\Theta; [0,\infty[)$ of edge weight systems valued in $[0,\infty[$ for Θ. Consider an edge weight system $a \in \mathcal{W}(\Theta; [0,\infty[)$. In a hexagon H touching the arcs k_i, k_j and k_l, consider the corresponding edges e_i, e_j and e_l. From the 3 switch relations and the fact that the weights of the other edges have to be non-negative, we see that the weights $a(e_i)$, $a(e_j)$, $a(e_l)$ have to satisfy the inequality $a(e_i) \leqslant a(e_j) + a(e_l)$ and the other two similar triangular inequalities obtained by cyclic permutations of the indices.

Conversely, given a family of non-negative numbers $a_1, \ldots, a_n$ that satisfy the triangular inequality $a_i \leqslant a_j + a_l$ whenever the arcs k_i, k_j, k_l touch a common hexagon, there is a unique way to assign weights on the remaining 3 edges of each hexagon, in order to obtain an edge weight system $a \in \mathcal{W}(\Theta; [0,\infty[)$ such that $a(e_i) = a_i$ for $i = 1, \ldots, n$.

Therefore, $\mathcal{W}(\Theta; [0,\infty[)$ is isomorphic to the set of points $(a_1, \ldots, a_n) \in [0,\infty[^n$, with $n = -3\chi(S)$, such that $a_i \leqslant a_j + a_l$ whenever the arcs k_i, k_j, k_l touch a common hexagon. Note that this set is a closed convex positively homogeneous subset $\mathbb{R}^n$, bounded by finitely many faces. Its interior is non-empty since it contains the point $(1, \ldots, 1)$.

Since all these sets are invariant by multiplication by positive numbers, it is convenient to projectivize the situation. Let $\mathcal{PML}(S)$ by the space quotient of $\mathcal{ML}(S) - \{0\}$ by the euivalence relation which identifies α to β whenever there is a number $\lambda > 0$ such that $\alpha = \lambda\beta$. The elements of $\mathcal{PML}(S)$ are called ***projective measured geodesic laminations***. We can clearly reconstruct $\mathcal{ML}(S)$ from $\mathcal{PML}(S)$ by a coning construction. The piecewise homogeneous integral linear structure of $\mathcal{ML}_0(S)$ induces a piecewise linear structure on $\mathcal{PML}(S)$.

The above discussion of $\mathcal{W}(\Theta; [0,\infty[)$ immediately gives:

THEOREM. *If S has non-empty boundary, then the space $\mathcal{PML}(S)$ of projective measured geodesic laminations is piecewise linearly isomorphic to a closed ball of dimension $-3\chi(S) - 1$.* □

Since $\mathcal{ML}(S)$ is a cone over $\mathcal{PML}(S)$, it follows $\mathcal{ML}(S) - \{0\}$ is piecewise linearly isomorphic to $\mathbb{R}^{-3\chi(S)-1} \times [0,\infty[- \{0\}$, and the piecewise homogeneous integral linear structure of $\mathcal{ML}(S) - \{0\}$ extends to a piecewise homogeneous integral linear structure on $\mathcal{ML}(S)$ making it isomorphic to $\mathbb{R}^{-3\chi(S)-1} \times [0,\infty[$.

We now determine the structure of the similarly defined projective space $\mathcal{PML}_0(S)$, consisting of projective measured geodesic laminations disjoint from the boundary.

The above embedding of $\mathcal{W}(\Theta; [0,\infty[)$ in $\mathbb{R}^n$ gives an embedding of $\mathcal{PML}(S)$ in $\mathbb{R}^{n-1}$, identified to the set of $(a_1, \ldots, a_n) \in \mathbb{R}^n$ such that $\sum_i a_i = 1$. Each boundary component of S determines a non-trivial element of $\mathcal{ML}(S)$, and therefore a point in $\mathbb{R}^{n-1}$. Let $x_1, \ldots, x_d \in \mathbb{R}^{n-1}$ be the points corresponding in this way to the boundary components of S. Note that these d points are linearly independent in $\mathbb{R}^{n-1}$.

Identify $\mathcal{PML}(S)$ to its image in $\mathbb{R}^{n-1}$. In this context, the decomposition of $\mathcal{ML}(S)$ as $\mathcal{ML}_0(S) \times [0,\infty[^d$ provided by Proposition ?? can be translated in the following way: For every $x \in \mathcal{PML}(S) \subset \mathbb{R}^{n-1}$ which is not in the affine $(d-1)$–dimensional space V spanned by $x_1, \ldots, x_d$, there is a unique point $y \in \mathcal{PML}_0(S) \subset \mathcal{PML}(S)$ such that x is a linear combination with non-negative coefficients of $x_1, \ldots, x_d$ and y. In particular, if W_x is the affine d–dimensional half-space in $\mathbb{R}^{n-1}$ which is bounded by V and contains x, then $W_x \cap \mathcal{PML}(S)$ is the d–dimensional simplex spanned by $x_1, \ldots, x_d$ and y, and $W_x \cap \mathcal{PML}_0(S)$ is equal to $\{y\}$.

Conversely, consider a d-dimensional half-space W in $\mathbb{R}^{n-1}$ which is bounded by the space V spanned by $x_1, \ldots, x_d$. Because the barycenter $\left(\frac{1}{d}, \ldots, \frac{1}{d}\right) \in \mathbb{R}^{n-1} \subset \mathbb{R}^n$ of $x_1, \ldots, x_d$ in in the interior of $\mathcal{PML}(S)$, the intersection $W \cap \mathcal{PML}(S)$ cannot be contained in V; by the above observation, it therefore contains a unique element $y \in W \cap \mathcal{PML}_0(S)$, defined by the property that $W \cap \mathcal{PML}(S)$ is the simplex spanned by $x_1, \ldots, x_d$ and y.

This proves that $\mathcal{PML}_0(S)$ is a piecewise linear submanifold of $\mathcal{PML}(S)$, and is isomorphic to the space of all d-dimensional half-spaces W in $\mathbb{R}^{n-1}$ which are bounded by V. This proves:

THEOREM. *If S has $d > 0$ boundary components, then the space $\mathcal{PML}_0(S)$ of projective measured geodesic laminations disjoint from the boundary is piecewise linearly isomorphic to a sphere of dimension $-3\chi(S) - d - 1$.* □

As a consequence, $\mathcal{ML}_0(S) - \{0\}$ is piecewise linearly isomorphic to $\mathbb{R}^{-3\chi(S)-d} - \{0\}$, and the piecewise homogeneous integral linear structure of $\mathcal{ML}_0(S) - \{0\}$ extends to a piecewise homogeneous integral linear structure on $\mathcal{ML}_0(S)$ making it isomorphic to $\mathbb{R}^{-3\chi(S)-d}$.

3.7. The global structure of $\mathcal{ML}(S)$ when S has empty boundary

■unchanged since modification of train tracks

Now, consider the case where the boundary ∂S is empty.

Since the Euler characteristic $\chi(S)$ is negative, the surface S contains a punctured torus. Considering the geodesics homotopic to a meridian and longitude for this punctured torus, we obtain two 2–sided simple closed m_0–geodesics k_1 and k_2 of S which meet in exactly one point x_0. We can then extend k_1, k_2 to a family $k_1, \ldots, k_n$ of arcs which go from x_0 to itself, whose interiors ar embedded and pairwise disjoint, and which decompose S into a family of triangles. A counting argument shows that the number n of the k_i is equal to $3(1-\chi(S))$.

We impose an additional condition for the k_i. We select a side of k_1 and we require that k_2 is the only arc hitting k_1 on this preferred side. This condition will automatically be realized if we choose for k_3 the arc defined as follows: Start from x_0 on the side opposite to the preferred side, follow an arc parallel to k_2, then, when approaching x_0 again, go once around an arc parallel to k_1 and located on its preferred side, and finally return to x_0 by following again an arc parallel to k_2 (but on the other side of k_2, this time).

In addition, we can isotop the arcs k_i so that they are m_0–geodesic. Note that the two branches of each k_i necessarily form an angle at x_0, except for $i = 1, 2$.

Let S_0 be the surface with boundary obtained from S by removing a small disk around the point x_0. Each arc k_i gives an arc in S_0, which we will also denote by k_i to alleviate the notation. These k_i decompose S_0 into hexagons, as in §3.6. Let Θ be the train track in S_0 associated to the k_i as in §3.6.

Consider a measured m_0–geodesic lamination $\alpha \in \mathcal{ML}(S)$, and let λ be its support. The m_0–geodesic lamination λ in S can be slightly isotoped to a union λ_0 of simple disjoint curves (possibly bi-infinite) in S_0. Because the k_i are m_0–geodesic in S, we can easily choose λ_0 so that, in each hexagon of the decomposition, no leaf of λ_0 hits the same side twice in a row. Note that his guarantees that each leaf of λ_0 can be deformed to be carried by the train track Θ, and λ_0 therefore defines a geodesic lamination in S_0.

As we go around ∂S_0, we follow an edge path $\langle e_1, e_2', e_2, e_2', \bar{e}_1, e_3', e_4, e_4', \dots, e_{2n}, e_{2n}' \rangle$ where each oriented edge e_j contains an arc k_{i_j}, each e_j' contains no arc k_i, e_1 contains k_1, e_2 contains k_2, and $\bar{e}_j$ denotes e_j with the orientation reversed. Note that the edges e_1' and e_2' are both located on the preferred side of k_1, that there is an e_j with $j \geqslant 4$ which is equal to $\bar{e}_2$, and that the remaining e_j with $j \geqslant 4$ can be grouped into pairs $\{e_j, e_{j'}\}$ where $e_{j'}$ is equal to $\bar{e}_j$ or e_j, according to whether the corresponding arc k_{i_j} is orientation-preserving or -reversing.

An important observation is that, because λ is m_0–geodesic and k_1 is a closed m_0–geodesic (with no angle at x_0), we can choose the deformation from λ to λ_0 so that no arc contained in a leaf of λ_0 follows the edge paths $\langle e_1, e_1', e_2, e_2', \bar{e}_1 \rangle$ or $\langle \bar{e}_1, e_3', e_4, e_4', \dots, e_{2n}, e_{2n}', e_1 \rangle$, namely crosses k_1 twice in a row while turning around ∂S_0.

Now, consider a leaf g of λ_0 which crosses the edge e_1'. It necessarily contains an arc a which follows an edge path $\left\langle f, e_k'', e_{k+1}, e_{k+1}', \dots, e_1, e_1', e_2, e_2'', f' \right\rangle$ where $4 \leqslant k \leqslant 2n$, $e_k'' \neq e_k'$ and $e_2'' \neq e_2'$. This arc a can be deformed, fixing a neighborhood of its ends, across the disk $S - S_0$ in S to an arc which follows an edge path $\left\langle f, e_k''', \bar{e}_k, \bar{e}_{k-1}', \bar{e}_{k-1}, \dots, e_1, e_2''', f' \right\rangle$. The arcs in leaves of λ_0 that realize such an edge path $\left\langle f, e_k'', e_{k+1}, e_{k+1}', \dots, e_1, e_1', e_2, e_2'', f' \right\rangle$ area all leaves of a foliated rectangle in θ. By successively deforming these foliated rectangles across $S - S_0$, starting with the one which is closest to ∂S_0 in e_1', we can therefore isotop λ_0 in S so that it is still carried by Θ but does not cross e_1' any more.

Performing a similar modification of the leaves of λ_0 crossing e_2', we can arrange that λ_0 crosses neither e_1' nor e_2'. Note that the fact that no arc contained in a leaf of λ_0 follows the edge path $\langle \bar{e}_1, e_3', e_4, e_4', \dots, e_{2n}, e_{2n}', e_1 \rangle$ also guarantees that λ_0 avoids at least one edge e_j' with $3 \leqslant j \leqslant 2n$.

Let $\mathcal{ML}_1(\Theta)$ denote the subset of $\mathcal{ML}(\Theta) \subset \mathcal{ML}(S_0)$ that consists of those measured geodesic laminations which are weakly carried by Θ in a way avoiding the edges e_1' and e_2' as well as some edge e_j' with $4 \leqslant j \leqslant 2n$. More precisely, for $3 \leqslant j \leqslant 2n$, let Θ_i be obtained from $\Theta - \left(e_1' \cup e_2' \cup e_j' \right)$ by pushing it slightly inside of S_0 and by perturbing its boundary so that it becomes transverse to the ties (so as to make Θ_j a train track in S_0). Then, $\mathcal{ML}_1(\Theta)$ is the union of all the subsets $\mathcal{ML}(\Theta_j) \subset \mathcal{ML}(S_0)$ with $4 \leqslant j \leqslant 2n$.

Note that the component of $S - \Theta_j$ that contains the disk $S - S_0$ is a disk with 3 spikes. Therefore, Θ_j is also a train track in S. In particular, we can consider $\mathcal{ML}(\Theta_j)$ as a subset of both $\mathcal{ML}(S_0)$ and $\mathcal{ML}(S)$. This defines a natural map $\Gamma : \mathcal{ML}_1(\Theta) \to \mathcal{ML}(S)$.

LEMMA. *The map $\Gamma : \mathcal{ML}_1(\Theta) \to \mathcal{ML}(S)$ is a piecewise linear isomorphism.*

PROOF. The discussion beginning of this section shows that the map Γ is surjective. Its definition through train tracks also shows that Γ is piecewise linear for the piecewise linear structures of $\mathcal{ML}(S_0)$ and $\mathcal{ML}(S)$.

It remains to check that Γ is injective. Let Θ' be obtained from $\Theta - (e_1' \cup e_2')$ by pushing it slightly inside of S_0 and by perturbing its boundary so that it becomes transverse to the ties. This Θ' is a train track in S_0, but not in S because one of the components of $S - \Theta'$ is a digon, namely a disk with 2 ties. In $\mathcal{ML}(S_0)$, the subset $\mathcal{ML}_1(\Theta)$ is clearly the same as the set of those $\alpha \in \mathcal{ML}(\Theta')$ such that $\alpha\left(e_j'\right) = 0$ for some $3 \leqslant j \leqslant 2n$ (denoting each edge of Θ' by the same symbol as the corresponding edge of Θ).

We will say that a curve is ***tightly carried*** by Θ' if it is carried by Θ' and does not contain any arc following the edge path $\langle \bar{e}_1, e_3', e_4, e_4', \dots, e_{2n}, e_{2n}', e_1 \rangle$. The definition is designed so that, if $\alpha \in \mathcal{ML}_1(\Theta) \subset \mathcal{ML}(\Theta')$, every leaf of α is asymptotic to a curve that is tightly carried by Θ'.

In S, the foliation of Θ' by its ties extends to a foliation of S with only saddle-type singularities, and with a closed leaf that meets the digon of $S - \Theta'$ and is homotopic to the closed geodesic k_2 (which implies that it is not homotopic to 0). The proofs of Lemmas ??, ?? and ?? in §1.8 easily extend to prove the following: If we lift a curve tightly carried by Θ' to the universal covering $\widetilde{S}$, this lift is quasi-geodesic; and two bi-infinite curves tightly carried by Θ' lift to curves which are asymptotic to the same geodesic of $\widetilde{S}$ if and only if the two curves follow the same bi-infinite edge path in Θ'.■check, hints?

The above observations immediately show the injectivity of the map $\Gamma : \mathcal{ML}_1(\Theta) \to \mathcal{ML}(S)$. □

In view of Lemma ??, we only have to determine the topology of $\mathcal{ML}_1(\Theta)$, namely of the set of those $\alpha \in \mathcal{ML}(\Theta')$ such that $\alpha\left(e_j'\right) = 0$ for some $3 \leqslant j \leqslant 2n$.

By ??, $\mathcal{ML}(\Theta')$ is isomorphic to the set $\mathcal{W}(\Theta'; [0,\infty[)$ of non-negative edge weight systems for Θ'. The edges of Θ' fall into two distinct categories: the 'e–edges', corresponding to edges of Θ that contain one of the arcs k_i; and the 'e'–edges', corresponding to edges of θ containing no such k_i. For an edge weight system, the switch relations are all requiring that the weight of an e–edge should be the sum of one or two weights of e'–edges. Therefore, an element of $\mathcal{W}(\Theta'; [0,\infty[)$ is completely determined by the weights it assigns to the $2n-2$ e'–edges, and $\mathcal{W}(\Theta'; [0,\infty[)$ is isomorphic to the intersection of the linear subspace L of $\mathbb{R}^{2n-2}$ defined by the switch relations with the cone C consisting of all $x \in \mathbb{R}^{2n-2}$ with non-negative coordinates. If we projectivize the situation, this gives an isomorphism between the projectivized space $\mathcal{PML}(\Theta')$ and $L \cap H \cap C$, where H is the hyperplane consisting of those $x \in \mathbb{R}^{2n-2}$ for which the sum of the coordinates is equal to 1.

Note that the linear space L is isomorphic to $\mathcal{W}(\Theta'; \mathbb{R})$. The train track Θ' is connected, and non-orientable since at least one component of $S_0 - \Theta'$ is a disk with an odd number (=3) of spikes. By Corollary ??, it follows that L has dimension $-\chi(\Theta') = -3\chi(S) + 1$.

By definition, an element of $\mathcal{ML}(\Theta')$ is in $\mathcal{ML}_1(\Theta)$ exactly when it assigns weight 0 to some e'–edge, namely when the corresponding edge weight system in $\mathcal{W}(\Theta'; [0, \infty[) \cong L \cap C$ is in the boundary ∂C. Therefore, $\mathcal{PML}(S) \cong \mathcal{PML}_1(\Theta)$ is isomorphic to the intersection of L with the boundary of the simplex $H \cap C$. By §3.5, $\mathcal{PML}(S)$ is a piecewise linear manifold of dimension $-3\chi(S) - 1$. It follows that L must meet the interior of the simplex $H \cap C$, and that $L \cap H \cap \partial C$ is piecewise linear isomorphic to a sphere of dimension $-3\chi(S) - 1$ (namely to the boundary of a simplex of dimension $-3\chi(S)$).

This proves that $\mathcal{PML}(S)$ is piecewise linear isomorphic to a sphere of dimension $-3\chi(S) - 1$. As a consequence, $\mathcal{ML}(S)$ is isomorphic to the cone over that same sphere, namely to $\mathbb{R}^{-3\chi(S)}$. In particular, $\mathcal{ML}(S)$ is a piecewise linear manifold even at the origin (compare Proposition ??).

Theorem. *When the boundary of the surface S is empty, the space $\mathcal{PML}(S)$ of projective measured geodesic laminations is piecewise linear isomorphic to a sphere of dimension $-3\chi(S) - 1$, and the space $\mathcal{ML}(S)$ of measured geodesic laminations is piecewise linear isomorphic to $\mathbb{R}^{-3\chi(S)}$.* □

3.8. Natural structures on the space $\mathcal{ML}_0(S) - \{0\}$

the PL structure is the only one for which Dehn Twists are PL (see Casson)

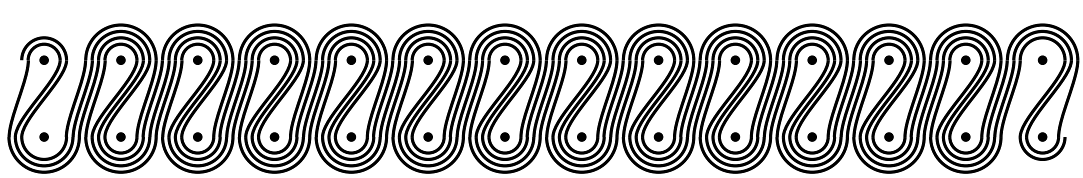

CHAPTER IV

GEODESIC CURRENTS

4.1. Measure geodesic currents

In the previous chapters, we used measured geodesic laminations to complete the space of homotopy classes of simple closed curves on a surface S. We now extend some of these ideas to closed curves which are not necessarily simple. We will use two distinct points of view to do this. The first one, developed in this section, is metric independent and is a natural generalization of our metric independent construction of measured geodesic laminations. The second point of view will be developed in the following section, where we will consider a more geometric presentation of the same construction which is probably more intuitive and easier to handle in applications, although metric dependent.

As usual, we endow S with a negatively curved metric m_0 for which the boundary ∂S is geodesic. Then, each closed curve c in S is, either homotopic to 0, or homotopic to a unique closed m_0–geodesic c^*. The geodesic c^* may wrap several times around another closed geodesic γ^*. This happens exactly when the homotopy class of c is ***divisible***, in the sense that it is represented by a non-trivial multiple γ^n, with $n \geqslant 2$, of another curve γ. We can choose γ^* so that it is ***indivisible***, namely so that it does not wrap several times around another smaller geodesic. By ??, this is equivalent to the property that the homotopy class of γ is indivisible.

In this way, we obtain a natural correspondence between, on one hand, the set of non-trivial homotopy classes of (unoriented) closed curves on S, and the set of indivisible (unoriented) closed m_0–geodesics γ^* endowed with positive integer weights $n \geqslant 2$ on the other hand. As in §1.2, where we used the same trick for simple closed curves, the main benefit of this construction is that we do not have to worry any more.

Now, we can go one step further, and consider the pre-image of the closed geodesic γ^* in the universal covering. This pre-image $\widetilde{\gamma}^*$ the union of an infinite family of geodesics of $\widetilde{S}$, and therefore defines a subset of $G\left(\widetilde{S}\right)$, which we we stll denote by $\widetilde{\gamma}^*$. Note that this subset in invariant under the action of $\pi_1(S)$.

LEMMA. *The subset $\widetilde{\gamma}^* \subset G\left(\widetilde{S}\right)$ associated to an indivisible closed geodesic γ^* of S is a closed discrete subset of $G\left(\widetilde{S}\right)$.*

PROOF. Since $G\left(\widetilde{S}\right)$ is locally compact, it suffices to show that every element $g \in G\left(\widetilde{S}\right)$ has a neighborhood which contains only finitely many elements of $\widetilde{\gamma}^*$. If x is a point of the m_0–geodesic g, and if $\widetilde{U}$ is a neighborhood of $\widetilde{x}$ in $\widetilde{S}$, the set of those m_0–geodesics which meet $\widetilde{U}$ is a neighborhood of g in $G\left(\widetilde{S}\right)$. Therefore, it suffices to show that every $\widetilde{x} \in \widetilde{S}$ admits a neighborhood $\widetilde{U}$ which meets only finitely many components of the pre-image of the m_0–geodesic γ. Any m_0–ball $\widetilde{U}$ centered at $\widetilde{x}$ and small enough that the projection $\widetilde{S} \to S$ is injective on $\widetilde{U}$ clearly has this property. □

At this point, to each non-trivial homotopy class $[c]$ of closed curves in S, we have associated an indivisible closed geodesic γ^* with a positive integer multiplicity n and consequently, by Lemma ??, a $\pi_1(S)$–invariant closed discrete subset $\widetilde{\gamma}^*$ of $G\left(\widetilde{S}\right)$ which is still endowed with a positive integer multiplicity n. Note that this construction is completely metric independent. Indeed, $\widetilde{\gamma}^*$ consists of all elements of $G\left(\widetilde{S}\right)$ which are fixed by some element of $\pi_1(S)$ that is represented by a curve freely homootpic to c. Also, n is the largest integer such that c is homotopic to a curve of the type γ^n for some closed curve γ.

Given a closed discrete subset A of a locally compact space x and a positive multiplicity a, it is natural to associate to this data the corresponding ***counting measure*** or ***Dirac measure*** α ***of weight*** a ***supported by*** A defined by the property that $\alpha(Z) = a\#A \cap Z \in [0, \infty]$ for every subset $Z \subset X$, where $\#A \cap Z$ denotes the cardinal of $A \cap Z$. Similarly, for every continuous function $\varphi : X \to \mathbb{R}$ with compact support, $\alpha(\varphi) = a \sum_{x \in A} \varphi(x)$. Note that this sum is finite since A is closed and discrete, so that α really defines a Radon measure on X.

In particular, to the discrete closed subset $\widetilde{\gamma}^* \subset G\left(\widetilde{S}\right)$ and the multiplicity $n \geqslant 1$, we can associate the counting measure of weight n supported by $\widetilde{\gamma}^*$ on $G\left(\widetilde{S}\right)$. Note that his measure is invariant under the action of $\pi_1(S)$, since $\widetilde{\gamma}^*$ is invariant under this action.

In this way, we have associated to each non-trivial homotopy class of closed curves on S a $\pi_1(S)$–invariant measure on $G\left(\widetilde{S}\right)$. Note that the trivial homotopy class is the class of γ^0 for an arbitrary γ; it is therefore natural to extend this construction by associating the 0 measure to the trivial homotopy class.

This leads us to consider measures on $G\left(\widetilde{S}\right)$ which are invariant under the action of $\pi_1(S)$. By definition, a ***(measure) geodesic current*** on S is a $\pi_1(S)$–invariant measure on $G\left(\widetilde{S}\right)$. We will explain the origin of the terminology in the next section.

Let $\mathcal{C}(S)$ denote the space of geodesic currents on S. We endow $\mathcal{C}(S)$ with the weak* topology already considered in §2.??.

What we accomplished in this section is to construct an embedding of the space of homotopy classes of simple closed curves in S into the space $\mathcal{C}(S)$. We observed that this construction is independent of the choice of the auxiliary metric m_0 on S.

4.2. Geodesic currents as transverse measures to the geodesic foliation

We now give another description of geodesic currents and of the embedding of the set of homotopy classes of closed curves in $\mathcal{C}(S)$. This construction (apparently) depends on the choice of negatively curved metric m_0 but, because it is much more geometric than the one we introduced in the previous section, it gives a better intuitive picture of geodesic currents and is often easier to handle.

Consider an unoriented indivisible closed curve γ in S. The corresponding m_0–geodesic γ^* may have self-intersection points, namely can pass twice through the same point x before closing up. We will make use of a key property of geodesics: the geodesics passing through a given point are completely determined by their tangent at that point. In particular, γ^* must pass through the self-intersection point x each timne with a different tangent. This leads us to a natural way of de-singularizing γ^*, by lifting it to the projective tangent bundle $PT(S)$.

▮ faire un dessin pour illustrer l'idée de passer au fibré tangent

Recall that $PT(S)$ is the 3–dimensional manifold consisting of all pairs (x, l) where $x \in S$ and where l is a line through the origin in the tangent space T_xS. A m_0–geodesic g lifts to a curve $\widehat{g}$ in $PT(S)$ by considering the tangent line l at each point $x \in g$. Since the m_0–geodesic passing through a point with a given direction is unique, the curve $\widehat{g}$ is embedded. Similarly, distinct complete m_0–geodesics g, g' lift to disjoint curves $\widehat{g}$, $\widehat{g}'$ in $PT(S)$. Recall that a geodesic is ***complete*** if it cannot be extended to a longer geodesics; in our case, an m_0–geodesic g is complete when it is closed or when each end of g, either has infinite length, or terminates at a point where g transversely hits the boundary ∂S. In particular, the lifts of the complete m_0–geodesics partition $PT(S)$ into a family of embedded curves. This partition is the ***geodesic foliation*** $\mathcal{G}$ of $PT(S)$.

Let us check that the geodesic foliation is really a foliation of the interior of $PT(S)$, namely that each point (x, l) in the interior of $PT(S)$ admits a neighborhood U which is diffeomorphic to $W \times [a, b]$, where W is an open subset of $\mathbb{R}^2$, wher $[a, b]$ is an interval in $\mathbb{R}$, and where each $* \times [a, b]$ is contained in a leaf of the geodesic foliation. Let k be a small arc in S which passes through x and is orthogonal to the line $l \subset T_xS$ at x. For each $y \in k$ and each angle $\theta \in \left]\frac{\pi}{2} + \varepsilon, \frac{\pi}{2} + \varepsilon\right[$, let $g_{(y,\theta)}$ be the m_0–geodesic arc which passes through y, makes an angle of θ with the oriented arc k, and has length ε' on either side of y, for fixed ε, ε' small. Lifting the $g_{(y,\theta)}$ to arcs $\widehat{g}_{(y,\theta)}$ in $PT(S)$ and using the obvious parametrizations, we obtain a map $k \times \left]\frac{\pi}{2} + \varepsilon, \frac{\pi}{2} + \varepsilon\right[\times [-\varepsilon', \varepsilon'] \to PT(S)$, which sends $(x, 0, 0)$ to (x, l). An application of the Inverse Function Theorem shows that this map is a local diffeomorphism at $(x, 0, 0)$. In other words, if we choose k, ε, ε' sufficiently small, we can assume that it is a diffeomorphism onto its image, which is a neighborhood of (x, λ). Since, by construction, the image of $\{x\} \times \{\theta\} \times [-\varepsilon', \varepsilon']$ is the arc $\widehat{g}_{(y,\theta)}$, this provides a diffeomorphism of a neighborhood U of (x, l) and $W \times [-\varepsilon', \varepsilon']$ where $W = (k - \partial k) \times \left]\frac{\pi}{2} + \varepsilon, \frac{\pi}{2} + \varepsilon\right[$ is diffeomorphic to an open subset of $\mathbb{R}^2$ and where each $\{*\} \times [-\varepsilon', \varepsilon']$ is contained in a leaf of the geodesic foliation. This proves that the geodesic foliation is really a foliation on the interior of $PT(S)$.

When S has non-empty boundary, we could wonder whether the geodesic foliation $\mathcal{F}$ is also a foliation near the boundary $\partial PT(S)$. In the literature, there are two competing definitions for a foliation on a manifold with boundary: One imposes that the boundary be a union of leaves, the other one that all leaves be transverse to the boundary. The geodesic foliation satisfies none: The components of ∂S are m_0–geodesic, and lift to leaves of $\mathcal{F}$ which are contained in $\partial PT(S)$; on the other hand, all other leaves which hit the boundary are transverse to it. In other word, the geodesic foliation is singular near $\partial PT(S)$, for any of the above two definitions of a foliation on a manifold with boundary.

Exercise. Carefully analyze the topology of the geodesic foliation near the boundary $\partial PT(S)$.

The reader may be more familiar with the ***geodesic flow*** of m_0, when the surface S has no boundary. This geodesic flow is defined on the unit tangent bundle T^1S consisting of all pairs (x, v), where $x \in S$, $v \in T_xS$, and v has norm 1 for the Riemannian metric m_0 (see ??). Recall that a ***flow*** on a manifold M is a family of diffeomorphisms $\Phi_t : M \to M$ where Φ_t which depend differentiably on $t \in \mathbb{R}$ and define a group action of the additive group $\mathbb{R}$, in the sense that $\Phi_{t+t'} = \Phi_t \circ \Phi_{t'}$ for every t, t'. The geodesic flow consists of the deffeomorphisms $\Phi_t : T^1S \to T^1S$ defined by the property that $\Phi_t(x, v) = (g_v(t), g'_v(t))$ where g_v is the unique parametrized m_0–geodesic such that $g_v(0) = x$ and $g'_v(0) = v$, using the notation $g'_v(t) = dg_v(t)/dt$. The projective tangent bundle $PT(S)$ is the quotient of T^1S under the free $\mathbb{Z}/2$–action defined by $(x, v) \mapsto (x, -v)$ and, by construction, the leaves of the geodesic foliation $\mathcal{G}$ of $PT(S)$ are exactly the images of the orbits of the geodesic flow of T^1S.

We will describe geodesic currents in terms of transverse measures for the geodesic foliation $\mathcal{G}$. A ***transverse measure*** α for $\mathcal{G}$ is the data of a (Radon) measure α_V defined on each surface V in $PT(S)$ transverse to $\mathcal{G}$, and which is invariant under homotopy respecting $\mathcal{G}$. ▮improve? locally same as measure on space of leaves

The Dirac transverse measure of weight $a > 0$ based at the closed leaf $\widehat{\gamma}^*$ is the transverse measure which, to each Borelian A contained in a submanifold V transverse to $\mathcal{G}$, associates the number $a\#(A \cap \widehat{\gamma}^*)$ where $\#X$ denotes the cardinal of X.

In this way, we associate a transverse measure α_γ for $\mathcal{G}$ to each indivisible closed curve γ in S.

We can extend this construction in a natural way to divisible curves. Indeed, if the closed curve γ is divisible but not homotopic to 0, its associated geodesic γ^* wraps $n \geqslant 2$ times around the closed geodesic γ_1^* associated to an indivisible closed curve γ_1. We can then think of γ^* as equal to γ_1^* with a multiplicity of n. It is therefore natural to associate to γ the transverse measure which is n times the transverse measure associated to γ_1^*, namely the Dirac transverse measure α_γ of weight n based at the closed leaf $\widehat{\gamma}_1^*$ of $\mathcal{G}$ that lifts γ_1^*. Finally, it is natural to associate the trivial transverse measure to a curve γ that is homotopic to 0.

In this way, we have defined an embedding of the set $C(S)$ of free homotopy classes of closed curves into the space of all transverse measures for the geodesic foliation $\mathcal{G}$.

Let the ***recurrent part*** $\mathcal{G}_0 \subset PT(S)$ denote the union of all closed or bi-infinite leaves of $\mathcal{G}$. Namely, $\mathcal{G}_0$ consists of the lifts of those geodesics of S which do not transversely hit the boundary. In particular, $\mathcal{G}_0$ contains the lifts of the boundary components of S. It is immediate that $\mathcal{G}_0$ is closed in $PT(S)$. ■ explain recurrent

By construction, the transverse measure α_γ associated to every closed curve γ has support contained in $\mathcal{G}_0$.

By definition, a ***measure*** m_0–***geodesic current*** is a transverse measure for the geodesic foliation $\mathcal{G}$ whose support is contained in $\mathcal{G}_0$. We will usually abbreviate "measure geodesic current" to "geodesic current". We will later encounter more general geodesic currents which are transverse distributions for $\mathcal{G}$.

Let $\mathcal{C}(S)$ denote the space of such m_0–geodesic currents, which we endow with the weak* topology, defined by the family of semi-distances $d_\varphi(\alpha, \beta) = |\alpha(\varphi) - \beta(\varphi)|$, where φ ranges over all continuous functions $\varphi : V \to \mathbb{R}$ which are defined on surfaces V transverse to $\mathcal{G}$ (■transversality near boundary?) and whose support is compact. In other words, the geodesic current $\alpha \in \mathcal{C}(S)$ converges to α_0 if and only if, for every continuous function $\varphi : V \to \mathbb{R}$ defined on a surface V transverse to $\mathcal{G}$ and such that the support of φ is compact and contained in the interior of V, the integral $\alpha(\varphi)$ converges to $\alpha_0(\varphi)$. Note that the condition on the support of φ is crucial to avoid problems when the boundary of V has non-trivial α_0–measure.

The semi-distances d_φ also define a uniform structure on $\mathcal{C}(S)$. Standard result on this weak* uniform structure (see for instance ??) give the following completeness result.

Theorem. *The uniform space $\mathcal{C}(S)$ is complete.* □

Remark. It is not hard to see that we can restrict attention to a countable number of functions $\varphi : V \to \mathbb{R}$ when defining the topology and uniform structure of $\mathcal{C}(S)$. It follows that this topology and this uniform structure are metrizable, and are associated to a metric which is complete, by Theorem ??.

Sometimes, the topology of $\mathcal{C}(S)$ will be easier to understand in terms of flow boxes for the recurrent part $\mathcal{G}_0$ of the geodesic foliation $\mathcal{G}$. A ***flow box*** for $\mathcal{G}_0$ is a compact subset of $\mathcal{G}_0$ which is homeomorphic to a product $X \times [0,1]$, where each arc $* \times [0,1]$ is contained in a leaf of $\mathcal{G}_0$. A typical example is obtained by considering a surface V transverse to the geodesic foliation $\mathcal{G}$ and a small subset X of $V \cap \mathcal{G}_0$; if we associate to $(x,t) \in [0,\varepsilon]$ the point of $\mathcal{G}_0$ image of x by the time t geodesic flow map (for a choice of transverse orientation for V), the map $X \times [0,\varepsilon] \to \mathcal{G}_0$ so defined is a homeomorphism onto its image for ε sufficiently small, and its image therefore is a flow box. In particular, this shows that every point of $\mathcal{G}_0$ has a neighborhood in $\mathcal{G}_0$ which is a flow box.

If $\alpha \in \mathcal{C}(S)$ is a geodesic current and if $B \cong X \times [0,1]$ is a flow box, then α defines a measure on the space X. We will denote by $\alpha(B) = \alpha(X)$ the total mass of this measure.

The topological frontier of a flow box $B \cong X \times [0,1]$ in $\mathcal{G}_0$ consists of $X \times \{0,1\}$ and of a subset of the form $(\delta X) \times [0,1]$ for some subset δX of X. For instance, in the above example where X is a subset of a transverse surface V to $\mathcal{G}$, the subset δX is just the topological frontier of X in $V \cap \mathcal{G}_0$. We will call this subset $\delta_{\mathcal{G}_0} B = (\delta X) \times [0,1]$ the ***tangential frontier*** of the flow box B. Note that $\delta_{\mathcal{G}_0} B$ is a flow box, so that we can define $\alpha(\delta_{\mathcal{G}_0} B)$ for every geodesic current $\alpha \in \mathcal{C}(S)$.

LEMMA. *Given a geodesic current $\alpha \in \mathcal{C}(S)$, every point of the recurrent part $\mathcal{G}_0$ admits arbitrarily small neighborhoods in $\mathcal{G}_0$ which are flow boxes B with $\alpha(\delta_{\mathcal{G}_0} B) = 0$.*

PROOF. We saw how to construct a flow box neighborhood of $x \in \mathcal{G}_0$ by considering a small subset X of $V \cap \mathcal{G}_0$ for some transverse surface V, in which case $\delta_{\mathcal{G}_0} B = (\delta X) \times [0,1]$ where δX is the frontier of X in $V \cap \mathcal{G}_0$. In particular, we can take X to be a ball of radius r in $V \cap \mathcal{G}_0$. Then, there are at most countably many values of r for which $\alpha(\delta X) \neq 0$. □

PROPOSITION. *For a geodesic current $\alpha \in \mathcal{C}(S)$, the geodesic current $\beta \in \mathcal{C}(S)$ converges to α if and only if the mass $\beta(B)$ converges to $\alpha(B)$ for every flow box B for $\mathcal{G}_0$ with $\alpha(\delta_{\mathcal{G}_0} B) = 0$.*

PROOF. To explain the restriction to flow boxes B with $\alpha(\delta_{\mathcal{G}_0} B) = 0$, we should probably remind the reader of a subtlety of the weak* topology: If a sequence of measures μ_i on a locally compact space X converges to μ for the weak* topology and if A is a measurable subset of X, the masses $\mu_i(A)$ converge to $\mu(A)$ if the mass $\mu(\delta A)$ of the frontier δA is equal to 0 (see for instance [**Bou**, Chap IV, §5, n° 12] ??), and this condition is absolutely necessary. For instance, on the real line $\mathbb{R}$, the Dirac measure D_x based at x converges to D_0 for the weak* topology as $x > 0$ tends to 0, but $D_x([-1,0]) = 0$ does not converge to $D_0([-1,0]) = 1$.

Assume that $\beta \in \mathcal{C}(S)$ converges to α for the topology of $\mathcal{C}(S)$, and let B be a flow box with $\alpha(\delta_{\mathcal{G}_0} B) = 0$. We want to show that $\beta(B)$ converges to $\alpha(B)$.

For every leaf g of the restriction of $\mathcal{G}_0$ to B, there is a surface V_g transverse to $\mathcal{G}_0$ and a compact subset A_g of V_g such that those leaves of B that meet A_g form a neighborhood of g in B, and so that no leaf of B cuts A_g more than once; by the same argument as in the proof of Lemma ??, we can require in addition that $\alpha(\delta A_g) = 0$. By compactness, we can find finitely many such A_{g_1}, A_{g_2}, … , A_{g_n} such that each leaf of B cuts at least one of the A_{g_i}.

For each subset I of $\{1, 2, \dots, n\}$, let B_I be the union of those leaves of B that meet all the A_{g_i} with $i \in I$ and no other A_{g_j}. Choose an arbitrary element i of I, let A_I be the intersection of B_I with $A_{g_i} \subset V_{g_i}$ and let V_I denote the transverse surface V_{g_i}. Each leaf of B that meets the frontier δA_I of A_I in V_I must also meet $\delta_{\mathcal{G}_0} B$ or one of the frontiers δA_{g_j} with $j \in I$. It follows that $\alpha(\delta A_I)$.

Since the B_I are pairwise disjoint, $\alpha(B)$ is equal to the sum of the $\alpha(A_I)$ as I ranges over all subsets of $\{1, 2, \dots, n\}$. Similarly, $\beta(B)$ if equal to the sum of the $\beta(A_I)$. As β tends to α, each $\beta(A_I)$ converges to $\alpha(A_I)$ since $\alpha(\delta A_I) = 0$

and since the measure defined by β on V_I weak* converges to that defined by α. Therefore, $\beta(B)$ converges to $\alpha(B)$.

Conversely, assume that $\beta(B)$ converges to $\alpha(B)$ for every flow box B for $\mathcal{G}_0$ with $\alpha(\delta_{\mathcal{G}_0}B) = 0$. Let $\varphi : V \to \mathbb{R}$ be a continuous function with compact support defined on a surface V transverse to $\mathcal{G}$. We want to show that $\beta(\varphi)$ converges to $\alpha(\varphi)$.

For every $\varepsilon > 0$, the function φ is uniformly ε–close to a step function, namely to a linear combination $\sum a_i \chi_{A_i}$ of the characteristic functions of finitely many pairwise disjoint subsets A_i of V. In addition, we can choose these A_i so that $\alpha(\delta A_i) = 0$ for every i. As above Lemma ??, the closure of each A_i can be used to define a flow box B_i with $\alpha(B_i) = \alpha(\bar{A}_i) = \alpha(A_i)$, $\beta(B_i) = \beta(\bar{A}_i)$ and $\alpha(\delta_{\mathcal{G}_0}B_i) = \alpha(\delta A_i) = 0$; in particular, $\beta(\bar{A}_i)$ converges to $\alpha(\bar{A}_i) = \alpha(A_i)$, and $\beta(\delta\bar{A}_i)$ converges to $\alpha(\delta\bar{A}_i) = 0$. Then, $\beta(\varphi)$ is within $\varepsilon\sum\beta(\bar{A}_i)$ of $\sum a_i\beta(A_i)$, which itself is within $\sum|a_i|\,\beta(\delta\bar{A}_i)$ of $\sum a_i\beta(\bar{A}_i)$; similar estimates hold for α. The convergence of $\beta(\varphi)$ to $\alpha(\varphi)$ easily follows from these considerations.

□

We can rephrase Proposition ?? in the following way. Given a geodesic current $\alpha \in \mathcal{C}(S)$, flow boxes $B_1, B_2, \ldots, B_n$ for $\mathcal{G}_0$, and a number $\varepsilon > 0$, define

$$\mathcal{U}(\alpha; B_1, B_2, \ldots, B_n) = \{\beta \in \mathcal{C}(S)\,; |\beta(B_i) - \alpha(B_i)| < \varepsilon \text{ for } i = 1, \ldots, n\}.$$

Then, Proposition ?? just says that, as $B_1, B_2, \ldots, B_n$ range over all finite sets of flow boxes B with $\alpha(\delta_{\mathcal{G}_0}B) = 0$, the corresponding $\mathcal{U}(\alpha; B_1, B_2, \ldots, B_n)$ form a basis of neighborhoods of α for the topology of $\mathcal{C}(S)$.

4.2. Geodesic currents from a metric independent viewpoint

As usual, we can use the space $G(\widetilde{S})$ of bi-infinite geodesics in $\widetilde{S}$ to make the notion of geodesic current independent of the choice of the negatively curved metric m_0.

In the projective tangent space $PT(\widetilde{S})$ of the universal covering $\widetilde{S}$, consider the foliation $\widetilde{\mathcal{G}}$ lifting the geodesic foliation $\mathcal{G}$ of $PT(S)$, and the preimage $\widetilde{\mathcal{G}}_0$ of the recurrent part $\mathcal{G}_0$ of $\mathcal{G}$. The space $\widetilde{\mathcal{G}}_0$ is the union of the lifts to $PT(\widetilde{S})$ of all bi-infinite geodesics of $\widetilde{S}$. Because no geodesic of $\widetilde{S}$ is recurrent, the space of leaves of $\widetilde{\mathcal{G}}_0$ is homeomorphic to the space $G(\widetilde{S})$.

Let α is a geodesic current. Given a geodesic $g \in G\left(\widetilde{S}\right)$, consider a codimension 1 submanifold V of $PT(\widetilde{S})$ which is transverse to the geodesic foliation $\widetilde{\mathcal{G}}$ and which crosses the lift of g to $PT(\widetilde{S})$. The elements of $G\left(\widetilde{S}\right)$ whose lifts cross V form a neighborhood U_V of g, and are in one-to-one correspondence with the points of $V \cap \widetilde{\mathcal{G}}_0$ if V is small enough. The transverse measure α for $\mathcal{G}$ lifts to a transverse measure for $\widetilde{\mathcal{G}}$, which we will also denote by α. This transverse measure for $\widetilde{G}_0$ determines a measure on V, with support contained in $V \cap \widetilde{G}_0$, and therefore induces a measure on the neighborhood U_V of g in $G\left(\widetilde{S}\right)$. The fact that α is invariant under holonomy shows that, locally near g, this measure is independent on the choice of the transverse submanifold V. Therefore, α defines a measure on $G\left(\widetilde{S}\right)$. By construction, this measure is invariant under the action of the fundamental group $\pi_1(S)$ on $G\left(\widetilde{S}\right)$.

Conversely, let α be a measure on $G\left(\widetilde{S}\right)$ which is invariant under the action of $\pi_1(S)$. If V is a codimension 1 submanifold of $PT(\widetilde{S})$ transverse to the geodesic foliation $\widetilde{G}$, the set U_V of those bi-infinite geodesics which cross V forms an open subset of $G\left(\widetilde{S}\right)$. In addition, if V is sufficiently small, the consideration of intersection points defines a homeomorphism between U_V and $V \cap \widetilde{\mathcal{G}}_0$. Consequently, the restriction of α to U_V defines a measure on $V \cap \widetilde{\mathcal{G}}_0$, and therefore defines a measure on V with support contained in $V \cap \widetilde{\mathcal{G}}_0$. Moving V by a homotopy respecting $\widetilde{\mathcal{G}}$ does not chanve the set U_V, and this measure on V is therefore independent under homotopy respecting $\widetilde{\mathcal{G}}$. In this way, we have defined a transverse measure for $\widetilde{\mathcal{G}}$ with support contained in $\widetilde{\mathcal{G}}_0$. Since α is $\pi_1(S)$–invariant, so is this trqansverse measure, and it descends to a transverse measure for the geodesic foliation $\mathcal{G}$ with support containedin $\mathcal{G}_0$, namely to a geodesic current.

This establishes a one-to-one correspondence between geodesic currents and $\pi_1(S)$–invariant measures on $G\left(\widetilde{S}\right)$. The same type of argument shows that this one-to-one correspondence is a homeomorphism if we endow the space of measures on $G\left(\widetilde{S}\right)$ with the weak* topology.

Proposition. *The space $\mathcal{C}(S)$ of (measure) geodesic currents is naturally homeomorphic to the space of $\pi_1(S)$–invariant measures on the space $G\left(\widetilde{S}\right)$ of all bi-infinite geodesics of $\widetilde{S}$, endowed with the weak* topology.* □

4.3. Weighted closed curves are dense in $\mathcal{C}(S)$

In §4.1, we associated a geodesic current to each (free homotopy class of) closed curve on the surface S. If we add a positive weight $a > 0$ to such a closed curve γ, we can associate to this weighted closed curve a times the geodesic current defined by γ. In particular, if γ^n is a closed curve which wraps n times around γ, for a positive integer n, then the same geodesic current is associated to γ with weight a and to γ^n with weight a/n.

Theorem. *The geodesic currents associated to weighted closed curves are dense in $\mathcal{C}(S)$.*

Proof. Theorem ?? is a special case of a more general result of K. Sigmund on Axiom A dynamical systems [**Sig**].

We first show that the sums of geodesic currents associated to weighted closed curves are dense in $\mathcal{C}(S)$.

Let $\alpha \in \mathcal{C}(S)$. We want to approximate α by a sum β of geodesic currents associated to weighted closed curves.

Recall that the topology of $\mathcal{C}(S)$ is defined using continuous functions defined on codimension 1 submanifolds V of $PT(S)$ which are transverse to the geodesic foliation.

□

4.4. The geometric intersection form

4.5. The length function

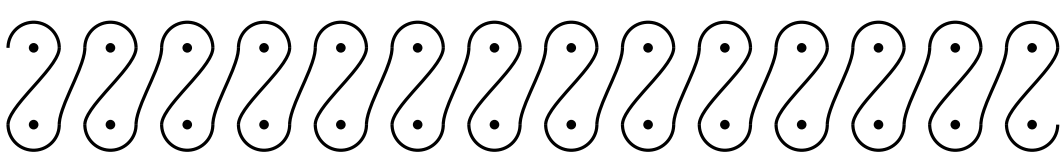

CHAPTER V

TANGENT VECTORS TO $\mathcal{ML}(S)$

5.1. Tangent vectors to piecewise linear manifolds

In §4.??, we showed that the length function $l_m : \mathcal{ML}(S) \to \mathbb{R}$ is continuous. If we want to push the analysis of this function one step further, a natural question is to ask whether it is differentiable. Because $\mathcal{ML}(S)$ does not have any natural differentiable structure (§3.??), we cannot use here the usual notion of differentiability. Nevertheless, the piecewise linear structure of $\mathcal{ML}(S)$ provides a notion of vector tangent to this manifold, as well as well as a notion of differential (although the terminology of "tangent map" seems to be more commonly used in this context) for functions defined on $\mathcal{ML}(S)$. The corresponding definitions are almost identical to those in the differentiable context.

For a function $f : U \to V$ going from an open subset U of $\mathbb{R}^m$ to an open subset V of $\mathbb{R}^n$, the ***tangent map*** of f at $x \in U$ is, if it exists, the map $T_x f : \mathbb{R}^m \to \mathbb{R}^n$ such that $T_x f(v) = (f \circ \sigma)'(0)$ for every curve $\sigma : [0, \varepsilon[\to U$ with $\sigma(0) = x$ and $\sigma'(0) = v$.

EXERCISE. Show that $T_x f : \mathbb{R}^m \to \mathbb{R}^n$ is the tangent map of f at x if and only if, as $t > 0$ tends to 0, the map $v \mapsto (f(x + tv) - f(x))/t$ converges to $T_x f$ uniformly on compact subsets of $\mathbb{R}^m$.

Clearly, the tangent map of a differentiable map is equal to its differential. Another example, which will be fundamental here, is provided by piecewise linear maps: For $f : U \to V$ piecewise linear, the restriction of $T_x f$ to each compact $K \subset \mathbb{R}^m$ coincides with the piecewise linear map $v \mapsto (f(x + tv) - f(x))/t$ for t sufficiently small (depending on K).

It immediately follows from the definition that the tangent map $T_x f$ is positive homogeneous, in the sense that $T_x f(\lambda v) = \lambda T_x(v)$ for every $\lambda \geqslant 0$ and $v \in \mathbb{R}^n$. However, the case of piecewise linear maps shows that $T_x f$ is not necessarily linear, and even that $T_x f(-v)$ is not necessarily equal to $-T_x f(v)$.

Given a piecewise linear manifold M and $x \in M$, we can define tangent vectors to M at x by mimicking the definition of tangent vectors to an abstract differentiable manifold. For instance, consider the set of all pairs (φ, v) where $\varphi : U \to \mathbb{R}^m$ is a chart of the piecewise linear structure of M such that the open subset $U \subset M$ contains x, where v is an element of $\mathbb{R}^n$. Then, a ***tangent vector*** to M at x is an equivalence class of such pairs, where (φ_1, v_1) and (φ_2, v_2) are equivalent exactly when the tangent map $T_{\varphi_1(x)} \varphi_2^{-1} \circ \varphi_1$ of the (piecewise linear) change of charts $\varphi_2^{-1} \circ \varphi_1$ sends v_1 to v_2.

In the particular case of $\mathcal{ML}(S)$ and $\mathcal{ML}_0(S)$, the piecewise linear structures of these manifolds are defined by charts $\Phi_\Theta : \mathcal{ML}^{\mathrm{pw}}(\Theta) \to \mathcal{W}(\Theta;]0,\infty[)$ associated to train tracks Θ; see §3.5. Given two such train tracks Θ and Θ', the proofs of Proposition ?? and Lemma ?? show that, near a measured geodesic lamination α, the piecewise linear change of charts $\Phi_{\Theta'} \circ \Phi_\Theta^{-1}$ is completely determined by the way the leaves of α cross the edges of Θ and Θ'. Therefore, a tangent vector to $\mathcal{ML}(S)$ or $\mathcal{ML}_0(S)$ can be described in terms of train tracks Θ and of the combinatorics of the changes of charts $\Phi_{\Theta'} \circ \Phi_\Theta^{-1}$. This combinatorial description of tangent vectors to $\mathcal{ML}(S)$ and $\mathcal{ML}_0(S)$ is unfortunately not very convenient to work with in practice, and we will focus on a more geometric approach.

The example investigated in the next section will illustrate this approach.

However, before we go any further, it is useful to introduce the following notation. A tangent vector is best described as the derivative of a curve and, throughout the rest of this monograph, we will consequently often have to take the right derivative with respect to t of a quantity X_t depending on a real parameter t, usually at $t = 0$. To alleviate some otherwise cumbersome expressions, we will write $\dot{X}_0$ for the right derivative $\frac{\partial}{\partial t^+} X_{t|t=0}$ of X_t at 0. For instance, given a train track Θ, consider a 1–parameter family of edge weight systems $a_t \in \mathcal{W}(\Theta; [0,\infty[)$ and the corresponding measured geodesic laminations $\alpha_t \in \mathcal{ML}^{\mathrm{w}}(\Theta) \subset \mathcal{ML}(S)$. Then $\dot{a}_0 \in \mathcal{W}(\Theta;\mathbb{R})$ represents the right derivative $\frac{\partial}{\partial t^+} a_{t|t=0}$ and, for every edge e of Θ, the right derivative of the edge weight $a_t(e)$ is denoted by $\dot{a}_0(e) = \frac{\partial}{\partial t^+} a_t(e)_{|t=0}$ (if these derivatives exist). Similarly, if $\psi : k \to \mathbb{R}$ is a function defined on an arc k that is transverse to all the α_t, the right derivative of the α_t–integral $\alpha_t(\psi)$ is $\dot{\alpha}_0(\psi) = \frac{\partial}{\partial t^+} \alpha_t(\psi)_{|t=0}$. In theory, there is a certain amount of ambiguity in a composite expression such as $\dot{a}_0(e)$, which could be read as $\frac{\partial}{\partial t^+}(a_t(e))_{|t=0}$ or as $\left(\frac{\partial}{\partial t^+} a_{t|t=0}\right)(e)$. However, for each occurrence of such an ambiguity where the two interpretations make sense, the reader will easily check that the two interpretations actually represent the same quantity.

5.2. An example

In a punctured torus contained in the surface S, consider the train track Θ represented in Figure 2, with one edge e_1 going around a meridian of the punctured torus and the other edge e_2 going around a longitude. For every $t \geqslant 0$, let $\alpha_t \in \mathcal{ML}(S)$ be the measured geodesic lamination which is weakly carried by Θ and is associated to the edge weight system $a_t \in \mathcal{W}(\Theta; [0,\infty[)$ such that $a_t(e_1) = 1$ and $a_t(e_2) = t$. Since the right derivative $\dot{a}_0 = \frac{\partial}{\partial t^+} a_{t|t=0} \in \mathcal{W}(\Theta;\mathbb{R})$ exists, the curve $t \mapsto \alpha_t$ in $\mathcal{ML}(S)$ admits a right derivative $\dot{\alpha}_0 = \frac{\partial}{\partial t^+} \alpha_{t|t=0}$, which is a vector tangent to $\mathcal{ML}(S)$ at α_0.

We want to give a geometric interpretation of this tangent vector $\dot{\alpha}_0$. In this heuristic section, we will restrict ourselves to formal computations. Using this example as a guide, we will later generalize there computations to a more formal setting, in the rest of this chapter and in the next two ones.

A measured geodesic lamination α is completely determined by the mass $\alpha(k) \in [0, \infty[$ it assigns to each arc k which is generic with respect to simple geodesics. Therefore, as a first step towards a geometric analysis of the tangent vector $\dot{\alpha}_0$, it is natural to determine the right derivative $\dot{\alpha}_0(k) = \frac{\partial}{\partial t^+}\alpha_t(k)_{|t=0}$ for every such generic arc k.

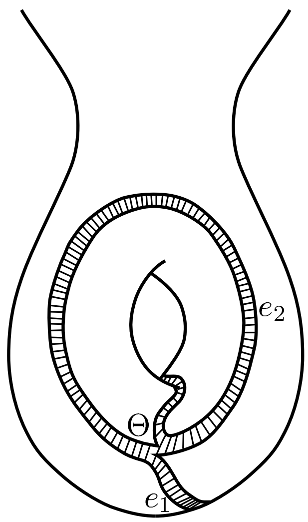

FIGURE 5.1

The geodesic lamination λ_0 underlying α_0 is just the closed geodesic which goes once around e_1, and the measure deposited by α_0 on a transverse arc k is the Dirac measure of weight 1 based at $k \cap \lambda_0$; in particular, $\alpha_0(k) = \#k \cap \lambda_0$ where $\#A$ denotes the cardinal of the set A. Although the structure of the α_t with t irrational is quite complex, the geodesic lamination $\lambda_{1/n}$ underlying $\alpha_{1/n}$ is the simple closed geodesic which wraps once around e_2 and n times around e_1; compare §1.??. Figure 5.2(b) illustrates an example for $n = 4$. In addition, the measure deposited by $\alpha_{1/n}$ on a transverse arc k is the Dirac measure of weight $1/n$ based at $k \cap \lambda_{1/n}$; in particular, $\alpha_{1/n}(k) = \left(\#k \cap \lambda_{1/n}\right)/n$. Therefore, if the derivative $\dot{\alpha}_0(k)$ exists for the arc k, it must be equal to

$$\begin{aligned}\dot{\alpha}_0(k) &= \lim_{t\to 0^+}\left(\alpha_t(k) - \alpha_0(k)\right)/t = \lim_{n\to\infty} n\left(\alpha_{1/n}(k) - \alpha_0(k)\right) \\ &= \lim_{n\to\infty}\left(\#k \cap \lambda_{1/n} - n\#k \cap \lambda_0\right).\end{aligned}$$

As n tends to ∞, the geodesic lamination $\lambda_{1/n}$ tends, for the Hausdorff topology on closed subsets of S, to the union λ_0^+ of the closed geodesic λ_0 and of an infinite geodesic spiraling on both sides of λ_0, as illustrated in Figure 5.2(c); compare the proof of ??.

Let us focus attention on an arc k' which is contained in the arc k represented on Figure 5.2. In this case, $k \cap \lambda_{1/n}$ consists of n points x_1^n, x_2^n, … , x_n^n. Similarly, $k \cap \lambda_0^+$ consists of the point $x_\infty^\infty = k \cap \lambda_0$ together with countably many points x_i^∞, $i \in \mathbb{N}$, accumulating on x_∞^∞. In addition, it is possible to

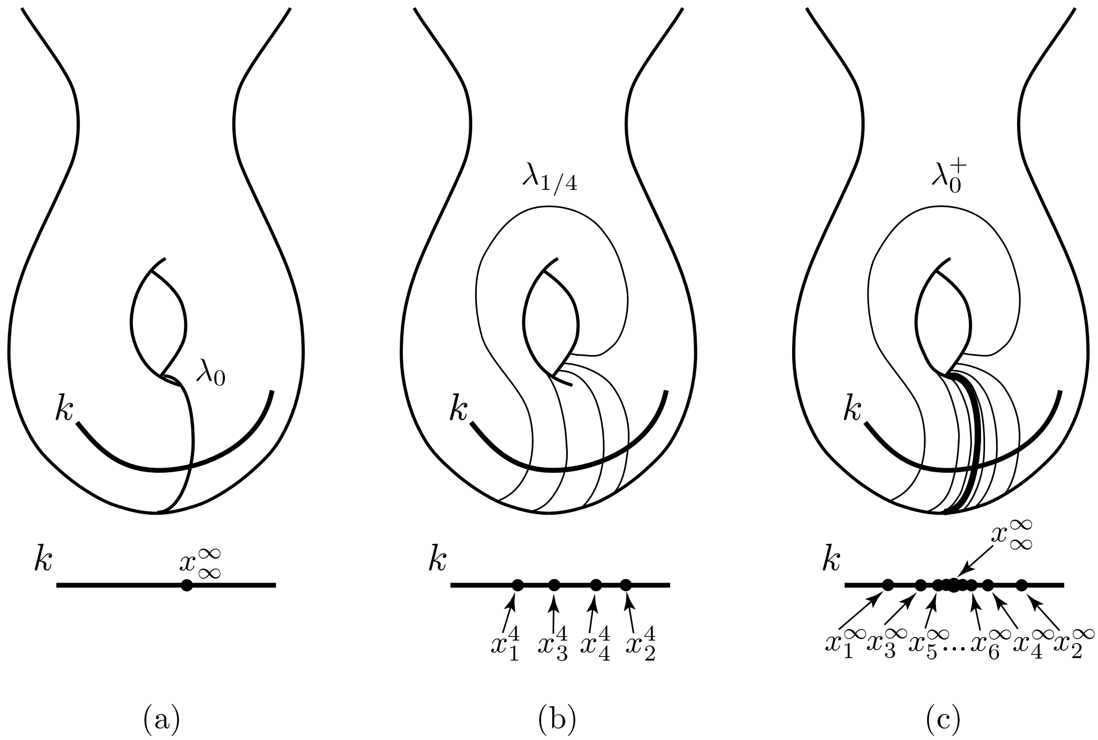

Figure 5.2

choose the indexing so that, for each i, the point x_i^n converges to x_i^∞ as n tends to infinity. We have to distinguish cases, according to whether k' contains the limit point x^∞_∞ or not. We assume k' generic enough so that its end points are not in λ_0^+.

If k' does not contain x^∞_∞, then $\# k' \cap \lambda_0$ is equal to 0 and, for n large enough, $\# k' \cap \lambda_{1/n}$ is equal to the finite number $\# k' \cap \lambda_0^+$. It follows that, if $\dot\alpha_0(k')$ exists, it must be equal to $\# k' \cap \lambda_0^+$.

If k' contains the point x^∞_∞, then $\# k' \cap \lambda_0$ is equal to 1 and, for n large enough, $\# k' \cap \lambda_{1/n}$ is equal to $n - \#(k - k') \cap \lambda_0^+$. It follows that, if $\dot\alpha_0(k')$ exists, it must be equal to $-\#(k-k') \cap \lambda_0^+$.

A similar analysis can easily be carried out for every arc k transverse to λ_0^+. It shows that the limit of $n\big(\alpha_{1/n}(k) - \alpha_0(k)\big)$ as n tends to ∞ is an integer, determined by the combinatorics of the intersection of k with λ_0^+. This putative map $k \mapsto \dot\alpha_0(k)$ has two fundamental properties: It is invariant under homotopy respecting λ_0^+, namely $\dot\alpha_0(k) = \dot\alpha_0(k')$ whenever the two transverse arcs k and k' are homotopic respecting λ; and it is finitely additive in the sense that $\dot\alpha_0(k) = \dot\alpha_0(k_1) + \dot\alpha_0(k_2)$ if we split the arc k as the union of two arcs k_1 and k_2 meeting in exactly one point. We will later summarize these two properties by saying that the rule $k \mapsto \dot\alpha_0(k)$ defines a transverse cocycle for the geodesic lamination λ_0^+.

In the rest of this chapter, we will rigorously prove that what we observed in this specific example is actually a general phenomenon. Given a curve $t \mapsto \alpha_t \in \mathcal{ML}(S)$ whose right derivative at $t = 0$ is the vector $\dot\alpha_0$ tangent to $\mathcal{ML}(S)$ at α_0, we will associate to this curve a geodesic lamination λ_0^+ containing the geodesic lamination λ_0 underlying α_0, called the essential support of α_t as t tends to 0^+. We will also show that, for every arc k that is generic with respect to simple geodesics, the right derivative $\dot\alpha_0(k)$ exists, and that the corresponding map $k \mapsto \dot\alpha_0(k)$ defines a transverse cocycle for the geodesic lamination λ_0^+ (however, in contrast to the example we just analyzed, this transverse cocycle may take non-integer values). In addition, we will see that the essential support λ_0^+ and its transverse cocycle depend only on the tangent vector $\dot\alpha_0$, and not on the specific curve $t \mapsto \alpha_t$ whose right derivative at $t = 0$ is equal to $\dot\alpha_0$.

Going back to the above example, we can be more ambitious and try to compute the right derivative of the measure defined by α_t on each arc k (and not just the right derivative of its total mass $\alpha_t(k)$). Namely, given a continuous function $\psi : k \to \mathbb{R}$ defined on a generic arc k, we would like to determine the right derivative $\dot\alpha_0(\psi)$ of the integral $\alpha_t(\psi)$ of ψ with respect to α_t, if this derivative exists.

Let us consider the case where k is the arc of Figure 5.2, and let the points x_∞^∞, x_i^∞, x_i^n be defined as before. The Dirac measures defined by $\alpha_{1/n}$ and α_0 on k respectively associate to a continuous function $\psi : k \to \mathbb{R}$ the numbers $\alpha_{1/n}(\psi) = \frac{1}{n}\sum_{i=1}^n \psi(x_i^n)$ and $\alpha_0(\psi) = \psi(x_\infty^\infty)$. From a purely formal point of view it follows that, if the derivative exists,

$$\begin{aligned}\dot\alpha_0(\psi) &= \lim_{t\to 0^+} \frac{\alpha_t(\psi) - \alpha_0(\psi)}{t} = \lim_{n\to\infty} n\left(\alpha_{1/n}(\psi) - \alpha_0(\psi)\right) \\ &= \lim_{n\to\infty} \sum_{i=1}^n \psi(x_i^n) - n\psi(x_\infty^\infty) = \lim_{n\to\infty} \sum_{i=1}^n \left(\psi(x_i^n) - \psi(x_\infty^\infty)\right) \\ &= \sum_{i=1}^\infty \left(\psi(x_i^\infty) - \psi(x_\infty^\infty)\right)\end{aligned}$$

since, for every i, the point x_i^n converges to x_i^∞ as n tends to ∞.

Observe that we cannot justify this formal reasoning for every ψ, because the above series does not even converge for an arbitrary continuous function $\psi : k \to \mathbb{R}$. However, all geodesic laminations considered here are geodesic with respect to a fixed metric m of negative curvature. A geometric estimate then shows that the series $\sum d(x_i^\infty, x_\infty^\infty)$ is equivalent to a geometric series. Therefore, the series $\sum_{i=1}^\infty \left(\psi(x_i^\infty) - \psi(x_\infty^\infty)\right)$ converges if we assume that the function ψ is, say, continuously differentiable. This suggests that the right derivative at $t = 0$ of the measure deposited by α_t on k is the distribution which associates $\sum_{i=1}^\infty \left(\psi(x_i^\infty) - \psi(x_\infty^\infty)\right)$ to each continuously differentiable function $\psi : k \to \mathbb{R}$ with support contained in the interior of k. (Recall that a distribution is a continuous linear form on the space of infinitely differentiable functions on k). This distribution is not a measure, in the sense that it does not extend to a continuous form on the space of all continuous functions on k.

We can clearly perform a similar analysis for every arc k transverse to the geodesic lamination λ_0^+. This seems to associate to $\dot\alpha_0$ a transverse distribution for the essential support λ_0^+ (although, as we will see in §6.??, we need to be careful in the definition of the transverse distribution for a geodesic lamination).

In Chapter 6, we will rigorously prove that, in general, a vector $\dot\alpha_0$ tangent to $\mathcal{ML}(S)$ determines a geodesic lamination λ_0^+ (its essential support) and a certain transverse distribution for λ_0^+. In Chapter 7, we will prove the stronger property that, if $\dot\alpha_0$ is the right derivative at $t = 0$ of a curve $t \mapsto \alpha_t \in \mathcal{ML}(S)$, the distribution defined by $\dot\alpha_0$ on an arc k transverse to λ_0^+ is equal to the derivative of the measure induced by α_t on k. This property will then lead to several practical applications.

5.2. The essential support of a family of measured geodesic laminations

case where $\alpha_0 = 0$ or α_0 is contained in the boundary

Consider a 1–parameter family of measured geodesic laminations $\alpha_t \in \mathcal{ML}(S)$, with $t \in [0, \varepsilon[$.

It is convenient to lift the situation to the universal covering $\widetilde{S}$, and to interpret the α_t as measures on $G(\widetilde{S})$. Consider the subset $\widetilde{\lambda}_{0^+} \subset G(\widetilde{S})$ consisting of those geodesics $\widetilde{g} \in G(\widetilde{S})$ which admit arbitrarily small neighborhoods U such that $\liminf_{t\to 0^+} \frac{1}{t}\alpha_t(U) > 0$. Clearly, $\widetilde{\lambda}_{0^+}$ is closed in $G(\widetilde{S})$, and is invariant under the action of $\pi_1(S)$. Also, no two geodesics of $\widetilde{\lambda}_{0^+}$ can cross each other since no two geodesics of the support of α_t cross each other. By ??, it follows that $\widetilde{\lambda}_{0^+}$ projects to a geodesic lamination λ_{0^+} of S, called the ***essential support of*** α_t ***as*** t ***tends to*** 0^+.

Observe that the geodesic lamination λ_0 underlying α_0 is contained in λ_{0^+}. Also, if t_n is a sequence converging to 0^+ such that the geodesic lamination λ_{t_n} underlying α_{t_n} converges to a geodesic lamination λ'_{0^+} for the Hausdorff topology as n tends to ∞, this limit λ'_{0^+} clearly contains λ_{0^+}. The inclusions $\lambda_0 \subset \lambda_{0^+} \subset \lambda'_{0^+}$ can all be strict; for instance, this occurs when the geodesic lamination λ_t underlying α_t is a fixed union of three disjoint simple closed geodesics, and when the transverse measure of α_t is the Dirac transverse measure which has weight 1 on the first closed geodesic, t on the second one and t^2 on the third one.

The following elementary test will often be useful to search for geodesics of $\widetilde{\lambda}_{0^+}$.

Lemma (Tracking Lemma). *Let K be a compact subset of $G(\widetilde{S})$ such that $\liminf_{t\to 0^+} \alpha_t(K)/t > 0$. Then K contains at least one geodesic of $\widetilde{\lambda}_{0^+}$.*

Proof. By compactness, K can be written as the union of finitely many compact subsets of diameter less than $\frac{1}{2}$. For at least one of these, say K_1, we must have $\liminf_{t\to 0^+} \alpha_t(K_1)/t > 0$. Reapplying the same process to K_1, we can construct a sequence of nested compact subsets $K \supset K_1 \supset \dots \supset K_n \supset \dots$ such that the diameter of K_n is less than 2^{-n} and such that $\liminf_{t\to 0^+} \alpha_t(K_n)/t > 0$ for every n. The intersection of the K_n consists of a point $g \in G(\widetilde{S})$ which, by construction, must be in $\widetilde{\lambda}_{0^+}$. □

Although easy to state, the above definition of λ_{0^+} is not very easy to handle in practice. We want to give a more convenient description of λ_{0^+} in terms of train tracks.

By Proposition ??, there is a train track Θ which weakly carries the α_t for t sufficiently close to 0.

Assume in addition that the curve $t \mapsto \alpha_t$ has a right derivative $\dot{\alpha}_0$ at $t = 0$. In this case, this is equivalent to the property that, for every edge e of Θ, the map $t \mapsto \alpha_t(e)$ has a right derivative $\dot{\alpha}_0(e)$ at $t = 0$. For every finite edge path γ of Θ, Proposition ?? shows that $\alpha_t(\gamma)$ is a certain piecewise linear function of the edge weights $\alpha_t(e)$. It follows that the map $t \mapsto \alpha_t(\gamma)$ has a right derivative $\dot{\alpha}_0(\gamma)$ at $t = 0$ for every finite edge path γ of Θ.

LEMMA. *Consider a curve $t \mapsto \alpha_t \in \mathcal{ML}(S)$, $t \in [0, \varepsilon[$, which has a right derivative $\dot{\alpha}_0$ at $t = 0$. Assume that the α_t are weakly carried by a train track Θ. Then, the essential support λ_{0^+} of α_t as t tends to 0^+ consists of those geodesics g of S which satisfy the following two conditions:*

(i) g is weakly carried by Θ; in particular, g is associated to a bi-infinite edge path γ_g of Θ, by Proposition ??;

(ii) for every finite edge path γ of Θ that is contained in γ_g, at least one of $\alpha_0(\gamma)$ and $\dot{\alpha}_0(\gamma)$ is different from 0.

PROOF. As usual, it is convenient to lift the situation to the universal covering $\widetilde{S}$. Let $\widetilde{\lambda}_{0^+}$ denote the preimage of λ_{0^+} in $\widetilde{S}$, and let $\widetilde{\lambda}_t$ be the preimage of the geodesic lamination λ_t underlying λ_t. By construction, every leaf of $\widetilde{\lambda}_{0^+}$ can be arbitrarily approximated by a leaf of some $\widetilde{\lambda}_t$. By Lemma 1.??, the set $G(\widetilde{\Theta})$ of those geodesics of $\widetilde{S}$ which are weakly carried by $\widetilde{\Theta}$ is closed in $G(\widetilde{S})$. It follows that $\widetilde{\lambda}_{0^+}$ is weakly carried by $\widetilde{\Theta}$, and therefore that λ_{0^+} is weakly carried by Θ.

As in §§1.?? and 3.??, for a finite edge path γ in $\widetilde{\Theta}$, let $G(\gamma)$ denote the open subset of $G(\widetilde{\Theta}) \subset G(\widetilde{S})$ consisting of those geodesics which realize the edge path γ. Consider a geodesic g which is weakly carried by $\widetilde{\Theta}$, and let γ_g be the bi-infinite edge path in $\widetilde{\Theta}$ associated to g by Lemma 1.??. By Lemma 1.??, as γ ranges over all finite edge paths contained in γ_g, the corresponding $G(\gamma)$ form a basis of neighborhoods of g in $G(\widetilde{\Theta})$. As a consequence, the geodesic $g \in G(\widetilde{\Theta})$ is in $\widetilde{\lambda}_{0^+}$ if and only if $\liminf_{t \to 0^+} \frac{1}{t}\alpha_t(\gamma) > 0$ for every finite edge path γ realized by g. Since $\dot{\alpha}_0(\gamma)$ exists, the condition that $\liminf_{t \to 0^+} \frac{1}{t}\alpha_t(\gamma) > 0$ is equivalent to the property that $\alpha_0(\gamma) > 0$ or $\dot{\alpha}_0(\gamma) > 0$, which concludes the proof. □

COROLLARY. *Consider a curve $t \mapsto \alpha_t \in \mathcal{ML}(S)$, $t \in [0, \varepsilon[$, which has a right derivative $\dot{\alpha}_0$ at $t = 0$. The essential support λ_{0^+} of α_t as t tends to 0^+ depends only on α_0 and on the vector $\dot{\alpha}_0$ tangent to $\mathcal{ML}(S)$ at α_0.* □

PROPOSITION 4. *If the path $t \mapsto \alpha_t \in \mathcal{ML}(S)$ is piecewise linear, the essential support λ_{0^+} is equal to the Hausdorff limit as t tends to 0^+ of the geodesic laminations λ_t underlying the α_t.*

PROOF. Let Θ be a train track which weakly carries the α_t for t small enough. It clearly suffices to prove that $\widetilde{\lambda}_{0^+} \subset G(\widetilde{S})$ is the Hausdorff limit of the supports

$\widetilde{\lambda}_t$ of the α_t, interpreting α_t as a measure on $G(\widetilde{S})$. Actually, it even suffices to prove that, for every generic tie k_0 of the preimage $\widetilde{\Theta} \subset \widetilde{S}$, the set of those geodesics of $\widetilde{\lambda}_t$ crossing k_0 converges to the set of those geodesics of $\widetilde{\lambda}_{0^+}$ crossing k_0.

For $r \geqslant 0$, let Γ_r denote the set of edge paths $\gamma = \langle e_{-r}, e_{-r+1}, \ldots, e_{r-1}, e_r\rangle$ in $\widetilde{\Theta}$ such that e_0 is the edge containing k_0. If $G(\widetilde{S})$ is endowed with a metric compatible with its Hölder structure, an easy geometric estimate provides two constants $A > 0$, $B > 0$, depending only on k_0 and on the lengths and widths of the edges of τ, such that the diameter of the set $G(\gamma)$ of geodesics realizing $\gamma \in \Gamma_r$ is bounded by Ae^{-Br}; indeed, it suffices to check this when $G(\widetilde{S}) = \left(\widetilde{S}_\infty \times \widetilde{S}_\infty - \Delta\right)/\mathbb{Z}_2$ is endowed with the product metric coming from the angle metric on $\widetilde{S}_\infty$ based at some point of k_0, in which case it follows from the negativity of the curvature of the metric of S.

By Lemma ??, $\widetilde{\lambda}_{0^+}$ consists of those geodesics g which are weakly carried by $\widetilde{\Theta}$ and such that $\alpha_0(\gamma) > 0$ or $\dot{\alpha}_0(\gamma) > 0$ for every finite edge path γ realized by g. By Lemma ??, the map $t \mapsto \alpha_t(\gamma)$ is piecewise linear. It follows that the above condition is equivalent to the condition that $\alpha_t(\gamma) > 0$ for every t sufficiently close to 0.

Since Γ_r is finite, this proves that, for t sufficiently close to 0, the edge paths $\gamma \in \Gamma_r$ for which $G(\gamma)$ meets $\widetilde{\lambda}_{0^+}$ are exactly those for which $G(\gamma)$ meets $\widetilde{\lambda}_t$. As a consequence, the Hausdorff distance between the set of geodesics of $\widetilde{\lambda}_{0^+}$ crossing k_0 and the set of geodesics of $\widetilde{\lambda}_t$ crossing k_0 is at most Ae^{-Br} for t sufficiently close to 0.

Letting r tend to ∞ then proves the result we wanted. □

5.4. Tangent vectors to $\mathcal{ML}(S)$ as geodesic laminations with transverse cocycles

If k is an arc transverse to the support of a measured m_0–geodesic lamination $\alpha \in \mathcal{ML}(S)$, recall that $\alpha(k)$ denotes the mass of the measure deposited by α on k. In theory, the notation should reflect the fact that this number $\alpha(k)$ depends on the negatively curved metric m_0. We assume m_0 fixed so that there is no ambiguity.

Proposition. *Consider a curve $t \mapsto \alpha_t \in \mathcal{ML}(S)$, $t \in [0, \varepsilon[$, which has a right derivative $\dot{\alpha}_0$ at $t = 0$, and let k be an arc which is in general position with respect to m_0–geodesics. Then, the right derivative $\dot{\alpha}_0(k)$ exists and depends only on the tangent vector $\dot{\alpha}_0$. In addition, consider the essential support λ_{0^+} of α_t as t tends to O^+; if k' is another arc which is in general position with respect to m_0–geodesic arcs and if k' is homotopic to k by a homotopy respecting λ_{0^+}, then $\dot{\alpha}_0(k') = \dot{\alpha}_0(k)$.*

Proof. By Proposition ??, there is a train track Θ which weakly carries the α_t for t sufficiently small. As usual, lewt $\widetilde{\Theta}$ be the pre-image of k in the universal covering $\widetilde{S}$, and lift k to an arc in $\widetilde{S}$, which we will also denote by k. For an m_0–geodesic $g \in G\left(\widetilde{S}\right)$, let $n_k(g)$ be the number of times it crosses

$k \subset \widetilde{S}$; note that, by transversality, $n_k(g)$ is finite and uniformly bounded if g be projects to a simple geodesic of S. Then, by definition,

$$\begin{aligned}\alpha_t(k) &= \int_{G(\widetilde{S})} n_k(g)\, d\alpha_t(g) \\ &= \sum_{n=1}^{\infty} n\alpha_t\left(\left\{g \in G^{\mathrm{s}}\left(\widetilde{S}\right); n_k(g) = n\right\}\right) \\ &= \sum_{n=1}^{\infty} n\alpha_t\left(\left\{g \in G\left(\widetilde{\Theta}\right) \cap G^{\mathrm{s}}\left(\widetilde{S}\right); n_k(g) = n\right\}\right)\end{aligned}$$

where $G^{\mathrm{s}}\left(\widetilde{S}\right)$ denotes the set of those geodesics of $\widetilde{S}$ which project to simple geodesics in S. Note that the sum is finite.

For every integer $n > 0$, the set $\left\{g \in G\left(\widetilde{\Theta}\right) \cap G^{\mathrm{s}}\left(\widetilde{S}\right); n_k(g) = n\right\}$ is compact since k is in general position with respect to simple m_0–geodesics (and since $G\left(\widetilde{\Theta}\right)$ is closed in $G\left(\widetilde{S}\right)$ by Lemma ??). Using Lemmas ?? and ??, we can therefore decompose this set as the disjoint union of finitely many subsets of the form $G(\gamma)$, where γ is a finite edge path in $\widetilde{\Theta}$. For each such γ, $\alpha_t(\gamma) = \alpha_t(G(\gamma))$ is a picewise linear function of $\alpha_t \in \mathcal{ML}(S)$. It follows that $\alpha_t(k)$ is a piecewise linear function of α_t. ■put this part in an earlier chapter?

In particular, if the curve $t \mapsto \alpha_t$ admits a tangent vector $\dot{\alpha}_0$ at $t = 0$, the function $t \mapsto \alpha_t(k)$ admits a derivative $\dot{\alpha}_0(k)$ which depends only on the tangent vector $\dot{\alpha}_0$. Namely, $\dot{\alpha}_0(k)$ is the image of $\dot{\alpha}_0$ under the tangent map of the piecewise linear function $\alpha \mapsto \alpha(k)$.

We now prove the second statement of Proposition ??. Let k and k' be two arcs which are in general position with respect to simple m_0–geodesics and which are homotopic by a homotopy respecting the essential support λ_{0^+}. In the universal covering $\widetilde{S}$, let $\widetilde{\lambda}_{0^+}$ be the pre-image of λ_{0^+}, and lift k and k' to two arcs, also denoted by k and k', that are homotpic respecting $\widetilde{\lambda}_{0^+}$.

Consider the set A of those geodesics $g \in G^{\mathrm{s}}\left(\widetilde{S}\right)$ which cut k and k' with a different multiplicity, namely such that $n_k(g) \neq n_{k'}(g)$. The set A is compact, and the term $|\alpha_t(k) - \alpha_t(k')|$ is bounded by $N\alpha_t(A)$, where N is any upper bound for $n_k(g)$ and $n_{k'}(g)$ with $g \in G^{\mathrm{s}}\left(\widetilde{S}\right)$. Because k and k' are homotopic respecting $\widetilde{\lambda}_{0^+}$, every geodesic of the essential support $\widetilde{\lambda}_{0^+}$ crosses k and k' with the same multiplicity. ■ Beware of limsup and liminf! With the Tracking Lemma ??, this implies that $\liminf_{t\to 0^+} \alpha_t(A)/t = 0$ and proves that $\dot{\alpha}_0(k) = \dot{\alpha}_0(k')$. □

Given a tangent vector $\dot{\alpha}$ of $\mathcal{ML}(S)$ at α_0, we have associated to $\dot{\alpha}$ a geodesic lamination $\lambda_{\dot{\alpha}}$ and a certain transverse structure to $\lambda_{\dot{\alpha}}$ as follows: Choose a curve $t \mapsto \alpha_t \in \mathcal{ML}(S)$, $t \in [0, \varepsilon[$, passing through α_0 and tangent to $\dot{\alpha}$ at α_0. Then, $\lambda_{\dot{\alpha}}$ is the essential support of α_t as t tends to 0^+. For each arc k transverse to $\lambda_{\dot{\alpha}}$, Proposition ?? enables us to approximate k in the C^1 topology by an arc k' which is in general position with respect to m_0–geodesics; any two such C^1–approximation of k are homotopic respecting $\lambda_{\dot\alpha}$. Then, Proposition ?? associates to k a well-defined number $\dot\alpha(k) = \dot\alpha_0(k') \in \mathbb{R}$. Propositions ?? and ?? show that the geodesic lamination $\lambda_{\dot\alpha}$ and the rule $k \mapsto \dot\alpha(k) \in \mathbb{R}$ depend only on the tangent vector $\dot\alpha$, and not on the choice of the curve $t \mapsto \alpha_t$.

Note that the rule $k \mapsto \dot\alpha(k)$ satisfies the following properties: It is invariant under homotopy of k respecting $\lambda_{\dot\alpha}$; it is ***finitely additive*** in the sense that, if we split k as the union of two sub-arcs k_1 and k_2 meeting in one point of $k - \lambda_{\dot\alpha}$, then $\dot\alpha(k) = \dot\alpha(k_1) + \dot\alpha(k_2)$. We summarize these two properties by saying that this defines an $\mathbb{R}$–***valued transverse cocycle*** for the geodesic lamination $\lambda_{\dot\alpha}$.

In the next chapter, we will study transverse cocycles for geodesic laminations in more details. In particular, we will characterize which geodesic laminations with transverse cocycles are associated to tangent vectors of $\mathcal{ML}(S)$ in this way.

CHAPTER VI

TRANSVERSE COCYCLES FOR GEODESIC LAMINATIONS

In the preceding chapter, we encountered a certain type of transverse structures for a geodesic lamination, namely transverse cocycles. This chapter is devoted to an extensive study of the combinatorial properties of these transverse cocycles. In particular, in §6.2, we completely classify the cocycles transverse to a given geodesic lamination λ, in terms of weights on a train track carrying λ. When we encountered transverse cocycles in Chapter 5, these occurred as tangent vectors to the space $\mathcal{ML}(S)$ of measured geodesic laminations. In §6.5, we determine exactly which transverse cocycles correspond to tangent vectors to $\mathcal{ML}(S)$. When the surface S is oriented, the space of all $\mathbb{R}$–valued transverse cocycles for the geodesic lamination λ inherits a certain anti-symmetric bilinear form, called the Thurston intersection form. This intersection is studied in detail in §6.5. A particularly inportant property is Proposition 6.22, which shows that the Thurston intersection form is non-degenerate when thegeodesic lamination λ is maximal.

6.1. Transverse cocycles

Consider an m_0–geodesic lamination λ and let $\mathbb{M}$ be a module over a commutative unitary ring R. In most cases of interest here, $\mathbb{M}$ will be $\mathbb{R}$, considered as a vector space over itself, and occasionally the abelian groups (= $\mathbb{Z}$–modules) $\mathbb{R}/2\pi\mathbb{Z}$, $\mathbb{C}/2\pi\mathbb{Z}$ or $\mathbb{Z}/2\mathbb{Z}$. An $\mathbb{M}$–***valued transverse cocycle*** for λ is a map which associates an element $\alpha(k) \in \mathbb{M}$ to each unoriented arc k transverse to λ, and which satisfies the following properties: α is ***(finitely) additive*** in the sense that $\alpha(k) = \alpha(k_1) + \alpha(k_2)$ if we split k into two subarcs k_1 and k_2 with disjoint interiors; and α is λ–***invariant*** in the sense that $\alpha(k) = \alpha(k')$ whenever the arc k can be deformed to the arc k' by a homotopy respecting λ.

Let $\mathcal{H}(\lambda;\mathbb{M})$ denote the space of transverse cocycles for the geodesic lamination λ, valued in the R–module $\mathbb{M}$. This space $\mathcal{H}(\lambda;\mathbb{M})$ has a natural structure of R–module.

Transverse (positive) measures α for the geodesic lamination λ provide examples of $\mathbb{R}$–valued transverse cocycles for λ, by associating to each transverse arc k the mass $\alpha(k)$ of the measure deposited by α on k. In §5.2, we encountered an example of a transverse cocycle which is not associated to a transverse measure. Here is a different characterization of the transverse cocycles associated to transverse measures.

Proposition 6.1. *An $\mathbb{R}$–valued transverse cocycle $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ for the geodesic lamination λ corresponds to a transverse measure for λ if and only if*

$\alpha(k) \geqslant 0$ for every arc k transverse to λ.

Proof. The condition is clearly necessary.

Conversely, assume that $\alpha(k) \geqslant 0$ for every arc k transverse to λ. Given such a transverse arc k and a continuous function $\psi : k \to \mathbb{R}$, decompose k into small arcs $k_1, k_2, \ldots, k_n$ with disjoint interiors, choose an arbitrary point x_i in each k_i, and consider the Riemann sum $\sum_{i=1}^{n} \alpha(k_i)\,\psi(x_i)$. Because the $\alpha(k_i)$ are non-negative and by uniform continuity of ψ, it is easy to check that this Riemann sum converges to some number $\alpha'(\psi)$ as the size of the k_i tends to 0. In addition $\alpha'(\psi) \leqslant \alpha(k)\,\|\psi\|_\infty$. Therefore, α' defines a positive Radon measure α' on k. Since this measure is clearly invariant under homotopy of k respecting λ, this defines a transverse measure α' for λ. By construction the mass of the measure defined by α' on every transverse arc k is equal to the number $\alpha(k) \in \mathbb{R}$ defined by the cocycle α. It follows that the transverse cocycle α is the one associated to the transverse measure α'. □

This Riemann sum argument also shows:

Proposition 6.2. *A transverse measure for the geodesic lamination λ is completely determined by its associated transverse cocycle. In particular, the space $\mathcal{M}(\lambda)$ of all transverse measures for λ is naturally identified to its image in $\mathcal{H}(\lambda; \mathbb{R})$.* □

A priori, our definition of transverse cocycles depends on the choice of a negatively curved metric m_0. We want to show that this is not the case, by giving a metric independent description of these transverse cocycles. We will accomplish this by modifying the framework which we used in §2.?? for transverse measures, giving it a more algebraic and less analytical flavor.

As usual, let $\widetilde{\lambda}$ denote the geodesic lamination of $\widetilde{S}$ which is the pre-image of λ in the universal covering $\widetilde{S}$. In the space $G\left(\widetilde{S}\right)$ of geodesics of $\widetilde{S}$, we can also identify $\widetilde{\lambda}$ to the closed subset of $G\left(\widetilde{S}\right)$ consisting of all leaves of $\widetilde{\lambda}$, as in §1.??. In §2.??, we saw that a transverse measure for λ is the same thing as a $\pi_1(S)$–invariant measure on $\widetilde{\lambda} \subset G\left(\widetilde{S}\right)$. If $C_0\left(\widetilde{\lambda}; \mathbb{R}\right)$ denote the space of all continuous functions $\varphi : \widetilde{\lambda} \to \mathbb{R}$ with compact support, a transverse measure for λ is therefore equivalent to a linear form $\alpha : C_0\left(\widetilde{\lambda}; \mathbb{R}\right) \to \mathbb{R}$ which is invariant under the action of $\pi_1(S)$ and which satisfies the positivity condition that $\alpha(\varphi) \geqslant 0$ whenever the function φ is non-negative, by the Riesz Representation Theorem.

Since an $\mathbb{M}$–valued transverse cocycle is more combinatorial and algebraic in nature, it is now natural to substitute to $C_0\left(\widetilde{\lambda}; \mathbb{R}\right)$ the space $LC_0\left(\widetilde{\lambda}; R\right)$ of all functions $\varphi : \widetilde{\lambda} \to R$ which are locally constant (the algebraic version of continuity!) and with compact support; here R is the ring over which $\mathbb{M}$ is a module. We can then consider linear maps $\alpha : LC_0\left(\widetilde{\lambda}; R\right) \to \mathbb{M}$ which are invariant under the natural action of $\pi_1(S)$.

Such a $\pi_1(S)$–invariant linear $\alpha : LC_0\left(\widetilde{\lambda}; R\right) \to \mathbb{M}$ defines a transverse cocycle as follows. If k is an arc transverse to λ, lift it to an arc $\widetilde{k}$ transverse to $\widetilde{\lambda}$ in $\widetilde{S}$. The arc $\widetilde{k}$ defines a locally constant function $\varphi_{\widetilde{k}} : \widetilde{\lambda} \to R$ which just counts the number of times each leaf of $\widetilde{\lambda}$ intersects $\widetilde{k}$ (identifying the integer n to n times the unit element in R); note that the support of $\varphi_{\widetilde{k}}$ is contained in the set of those elements of $G\left(\widetilde{S}\right)$ which meet k, and is therefore compact. Then, define $\alpha(k) = \alpha\left(\varphi_{\widetilde{k}}\right) \in \mathbb{M}$. This $\alpha(k)$ is independent of the choice of the lift $\widetilde{k}$ by $\pi_1(S)$–invariance of α. If the transverse arc k' is homotopic to k by a homotopy respecting λ, we can lift it to an arc $\widetilde{k}'$ that is homotopic to $\widetilde{k}$ by a homotopy respecting $\widetilde{\lambda}$; since $\varphi_{\widetilde{k}'} = \varphi_{\widetilde{k}}$, it follows that $\alpha(k') = \alpha(k)$. Since the additivity is immediate, this proves that we have indeed associated an $\mathbb{M}$–valued transverse cocycle to the $\pi_1(S)$–invariant linear function $\alpha : LC_0\left(\widetilde{\lambda}; R\right) \to \mathbb{M}$.

Proposition 6.3. *The above construction defines a one-to-one correspondence between $\mathbb{M}$–valued transverse cocycles for the geodesic lamination λ and $\pi_1(S)$–invariant linear functions $\alpha : LC_0\left(\widetilde{\lambda}; R\right) \to \mathbb{M}$.*

Proof. We need to show that every $\mathbb{M}$–valued transverse cocycle α is associated to a unique such linear $LC_0\left(\widetilde{\lambda}; R\right) \to \mathbb{M}$.

A geodesic $g \in \widetilde{\lambda} \subset G\left(\widetilde{S}\right)$ admits arbitrarily small neighborhoods in $\widetilde{\lambda}$ of the form $U_{\widetilde{k}} = \left\{h \in \widetilde{\lambda}; h \cap \widetilde{k} \neq \emptyset\right\}$, where $\widetilde{k}$ is a small arc transverse to $\widetilde{\lambda}$, small enough that it intersects each $h \in \widetilde{\lambda}$ in at most one point. Also, for any two such $U_{\widetilde{k}}$ and $U_{\widetilde{k}'}$, there is an arc $\widetilde{k}'' \subset \widetilde{k}$ such that $U_{\widetilde{k}''} = U_{\widetilde{k}} \cap U_{\widetilde{k}'}$. Also, $U_{\widetilde{k}} = U_{\widetilde{k}'}$ precisely when $\widetilde{k}$ and $\widetilde{k}'$ are homotopic by a homotopy respecting $\widetilde{\lambda}$.

It follows from these observations that every $\varphi \in LC_0\left(\widetilde{\lambda}; R\right)$ can be written as a linear combination $\varphi = \sum_{i=1}^{n} a_i \varphi_{\widetilde{k}_i}$, where the coefficients a_i are in R, where the $\varphi_{\widetilde{k}_i} \in LC_0\left(\widetilde{\lambda}; R\right)$ are the counting functions associated to transverse arcs $\widetilde{k}_i$ as above, and where each geodesic $g \in \widetilde{\lambda}$ meets at most one $\widetilde{k}_i$. Set $\alpha(\varphi) = \sum_{i=1}^{n} \alpha(k_i) \in \mathbb{M}$, where k_i is the arc transverse to λ which is the image of $\widetilde{k}_i$ in S. The same observations easily prove that $\alpha(\varphi)$ is independent of the choice of the linear combination decomposition $\varphi = \sum_{i=1}^{n} a_i \varphi_{\widetilde{k}_i}$.

The linear map $LC_0\left(\widetilde{\lambda}; R\right) \to \mathbb{M}$ defined by $\varphi \mapsto \alpha(\varphi)$ is clearly invariant under the action of $\pi_1(S)$. Also, it immediately follows from the definitions that this is the only such map $LC_0\left(\widetilde{\lambda}; R\right) \to \mathbb{M}$ whose associated transverse cocycle is equal to α, which completes the proof of Proposition 6.3. □

Since the topological space $G\left(\widetilde{S}\right)$ was constructed in a metric independent way in §1.?? and since $\widetilde{\lambda} \subset G\left(\widetilde{S}\right)$ is endowed with the subspace topology, an immediate corollary of Proposition 6.3 is that the notion of transverse cocycle for the geodesic lamination λ can be defined in a metric independent way.

6.2. Transverse cocycles to a given geodesic lamination

Let λ be a geodesic lamination, and let $\mathbb{M}$ be an R–module. We want to determine the R–module $\mathcal{H}(\lambda;\mathbb{M})$ of all $\mathbb{M}$–valued transverse cocycles for λ.

By Proposition 1.??, λ is strongly carried by a fattened train track Φ. As usual, let $\widetilde{\lambda}$ and $\widetilde{\Phi}$ denote the pre-images of λ and Φ in the universal covering $\widetilde{S}$.

For every generic tie k of Φ, we can consider the element $\alpha(k) \in \mathbb{M}$ associated by α to k. Since every two ties of an edge are homotopic respecting λ, this $\alpha(k)$ depends only on the edge e containing k, and will also be denoted by $\alpha(e)$.

We first show that α is uniquely determined by these finitely many edge weights $\alpha(e) \in \mathbb{M}$.

Consider a fixed generic tie k_0 of Φ, and choose an orientation for k_0.

For a component d of $k_0 - \lambda$, pick a point $x_d \in d$ and consider the component k_d^- of $k_0 - \{x_d\}$ that is below x_d for the orientation of k_0. By invariance of α under homotopy respecting λ, the element $\alpha\left(k_d^-\right) \in \mathbb{M}$ is independent of the choice of x_d, and is called the α–***height*** $h_\alpha(d)$ of d with respect to α. This α–height will also play an important rôle in the next chapter, when we consider transverse distributions for geodesic laminations.

We first show how to compute $h_\alpha(d) \in \mathbb{M}$ from the edge weights $\alpha(e)$ associated by α to the edges of Φ.

When d is one of the two components of $k_0 - \lambda$ that contain an end point of k_0, then $h_\alpha(d)$ is equal to 0 for the component containing the negative end point, and is equal to the weight $\alpha(e_0)$ of the edge e_0 containing k_0 for the one containing the positive end point. Therefore, we can restrict attention to the case where d contains no end point of k_0.

Identify k_0 to one of its lifts as a generic tie of $\widetilde{\Phi}$. Then, there are two leaves g_d^+ and g_d^- of $\widetilde{\lambda}$ which pass through the end points of d, where $k_0 \cap g_d^-$ is below $k_0 \cap g_d^+$ for the orientation of k_0.

Since g_d^- and g_d^+ are distinct, they cannot realize the same bi-infinite edge path in $\widetilde{\Phi}$. Therefore, g_d^+ and g_d^- respectively realize some edge paths $\langle e_0, e_1, \dots, e_r, e_{r+1}\rangle$ and $\left\langle e_0, e_1, \dots, e_r, e'_{r+1}\right\rangle$ with $e_{r+1} \neq e'_{r+1}$, where e_0 is the edge containing k_0.

■Figure with Φ, g_d^+ and g_d^-, edges branching in and out, etc...

Any leaf g of $\widetilde{\lambda}$ that hits k_0 below d is disjoint from g_d^+, and therefore must realize an edge path $\langle e_0, e_1, \dots, e_i, f\rangle$, with $0 < i \leqslant r$, where f is different from e_{i+1} and branches out on the negative side of $\langle e_0, e_1, \dots, e_r, e_{r+1}\rangle$, for the transverse orientation of this edge path determined by the orientation of k_0. Let f_1, f_2, … , f_p be the collection of these edges f which branch out on the negative side of $\langle e_0, e_1, \dots, e_r, e_{r+1}\rangle$, including $f_1 = e'_{r+1}$.

Conversely, every leaf which realizes such an $\langle f_k\rangle$ must, either hit k_0 below d, or realize an edge path $\langle f, e_j, \dots, e_{i-1}, e_i, f_k\rangle$ where f is different from e_{j-1} and branches in on the negative side of $\langle e_0, e_1, \dots, e_r, e_{r+1}\rangle$. Let f_{p+1}, f_{p+2}, … , f_q be the collection of these edges f which branch in on the negative side of $\langle e_0, e_1, \dots, e_r, e_{r+1}\rangle$.

In addition, any leaf which realizes $\langle f_i\rangle$ with $p+1 \leqslant i \leqslant q$ must exit through one of the f_j with $1 \leqslant j \leqslant p$.

By additivity of α, we obtain from these observations that $h_\alpha(d) = \sum_{i=1}^{p} \alpha(f_i) - \sum_{i=p+1}^{q} \alpha(f_i)$. This proves:

need to introduce $\alpha(\gamma)$ for an edge path γ, to be able to use additivity

LEMMA 6.4. *For every component d of $k_0 - \lambda$, the α–height $h_\alpha(d)$ is a linear function of the $\alpha(e)$ associated to the edges of Φ. More precisely, if the component d of $k_0 - \widetilde{\lambda}$ contains no end point of k_0,*

$$h_\alpha(d) = \sum_f \varepsilon(f)\,\alpha(f)$$

where the sum is over all edges f branching in or out on the negative side of the edge path $\langle e_0, e_1, \ldots, e_r, e_{r+1}\rangle$ defined above, and where $\varepsilon(f) = -1$ or $+1$ according to whether f branches in or out. If d contains one of the end points of k_0, then $h_\alpha(d)$ is equal to $\alpha(e_0)$ for the positive end point, and to 0 for the negative end point. □

As a corollary, we get:

PROPOSITION 6.5. *For every arc k transverse to λ, the number $\alpha(k)$ is uniquely determined by the weights $\alpha(e)$ associated by α to the edges e of the train track Φ.*

PROOF. If k is contained in a generic tie k_0 of Φ, this immediately follows from Lemma 6 since $\alpha(k) = h_\alpha\left(d_k^+\right) - h_\alpha\left(d_k^-\right)$, where d_k^+ and d_k^- are the components of $k_0 - \lambda$ respectively containing the positive and negative end points of k.

In the general case, k can be decomposed into finitely many small arcs k_i with disjoint interiors such that each k_i can be homotoped respecting α to an arc contained in a generic tie of Φ. By additivity and invariance of α under homotopy respecting λ, the result then follows by applying the previous case to each k_i. □

We can rephrase Proposition 6.5 in the following way. Recall that $\mathcal{H}(\lambda;\mathbb{M})$ denotes the R–module of all $\mathbb{M}$–valued transverse cocycles for λ. If λ is carried by the fattened train track Φ, the edge weights $\alpha(e)$ associated to the edges e of Φ by a transverse cocycle α clearly satisfy the switch relations. This defines an R–linear map $\mathcal{H}(\lambda;\mathbb{M})$ to the module $\mathcal{W}(\Phi;\mathbb{M})$ of all $\mathbb{M}$–valued edge weight systems for Φ. What Proposition 6.5 says is that this linear map $\mathcal{H}(\lambda;\mathbb{M}) \to \mathcal{W}(\Phi;\mathbb{M})$ is injective.

There is no reason for the map $\mathcal{H}(\lambda;\mathbb{M}) \to \mathcal{W}(\Phi;\mathbb{M})$ to be surjective without further hypothesis on the fattened train track Φ. For instance, consider the case where λ consists of only a simple closed geodesic. A transverse cocycle $\alpha \in \mathcal{H}(\lambda;\mathbb{M})$ is completely determined by the weight $\alpha(k)$ it assigns to an arbitrary k transverse to λ and cutting λ in exactly one point; this immediately follows from the additivity of α and from its invariance under homotopy respecting λ. This provides a natural isomorphism between $\mathcal{H}(\lambda;\mathbb{M})$ and $\mathbb{M}$. On the other hand, Propositions 3.?? and 3.?? show that, for some negatively curved metric,

λ is strongly carried by a fattened train track Φ for which $\mathcal{W}(\Phi;\mathbb{M})$ is much larger.

We will say that the m_0–geodesic lamination λ is ***snugly carried*** by the fattened train track Φ if λ is carried by Φ, if λ meets every tie of Φ, and if there is no curve carried by Φ which is disjoint from λ and which joins an end point of a spike of $S-\Phi$ to another one. If λ is carried but not snugly carried by a fattened train track Φ, cutting Φ open along appropriate curves gives a train track Φ' which carries λ and is topologically simpler than Φ. Applying a succession of such reductions, we eventually reach a train track which snugly carries λ. Also, note that the train track constructed in the proof of Proposition 1.?? snugly carries λ.

THEOREM 6.6. *Consider a geodesic lamination λ which is snugly carried by the fattened train track Φ. Then, every $\mathbb{M}$–valued edge weight system for Φ can be realized by a transverse cocycle for λ. As a consequence, the natural map $\mathcal{H}(\lambda;\mathbb{M}) \to \mathcal{W}(\Phi;\mathbb{M})$ is an isomorphism.*

PROOF. If E denotes the set of edges of Φ, let $a : E \to \mathbb{M}$ be an $\mathbb{M}$–valued edge weight system for Φ, satisfying the switch relations. We want to construct a transverse cocycle α for λ such that $\alpha(e) = a(e)$ for every edge e of Φ.

Consider an oriented generic tie k_0 of Φ.

To each component d of $k_0 - \lambda$, we associate a number $h_a(d)$ by the formula of Lemma 6.4 applied to a instead of α. Namely, if d does not contain either of the end points of k_0, we lift k_0 and d to the universal covering $\widetilde{S}$ and consider the two leaves g_d^+ and g_d^- of $\widetilde{\lambda}$ passing through its end points. There is an edge path $\left\langle e_0, e_1, \dots, e_{r(d)}\right\rangle$, starting at the edge e_0 of $\widetilde{\Phi}$ containing k_0, such that g_d^+ and g_d^- respectively realize edge paths $\left\langle e_0, e_1, \dots, e_{r(d)}, e_{r(d)+1}\right\rangle$ and $\left\langle e_0, e_1, \dots, e_{r(d)}, e'_{r(d)+1}\right\rangle$ with $e_{r(d)+1} \neq e'_{r(d)+1}$. Note that this edge path is uniquely determined because Φ snugly carries λ. Then define

$$h_\alpha(d) = \sum_f \varepsilon(f)\,\alpha(f)$$

where the sum is over all edges f branching in or out on the negative side of the edge path $\langle e_0, e_1, \dots, e_r, e_{r+1}\rangle$, and where $\varepsilon(f) = -1$ or $+1$ according to whether f branches in or out. When d contains one of the end points of k_0, define $h_a(d) = 0$ if d contains the negative end point, and $h_a(d) = a(e_0)$ if d contains the positive end point, where e_0 is the edge containing k_0.

If we reverse the orientation of k_0, then the number $h_a(d)$ is replaced by another number $\bar{h}_a(d)$ for every component d of $k_0 - \lambda$.

LEMMA 6.7. *For every component d of $k_0 - \widetilde{\lambda}$,*

$$h_a(d) + \bar{h}_a(d) = a(e_0).$$

PROOF. This is only the point in the proof of Theorem 6.6 where we use the switch relations.

The result is immediate when d contains an end point of k_0, so we can assume that this is not the case. Then, the leaves g_d^+ and g_d^- passing through the end points of d respectively realize edge paths $\left\langle e_0, e_1, \ldots, e_{r(d)}, e_{r(d)+1}\right\rangle$ and $\left\langle e_0, e_1, \ldots, e_{r(d)}, e'_{r(d)+1}\right\rangle$ with $e_{r(d)+1} \neq e'_{r(d)+1}$.

By definition, $h_a(d)$ is equal to the sum of the weights of those edges which branch out on the negative side of $\left\langle e_0, e_1, \ldots, e_{r(d)}, e_{r(d)+1}\right\rangle$ minus the sum of the weights of those edges which branch in on the negative side of $\left\langle e_0, e_1, \ldots, e_{r(d)}, e'_{r(d)+1}\right\rangle$. Similarly, $\bar{h}_a(d)$ is equal to the sum of the weights of those edges which branch out on the positive side of $\left\langle e_0, e_1, \ldots, e_{r(d)}, e'_{r(d)+1}\right\rangle$ minus the sum of the weights of those edges which branch in on the positive side of $\left\langle e_0, e_1, \ldots, e_{r(d)}, e_{r(d)+1}\right\rangle$ (still using the transverse orientation of edge paths defined by the orientation of k_0).

Using the fact that the weights $a(e)$ satisfy the switch relations, we conclude that the contribution to $h_a(d) + \bar{h}_a(d)$ of those edges which branch in or out at the switch separating e_k from e_{k+1} is equal to $a(e_k) - a(e_{k+1})$ if $0 \leqslant k < r(d)$, and to $a\left(e_{r(d)}\right)$ if $k = r(d)$. It follows that $h_a(d) + \bar{h}_a(d) = a(e_0)$. □

If k is an arc transverse to λ and contained in the oriented generic tie k_0, we define

$$\alpha(k) = h_a\left(d_k^+\right) - h_a\left(d_k^-\right)$$

where d_k^+ and d_k^- are the components of $k_0 - \lambda$ that contain the positive and negative end points of k, respectively. By Lemma 6.7, this $\alpha(k)$ is independent of the choice of an orientation for k_0. Also, $\alpha(k)$ is clearly an additive function of the arc k.

Let us show that, if k' is another arc transverse to λ, contained in a generic tie k'_0 of Φ, and homotopic to k respecting λ, then $\alpha(k) = \alpha(k')$. Subdividing k and the homotopy to k' if necessary, we can restrict attention to the case where the homotopy from k to k' crosses exactly one switch s of Φ. In particular, the edges e_0 and e'_0 of Φ that respectively contain k_0 and k'_0 meet along s. Orient k_0 and k'_0 so that they induce the same orientation for s.

We first compare $h_a\left(d_k^+\right)$ and $h_a\left(d_{k'}^+\right)$. If the end points of d_k^+ and $d_{k'}^+$ are both contained in λ, then d_k^+ and $d_{k'}^+$ are adjacent to the same leaves of λ since the homotopy respecting λ from k to k' crosses only the switch s. From the definition of $h_a\left(d_k^+\right)$ and $h_a\left(d_{k'}^+\right)$, it follows that $h_a\left(d_k^+\right) - h_a\left(d_{k'}^+\right)$ is equal to $A - A'$ where: A is the sum of the weights of those edges which branch in at s on the same side as e_0 and on the negative side of e_0, for the transverse orientation of the edge e_0 defined by the orientation of its tie k_0; A' is the sum of the weights of those edges which branch out at s on the same side as e'_0 and on the negative side of e'_0. When one of the end points of d_k^+ is an end point of k_0, or when one of the end points of $d_{k'}^+$ is an end point of k'_0, a similar analysis shows that the equality $h_a\left(d_k^+\right) - h_a\left(d_{k'}^+\right) = A - A'$ still holds in this case.

Similarly, $h_a\left(d_k^-\right) - h_a\left(d_{k'}^-\right) = A - A'$. We conclude that

$$\alpha(k) = h_a\left(d_k^+\right) - h_a\left(d_k^-\right) = h_a\left(d_{k'}^+\right) - h_a\left(d_{k'}^-\right) = \alpha(k')$$

as required.

Consider now an arbitrary arc k transverse to λ. This arc k can be subdivided as the union of finitely many arcs $k_1, \ldots, k_n$ with disjoint interiors such that each k_i is homotopic respecting λ to a transverse arc k'_i which is contained in a generic tie of Φ. We then define

$$\alpha(k) = \sum_{i=1}^{n} \alpha(k'_i)$$

where, for the transverse arc k'_i contained in a generic tie, $\alpha(k'_i)$ is defined as above. The above analysis of the case of subarcs of generic ties immediately shows that this $\alpha(k)$ is independent of the choice of the k'_i, is invariant under homotopy respecting λ, and is additive. Therefore, this defines a transverse cocycle $\alpha \in \mathcal{H}(\lambda; \mathbb{M})$, whose associated edge weight system is $a \in \mathcal{W}(\Phi; \mathbb{M})$ by construction.

This proves that the map $\mathcal{H}(\lambda; \mathbb{M}) \to \mathcal{W}(\Phi; \mathbb{M})$ is surjective. By Proposition 6.5, it is consequently an R–linear isomorphism. □

If Φ is a fattened train track which snugly carries λ, Theorem 6.6 provides an isomorphism between $\mathcal{H}(\lambda; \mathbb{M})$ and $\mathcal{W}(\Phi; \mathbb{M})$, and we computed $\mathcal{W}(\Phi; \mathbb{M})$ in Theorem 3.??. We can rephrase this computation in a train track independent way in the following manner.

Let the ***Euler characteristic*** $\chi(\lambda)$ of the geodesic lamination λ be the alternating sum of the ranks of its Čech cohomology groups (considering λ as a subset of S).

Lemma 6.8. *The Euler characteristic of a geodesic lamination λ is equal to the Euler characteristic of any fattened train track Φ that snugly carries λ.*

Proof. Splitting up Φ along arcs that are carried by Φ, originate from the spikes of $S - \Phi$ and are disjoint from λ, we can create a family of fattened train tracks Φ_i which form a nested basis of neighborhoods for λ and such that the inclusion maps $\Phi_i \to \Phi$ are homotopy equivalences. The inclusion map $\lambda \to \Phi$ then induces an isomorphism of Čech cohomology, and therefore $\chi(\lambda) = \chi(\Phi)$. □

It easily follows from Lemma 6.8 and from the definition of fattened train tracks that the Euler characteristic of a geodesic lamination is always nonpositive.

A geodesic lamination λ is ***orientable*** if it is possible to continuously orient its leaves. In practice, to decide whether a geodesic lamination is orientable, it is useful to consider a fattened train track which snugly carries it. A fattened train track is ***orientable*** if it is possible to continuously choose a transverse orientation for its ties. There is an obvious finite algorithm to decide whether or not a given fattened train track is orientable.

Lemma6.9. *If the geodesic lamination λ is snugly carried by fattened the train track Φ, then λ is orientable if and only if Φ is orientable.*

Proof. Clearly, an orientation for Φ determines an orientation for λ.

Conversely, assume that λ is oriented. We claim that the orientation of the leaves of λ induces the same transverse orientation on each tie k of Φ. Indeed, there would otherwise be a component d of $k - \lambda$ such that the leaves g_d^+ and g_d^- of λ passing through the end points of d induce opposite transverse orientations on k. Because Φ snugly carries λ, these two leaves have to follow the same infinite edge path in one direction. In particular, they are asymptotic on one side, contradicting the continuity of the orientation of λ since they carry opposite orientations.

Therefore, the orientation of λ induces a well defined transverse orientation for the ties of Φ, and defines an orientation for Φ. □

Lemma 6.9 makes it possible to decide whether or not a given geodesic lamination is orientable.

THEOREM 6.10. *For a geodesic lamination λ, the module $\mathcal{W}(\lambda;\mathbb{M})$ of all transverse $\mathbb{M}$–valued cocycles for λ is isomorphic to*

$$\mathbb{M}^{-\chi(\lambda)} \oplus \mathbb{M}^{N_{\mathrm{o}}} \oplus \{m \in \mathbb{M}; 2m = 0\}^{N_{\mathrm{no}}}$$

where $\chi(\lambda)$ is the Euler characteristic, N_{o} is the number of orientable components of λ and N_{no} is the number of its non-orientable components.

PROOF. Let Φ be a fattened train track which snugly carries λ. In view of Theorem 6.6, Proposition 3.??, and Lemmas 6.8 and 6.9, the only thing to check is that each component of Φ contains exactly one component of λ, but this follows by an argument which is almost identical to the proof of Lemma 6.9. □

COROLLARY 6.11. *The transverse $\mathbb{R}$–valued cocycles for a geodesic lamination λ form a vector space $\mathcal{H}(\lambda;\mathbb{R})$ of dimension $-\chi(\lambda) + N_{\mathrm{o}}$, where N_{o} is the number of orientable components of λ.* □

EXERCISE. We want to investigate the behavior of transverse cocycles with respect to the decomposition of the geodesic lamination into minimal sub-laminations and infinite isolated leaves, as in 1.??. If g is an infinite isolated leaf of the geodesic lamination λ and if $\alpha \in \mathcal{H}(\lambda;\mathbb{M})$ is a transverse cocycle for λ, pick a small arc k_g transverse to λ such that $k_g \cap \lambda$ consists of a single point and is contained in g, and define the α–***mass*** $\alpha(g)$ of g with respect to α as the element $\alpha(k_g) \in \mathbb{M}$.

a) Show that this α–mass $\alpha(g)$ is independent of the choice of k_g.
b) Let λ_1 be an orientable minimal sublamination of λ. For each orientation of λ_1, consider those infinite isolated half-leaves of λ which spiral along λ_1 in the direction of the orientation considered, and consider the sum of the α–masses of these half-leaves. Show that the sums corresponding to the two orientations of λ_1 are equal.
c) Let λ_{m} denote the union of the minimal sub-laminations of λ. Also, let $\mathcal{W}(\lambda - \lambda_{\mathrm{m}};\mathbb{M})$ be the R–module consisting of all assignments a of a weight $a(g) \in \mathbb{M}$ to each infinite isolated leaf of λ such that, for every orientable minimal sub-lamination of λ, a satisfies the Orientable Sub-lamination Condition of b). Show that there is an exact sequence

$$0 \to \mathcal{H}(\lambda_{\mathrm{m}};\mathbb{M}) \to \mathcal{H}(\lambda;\mathbb{M}) \to \mathcal{W}(\lambda - \lambda_{\mathrm{m}};\mathbb{M}) \to 0$$

As a corollary of the fact that $\mathcal{H}(\lambda;\mathbb{R})$ is finite dimensional, we can obtain the following result of A. Katok [**Ka**]. (Katok's proof of this result can actually be seen as a precursor to Proposition 6.5).

PROPOSITION 6.12. *The space $\mathcal{M}(\lambda)$ of all transverse measures for a geodesic lamination λ is the cone over a certain linear simplex in the vector space $\mathcal{H}(\lambda;\mathbb{R})$. Namely, there exists finitely many transverse measures $\mu_1, \dots \mu_n$ such that every transverse measure μ for λ can be uniquely decomposed as a linear combination $\mu = \sum_{i=1}^n a_i\mu_i$ with non-negative coefficients $a_i \geqslant 0$.*

PROOF. We will prove Proposition 6.12 with the additional conclusion that each transverse measure μ_i is ***ergodic***, namely is non-trivial and cannot be decomposed as the sum $\mu_i = \mu_i' + \mu_i''$ of two transverse measures μ_i' and μ_i'' which are not real multiples of μ_i.

Step 1. Up to multiplication by real numbers, the geodesic lamination λ admits only finitely many ergodic transverse measures.

Suppose otherwise. Since the vector space $\mathcal{H}(\lambda;\mathbb{R})$ is finite dimensional (Theorem 6.10) and since a transverse measure is completely determined by its associated transverse cocycle (Proposition 6.2), we can then find a transverse measure μ which admits two distinct decompositions $\mu = \sum_{i=1}^m a_i\mu_i$ and $\mu = \sum_{j=1}^n b_j\nu_j$, where the transverse measures μ_i and ν_j are ergodic, where the coefficients a_i, b_j are positive, and where no ν_j is a real multiple of μ_1.

Consider an arc k transverse to λ, and the measures defined on k by μ_1 and ν_j. By a classical result of measure theory (see for instance ??), we can uniquely decompose the measure defined by ν_j as the sum of two measures μ_1^{ac} and μ_1^{s} which respectively are absolutely continuous and singular with respect to the measure defined by μ_1. Recall that this means that there exists a measurable function $f_k : k \to [0,\infty[$ such that $\mu_1^{\mathrm{ac}}(A) = \int_A f_k\,d\mu_1$ for every measurable subset A of k and a measurable subset B of k with $\mu_1(B) = 0$ such that $\mu_1^{\mathrm{s}}(A) = \mu_1^{\mathrm{s}}(A\cap B)$ for every measurable $A \subset k$. By uniqueness of this decomposition, μ_1^{ac} and μ_1^{s} actually define transverse measures for λ. Since ν_j is ergodic, we conclude that either $\nu_j = \mu_1^{\mathrm{ac}}$ or $\nu_j = \mu_1^{\mathrm{s}}$, namely that ν_j is either absolutely continuous or singular with respect to μ_1.

If ν_j is abolutely continuous with respect to μ_1 consider, for each transverse arc k, the Radon-Nykodym derivative $f_k : k \to [0,\infty[$ defined by the property that $\nu_j = \int_A f_k\,d\mu_1$ for every measurable $A \subset k$. This f_k is uniquely determined μ_1–almost everywhere. It follows that every homotopy respecting λ which sends k to another transverse arc k' sends f_k to $f_{k'}$ μ_1–almost everywhere. If we now use the fact that μ_1 is ergodic, we conclude that f_k is constant μ_1–almost everywhere; otherwise, if we write μ_1 as the sum of its restriction to $\{x; f_k(x) > c\}$ and of its restriction to $\{x; f_k(x) \leqslant c\}$, this decomposition would be non-trivial for some constant c, contradicting the ergodicity of μ_1. But the fact that f_k is constant means that ν_j is a real multiple of μ_1, contradicting the definition of these transverse measures. Therefore ν_j cannot be absolutely continuous with respect to μ_1.

This proves that ν_j is singular with respect to μ_1. Noting that the relation "is singular with respect to" is symmetric, we conclude that μ_1 is singular with respect to each ν_j, and therefore with respect to $\mu = \sum_{j=1}^{n} a_j \nu_j$. But this contradicts the hypothesis that $\mu = \sum_{i=1}^{n} a_i \mu_i$ with $a_1 > 0$.

This concludes the proof of Step 1.

Let $\mu_1, \ldots, \mu_n$ be a list of ergodic transverse measures such that every ergodic transverse measure is a real multiple of exactly one μ_i.

Step 2. Every transverse measure μ can be written as a linear combination $\mu = \sum_{i=1}^{n} a_i \mu_i$ with $a_i \geqslant 0$.

Because $\mathcal{H}(\lambda; \mathbb{R})$ is finite dimensional, this is essentially a matter of linear algebra. By Proposition 6.1, $\mathcal{M}(\lambda)$ is a closed convex cone in $\mathcal{H}(\lambda; \mathbb{R})$. Define the ***contact order*** of μ in $\mathcal{M}(\lambda)$ as the largest dimension p of a linear subspace L of $\mathcal{H}(\lambda; \mathbb{R})$ (passing through the origin) such that μ is in the interior of $L \cap \mathcal{M}(\lambda)$ considered as a subspace of L. We will prove the property by induction on this contact order p. Note that $p = 0$ only for $\mu = 0$.

If $p = 1$, it immediately follows from definitions that μ is ergodic, and the property is trivial.

Assume that μ has contact order $p \geqslant 2$, and that the property is proved for every transverse measure of contact order less than p. Let L be a linear subspace of dimension p of $\mathcal{H}(\lambda; \mathbb{R})$ such that μ is in the interior of $L \cap \mathcal{M}(\lambda)$.

We claim that there is an affine line L' (avoiding the origin) in L such that $L' \cap \mathcal{M}(\lambda)$ is compact. The easier way to see this is to choose a fattened train track Φ carrying λ and to consider the linear form $u : \mathcal{H}(\lambda; \mathbb{R}) \to \mathbb{R}$ defined by $u(\alpha) = \sum_e \alpha(e)$, where the sum is over all edges e of Φ. The affine hyperplane $H = u^{-1}(u(\mu))$ has the property that $H \cap \mathcal{M}(\lambda)$ is compact, since $0 \leqslant \alpha(e) \leqslant u(\mu)$ for every $\alpha \in H \cap \mathcal{M}(\lambda)$ and every edge e. Also, $H \cap L$ has dimension $p - 1 \geqslant 1$ by transversality. Any affine line L' contained in $H \cap L$ and passing through μ will then have compact intersection with $\mathcal{M}(\lambda)$.

By convexity, $L' \cap \mathcal{M}(\lambda)$ is a compact interval containing μ, with end points μ' and μ''. These two elements μ', μ'' must have contact order less than p, because the order of μ would otherwise be at least $p + 1$. By induction hypothesis, it follows that μ' and μ'' are both linear combination of the μ_i with non-negative coefficients, and therefore so is μ since it is a linear combination of μ' and μ'' with non-negative coefficients.

This concludes the proof of Step 2.

Step 3. The coefficients $a_i \geqslant 0$ of the linear decomposition $\mu = \sum_{i=1}^{n} a_i \mu_i$ are unique.

This is proved by the same arguments as Step 1, and completes the proof of Proposition 6.12. □

At this point, we only know that the dimension of $\mathcal{M}(\lambda)$ is less than or equal to the dimension of $\mathcal{H}(\lambda; \mathbb{R})$, as determined by Proposition 6.12. This estimate will be greatly improved in §??.

6.3. Transverse cocycles as cocycles

A ***local transverse orientation*** of the geodesic lamination λ at the point $x \in \lambda$ is an orientation of the normal at x of the leaf of λ passing through x. A ***transverse***

orientation for λ is a continuous choice of a local transverse orientation at each $x \in \lambda$.

A geodesic lamination does not necessarily admit a transverse orientation. However, the geodesic lamination λ always admits a ***transverse orientation*** 2***–fold covering*** $\widehat{\lambda} \to \lambda$, formally defined as the space of all pairs (x, o) where $x \in \lambda$ and where o is a local transverse orientation of λ at x, with the obvious topology. By continuity, this covering $\widehat{\lambda} \to \lambda$ extends to a 2–fold covering $\widehat{U} \to U$ where U is a small neighborhood of λ. Note that the geodesic lamination $\widehat{\lambda} \subset \widehat{U}$ is canonically endowed with a transverse orientation, defined by the choice of the transverse orientation o at $(x, o) \in \widehat{\lambda}$.

A transverse cocycle $\alpha \in \mathcal{H}(\lambda; \mathbb{M})$ defines an $\mathbb{M}$–valued cocycle on $\widehat{U}$ as follows.

Consider all differentiable simplices in $\widehat{U}$ which are transverse to the leaves of $\widehat{\lambda}$. Since every singular simplex can be approximated by such a simplex, the chain complex consisting of linear combinations of such transverse simplices can be used to comput the homology and cohomology of $\widehat{U}$, with various coefficients. Any oriented arc (= 1–simplex) $\widehat{k}$ in $\widehat{U}$ transverse to $\widehat{\lambda}$ can be split into finitely many small arcs $\widehat{k}_1, \ldots, \widehat{k}_n$ with disjoint interiors in such a way that, on each $\widehat{k}_i$, the two orientations respectively defined by the orientation of $\widehat{k}$ and by the transverse orientation of $\widehat{\lambda}$ either agree everywhere on $\widehat{k}_i \cap \widehat{\lambda}$ or disagree everywhere on $\widehat{k}_i \cap \widehat{\lambda}$. Then, define $\widehat{\alpha}(k) = \sum_{i=1}^{n} \varepsilon\big(\widehat{k}_i\big)\alpha\big(k_i\big) \in \mathbb{M}$, where $\varepsilon\big(\widehat{k}_i\big) = +1$ or -1 according to whether the orientations defined by the orientation of $\widehat{k}$ and by the transverse orientation of $\widehat{\lambda}$ agree or disagree on $\widehat{k}_i \cap \widehat{\lambda}$, and where k_i is the projection of $\widehat{k}_i$ to U. The fact that α is additive immediately shows that $\widehat{\alpha}\big(\widehat{k}\big)$ is independent of the decomposition of $\widehat{k}$ into the $\widehat{k}_i$. In this way, we define a cochain $\widehat{\alpha}$ of the chain complex of simplices transverse to $\widehat{\lambda}$, with coefficients in $\mathbb{M}$.

Using the invariance of α under homotopy respecting λ, one easily shows that the cochain $\widehat{\alpha}$ is actually a cocycle, namely that $\widehat{\alpha}$ sends to 0 the boundary of any 2–simplex transverse to $\widehat{\lambda}$.

Note that, if the boundary of U is piecewise differentiable, it makes sense to talk of the restriction of $\widehat{\alpha}$ to the boundary $\partial\widehat{U}$, and this restriction is equal to 0. Therefore, $\widehat{\alpha}$ defines a relative cohomology class $\big[\widehat{\alpha}\big] \in H^1\big(\widehat{U}, \partial\widehat{U}; \mathbb{M}\big)$.

In particular, we can consider the case where U is a fattened train track Φ which snugly carries λ. Then, by Lemma 6.9, the covering $\widehat{\lambda} \to \lambda$ uniquely extends to a covering $\widehat{\Phi} \to \Phi$, where $\widehat{\Phi}$ is a fattened train track which snugly carries $\widehat{\lambda}$.

Proposition 6.13. *Let Φ be a fattened train track which snugly carries the geodesic lamination λ. The transverse cocycle $\alpha \in \mathcal{H}(\lambda; \mathbb{M})$ is completely determined by its associated cohomology class $[\widehat{\alpha}] \in H^1\big(\widehat{\Phi}, \partial\widehat{\Phi}; \mathbb{M}\big)$. Conversely, a cohomology class $c \in H^1\big(\widehat{\Phi}, \partial\widehat{\Phi}; \mathbb{M}\big)$ corresponds to a transverse cocycle $\alpha \in \mathcal{H}(\lambda; \mathbb{M})$ if and only if $\rho^*(c) = -c$, where ρ^* is the isomorphism induced by the covering automorphism ρ that exchanges the two sheets of the covering $\widehat{\Phi} \to \Phi$.*

PROOF. Pick a generic tie k_e in each edge e of Φ, and lift it to a generic tie $\widehat{k}_e$ of $\widehat{\Phi}$. The train fattened train track $\widehat{\Phi}$ snugly carries the geodesic lamination. An argument similar to that of Lemma 6.9 then shows that the transverse orientation of $\widehat{\lambda}$ induces the same orientation for $\widehat{k}_e$ at each point of $\widehat{k}_e \cap \widehat{\lambda}$. Orient $\widehat{k}_e$ with this orientation, and consider the relative homology class $\big[\widehat{k}_e\big] \in H_1\big(\widehat{\Phi}, \partial\widehat{\Phi}; R\big)$ (recall that R is the ring for which $\mathbb{M}$ is an R–module). Then, by definition of the cocycle $\widehat{\alpha}$, $\big\langle\big[\widehat{\alpha}\big], \big[\widehat{k}_e\big]\big\rangle = \alpha\,(e)$ where $\langle\ ,\ \rangle$ denote the evaluation pairing $H^1\big(\widehat{\Phi}, \partial\widehat{\Phi}; \mathbb{M}\big) \times H_1\big(\widehat{\Phi}, \partial\widehat{\Phi}; R\big) \to \mathbb{M}$. This proves that the cohomology class $\big[\widehat{\alpha}\big] \in H^1\big(\widehat{\Phi}, \partial\widehat{\Phi}; \mathbb{M}\big)$ completely determines the edge weights $\alpha\,(e) \in \mathbb{M}$, and therefore the transverse cocycle $\alpha \in \mathcal{H}\,(\lambda; \mathbb{M})$ by Proposition 6.5.

To prove the second statement, fist note that ρ reverses the transverse orientation of $\widehat{\lambda}$, and therefore transforms the cocycle $\widehat{\alpha}$ to $-\widehat{\alpha}$. As a consequence, $\rho^*\big(\big[\widehat{\alpha}\big]\big) = -\big[\widehat{\alpha}\big]$ for every $\alpha \in \mathcal{H}\,(\lambda; \mathbb{M})$.

Conversely, let $c \in H^1\big(\widehat{\Phi}, \partial\widehat{\Phi}; \mathbb{M}\big)$ be such that $\rho^*\,(c) = -c$. Pick a generic tie $\widehat{k}_{\widehat{e}}$ in each edge $\widehat{e}$ of $\widehat{\Phi}$, and orient it by the transverse orientation of $\widehat{\lambda}$. Set $\alpha\big(\widehat{e}\big) = \big\langle c, \big[\widehat{k}_{\widehat{e}}\big]\big\rangle \in \mathbb{M}$, where $\big[\widehat{k}_{\widehat{e}}\big]$ denotes the class of $\widehat{k}_{\widehat{e}}$ in $H_1\big(\widehat{\Phi}, \partial\widehat{\Phi}; R\big)$.

If $\widehat{s}$ is a switch of $\widehat{\Phi}$, let $\widehat{e}_1, \ldots, \widehat{e}_p$ be the edges of $\widehat{\Phi}$ that are adjacent to one side of $\widehat{s}$, and let $\widehat{e}_{p+1}, \ldots, \widehat{e}_q$ be the edges adjacent to the other side. Then,

$$\sum_{i=1}^{p} \big[\widehat{k}_{\widehat{e}_i}\big] = \big[\widehat{s}\big] = \sum_{j=p+1}^{q} \big[\widehat{k}_{\widehat{e}_j}\big]$$

in $H_1\big(\widehat{\Phi}, \partial\widehat{\Phi}; R\big)$. It follows that

$$\sum_{i=1}^{p} \alpha\big(\widehat{e}_i\big) = \sum_{j=p+1}^{q} \alpha\big(\widehat{e}_j\big)$$

namely that the edge weights $\alpha\,(\widehat{e}) \in \mathbb{M}$ form an edge weight system for the fattened train track $\widehat{\Phi}$.

By choice of orientations, $\rho^*\big(\big[\widehat{k}_{\widehat{e}}\big]\big) = -\big[\widehat{k}_{\rho(\widehat{e})}\big]$. Since $\rho^*\,(c) = -c$, we conclude that $\alpha\big(\rho\big(\widehat{e}\big)\big) = \alpha\big(\widehat{e}\big)$ for every edge $\widehat{e}$ of $\widehat{\Phi}$. Namely, the edge weights $\alpha\big(\widehat{e}\big)$ descend to an edge weight system for the fattened train track Φ in S. Let $\alpha \in \mathcal{H}\,(\lambda; \mathbb{M})$ be the transverse cocycle associated to this edge weight system by Theorem 6.6, and let $\big[\widehat{\alpha}\big] \in H^1\big(\widehat{\Phi}, \partial\widehat{\Phi}; \mathbb{M}\big)$ be the cohomology class associated to this α.

By construction, $\big\langle\big[\widehat{\alpha}\big], \big[\widehat{k}_{\widehat{e}}\big]\big\rangle = \alpha\big(\widehat{e}\big) = \big\langle c, \big[\widehat{k}_{\widehat{e}}\big]\big\rangle$ for every edge $\widehat{e}$ of $\widehat{\Phi}$. If we cut $\widehat{\Phi}$ along the arcs $\widehat{k}_{\widehat{e}}$, we obtain a family of topological disks, and it follows that the homology classes $\big[\widehat{k}_{\widehat{e}}\big]$ generate $H_1\big(\widehat{\Phi}, \partial\widehat{\Phi}; R\big)$. By the universal coefficient theorem, this implies that the two cohomology classes $\big[\widehat{\alpha}\big]$ and $c \in H^1\big(\widehat{\Phi}, \partial\widehat{\Phi}; \mathbb{M}\big)$ coincide. Therefore, c is the cohomology class associated to the transverse cocycle $\alpha \in \mathcal{H}\,(\lambda; \mathbb{M})$. ◻

In particular, Proposition 6.13 establishes an isomorphism between the space $\mathcal{H}(\lambda;\mathbb{M})$ of all $\mathbb{M}$–valued transverse cocycles and the submodule of $H^1(\widehat{\Phi},\partial\widehat{\Phi};\mathbb{M})$ consisting of those homology classes c for which $\rho^*(c)=-c$.

We can rephrase this differently. Note that the cocycle $\rho^*(\widehat{\alpha})$ is equal to $-\widehat{\alpha}$, which is a stronger property that just saying that they represent the same relative cohomology class. We can therefore interpret $\widehat{\alpha}$ as a cocycle on Φ valued in the coefficient bundle $\widehat{\mathbb{M}}=(\widehat{\Phi}\times\mathbb{M})/(\mathbb{Z}/2)$ over Φ, where $\mathbb{Z}/2$ acts by ρ on $\widehat{\Phi}$ and by multiplication by -1 on $\mathbb{M}$. In this framework, $\widehat{\alpha}$ defines a cohomology class in the twisted cohomology module $H^1(\Phi,\partial\Phi;\widehat{\mathbb{M}})$; see for instance ?? for a definition of twisted cohomology. An argument almost identical to the one used in the proof of Proposition 6.13 (or an elementary proof that $H^1(\Phi,\partial\Phi;\widehat{\mathbb{M}})$ coincides with the set of those $c\in H^1(\widehat{\Phi},\partial\widehat{\Phi};\mathbb{M})$ for which $\rho^*(c)=-c$) then shows that this provides an isomorphism between $\mathcal{H}(\lambda;\mathbb{M})$ and $H^1(\Phi,\partial\Phi;\widehat{\mathbb{M}})$.

6.4. Geodesic laminations with transverse cocycles as tangent vectors to $\mathcal{ML}(S)$

In Chapter 5 we associated to a vector $\dot{\alpha}_0$ tangent to $\mathcal{ML}(S)$ at α_0 its essential support λ_{0^+}, which is a geodesic lamination containing the support λ_0 of α_0, and an $\mathbb{R}$–valued transverse cocycle for λ_{0^+}. We now consider the converse problem. Given a geodesic lamination λ containing the supportof $\alpha_0\in\mathcal{ML}(S)$ and given a transverse cocycle $\alpha\in\mathcal{H}(\lambda;\mathbb{R})$, do λ and α correspond to a vector $\dot{\alpha}_0$ tangent to $\mathcal{ML}(S)$ at α_0?

Theorem 6.14. *Let $\alpha_0\in\mathcal{ML}(S)$ and let λ be a geodesic lamination which contains the support λ_0 of α_0. Then, a transverse cocycle $\alpha\in\mathcal{H}(\lambda;\mathbb{R})$ for λ corresponds to a tangent vector of $\mathcal{ML}(S)$ at α_0 if and only if $\alpha(k)\geqslant 0$ for every arc k transverse to λ and disjoint from λ_0.*

Proof. The condition is clearly necessary. Indeed, for such an arc k and if α corresponds to the vector $\dot{\alpha}_0$ tangent to a path $t\mapsto\alpha_t\in\mathcal{ML}(S)$ starting at α_0, then $\alpha(k)$ is the right derivative of $t\mapsto\alpha_t(k)$ at $t=0$. Since $\alpha_0(k)=0$ and $\alpha_t(k)\geqslant 0$ for every t, this right derivative must be non-negative.

Conversely, assume that $\alpha(k)\geqslant 0$ for every arc k transverse to λ and disjoint from λ_0.

Let Φ be a fattened train track which strongly carries λ. In §3.4 we saw that, for every $t\geqslant 0$, there is a measured geodesic lamination α_t which is weakly carried by Φ and such that $\alpha_t(e)=\alpha_0(e)+t\alpha(e)$ for every edge e of Φ.

The path $t\mapsto\alpha_t$ in $\mathcal{ML}(S)$ is piecewise linear, by definition of the piecewise linear structure. We want to show that λ contains the essential support of the tangent vector $\dot{\alpha}_0$, and that the transverse cocycle associated to $\dot{\alpha}_0$ is equal to α.

As usual, we lift the situation to the universal covering $\widetilde{S}$, and we let tildes $\sim$ denote preimages in $\widetilde{S}$. If γ is a finite edge path in $\widetilde{\Phi}$, there is an arc k_γ contained in a tie of $\widetilde{\Phi}$ such that the leaves of $\widetilde{\lambda}$ that realize γ are exactly those which cut k_γ. Since this number is clearly independent on the choice of k_γ, we set $\alpha(\gamma) = \alpha(k_\gamma) \in \mathbb{R}$.

Note that the condition that $\alpha(k) \geqslant 0$ for every arc disjoint from λ_0 implies that $\alpha(\gamma) \geqslant 0$ for every edge path γ such that $\alpha_0(\gamma) = 0$.

Because the α_t are weakly carried by Φ, they also associate a number $\alpha_t(\gamma)$ to each finite edge path γ of $\widetilde{\Phi}$.

LEMMA 6.15. *For every finite edge path γ of $\widetilde{\Phi}$, $\alpha_t(\gamma) = \alpha_0(\gamma) + t\alpha(\gamma)$ for t sufficient small.*

PROOF. This property is much less inocuous than it might appear at first glance. Indeed, because $\alpha_t(\gamma)$ has to be non-negative, it cannot hold without the hypothesis that $\alpha(k) \geqslant 0$ for every arc k disjoint from λ_0. It actually constitutes the core of the proof of Theorem ??.

We will prove the lemma by induction on the length n of the edge path $\gamma_n = \langle e_1, \ldots, e_n \rangle$ of $\widetilde{\Phi}$.

The property holds for $n = 1$ by definition of α_t. Assume as induction hypothesis that it holds for every edge path of length at most $n-1$.

As in the proof of Lemma ??, select a 'left' side and a 'right' side for γ_n. Let $G_n^{\mathrm{l}} \subset G(\widetilde{S})$ consist of those geodesics which realize some path $\langle e_i', e_{i+1}, \ldots, e_n \rangle$ with $e_i' \neq e_i$ branching in on the left side of γ_n. Similarly define G_n^{r} consisting of geodesics branching in on the right side of γ_n. Also, if s is the switch separating e_{n-1} from e_n, let A_-^{l} (resp. A_-^{r}, A_+^{l}, A_+^{r}) consist of those geodesics of $G(\widetilde{S})$ which realize an edge path consisting of one of the edges entering s on the same side as e_{n-1} (resp. e_{n-1}, e_n, e_n) and on the left (resp. right, left, right) side of γ_n. Finally, let $H_n^{\mathrm{l}} = G_{n-1}^{\mathrm{l}} \cup A_-^{\mathrm{l}} - A_+^{\mathrm{l}}$ and $H_n^{\mathrm{r}} = G_{n-1}^{\mathrm{r}} \cup A_-^{\mathrm{r}} - A_+^{\mathrm{r}}$; for instance, H_n^{l} consists of those geodesics which cross s after branching in on the left side of γ_n and which do not branch out on the left side of e_n. Then, as in the proof of Lemma 2,

$$\begin{aligned}
\alpha_t\left(H_n^{\mathrm{l}}\right) &= \max\left\{\alpha_t\left(G_{n-1}^{\mathrm{l}}\right) + \alpha_t\left(A_-^{\mathrm{l}}\right) - \alpha_t\left(A_+^{\mathrm{l}}\right), 0\right\} \\
\alpha_t\left(H_n^{\mathrm{r}}\right) &= \max\left\{\alpha_t\left(G_{n-1}^{\mathrm{r}}\right) + \alpha_t\left(A_-^{\mathrm{r}}\right) - \alpha_t\left(A_+^{\mathrm{r}}\right), 0\right\} \\
\alpha_t\left(G_n^{\mathrm{l}}\right) &= \min\left\{\alpha_t(e_n), \alpha_t\left(H_n^{\mathrm{l}}\right)\right\} \\
\alpha_t\left(G_n^{\mathrm{r}}\right) &= \min\left\{\alpha_t(e_n), \alpha_t\left(H_n^{\mathrm{r}}\right)\right\} \\
\alpha_t(\gamma_n) &= \alpha_t(e_n) - \alpha_t\left(G_n^{\mathrm{l}}\right) - \alpha_t\left(G_n^{+}\right).
\end{aligned}$$

Note that H_n^{l} is a disjoint union of G_γ associated to finite edge paths γ. Therefore, we can talk of $\alpha\left(H_n^{\mathrm{l}}\right)$, defined as the union of the corresponding $\alpha(\gamma)$. The same applies to H_n^{r}, G_n^{l}, G_n^{r}, A_-^{l}, A_-^{r}, A_+^{l}, A_+^{r}.

Our first goal is to prove that $\alpha_t\left(H_n^{\mathrm{l}}\right) = \alpha_0\left(H_n^{\mathrm{l}}\right) + t\alpha\left(H_n^{\mathrm{l}}\right)$ for t sufficiently small. For this, note that G_{n-1}^{l}, A_-^{l} and A_+^{l} are disjoint unions of G_γ associated to finite edge paths γ of length at most $n-1$. By induction hypothesis, is follows that

$$\alpha_t\left(H_n^{\mathrm{l}}\right) = \max\left\{\alpha_0\left(G_{n-1}^{\mathrm{l}}\right) + \alpha_0\left(A_-^{\mathrm{l}}\right) - \alpha_0\left(A_+^{\mathrm{l}}\right) + t\left(\alpha\left(G_{n-1}^{\mathrm{l}}\right) + \alpha\left(A_-^{\mathrm{l}}\right) - \alpha\left(A_+^{\mathrm{l}}\right)\right), 0\right\}.$$

The intersection points of geodesics of $\widetilde{\lambda} \cap \left(G^{\mathrm{l}}_{n-1} \cup A^{\mathrm{l}}_{-}\right)$ form an interval of $s \cap \widetilde{\lambda}$ adjacent to the left end of the tie s. The same property holds for the geodesics of $\widetilde{\lambda} \cap A^{\mathrm{l}}_{+}$. It follows that, either $\widetilde{\lambda} \cap \left(G^{\mathrm{l}}_{n-1} \cup A^{\mathrm{l}}_{-}\right)$ is contained in $\widetilde{\lambda} \cap A^{\mathrm{l}}_{+}$, or $\widetilde{\lambda} \cap A^{\mathrm{l}}_{+}$ is contained in $\widetilde{\lambda} \cap \left(G^{\mathrm{l}}_{n-1} \cup A^{\mathrm{l}}_{-}\right)$. We now distinguish cases.

If $\widetilde{\lambda} \cap \left(G^{\mathrm{l}}_{n-1} \cup A^{\mathrm{l}}_{-}\right)$ is contained in $\widetilde{\lambda} \cap A^{\mathrm{l}}_{+}$, then $\alpha_0\left(H^{\mathrm{l}}_n\right) = \alpha\left(H^{\mathrm{l}}_n\right) = 0$. Also, $\alpha_0\left(A^{\mathrm{l}}_{+}\right) - \alpha_0\left(G^{\mathrm{l}}_{n-1}\right) - \alpha_0\left(A^{\mathrm{l}}_{-}\right) \geqslant 0$. In addition, $\widetilde{\lambda} \cap A^{\mathrm{l}}_{+} - \widetilde{\lambda} \cap \left(G^{\mathrm{l}}_{n-1} \cup A^{\mathrm{l}}_{-}\right)$ is the disjoint union of finitely many $\widetilde{\lambda} \cap G_\gamma$. It follows that, if $\alpha_0\left(A^{\mathrm{l}}_{+}\right) - \alpha_0\left(G^{\mathrm{l}}_{n-1}\right) - \alpha_0\left(A^{\mathrm{l}}_{-}\right) = 0$, then $\alpha\left(A^{\mathrm{l}}_{+}\right) - \alpha\left(G^{\mathrm{l}}_{n-1}\right) - \alpha\left(A^{\mathrm{l}}_{-}\right) \geqslant 0$ since $\alpha(\gamma) \geqslant 0$ for every γ with $\alpha_0(\gamma) = 0$. As a consequence,

$$\begin{aligned}\alpha_t\left(H^{\mathrm{l}}_n\right) &= \max\left\{\alpha_0\left(G^{\mathrm{l}}_{n-1}\right) + \alpha_0\left(A^{\mathrm{l}}_{-}\right) - \alpha_0\left(A^{\mathrm{l}}_{+}\right) + t\left(\alpha\left(G^{\mathrm{l}}_{n-1}\right) + \alpha\left(A^{\mathrm{l}}_{-}\right) - \alpha\left(A^{\mathrm{l}}_{+}\right)\right), 0\right\}\\ &= 0\\ &= \alpha_0\left(H^{\mathrm{l}}_n\right) + t\alpha\left(H^{\mathrm{l}}_n\right)\end{aligned}$$

for t sufficiently small.

The other case is when $\widetilde{\lambda} \cap A^{\mathrm{l}}_{+}$ is contained in $\widetilde{\lambda} \cap \left(G^{\mathrm{l}}_{n-1} \cup A^{\mathrm{l}}_{-}\right)$. Then,

$$\begin{aligned}\alpha_0\left(H^{\mathrm{l}}_n\right) &= \alpha_0\left(G^{\mathrm{l}}_{n-1}\right) + \alpha_0\left(A^{\mathrm{l}}_{-}\right) - \alpha_0\left(A^{\mathrm{l}}_{+}\right)\\ \alpha\left(H^{\mathrm{l}}_n\right) &= \alpha\left(G^{\mathrm{l}}_{n-1}\right) + \alpha\left(A^{\mathrm{l}}_{-}\right) - \alpha\left(A^{\mathrm{l}}_{+}\right).\end{aligned}$$

In addition, $\widetilde{\lambda} \cap \left(G^{\mathrm{l}}_{n-1} \cup A^{\mathrm{l}}_{-}\right) - \widetilde{\lambda} \cap A^{\mathrm{l}}_{+}$ is the disjoint union of finitely many $\widetilde{\lambda} \cap G_\gamma$. It follows that, if $\alpha_0\left(G^{\mathrm{l}}_{n-1}\right) + \alpha_0\left(A^{\mathrm{l}}_{-}\right) - \alpha_0\left(A^{\mathrm{l}}_{+}\right) = 0$, then $\alpha\left(G^{\mathrm{l}}_{n-1}\right) + \alpha\left(A^{\mathrm{l}}_{-}\right) - \alpha\left(A^{\mathrm{l}}_{+}\right) \geqslant 0$ by hypothesis on α. As a consequence,

$$\begin{aligned}\alpha_t\left(H^{\mathrm{l}}_n\right) &= \max\left\{\alpha_0\left(G^{\mathrm{l}}_{n-1}\right) + \alpha_0\left(A^{\mathrm{l}}_{-}\right) - \alpha_0\left(A^{\mathrm{l}}_{+}\right) + t\left(\alpha\left(G^{\mathrm{l}}_{n-1}\right) + \alpha\left(A^{\mathrm{l}}_{-}\right) - \alpha\left(A^{\mathrm{l}}_{+}\right)\right), 0\right\}\\ &= \alpha_0\left(G^{\mathrm{l}}_{n-1}\right) + \alpha_0\left(A^{\mathrm{l}}_{-}\right) - \alpha_0\left(A^{\mathrm{l}}_{+}\right) + t\left(\alpha\left(G^{\mathrm{l}}_{n-1}\right) + \alpha\left(A^{\mathrm{l}}_{-}\right) - \alpha\left(A^{\mathrm{l}}_{+}\right)\right)\\ &= \alpha_0\left(H^{\mathrm{l}}_n\right) + t\alpha\left(H^{\mathrm{l}}_n\right)\end{aligned}$$

for t sufficiently small.

This proves that, in all cases, $\alpha_t\left(H^{\mathrm{l}}_n\right) = \alpha_0\left(H^{\mathrm{l}}_n\right) + t\alpha\left(H^{\mathrm{l}}_n\right)$ for t sufficiently small.

Similarly, if $E_n \subset G(\widetilde{S})$ consists of those geodesics which realize the edge path $\langle e_n \rangle$, either $\widetilde{\lambda} \cap E_n$ is contained in $\widetilde{\lambda} \cap H^{\mathrm{l}}_n$, or $\widetilde{\lambda} \cap H^{\mathrm{l}}_n$ is contained in $\widetilde{\lambda} \cap E_n$. In both cases, the same kind of argument as above, using the fact that $\alpha(\gamma) \geqslant 0$ for every γ with $\alpha_0(\gamma) = 0$, shows that

$$\alpha_t\left(G^{\mathrm{l}}_n\right) = \min\left\{\alpha_0(e_n) + t\alpha(e_n), \alpha_0\left(H^{\mathrm{l}}_n\right) + t\alpha\left(H^{\mathrm{l}}_n\right)\right\} = \alpha_0\left(G^{\mathrm{l}}_n\right) + t\alpha\left(G^{\mathrm{l}}_n\right)$$

for t sufficiently small.

Replacing 'left' by 'right' everywhere also shows that $\alpha_t\left(G^{\mathrm{r}}_n\right) = \alpha_0\left(G^{\mathrm{r}}_n\right) + t\alpha\left(G^{\mathrm{r}}_n\right)$ for t sufficiently small.

Finally, for every geodesic lamination $\widetilde{\mu}$ carried by $\widetilde{\Phi}$, $\widetilde{\mu} \cap E_n$ is the disjoint union of $\widetilde{\mu} \cap G_{\gamma_n}$, $\widetilde{\mu} \cap G^{\mathrm{l}}_n$ and $\widetilde{\mu} \cap G^{\mathrm{r}}_n$. We conclude that,

$$\begin{aligned}\alpha_t(\gamma_n) &= \alpha_t(e_n) - \alpha_t\left(G^{\mathrm{l}}_n\right) - \alpha_t\left(G^{+}_n\right)\\ &= \alpha_0(e_n) + t\alpha(e_n) - \alpha_0\left(G^{\mathrm{l}}_n\right) - t\alpha\left(G^{\mathrm{l}}_n\right) - \alpha_0\left(G^{+}_n\right) - t\alpha\left(G^{+}_n\right)\\ &= \alpha_0(\gamma_n) + t\alpha(\gamma_n)\end{aligned}$$

for t sufficiently small.

This concludes the proof of Lemma ??. □

We can now determine the essential support λ_{0^+} of the α_t as t tends to 0^+. By Lemma ??, a geodesic $g \in G(\widetilde{S})$ belongs to $\widetilde{\lambda}_{0^+}$ if and only if it is weakly carried by $\widetilde{\Phi}$ and if $\alpha_0(\gamma) > 0$ or $\dot{\alpha}_0(\gamma) > 0$ for every finite edge path γ that is contained in the bi-infinite edge path associated to g. Lemma ?? implies that $\dot{\alpha}_0(\gamma) = \alpha(\gamma)$ for every finite edge path γ. Therefore, for every finite edge path γ realized by a geodesic $g \in \widetilde{\lambda}_{0^+}$, G_γ meets the support of α_0 or the support of α. In particular, since these G_γ form a basis of neighborhoods for $g \in \widetilde{\lambda}_{0^+}$ and since $\widetilde{\lambda}$ is closed, every $g \in \widetilde{\lambda}_{0^+}$ belongs to $\widetilde{\lambda}$. As a consequence, the essential support λ_{0^+} is contained in λ, and the tangent vector $\dot{\alpha}_0$ determines a transverse cocycle for λ which we also denote by $\dot{\alpha}_0$.

By Theorem ??, a transverse cocycle for λ is determined by the weights it associates to the edges of Φ. By construction of the α_t, $\dot{\alpha}_0(e) = \alpha(e)$ for every edge of Φ. Therefore, α is exactly the transverse cocycle associated to the tangent vector $\dot{\alpha}_0$ of $\mathcal{ML}(S)$ at α_0.

This concludes the proof of Theorem ??. □

The criterion provided by Theorem ?? can be rephrased in a more practical way in terms of the decomposition of the geodesic lamination λ into minimal sublaminations and infinite isolated leaves provided by Proposition 1.??.

If g is an isolated leaf of λ and if α is an $\mathbb{R}$–valued transverse cocycle for λ, there exists an arc k_g transverse to λ such that $k_g \cap \lambda = k_g \cap g$ consists of exactly one point. Define the α–***mass*** $\alpha(g)$ of g as the number $\alpha(k_g) \in \mathbb{R}$. By invariance under homotopy respecting λ, this does not depend of the choice of k_g.

If λ_1 is a minimal sublamination of λ and if every infinite isolated leaf that spirals along λ_1 has α–mass 0, we can define the ***restriction*** $\alpha_1 \in \mathcal{H}(\lambda_1;\mathbb{R})$ of the transverse cocycle $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ to λ_1 as follows: If k is an arc transverse to λ_1, we can cover $k \cap \lambda_1$ by finitely many disjoint subarcs $k_1, \dots, k_n$ of k such that each k_i is transverse to λ and meets λ only along $k_i \cup \lambda_1$ and along some points contained in infinite isolated leaves spiralling along λ_1; we then define $\alpha_1(k) = \sum_{i=1}^n \alpha(k_i)$. The hypothesis that every infinite isolated leaf spiralling along λ_1 has α–mass 0 guarantees that this is independent of the choice of the k_i, and the invariance under homotopy respecting λ_1 easily follows.

THEOREM 6.16. *Let the geodesic lamination λ contain the geodesic lamination underlying $\alpha_0 \in \mathcal{ML}(S)$. Then, the transverse cocycle $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ for λ corresponds to a tangent vector of $\mathcal{ML}(S)$ at α_0 if and only if the following three conditions hold:*

(i) every infinite isolated geodesic of λ has non-negative α–mass;

(ii) every infinite isolated leaf of λ which is asymptotic to a minimal sublamination of λ that is not contained in the support of α_0 has α–mass 0;

(iii) for each minimal sublamination λ_1 of λ that is not contained in the support of α_0, the restriction α_1 of α to λ_1 (which makes sense because of Condition (ii)*) is associated to a transverse measure for λ_1.*

PROOF. First, let us show that Conditions (i), (ii) and (iii) are necessary for α to correspond to a tangent vector of $\mathcal{ML}(S)$ at α_0.

By ??, the support of the transverse measure α_0 contains no infinite isolated leaf. As a consequence, every infinite isolated leaf of λ has α_0–mass 0, and Condition (i) immediately follows from Theorem ??.

Now, consider an infinite isolated geodesic g which is asymptotic to a minimal sublamination λ_1 of λ that is not in the support of α_0. Let k be an arc transverse to λ which meets λ_1 and no other minimal sublamination of λ. Then, g hits k in infinitely many isolated points. Choose small arcs around n points of $g \cap k$. The complement of these small arcs consists of $n+1$ arcs k' such that $\alpha_0(k') = 0$, and therefore such that $\alpha(k') \geqslant 0$ by Theorem ??. In particular, $\alpha(k)$ is the sum of $n\alpha(g) \geqslant 0$ and of a non-negative number. Since this is true for every n, we conclude that $\alpha(g) = 0$. This proves that Condition (ii) is necessary.

Finally, consider the restriction α_1 of α to a minimal sublamination λ_1 of λ which is not in the support of α_0. By Theorem ?? and by definition of this restriction, $\alpha_1(k) \geqslant 0$ for every arc k transverse to λ_1. By Proposition! ??, it follows that α_1 is associated to a transverse measure for λ_1. This proves that Condition (iii) is necessary.

Conversely, assume that Conditions (i), (ii) and (iii) are satisfied. If k is an arc transverse to λ and disjoint from the support of α_0, we can split it into finitely many arcs k_i such that, either k_i intersects λ in finitely many points located on infinite isolated leaves, or k_i meets exactly one minimal sublamination λ_i of λ and possibly some infinite isolated leaves asymptotic to λ_i. In the first case, $\alpha(k_i) \geqslant 0$ by Condition (i). In the second case, $\alpha(k_i) \geqslant 0$ by Conditions (ii) and (iii). We conclude that $\alpha(k) \geqslant 0$ for every arc k transverse to λ and disjoint from the support of α_0. By Theorem ??, this proves that α corresponds to a tangent vector of $\mathcal{ML}(S)$ at α_0. □

THEOREM 6.17. *Let the geodesic lamination λ contain the geodesic lamination underlying $\alpha_0 \in \mathcal{ML}(S)$. Then, the transverse cocycles for λ that correspond to tangent vectors of $\mathcal{ML}(S)$ at α_0 form a closed convex cone in the vector space $\mathcal{H}(\lambda;\mathbb{R})$, and this cone is bounded by finitely many hyperplanes.*

PROOF. Let Φ be a fattened train track which strongly carries λ. By cutting Φ open along some arcs carried by Φ, originating from switch points and disjoint from λ, we can arrange that the following two conditions are met:

(a) if a tie k of Φ meets a minimal sublamination λ_1, every other leaf of λ that meets k is an isolated leaf that is asymptotic to λ_1;

(b) for every infinite isolated leaf g, there is a tie k of Φ such that $k \cap \lambda = k \cap g$ consists of exactly one point.

The transverse cocycle $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ is determined by the edge weights $\alpha(e)$ it associates to the edges e of Φ, by Theorem ??. By (b), Conditions (i) and (ii) of Theorem ?? are equivalent to the property that some of these $\alpha(e)$ are non-negative or are equal to 0. If λ_1 is a minimal sublamination which is not in the support of α_0, Proposition ?? translates Condition (iii) for λ_1 to the property that the set of those $\alpha(e)$ with e meeting λ_1 belongs to the cone over a certain linear simplex. This clearly concludes the proof. □

The representation of tangent vectors to $\mathcal{ML}(S)$ by geodesic laminations with $\mathbb{R}$–valued transverse cocycles provides a nice interpretation of the faces of the piecewise linear structure of $\mathcal{ML}(S)$. By definition, two tangent vectors at the same point of a piecewise linear n–manifold belong to the same ***face*** if, when we consider the two tangent vectors of $\mathbb{R}^n$ associated to these two vectors by a local chart, the differential of every change of chart is linear on the positive cone generated by these two vectors. ■need piecewise integral linear? In the case of $\mathcal{ML}(S)$, recall that the piecewise linear structure is defined by the maps $f_k : \mathcal{ML}(S) \to \mathbb{R}^+$ where k ranges over all generic arcs of S and where $f_k(\alpha_0)$ is the total mass of the measure deposited by α_0 on k. Therefore, two tangent vectors α and β at $\alpha_0 \in \mathcal{ML}(S)$ belong to the same face if and only if, for every $a, b \geqslant 0$, there is a third tangent vector γ such that $df_k(\gamma) = a\,df_k(\alpha) + b\,df_k(\beta)$ for every generic arc k. Note that, if we interpret the tangent vector α as a geodesic lamination with transverse Hölder distribution, the image $df_k(\alpha)$ of α under the differential of f_k is just the integral $\alpha(k)$ of the constant function 1 under the Hölder distribution deposited by α on k.

PROPOSITION 6.18. *Let α and β be two tangent vectors at $\alpha_0 \in \mathcal{ML}(S)$, considered as geodesic laminations with $\mathbb{R}$–valued transverse cocycles. Then, α and β belong to the same face of the piecewise linear structure of $\mathcal{ML}(S)$ if and only if their supports λ_α and λ_β are sublaminations of a same geodesic lamination, namely if and only if no geodesic of λ_α transversely crosses a geodesic of λ_β.*

PROOF. If λ_α and λ_β are sublaminations of a geodesic lamination λ, then α and β are transverse Hölder distributions for λ. If a and b are non-negative numbers, then $\gamma = a\alpha + b\beta$ is also a transverse Hölder distribution for λ which, by Theorem 22, is associated to a tangent vector at α_0. Since $\gamma(k) = a\alpha(k) + b\beta(k)$ for every generic arc k, this proves that α and β belong to the same face.

Conversely, assume that there is a geodesic of λ_α which crosses a geodesic of λ_β. Suppose that α and β belong to the same face. Then, for $a, b > 0$, there is a tangent vector γ such that $\gamma(k) = a\alpha(k) + b\beta(k)$ for every generic arc k. The support of γ in $G(\widetilde{S})$ contains at least the symmetric difference of the supports $\widetilde{\lambda}_\alpha$ and $\widetilde{\lambda}_\beta$. In particular, by assumption on λ_α and λ_β, there are two geodesics of the support of γ which cross each other, contradicting the fact that this support is a geodesic lamination of $\widetilde{S}$. Therefore, α and β cannot belong to the same face. □

■ tangent vectors which are not contained in a face of maximal dimension.

6.4. The Thurston intersection form

We now assume that the surface S is oriented.

The Thurston intersection form is an antisymmetric bilinear form τ on the space $\mathcal{H}(\lambda;\mathbb{R})$ of $\mathbb{R}$–valued transverse cocycle for the geodesic lamination λ, and corresponds to an algebraic intersection number.

When λ is orientable, the idea is very simple: Choosing an orientation for λ, we can associate to $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ a homology class $[\alpha] \in H_1(S;\mathbb{R})$ by cutting λ into arcs which are oriented by the orientation of λ, and using α as a multiplicity (see below). Then, the Thurston intersection form τ associates to $\alpha, \beta \in \mathcal{H}(\lambda;\mathbb{R})$ the algebraic intersection number $\tau(\alpha,\beta) = [\alpha]\cdot[\beta] \in \mathbb{R}$ of the homology classes $[\alpha]$, $[\beta] \in H_1(S;\mathbb{R})$. If we consider the factors of $[\alpha]$ and $[\beta]$ corresponding to a component λ_1 of λ, reversing the orientation of λ_1 multiplies these two factors by -1, so that the intersection number is unchanged; it follows that $\tau(\alpha,\beta)$ is independent of the choice of the orientation of λ.

When λ is non-orientable, one performs the same construction locally, by using the following idea: Consider a small C^1–perturbation K' of a 1–dimensional object K on the oriented surface S. If we choose a local orientation of K near a transverse intersection point x of K with K', this also defines a local orientation of K' near x, and therefore assigns a sign to the intersection point x; reversing the local orientation of K near x also reverses the local orientation of K', so that the sign of x is independent of the choice of the local orientation for K. Intuitively, $\tau(\alpha,\beta)$ is defined by perturbing λ to a λ' transverse to λ, and counting the signs of intersection points with the multiplicities defined by α and β. We can now be more precise.

The lamination λ admits an orientation 2–fold covering $\widehat{\lambda} \to \lambda$, formally defined as the space of all pairs (x,o) where $x \in \lambda$ and where o is a local orientation of the leaf of λ passing through x, with the obvious topology. By continuity, this covering $\widehat{\lambda} \to \lambda$ extends to a 2–fold covering $\widehat{U} \to U$ where U is a small neighborhood of λ. Note that $\widehat{U}$ carries an orientation induced by the orientation of S, and that the leaves of $\widehat{\lambda}$ are canonically oriented by the choice of the orientation o at $(x,o) \in \widehat{\lambda}$.

If $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$, it lifts to a transverse cocycle $\widehat{\alpha}$ for $\widehat{\lambda}$. We now associate to the oriented lamination $\widehat{\lambda}$ and to the transverse cocycle $\widehat{\alpha}$ a homology class $[\widehat{\alpha}] \in H_1(\widehat{U};\mathbb{R})$, defined as follows. Let $\widehat{\Phi}$ be a fattened train track which snugly carries $\widehat{\lambda}$ and which is contained in $\widehat{U}$, and let $\widehat{\Theta} \subset \widehat{U}$ be a train track obtained by collapsing each edge of $\widehat{\Theta}$ to an arc. Note that the orientation of $\widehat{\lambda}$ orients each edge e of $\widehat{\Phi}$ and $\widehat{\Theta}$, by Lemma ??. If we endow each edge e of $\widehat{\Theta}$ with the weight $\widehat{\alpha}(e) \in \mathbb{R}$ associated by $\widehat{\alpha}$ to the corresponding edge of $\widehat{\Phi}$, the sum $\sum_e \widehat{\alpha}(e)\, e$ can then be interpreted as a 1–chain in $\widehat{U}$, with coefficients in $\mathbb{R}$. The fact that the weights $\widehat{\alpha}(e)$ satisfy the switch relations implies that this chain is closed, and therefore defines a homology class $[\widehat{\alpha}] \in H_1(\widehat{U};\mathbb{R})$.

Lemma 6.19. *The homology class $[\widehat{\alpha}] \in H_1(\widehat{U};\mathbb{R})$ does not depend on the choice of the fattened train track $\widehat{\Phi}$.*

Proof. Every fattened train track $\widehat{\Phi}$ that carries $\widehat{\lambda}$ contains arbitrary small fattened train tracks $\widehat{\Phi}'$ snugly carrying λ, where a fattened train track is "small" if its ties are uniformly short. If $\widehat{\Theta}$ and $\widehat{\Theta}'$ are the train tracks respectively associated to $\widehat{\Phi}$ and $\widehat{\Phi}'$, it therefore suffices to show that, for $\widehat{\Phi}'$ sufficiently small, the inclusion map $\widehat{\Phi}' \to \widehat{\Phi}$ sends the class $[\sum_{e'} \alpha(e')\, e'] \in H_1(\widehat{\Theta}';\mathbb{R}) \cong H_1(\widehat{\Phi}';\mathbb{R})$ to $[\sum_e \alpha(e)\, e] \in H_1(\widehat{\Theta};\mathbb{R}) \cong H_1(\widehat{\Phi};\mathbb{R})$, where e' and e respectively range over all edges of $\widehat{\Theta}'$ and $\widehat{\Theta}$.

Since $\widehat{\Phi}$ collapses to the graph $\widehat{\Theta}$, a homology class $c \in H_1\big(\widehat{\Phi};\mathbb{R}\big)$ is completely determined by 'the amount by which it crosses' each edge e of $\widehat{\Phi}$. Namely, pick a tie k_e in each edge e of $\widehat{\Phi}$, orient it so that it has intersection number $+1$ with an (arbitrary) leaf of $\widehat{\lambda} \cap e$ for the canonical orientation of $\widehat{\lambda}$, and consider the relative class $[k_e] \in H_1\left(\widehat{\Phi}, \partial\widehat{\Phi};\mathbb{R}\right)$. Then, the class of $c \in H_1\big(\widehat{\Phi};\mathbb{R}\big)$ is completely (over)determined by the family of algebraic intersection numbers $[k_e]\cdot c \in \mathbb{R}$, where e ranges over all edges of $\widehat{\Phi}$.

Now, the additivity of $\widehat{\alpha}$ immediately shows that, for $\widehat{\Phi}'$ sufficiently small, the algebraic intersection number of k_e with $\sum_{e'} \alpha\,(e')\,e'$, with e' ranging over all edges of $\widehat{\Theta}'$, is equal to $\alpha\,(e)$ for every edge e of $\widehat{\Phi}$ (or $\widehat{\Theta}$). It follows that, as e and e' respectively range over all edges of $\widehat{\Theta}$ and $\widehat{\Theta}'$, the cycles $\sum_e \alpha\,(e)\,e$ and $\sum_{e'} \alpha\,(e')\,e'$ represent the same class in $H_1\big(\widehat{\Phi};\mathbb{R}\big)$. This concludes the proof. □

▮Add comments on De Rham currents in the next chapter, or when defining the length of a transverse cocycle.

▮characterize which homology classes correspond to transverse cocycles, when U is a snug train track?

Given two transverse cocycles α, $\beta \in \mathcal{H}\,(\lambda;\mathbb{R})$, consider $\tau\,(\alpha,\beta) = \frac{1}{2}\big[\widehat{\alpha}\big]\cdot\big[\widehat{\beta}\big]$, namely one half of the algebraic intersection number of the two classes $\big[\widehat{\alpha}\big]$, $\big[\widehat{\beta}\big] \in H_1\big(\widehat{U};\mathbb{R}\big)$. This defines an antisymmetric bilinear form $\tau : \mathcal{H}\,(\lambda;\mathbb{R}) \times \mathcal{H}\,(\lambda;\mathbb{R}) \to \mathbb{R}$, called the ***Thurston intersection form*** of $\mathcal{H}\,(\lambda;\mathbb{R})$.

Note that we can take $\widehat{\Phi}$ as the pre-image in $\widehat{U}$ of a fattened train track $\Phi \subset U$ carrying λ, and $\widehat{\Theta}$ as the pre-image of a train track $\Theta \subset U$ associated to Φ. If, by a small C^1–deformation, we slightly perturb Θ to a train track Θ' that is transverse to Θ and if $\widehat{\Theta}' \subset \widehat{U}$ denotes the preimage of Θ', note that each intersection point $x \in \Theta \cap \Theta'$ lifts to two intersection points of $\widehat{\Theta}$ with $\widehat{\Theta}'$ which have the same sign, for the orientation of $\widehat{\Theta}'$ coming from the orientation of $\widehat{\Theta}$. Indeed, the two points of the preimage of x correspond to two distinct local orientations for Θ and, since Θ' is obtained from Θ by a small C^1–deformation, to two distinct local orientations of Θ' as well, so that the two intersections have the same sign. Therefore,

$$\tau\,(\alpha,\beta) = \sum_x \varepsilon\,(x)\,\alpha\,(e_x)\,\beta\,(e'_x)$$

where the sum is over all intersection points $x \in \Theta \cap \Theta'$, where $\varepsilon\,(x) = \pm 1$ is the sign of the intersection of Θ with Θ' at the two points of the pre-image of x, and where e_x and e'_x are the edges of Θ and Θ' containing x, respectively.

This method of computing $\tau\,(\alpha,\beta)$ may be relatively intuitive but, in practive, it is better to rely on the following formula.

PROPOSITION 6.20. *Let Φ be a fattened train track which snugly carries the geodesic lamination λ and which is generic, in the sense that each switch is adjacent to exactly 3 edges. At each switch s of Θ, there consequently are one* *incoming edge* e_s^{i} *and two outgoing edges* e_s^{r} *and* e_s^{l} *where, as seen from the incoming edge* e_s^{i} *and for the orientation of* S*,* e_s^{r} *branches out to the right and* e_s^{l} *branches out to the left; see Figure ??(a). Then, for every* α*,* $\beta \in \mathcal{H}(\lambda;\mathbb{R})$*,*

$$\tau(\alpha,\beta) = \tfrac{1}{2}\sum_s \left(\alpha\left(e_s^{\mathrm{r}}\right)\beta\left(e_s^{\mathrm{l}}\right) - \alpha\left(e_s^{\mathrm{l}}\right)\beta\left(e_s^{\mathrm{r}}\right)\right)$$

where $\alpha(e)$*,* $\beta(e) \in \mathbb{R}$ *are the weights associated by* α *and* β *to the edge* e *of* Φ*.*

Proof. Let $\Theta \subset U$ be a train track associated to Φ, and let $\widehat{\Theta}$ be its preimage in $\widehat{U}$. Since $\widehat{\Theta}$ is oriented by the canonical orientation of $\widehat{\lambda}$, it makes sense to perturb it to a train track $\widehat{\Theta}'$ that is obtained by pushing each edge of $\widehat{\Theta}$ to the left (say); see Figure ??. In particular, we can do this in such a way that there is one intersection point of $\widehat{\Theta}$ and $\widehat{\Theta}'$ near each switch of $\widehat{\Theta}$, and no other intersection point. In particular, to each switch s of Θ correspond two intersection points x_s^+ and x_s^- of $\widehat{\Theta}$ with $\widehat{\Theta}'$, of respective signs $+1$ and -1, as in Figure ??. Let us realize the homology class $[\widehat{\alpha}] \in H_1\left(\widehat{U};\mathbb{R}\right)$ by $\sum_e \widehat{\alpha}(e)\,e$, where e ranges over all edges of $\widehat{\Theta}$, and the class $[\beta]$ by $\sum_{e'} \widehat{\beta}(e')\,e'$, where e' ranges over all edges of $\widehat{\Theta}'$. The formula for $\tau(\alpha,\beta) = \frac{1}{2}\left[\widehat{\alpha}\right]\cdot\left[\widehat{\beta}\right]$ then easily follows from the above considerations on the intersection points of $\widehat{\Theta}$ and $\widehat{\Theta}'$. □

Exercise. Let Φ be a fattened train track which snugly carries λ but which is not necessarily generic. Give an expression, analogous to that of Proposition ??, for the Thurston intersection form on $\mathcal{H}(\lambda;\mathbb{R})$ in terms to the weights associated to the edges of Φ by elements of $\mathcal{H}(\lambda;\mathbb{R})$.

■show τ metric independent

Note that the intersection form τ may be degenerate. An extreme case occurs when λ consists of finitely many disjoint simple closed geodesics, in which case the form τ is actually trivial. A more general example is provided by transverse cocycles which are carried by boundary leaves, as we now define.

Given a geodesic lamination λ, consider the completion $\widehat{S-\lambda}$ of the path metric of its complement $S-\lambda$, as in §1.4. (Beware that the hat $\widehat{\ }$ here does not mean that we are in an orientation covering). We saw in Proposition ?? that the boundary $\partial\widehat{S-\lambda}$ consists of finitely many simple geodesics. In addition, those geodesics of $\partial\widehat{S-\lambda}$ which are infinite come in 'chains', where any to consecutive geodesics are asymptotic and bound the same spike of $\widehat{S-\lambda}$. Let a ***boundary cocycle*** of $\widehat{S-\lambda}$ be a map which associates a weight $a(g) \in \mathbb{R}$ to each component g of $\partial\widehat{S-\lambda}$, in such a way that any two boundary components which are asymptotic have opposite weights. Note that the only boundary components g for which $a(g)$ can be non-trivial are those which, either are closed, or belong to a chain made of an even number of infinite geodesics.

A boundary cocycle a of $\widehat{S-\lambda}$ defines a transverse cocycle $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ as follows. If k is an arc transverse to λ, we set

$$\alpha(k) = \sum_d \left(a\left(x_d^+\right) + a\left(x_d^-\right)\right)$$

where d ranges over all components of $k - \lambda$, where x_d^+ and x_d^- are the end points of d, and where, if we lift d to an arc in $\widehat{S - \lambda}$, $a\left(x_d^\pm\right)$ is the weight assigned by a to the component of $\partial\widehat{S - \lambda}$ that passes through the point of this lift corresponding to $x_d^\pm$; by convention, $a\left(x_d^\pm\right) = 0$ when there is no such component of $\partial\widehat{S - \lambda}$ passing through $x_d^\pm$, namely when $x_d^\pm$ is an end point of k. The reader should beware that the notation is somewhat ambiguous when k crosses an isolated leaf of λ, so that there exists two distinct components d, d' of $S - \lambda$ which have a common end point $x_d^+ = x_{d'}^-$; in this case, the points x_d^+ and $x_{d'}^-$ lift to different points of $\widehat{S - \lambda}$, and we may have that $a\left(x_d^+\right) \neq a\left(x_{d'}^-\right)$ although $x_d^+ = x_{d'}^-$ as points in k. A formal way to get around this difficulty would be to use pairs $\left(x_d^\pm, d\right)$ instead of points $x_d^\pm$, but we will stick to this ambiguous notation to avoid any further cluttering of the formula.

Note that, for every $\varepsilon > 0$, there are at most finitely many components d of $S - \lambda$ of length larger than ε. By Proposition 1.??, it follows that all but finitely many d lift to arcs that are contained in a spike of $\widehat{S - \lambda}$ and join the two components of $\partial\widehat{S - \lambda}$ delimiting this spike. In particular, $a\left(x_d^+\right) + a\left(x_d^-\right) = 0$ for all but finitely many d, so that the sum in the definition of $\alpha(k) \in \mathbb{R}$ is actually finite. The additivity and invariance under homotopy respecting λ are immediate, so that α is a well defined transverse cocycle for λ.

By definition, a transverse cocycle $\alpha \in \mathcal{H}(\lambda; \mathbb{R})$ is ***carried by boundary leaves*** if it is associated in this way to a boundary cocycle of $\widehat{S - \lambda}$. Recall that a boundary leaf of λ is one which is in the image of the natural map $\partial\widehat{S - \lambda} \to S$.

From a combinatorial point of view, it is convenient use a fattened train track Φ which snugly carries λ, and to characterize those transverse cocycles which are carried by boundary leaves in terms of the isomorphism between $\mathcal{H}(\lambda; \mathbb{R})$ and the space $\mathcal{W}(\Phi; \mathbb{R})$ of edge weights systems for Φ. The boundary of Φ has a finite number of corners, located on switches of Φ. Let a ***smooth section*** of $\partial\Phi$ be a smooth arc in $\partial\Phi$ going from one corner to another, or a component of $\partial\Phi$ which does not have any corner. If we consider the foliation of $\Phi - \lambda$ by its intersections with the ties of Φ, the fact that Φ snugly carries λ implies that the leaves issued from a smooth section of $\partial\Phi$ all lead to the same boundary leaf of λ, and actually to the same component of $\partial\widehat{S - \lambda}$. For the same reason, the leaves which do not originate on a smooth section of $\partial\Phi$ form spikes in $\widehat{S - \lambda}$. In particular, there is a one-to-one correspondence between components of $\partial\widehat{S - \lambda}$ and smooth sections of $\partial\Phi$. As a consequence, a boundary cocycle for $\widehat{S - \lambda}$ is equivalent to the data of a ***boundary cocycle for*** Φ, namely a map a associating a weight $a(\sigma) \in \mathbb{R}$ to each smooth section σ of $\partial\Phi$ in such a way that, at each corner, the two smooth sections meeting at this corner have opposite weights. If $\alpha \in \mathcal{H}(\lambda; \mathbb{R})$ is the transverse cocycle associated to the boundary cocycle a, it immediately follows from definitions that, for every edge e of Φ, the weight $\alpha(e)$ is equal to the sum of the weights associated by a to the two (possibly equal) smooth section of $\partial\Phi$ touching the two sides of e. This completely characterizes the edge weight systems in $\mathcal{W}(\Phi; \mathbb{R})$ that correspond to transverse cocycles which are carried by boundary leaves of λ.

■Give an example for a lamination with 2 spiralling leaves

PROPOSITION 6.21. *A transverse cocycle $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ is carried by boundary leaves if and only if it belongs to the kernel of the Thurston intersection form τ, namely if and only if $\tau(\alpha,\beta)=0$ for every $\beta \in \mathcal{H}(\lambda;\mathbb{R})$.*

PROOF. Let Φ be a train track which snugly carries λ. As usual, we extend the orientation covering $\widehat{\lambda} \to \lambda$ to a 2–fold covering $\widehat{\Phi} \to \Phi$.

If α is carried by boundary leaves, it is associated to a boundary cocycle a for $\widehat{S-\lambda}$ or, equivalently, for Φ. Each smooth section σ of $\widehat{\Phi}$ carries a natural orientation coming from the orientation of $\widehat{\Phi}$, which itself comes from the cacnoical orientation of $\widehat{\lambda}$. At each corner of $\partial\widehat{\Phi}$, the two smooth sections meeting there have opposite orientations. By definition of boundary cocycles it follows that the chain $\sum_\sigma a(\sigma)\,\sigma$, where the sum is over all smooth sections of $\partial\widehat{\Phi}$ and where $a(\sigma) \in \mathbb{R}$ denotes the weight associated by a to the projection of σ in $\partial\Phi$, is a closed chain. By construction, this chain has the same intersection numbers with the ties of $\widehat{\Phi}$ as the chain $\sum_e \widehat{\alpha}(e)\,e$ of Lemma ??, and therefore represents the homology class $[\widehat{\alpha}] \in H_1(\widehat{\Phi};\mathbb{R})$ of Lemma ??. In particular, $[\widehat{\alpha}]$ is represented by a chain contained in the boundary of $\widehat{\Phi}$, and therefore must have intersection number 0 with every class of $H_1\left(\widehat{\Phi};\mathbb{R}\right)$. As a consequence, $\tau(\alpha,\beta) = \frac{1}{2}[\widehat{\alpha}]\cdot\left[\widehat{\beta}\right] = 0$ for every $\beta \in \mathcal{H}(\lambda;\mathbb{R})$.

Conversely, suppose that α is in the kernel of τ, namely that $\tau(\alpha,\beta) = \frac{1}{2}[\widehat{\alpha}]\cdot\left[\widehat{\beta}\right] = 0$ for every $\beta \in \mathcal{H}(\lambda;\mathbb{R})$.

LEMMA 6.22. *The homology class $[\widehat{\alpha}] \in H_1(\widehat{\Phi};\mathbb{R})$ associated to α is such that $[\widehat{\alpha}]\cdot b = 0$ for every $b \in H_1(\widehat{\Phi};\mathbb{R})$.*

PROOF. Suppose, in search of a contradiction, that there exists $b \in H_1(\widehat{\Phi};\mathbb{R})$ with $[\widehat{\alpha}]\cdot b \neq 0$.

We will use the covering involution $\rho : \widehat{\Phi} \to \widehat{\Phi}$, exchanging the two sheets of the covering $\widehat{\Phi} \to \Phi$. By construction, this involution reverses the orientation of the edges of $\widehat{\Phi}$, and it follows that $\rho_*([\widehat{\alpha}]) = -[\widehat{\alpha}]$ since $[\widehat{\alpha}] = [\sum_e \widehat{\alpha}(e)\,e]$. On the other hand, if we simply consider $\widehat{\Phi}$ as a surface, ρ respects the orientation of this surface defined by the orientation of S. In particular, ρ_* respects intersection numbers in $H_1\left(\widehat{\Phi};\mathbb{R}\right)$. Replacing b by $\frac{1}{2}(b-\rho_*(b))$ if necessary, we can therefore assume that $\rho_*(b) = -b$ without loss of generality.

As in the proof of Lemma ??, pick a tie k_e in each edge e of $\widehat{\Phi}$, orient it so that it has intersection number $+1$ with an (arbitrary) leaf of $\widehat{\lambda} \cap e$ for the canonical orientation of $\widehat{\lambda}$, and consider the relative class $[k_e] \in H_1\left(\widehat{\Phi},\partial\widehat{\Phi};\mathbb{R}\right)$. Then, set $\widehat{\beta}(e) = [k_e]\cdot b$, and note that the chain $\sum_e \widehat{\beta}(e)\,e$ represents b in $H_1\left(\widehat{\Phi};\mathbb{R}\right)$. The weights $\widehat{\beta}(e)$ clearly satisfy the switch relations, and therefore define an edge weight system $\widehat{\beta} \in \mathcal{W}\left(\widehat{\Phi};\mathbb{R}\right)$. Since $\rho_*(b) = -b$, since ρ reverses the orientation of the edges of $\widehat{\Phi}$ and since ρ_* respects intersection numbers,

$\widehat{\beta}(\rho(e)) = \widehat{\beta}(e)$ for every edge e, and $\widehat{\beta}$ descends to an edge weight system $\beta \in \mathcal{W}(\Phi; \mathbb{R})$.

Let $\beta \in \mathcal{H}(\lambda; \mathbb{R})$ denote the transverse cocycle associated to the edge weight system $\beta \in \mathcal{W}(\Phi; \mathbb{R})$. Since $\sum_e \widehat{\beta}(e)\, e$ represents b in $H_1\left(\widehat{\Phi}; \mathbb{R}\right)$, the homology class $\left[\widehat{\beta}\right] \in H_1\left(\widehat{\Phi}; \mathbb{R}\right)$ associated to β is equal to b, and $\tau(\alpha, \beta) = \frac{1}{2}[\widehat{\alpha}] \cdot \left[\widehat{\beta}\right] = \frac{1}{2}[\widehat{\alpha}] \cdot b \neq 0$, contradicting our hypothesis. □

By Lemma ??, $[\widehat{\alpha}] \cdot b = 0$ for every $b \in H_1\left(\widehat{\Phi}; \mathbb{R}\right)$. Since the intersection form $H_1\left(\widehat{\Phi}, \partial\widehat{\Phi}; \mathbb{R}\right) \times H_1\left(\widehat{\Phi}; \mathbb{R}\right) \to \mathbb{R}$ is non-degenerate (see for instance ??), we conclude that the image of $[\widehat{\alpha}] \in H_1\left(\widehat{\Phi}; \mathbb{R}\right)$ in $H_1\left(\widehat{\Phi}, \partial\widehat{\Phi}; \mathbb{R}\right)$ is 0, namely that $[\widehat{\alpha}]$ is in the image of the map $H_1\left(\partial\widehat{\Phi}; \mathbb{R}\right) \to H_1\left(\widehat{\Phi}; \mathbb{R}\right)$. In other words, as σ ranges over all smooth sections of $\partial\widehat{\Phi}$, we can find coefficients in $a(\sigma) \in \mathbb{R}$ such that $[\widehat{\alpha}] = \left[\sum_\sigma a(\sigma)\, \sigma\right]$ in $H_1\left(\widehat{\Phi}; \mathbb{R}\right)$.

For each corner of $\partial\widehat{\Phi}$, the two smooth sections σ and σ' meeting at this corner have opposite orientations. Since the chain $\sum_\sigma a(\sigma)\, \sigma$ is closed, it follows that $a(\sigma) + a(\sigma') = 0$ whenever the smooth sections meet at some corner. In other words, a defines a boundary cocycle for $\widehat{\Phi}$.

Note that we also have that $[\widehat{\alpha}] = -\rho_*([\widehat{\alpha}]) = -\rho_*\left(\left[\sum_\sigma a(\sigma)\, \sigma\right]\right) = \left[\sum_\sigma a(\sigma)\, \rho(\sigma)\right] = \left[\sum_{\sigma'} a(\rho(\sigma'))\, \sigma'\right]$ because ρ sends each smooth section σ to the smooth section $\rho(\sigma)$ with the orientation reversed. Replacing the boundary cocycle a by $\frac{1}{2}(a + a \circ \rho)$ if necessary, we can therefore assume that $a(\rho(\sigma)) = a(\sigma)$ for every smooth section σ, namely that a descends to a boundary cocycle for Φ, which we also denote by a.

Let $\alpha' \in \mathcal{H}(\lambda; \mathbb{R})$ be the transverse cocycle associated to the boundary cocycle a, and let $\widehat{\alpha}'$ be its lift to a transverse cocycle for $\widehat{\lambda}$. For every is an edge e of $\widehat{\Phi}$, the weight $\widehat{\alpha}'(e)$ is equal to the sum $a(\sigma') + a(\sigma'')$, where σ' and σ'' are the two (possibly equal) smooth section of $\partial\widehat{\Phi}$ that touch e. In other word, if we choose a tie k_e of e so that it has intersection number $+1$ with σ' and σ'',

$$\widehat{\alpha}'(e) = [k_e] \cdot \left[\sum_\sigma a(\sigma)\, \sigma\right] = [k_e] \cdot [\widehat{\alpha}] = \widehat{\alpha}(e).$$

Since this holds for every e, we conclude that $\widehat{\alpha}' = \widehat{\alpha}$ as transverse cocycles for $\widehat{\lambda}$, and therefore that $\alpha' = \alpha$.

As a consequence, $\alpha \in \mathcal{H}(\lambda; \mathbb{R})$ is associated to the boundary cocycle a, and is therefore carried by boundary leaves. This concludes the proof of Proposition ??. □

Let an ***even boundary chain*** of the completion $\widehat{S - \lambda}$ of the path metric of $S - \lambda$ be a chain of components of $\partial\widehat{S - \lambda}$ which has an even number of spikes. This includes chains with no spikes, namely components of $\partial\widehat{S - \lambda}$ which are closed geodesics. We already observed that a boundary cocycle for $\widehat{S - \lambda}$ can assign a non-zero weight only to those components of $\partial \widehat{S-\lambda}$ which are contained in an even boundary chain. Also, if Φ is a fattened train track which snugly carries λ, there is a natural one-to-one correspondence between even boundary chains of $\widehat{S-\lambda}$ and components of $\partial\Phi$ which have an even number of corners.

COROLLARY 6.23. *The kernel of the Thurston intersection form* $\tau : \mathcal{H}(\lambda; \mathbb{R}) \times \mathcal{H}(\lambda; \mathbb{R}) \to \mathbb{R}$ *has dimension* $N_{\mathrm{e}} - N_{\mathrm{o}}$ *where* N_{e} *is the number of even boundary chains of* $\widehat{S-\lambda}$*, and where* N_{o} *is the number of orientable components of* λ*.*

PROOF. By linearity, we can restrict attention to the case where λ is connected. The space $\mathcal{B}$ of boundary cocycles for $S - \lambda$ clearly has dimension N_{e} and, by Proposition ??, the kernel of τ is equal to the image of the map $\mathcal{B} \to \mathcal{H}(\lambda; \mathbb{R})$. Therefore, it suffices to show that the kernel of $\mathcal{B} \to \mathcal{H}(\lambda; \mathbb{R})$ has dimension 1 if λ is orientable and is 0 otherwise.

Let a be a non-zero boundary cocycle which is in the kernel of $\mathcal{B} \to \mathcal{H}(\lambda; \mathbb{R})$. If Φ is a fattened train track which snugly carries λ, this means that a assigns opposite weights to the two smooth sections of $\partial\Phi$ that touch each edge of Φ. By connectedness, it follows that a takes only two values w, $-w$ with $w \in \mathbb{R} - \{0\}$. Each tie k of Φ then touches two smooths sections of $\partial\Phi$ of respective weights w and $-w$, and we can orient k from the smooth section weighted by $-w$ to the smooth section weighted by w. By inspection near the switches of τ, we see that this defines a continuous orientation of the ties of Φ; see Figure ??. We can then combine this orientation of the ties of Φ with the orienation of S to define a continous orientation of the leaves of λ.

Therefore, if the kernel of $\mathcal{B} \to \mathcal{H}(\lambda; \mathbb{R})$ is non-trivial, then λ is orientable and this kernel has dimension at most 1.

Conversely, if λ is orientable, we can combine orientations of λ and S to continuously orient the ties of k. Given $w \in \mathbb{R}$, we can then assign a weight $+w$ or $-w$ to each smooth section of $\partial\Phi$, according to whether the oriented ties point in the direction of this smooth section or away from it. By inspection near the switches, we see that two smooth sections meeting at a corner then have opposite weights, so that this defines a boundary cocycle a for Φ. By construction, this boundary cocycle is in the kernel of the map $\mathcal{B} \to \mathcal{H}(\lambda; \mathbb{R})$. Since we were free to chose the weight $w \in \mathbb{R}$, we conclude that the kernel of $\mathcal{B} \to \mathcal{H}(\lambda; \mathbb{R})$ has dimension 1 when λ is orientable (and connected).

This concludes the proof of Corollary ??. □

When λ is maximal, every component of $\widehat{S-\lambda}$ is a triangle, with three spikes, so that $\widehat{S-\lambda}$ has no even boundary cocycle. This immediately proves the following:

COROLLARY 6.24. *If the geodesic lamination* λ *is maximal, then the Thurston intersection form* $\tau : \mathcal{H}(\lambda; \mathbb{R}) \times \mathcal{H}(\lambda; \mathbb{R}) \to \mathbb{R}$ *is non-degenerate.* □

6.5. The dual space $\mathcal{H}^*(\lambda; \mathbb{R})$

Consider the vector space $\mathcal{H}^*(\lambda; \mathbb{R})$ defined as follows: We start from the direct sum $E = \bigoplus_k \mathbb{R}_k$ where k ranges over all arcs transverse to λ and where

$\mathbb{R}_k$ denotes a copy of $\mathbb{R}$ asociated to each such k. In other words, E consists of all formal finite linear combinations $\sum_{i=1}^n a_i k_i$ where, for each $i = 1, \dots, n$, k_i is an arc transverse to λ and a_i is a real number. In E, consider the linear subspace F generated by all elements of the form $k - k'$ or $k - k_1 - k_2$, where k' is homotopic to k by a homotopy respecting λ, and where k_1 and k_2 are two transverse arcs obtained by splitting k at some point of $k - \lambda$. Then, $\mathcal{H}^*(\lambda;\mathbb{R})$ is defined as the quotient space E/F.

The definition of $\mathcal{H}^*(\lambda;\mathbb{R})$ is designed so that an $\mathbb{R}$–valued transverse cocycle is exactly a linear form on this space. In other words, $\mathcal{H}(\lambda;\mathbb{R})$ is the dual of $\mathcal{H}^*(\lambda;\mathbb{R})$. Since the dimension of $\mathcal{H}(\lambda;\mathbb{R})$ is finite, we conclude:

Proposition 6.25. *The above space $\mathcal{H}^*(\lambda;\mathbb{R})$ is the dual of $\mathcal{H}(\lambda;\mathbb{R})$.* □

In particular, the notation is consistent.

Exercise. Use Proposition 6.3 to express the space $\mathcal{H}^*(\lambda;\mathbb{R})$ in terms of the space $LC_0\left(\widetilde{\lambda};\mathbb{R}\right)$ of all functions $\varphi : \widetilde{\lambda} \to \mathbb{R}$ which are locally constant and with compact support.

If $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ and $A \in \mathcal{H}^*(\lambda;\mathbb{R})$, Proposition ?? enables us to indifferently write $\alpha(A)$ or $A(\alpha)$ for the image of α and A under the evaluation pairing $\mathcal{H}(\lambda;\mathbb{R}) \times \mathcal{H}^*(\lambda;\mathbb{R}) \to \mathbb{R}$.

From a practical point of view, the space $\mathcal{H}^*(\lambda;\mathbb{R})$ is probably best understood if we consider a fattened train track Φ which snugly carries λ, and the space $\mathcal{W}(\Phi;\mathbb{R})$ of all edge weight systems for Φ. The isomorphism $\mathcal{H}(\lambda;\mathbb{R}) \to \mathcal{W}(\Phi;\mathbb{R})$ given by Theorem ?? then provides an isomorphism of dual spaces $\mathcal{W}^*(\Phi;\mathbb{R}) \to \mathcal{H}^*(\lambda;\mathbb{R})$. By definition of edge weight systems and of the switch relations, the dual space $\mathcal{W}^*(\Phi;\mathbb{R})$ of $\mathcal{W}(\Phi;\mathbb{R})$ coincides with the quotient E'/F', where the vector space E' is the direct sum $\bigoplus_e \mathbb{R}_e$, with e ranging over all edges of Φ and with each $\mathbb{R}_e$ a copy of $\mathbb{R}$, and where the linear subspace F' is generated by the linear combinations $\sum_{i=1}^p e_i^{(s)} - \sum_{j=p+1}^q e_j^{(s)}$ associated to the switches s of Φ, where $e_1^{(s)}, \dots, e_p^{(s)}$ are the edges of Φ adjacent to one side of s and where $e_{p+1}^{(s)}, \dots, e_q^{(s)}$ are the edges adjacent to the other side.

We say that an element of $\mathcal{H}^*(\lambda;\mathbb{R})$ is ***positive*** if it can be represented by a linear combination $\sum_{i=1}^n a_i k_i$ of transverse arcs k_i where all the coefficients $a_i \in \mathbb{R}$ are positive and where every leaf of λ meets at least one of the k_i. These positive elements of $\mathcal{H}^*(\lambda;\mathbb{R})$ will play an important role in geometric applications in later chapters.

Proposition 6.26. *An element $A \in \mathcal{H}^*(\lambda;\mathbb{R})$ is positive if and only if $A(\mu) > 0$ for every non-zero transverse measure μ for the geodesic lamination λ.*

Proof. It is convenient to introduce the following notation. Given a family k of arcs transverse to λ and a locally constant function $\chi : k \cap \lambda \to \mathbb{R}$, subdivide k into small arcs $k_1, \dots, k_n$ such that χ is constant equal to $a_i \in \mathbb{R}$ on k_i. Then, define $[\chi k]$ as the element $\sum_{i=1}^n a_i k_i \in \mathcal{H}^*(\lambda;\mathbb{R})$. By definition of $\mathcal{H}^*(\lambda;\mathbb{R})$, this does not depend on the choice of the k_i. Also, if k is homotopic to k' respecting

λ and if $\chi' : k' \to \mathbb{R}$ is associated to χ by this homotopy, then $[\chi' k'] = \chi k$ in $\mathcal{H}^*(\lambda;\mathbb{R})$.

If A is positive, realize it by $\sum_{i=1}^n a_i k_i$ where all the coefficients $a_i \in \mathbb{R}$ are positive and where every leaf of λ meets at least one of the arcs k_i. If μ is a non-zero transverse measure, then $\mu(k_i) \geqslant 0$ for every i, and $\mu(k_i) > 0$ for at least one k_i meeting the support of μ. Therefore, $A(\mu) = \mu(A) = \sum_{i=1}^n a_i \mu(k_i)$ is positive.

Conversely, assume that $A(\mu) > 0$ for every transverse measure μ. Represent A by a linear combination $\sum_{i=1}^n a_i k_i$ of transverse arcs k_i.

We claim that there is a number n_0 such that, for every arc I which is contained in a leaf of λ and crosses at least n_0 times the arcs k_i, the sum $\sum_{i=1}^n a_i \#(I \cap k_i)$ is positive, where $\#X$ denotes the cardinal of X. Otherwise, we would be able to construct a sequence of such arcs I_p, cutting n_p times the k_i, such that n_p tends to infinity and $\sum_{i=1}^n a_i \#(I_p \cap k_i) \leqslant 0$. As in the proof of Proposition 1.??, passing to a subsequence if necessary, there is a non-trivial transverse measure μ such that $\mu(k) = \lim_{p\to\infty} \#(I_p \cap k)/n_p$ for every arc k transverse to λ. But then, $A(\mu) = \lim_{p\to\infty} \sum_{i=1}^n a_i \#(I_p \cap k_i)/n_p \leqslant 0$, contradicting the hypothesis on A.

Let k be the union of the k_i, and let $\chi : k \cap \lambda \to \mathbb{R}$ be the function which is equal to a_i on each $k_i \cap \lambda$. Then, with the notation introduced at the beginning of this proof, $A = [\chi k]$ in $\mathcal{H}^*(\lambda;\mathbb{R})$.

Every infinite half-geodesic which is contained in a leaf of λ must meet k infinitely many times. Otherwise, the set of limit points of this half-geodesic would provide a minimal sub-lamination λ_1 of λ which is disjoint from the k_i; Proposition 1.?? then provides a non-trivial transverse measure μ whose support λ_1 is disjoint from the k_i, and therefore such that $A(\mu) = 0$, contradicting our hypothesis on A. Choose a transverse orientation for each k_i, and therefore for k. For each $x \in k \cap \lambda$, follow the leaf of λ that passes through x in the direction of the transverse orientation of k, and let $h_1(x), \ldots, h_{n_0}(x)$ be the first n_0 intersection points of this leaf with k which we meet in this way.

By construction, the map $\chi \circ h_j$ is locally constant. Also, $[(\chi \circ h_j) k] = [\chi k] = A$ by invariance inder homotopy. Therefore, $A = \left[\frac{1}{n_0} \sum_{j=1}^{n_0} (\chi \circ h_j) k\right]$. By choice of n_0, $\sum_{j=1}^{n_0} \chi \circ h_j(x) > 0$ for every $x \in k \cap \lambda$. Therefore, we have represented A by a linear combination $\sum_{i'=1}^{n'} a'_{i'} k'_{i'}$ where the coefficients $a'_{i'}$ are positive. In addition, the union of the $k'_{i'}$ is equal to the union of the k_i, and therefore meets every leaf of A. This proves that A is a positive element of $\mathcal{H}^*(\lambda;\mathbb{R})$. □

Exercise. We can say that an element of $\mathcal{H}^*(\lambda;\mathbb{R})$ is ***weakly positive*** if it can be represented by a non-trivial linear combination $\sum_{i=1}^n a_i k_i$ of transverse arcs k_i where all the coefficients $a_i \in \mathbb{R}$ are positive. Namely, we drop the requirement that the k_i all the leaves of λ in the definition of positive elements of $\mathcal{H}^*(\lambda;\mathbb{R})$. Clearly, the two notions coincide when λ is minimal.

a) Give an example of:
   (i) a connected lamination with a weakly positive element of $\mathcal{H}^*(\lambda;\mathbb{R})$ which is not positive;

(ii) a connected non-minimal lamination where every weakly positive element of $\mathcal{H}^*(\lambda;\mathbb{R})$ is positive.

b) Let λ consist of a closed geodesic b_0 and of an infinite isolated leaf g_1 whose two ends spiral along g_0, in opposite directions, as in Figure ??. Let k be a small transverse arc such that $k\cap\lambda = k\cap g_1$ consists of a single point. Show that $-k$ defines an element of $\mathcal{H}^*(\lambda;\mathbb{R})$ which is not weakly positive, but which is in the closure of the subset consisting of all weakly positive elements of $\mathcal{H}^*(\lambda;\mathbb{R})$.

c) Show that $A\in\mathcal{H}^*(\lambda;\mathbb{R})$ is in the closure of the set of weakly positive elements of $\mathcal{H}^*(\lambda;\mathbb{R})$ if and only if $A(\mu)\geqslant 0$ for every transverse measure μ for λ.

6.6. Continuous transverse cocycles

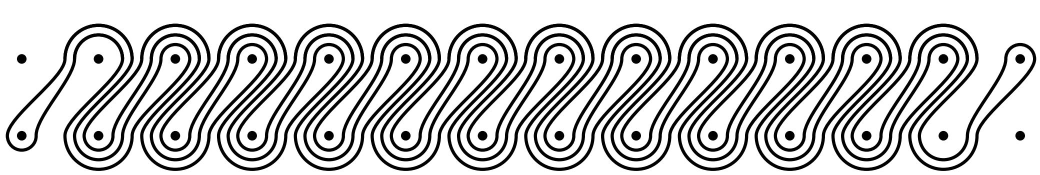

CHAPTER VII

THE NIELSEN-THURSTON CLASSIFICATION OF SURFACE HOMEOMORPHISMS

We consider diffeomorphisms $f : S \to S$ of the surface S, up to isotopy. In a series of papers [**Nie1**] around 1930, J. Nielsen developed a canonical decomposition of such an isotopy class of diffeomorphisms into building blocks which are either homotopically periodic or homotopically aperiodic in a sense defined below. This decomposition was re-discovered in the mid-seventies by W. Thurston, who also exhibited many quantitative and dynamical properties of the aperiodic building blocks. This chapter is devoted to this Nielsen-Thurston decomposition of surface diffeomorphisms.

Let us state the Nielsen decomposition theorem.

Recall that a homeomorphism $f : S \to S$ is ***periodic*** if there exists an integer $n \neq 0$ such that f^n is equal to the identity. We will say that a diffeomorphism f is ***homotopically periodic*** if there is an integer $n \neq 0$ such that f^n is *homotopic* to the identity.

The extreme opposite occurs when, for every essential simple closed curve γ on S and every $n \neq 0$, the curve $f^n(\gamma)$ is never homotopic to γ. In this case, we will say that $f : S \to S$ is ***homotopically aperiodic***.

Recall that a multicurve in S is a family of disjoint simple closed curves in S. Such a multicurve γ is ***essential*** if no component of $S-\gamma$ is a disk, an annulus or a Möbius strip. One easily sees that this is equivalent to the property that all components of γ are indivisible, that no two of them are parallel, and that none is parallel to a boundary component of S. If γ is an essential multicurve in S, we can consider the surface S_γ obtained by splitting S open along γ. Finally, the diffeomorphism $f : S \to S$ is ***reducible*** if there exists a non-empty family γ of pairwise disjoint essential simple closed curves, none of which is parallel to a boundary component of S, such that $f(\gamma)$ is homotopic to γ. In this case, removing components of γ if necessary, we can arrange that no component of $S-\gamma$ is an annulus. We can also homotope γ so that $f(\gamma)$ is actually equal to γ. Then, f induces a diffeomorphism f_γ of the surface S_γ obtained by splitting S open along γ. Formally, S_γ is the metric completion of $S-\gamma$ for the path metric induced by an arbitrary riemannian metric on S; it is also the union of $S-\gamma$ and of the unit normal bundle of γ in S, with the obvious topology. The compact surface S_γ may be disconnected, but it is topologically 'simpler' than S in the sense that it has the same Euler characteristic as S but has more boundary components.

The Nielsen classification theorem states:

THEOREM 7.1. *If $f : S \to S$ is a diffeomorphism of the surface S, then f can be isotoped so that it respects a (possibly empty) essential multicurve γ in such a way that, if $f_\gamma : S_\gamma \to S_\gamma$ denotes the diffeomorphism induced by f, the restriction of f_γ to each component of S_γ is either homotopically periodic or homotopically aperiodic.*

This result was later improved by Nielsen, who proved in [**Nie2**] that every homotopically periodic diffeomorphism is isotopic to a periodic diffeomorphism. We will prove this second result in Chapter ??, using hyperbolic geometry. This enables us to strengthen the conclusion of Theorem 7.1 by adding that the restriction of f_γ to each component S_1 of S_γ is, either homotopically aperiodic, or periodic outside of a small collar neighborhood of the boundary ∂S_1.

The conclusion of Theorem 7.1 for the homotopically aperiodic parts was greatly improved by Thurston. He showed that every homotopically aperiodic isotopy class of surface homeomorphisms contains a representative of a certain type, called pseudo-Anosov. This pseudo-Anosov representative is unique, in the sense that any two pseudo-Anosov homeomorphisms which are in the same isotopy class are conjugate by a homeomorphism which is isotopic to the identity. In addition, the dynamics of pseudo-Anosov homeomorphisms are simpler than that of any other homeomorphism in the same isotopy class.

We will prove Theorem 7.1 in §§7.1-??. The proof we give builds on ideas originally introduced by A. Casson [**Ca1**][**CaB**], and uses only basic facts on geodesic laminations from Chapter 1 and the combinatorial aspects of transverse cocycles developed in §6.1-2 and 6.6. We will then study further properties of pseudo-Anosov homeomorphism in §7.??.

7.1. Finding an invariant geodesic lamination

Let $f : S \to S$ be a diffeomorphism of the surface f. Considering geodesic laminations from the metric independent view-point developed in §1.??, we see that f induces a homeomorphism $f_* : \mathcal{L}(S) \to \mathcal{L}(S)$ on the space $\mathcal{L}(S)$ of geodesic laminations on S.

Since we aim at proving Theorem 7.1, we can assume that f is not homotopically periodic. Otherwise, we can take $\gamma = \emptyset$, and the result is proved.

PROPOSITION 7.2. *If the diffeomorphism $f : S \to S$ is homotopically nonperiodic, then there exists a geodesic lamination $\lambda \in \mathcal{L}(S)$ which is disjoint from the boundary ∂S and such that $f_*(\lambda)$.*

PROOF. We borrow the proof from [**CaB**].

LEMMA. *If $f : S \to S$ is not homotopically periodic, there exists a simple closed geodesic $\gamma \in \mathcal{L}(S)$ whose images $f_*^n(\gamma)$, $n \in \mathbb{Z}$, are all distinct in $\mathcal{L}(S)$.*

PROOF OF LEMMA ??. Note that the condition that f is not homotopically periodic is also necessary for the existence of such a simple closed geodesic γ.

To prove Lemma ?? suppose that, for every simple closed geodesic γ, there is an $n > 0$ such that $f_*^n(\gamma) = \gamma$. We want to conclude that f is homotopically periodic.

Take two simple closed geodesics γ_1 and γ_2 which 'fill' the surface S, namely such that each component of $S - \gamma_1 \cup \gamma_2$ is, either a disk, or an annulus containing a component of ∂S. ■Existence? Then, by hypothesis, there is an integer $n > 0$ such that $f_*^n(\gamma_1) = \gamma_1$ and $f_*^n(\gamma_2) = \gamma_2$. This means that the simple closed curves $f^n(\gamma_1)$ and $f^n(\gamma_2)$ are respectively isotopic to γ_1 and γ_2. In particular, there is a diffeomorphism $g : S \to S$ which is isotopic to f^n and such that $g(\gamma_1) = \gamma_1$. Now, the two simple closed curves and γ_2 and $g(\gamma_2)$ are both in taut position with respect to $\gamma_1 = g(\gamma_1)$, and are isotopic. Lemma ?? in the Appendix show that they are isotopic by an isotopy respecting γ_1. As a consequence, we can choose g so that $g(\gamma_1) = \gamma_1$ and $g(\gamma_2) = \gamma_2$.

Now, g respects the finite graph $\gamma_1 \cup \gamma_2$. As a consequence, there is a power $p > 0$ such that g^p respects each vertex and each oriented edge of this graph. We can therefore isotop g^p to a diffeomorphism $h : S \to S$ which fixes $\gamma_1 \cup \gamma_2$ pointwise. Since every component of $S - \gamma_1 \cup \gamma_2$ is a disk or an annulus containing a component of ∂S, we can additionally deform h to the identity by an isotopy fixing $\gamma_1 \cup \gamma_2$. In particular, f^{np} is isotopic to the identity, contadicting our hypothesis. □

Let the simple closed geodesic γ be provided by Lemma ??. Since $\mathcal{L}(S)$ is compact (Theorem 1.??), there is a sequence $n_i \in \mathbb{Z}$, tending to ∞ as i tends to ∞, such that $f_*^{n_i}(\gamma)$ converges to some geodesic lamination λ in $\mathcal{L}(S)$.

Note that γ cannot be a component of ∂S, since f permutes these boundary components. By Lemma 2.??, it follows that all $f_*^{n_i}(\gamma)$ stay away from a neighborhood of ∂S. As a consequence, λ is disjoint from ∂S.

At least one leaf of λ is non-compact. Indeed, by Proposition 1.??, λ would otherwise be the union of finitely many disjoint geodesics. For i large enough, $f_*^{n_i}(\gamma)$ would be contained in a small tubular neighborhood U of λ. Since the components of λ are the only geodesics that are completely contained in U, we conclude that λ is a single simple closed geodesic and that $f_*^{n_i}(\gamma) = \lambda$ for large enough. But this contradicts the fact that the $f_*^n(\gamma)$ are pairwise distinct.

Let g be such a non-compact leaf of λ. Let λ_1 be the set of cluster points of g, namely the set of those $x \in S$ which are limits in S of sequences $x_n \in g$ where the x_n tend to the end of g in g. Clearly, λ_1 is a sublamination of λ.

Lemma. *For every $p > 0$, no leaf of λ_1 has a non-trivial transverse intersection with the geodesic lamination $f_*^p(\lambda_1)$.*

Proof. Otherwise, there is an arc k contained in a leaf of λ_1 which has a non-trivial transverse intersection with $f_*^p(\lambda_1)$. As a consequence, the leaf $f_*^p(g)$ of $f_*^p(\lambda)$ meets k infinitely many often. In particular, if we fix an integer $N > 0$, there is an arc k' contained in $f_*^p(g)$ which transversely meets k at least N times.

Since $f_*^{n_i}(\gamma)$ converges to λ for the Hausdorff topology, there is an arc contained in $f_*^{n_i}(\gamma)$ which is arbitrarily close to k for i large. Similarly, since $f_*^{n_i+p}(\gamma)$ converges to $f_*^p(\lambda)$, there is an arc contained in $f_*^{n_i+p}(\gamma)$ which is arbitrarily close to k' for i large enough. Since the arcs k and k' meet in at least N points, we conclude that, for i large enough, the geodesics $f_*^{n_i}(\gamma)$ and $f_*^{n_i+p}(\gamma)$ also meet in at least N points.

However, the number of intersection points of the geodesics $f_*^{n_i}(\gamma)$ and

$f_*^{n_i+p}(\gamma)$ is the same as the number of intersection points of γ and $f_*^p(\gamma)$ ∎ obvious?, and is therefore constant. This leads to a contradiction if we choose the integer N larger than this constant. □

By Lemma ??, no leaf of $f_*^p(\lambda_1)$ transversely meets a leaf of $f_*^q(\lambda_1)$. If follows that $\lambda_p = \bigcup_{i=0}^{p-1} f_*^i(\lambda_1)$ is a geodesic lamination, for every $p \geqslant 1$. By Proposition 1.??, the number of minimal sublaminations and infinite isolated leaves of λ_p is uniformly bounded. Since $\lambda_{p+1} = \lambda_p \cup f_*(\lambda_p)$, we conclude that $f_*(\lambda_p) = \lambda_p$ for p large enough.

This concludes the proof of Proposition 7.??. □

We now assume in addition that f is ***irreducible***, in the sense that there exists no essential multicurve $\gamma \neq \emptyset$ such that $f(\gamma)$ is isotopic to γ. Otherwise, we can arrange that $f(\gamma) = \gamma$ by an isotopy, and f induces a diffeomorphism f_γ of the surface S_γ obtained by splitting S open along γ. Because the surface S_γ is topologically simpler than S, in a sense to be precised later on, this leads to an inductive procedure which essentially only needs to start. In the rest of this section and in the next ?? ones, we will therefore focus on the starting point of the induction, namely on the case where f is irreducible. The details of the induction will be developed in §7.??.

LEMMA. *If the surface diffeomorphism $f : S \to S$ is irreducible and non-periodic, there exists a geodesic lamination $\lambda \in \mathcal{L}(S)$ such that $f_*(\lambda) = \lambda$, such that λ is minimal and disjoint from the boundary ∂S, and such that every component of $S - \lambda$ is either an open disk or a semi-open annulus containing a component of ∂S.*

PROOF. Let λ be the geodesic lamination with $f_*(\lambda) = \lambda$ provided by Lemma 7.2. Since f_* permutes the finitely many isolated leaves of λ, we can remove those from λ, and arrange that every connected component of λ is minimal (see Proposition 1.??).

Because of our hypothesis that f is irreducible, no leaf of λ is closed. Indeed, f_* would otherwise respect the union γ of these compact leaves. Since λ can have only finitely many closed components by Proposition 1.??, γ would then provide an essential multicurve which is invariant under f up to isotopy, contradicting the irreducibility hypothesis.

In particular, no leaf of λ is isolated. We now prove that every component of $S - \lambda$ is either an open disk or a semi-open annulus containing a boundary component.

Consider the complement $S - \lambda$ and its completion $\widehat{S - \lambda}$, as in §1.??. By Proposition 1.??, $\widehat{S - \lambda}$ is the union of a compact surface and of finitely many spikes properly diffeomorphic to $[0,1] \times [0,\infty[$, where the intersection of a spike with the boundary $\partial\widehat{S - \lambda}$ corresponds to $\{0,1\} \times [0,\infty[$. If we remove from $\widehat{S - \lambda}$ the topological interior of the spikes (corresponding to $[0,1] \times]0,\infty[$ on each spike) and suitably round the corners, we get a compact surface whose boundary projects to a multicurve γ in S. Let γ' consist of those components of γ which are not homotopic to 0.

Let us show that $f(\gamma')$ is homotopic to γ'. This is not completely obvious because we only know that $f_*(\lambda) = \lambda$ in $\mathcal{L}(S)$. We will encode the components of γ' in terms of $\widetilde{\lambda}$ where, as usual, $\widetilde{\lambda}$ denotes both the preimage of λ in the universal covering $\widetilde{S}$ and the subset of $G(\widetilde{S})$ consisting of the leaves of this geodesic lamination. A component of γ' can be represented by an element of $\pi_1(S)$ which respects an infinite chain of boundary leaves of $\widetilde{\lambda}$, where any two consecutive geodesics have an end point in common on $\partial_\infty \widetilde{S}$. Conversely, because no leaf of λ is isolated, any such infinite chain of boundary leaves must correspond to a component of γ which is not homotopic to 0. Now, a leaf of $\widetilde{\lambda}$ is a boundary leaf if and only if it is isolated on one side, a property which is easily encoded in terms of the topology and the dihedral order of the boundary at infinity $\partial_\infty \widetilde{S}$. It follows that the map $f_* : G(\widetilde{S}) \to G(\widetilde{S})$ induced by f permutes the boundary leaves of $\widetilde{\lambda}$, and therefore permutes the infinite chains of boundary leaves as above. This shows that the isomorphism of $\pi_1(S)$ induced by f permutes those conjugacy classes which are represented by elements of γ'. As a consequence, $f(\gamma')$ is homotopic to γ'.

Let γ'' be obtained by removing from γ' those components which are parallel to components of ∂S. By construction, the multicurve is essential and respected by f up to homotopy. Since f is irreducible, γ'' must be empty. In other words, every component of γ is either homotopic to 0 or parallel to a component of ∂S. It easily follows that every component of $S - \lambda$ is, either an open disk, or a semi-open annulus containg a component of ∂S.

From this property we conclude that λ is connected. Since we had arranged that every component of λ is minimal, this proves that λ is minimal. □

■ say that λ fills the surface and define

For future reference, the proof immediately gives:

Complement. *In the conclusion of Lemma 7.??, we can choose the geodesic lamination λ so that, in addition, it is contained in a Hausdorff limit $\lim_{i\to\infty} f_*^{n_i}(\gamma)$, where γ is a simple closed geodesic and where n_i tends to ∞ as i tends to ∞.* □

7.2. Finding invariant projective measured laminations

As in §7.1, let $f : S \to S$ be a diffeomorphism which is irreducible and homotopically non-periodic. Lemmas 7.?? provides a geodesic lamination $\lambda \in \mathcal{L}(S)$ which is invariant under the homeomoprhism $f_* : \mathcal{L}(S) \to \mathcal{L}(S)$ induced by f, which is minimal, and whose complement $S - \lambda$ consists of open disks and semi-open annuli containing boundary components of S. In addition, we require as in Complement 7.?? that λ is contained in a Hausdorff limit $\lim_{k\to\infty} f_*^{n_k}(\gamma)$, where γ is a simple closed geodesic and where n_k tends to ∞ as k tends to ∞.

Lift f to a homeomorphism $\widetilde{f} : \widetilde{S} \to \widetilde{S}$, and consider the homeomorphism $\widetilde{f}_* : G(\widetilde{S}) \to G(\widetilde{S})$ induced by $\widetilde{f}$ on the space of geodesics $G(\widetilde{S})$. As usual, let $\widetilde{\lambda}$ denote both the preimage of λ in the universal covering $\widetilde{S}$ and the subset of $G(\widetilde{S})$ consisting of the leaves of this geodesic lamination. The fact that $f_*(\lambda) = \lambda$ means that $\widetilde{f}_*$ respects $\widetilde{\lambda} \subset G(\widetilde{S})$. As a consequence, $\widetilde{f}_*$ acts on the space of $\pi_1(S)$–invariant measures on $G(\widetilde{S})$ with support in $\widetilde{\lambda}$, namely on the space $\mathcal{M}(\lambda)$ of transverse measures for λ. To be more precise, if $\alpha \in \mathcal{M}(\lambda)$, the measure $\widetilde{f}_*(\alpha)$ is defined by the property that $\widetilde{f}_*(\alpha)(\varphi) = \alpha(\varphi \circ \widetilde{f}_*)$ for every continuous function $\varphi : G(\widetilde{S}) \to \mathbb{R}$ with compact support; similarly, $\widetilde{f}_*(\alpha)(X) = \alpha(\widetilde{f}_*^{-1}(X))$ for every relatively compact subset $X \subset G(\widetilde{S})$. By $\pi_1(S)$–invariance, the transverse measure $\widetilde{f}_*(\alpha)$ in independent of the choice of the lift $\widetilde{f}$, and will therefore be denoted by $f_*(\alpha)$.

The map $f_* : \mathcal{M}(\lambda) \to \mathcal{M}(\lambda)$ so defined is clearly linear. By Proposition 6.??, $\mathcal{M}(\lambda)$ is the cone over a certain simplex. By the Brouwer Fixed Point Theorem or by an elementary linear algebra argument, it follows that there exists a transverse measure $\mu \in \mathcal{M}(\lambda)$ and a number $h > 0$ such that $f_*(\mu) = h\mu$.

We will show that $h > 1$ because of the condition of Complement 7.??. As a first step, we first prove that $h \geqslant 1$.

Lemma. *The number h is greater than or equal to 1.*

Proof. For this, suppose that $h < 1$ and look for a contradiction.

By Proposition 1.??, λ has only finitely many boundary leaves. Therefore, $\widetilde{\lambda}$ has only finitely many boundary leaves modulo the action of $\pi_1(S)$. It follows that there is a lift $\widetilde{f}' : \widetilde{S} \to \widetilde{S}$ of some positive power f^p of f, with $p > 0$, such that the homeomorphism $\widetilde{f}'_* : G(\widetilde{S}) \to G(\widetilde{S})$ induced by $\widetilde{f}'$ fixes a boundary leaf g of $\widetilde{\lambda}$.

We use the simple closed geodesic $\gamma \in \mathcal{L}(S)$ and the sequence n_i such that, as i tends to ∞, $f_*^{n_k}(\gamma)$ tends to a geodesic lamination λ' containing λ in $\mathcal{L}(S)$. Let $\widetilde{\gamma}$ denote the pre-image of γ in $\widetilde{S}$, as well as the subset of $G(\widetilde{S})$ consisting of the leaves of $\widetilde{\gamma}$, as usual. Also, passing to a subsequence if necessary, we can assume that all the n_i are congruent modulo p to a fixed integer q between 0 and $p-1$. Replacing γ by $f_*^q(\gamma)$, we can then arrange that $q = 0$ namely that each $m_i = n_i/p$ is an integer.

Because the complement $S - \lambda$ consists of open disks and semi-open annuli and because γ is not a component of ∂S, the closed geodesic γ cannot be disjoint from λ. Since λ has no closed leaf, it follows that γ transversely cuts λ. Because λ is minimal, we conclude that γ transversely meets the leaf of λ that is the projection of the boundary leaf g of $\widetilde{\lambda}$. In particular, g transversely meets a leaf l of $\widetilde{\gamma}$.

In $\widetilde{\lambda}$, the boundary leaf g is isolated on one side but not on the other side since λ has no isolated leaf. Therefore, there are leaves $g' \neq g$ of $\widetilde{\lambda}$ which are arbitrarily close to g in $G(\widetilde{S})$. Pick such a g', and let $X \subset \widetilde{\lambda} \subset G(\widetilde{S})$ consist of those leaves of $\widetilde{\lambda}$ which separate g from g'. The set X is open in $\widetilde{\lambda}$, and non-empty since g is not isolated in $\widetilde{\lambda}$ on the side facing g'. Because λ is minimal, it is equal to the support of μ, and it follows that $\mu(X) > 0$. As a consequence, our hypothesis that $h < 1$ implies that $\mu(\widetilde{f}'_*(X)) = f_*^{-p}(\mu)(X) = h^{-p}\mu(X) > \mu(X)$ since $\widetilde{f}'$ lifts f^p.

If g' is close enough to g, the three geodesics $g = \widetilde{f}'(g)$, g' and $\widetilde{f}'(g')$ are close

Figure 7.1

to each other and do not transversely meet in $\widetilde{S}$ since they are all leaves of $\widetilde{\lambda}$. We conclude that, either g' separates g from $\widetilde{f}'(g')$, or $\widetilde{f}'(g')$ separates g from g', or $g' = \widetilde{f}'(g')$. In the last two cases, $\widetilde{f}'_*(X)$ is contained in X, contradicting our conclusion that $\mu\big(\widetilde{f}'_*(X)\big) > \mu(X)$. Therefore, g' separates g from $\widetilde{f}'(g')$. As a consequence, g' separates g from any $\big(\widetilde{f}'\big)^m_*(g')$ with $m > 0$.

If g' is sufficiently close to g, it will also cross the leaf l of $\widetilde{\gamma}$. From the above, we conclude that the leaf $\big(\widetilde{f}'\big)^{m_i}_*(l)$ of $\big(\widetilde{f}'\big)^{m_i}_*(\widetilde{\gamma})$ also crosses g and g'. By construction, $\big(\widetilde{f}'\big)^{m_i}$ lifts $(f^p)^{m_i} = f^{n_i}$, so that $\big(\widetilde{f}'\big)^{m_i}_*(\widetilde{\gamma})$ is also the pre-image of $f^{n_i}_*(\gamma)$.

The end points of g and g' are disjoint. Indeed, suppose that g and g' have an end point $x \in \partial_\infty \widetilde{S}$ in common. Since g is not isolated on the side facing g', there would also be infinitely many leaves of $\widetilde{\lambda}$ with x as an end point, and actually infinitely many such leaves which are boundary leaves since these are dense. However, we know that $\widetilde{\lambda}$ has only finitely many boundary leaves modulo the action of $\pi_1(S)$. We conclude that there are two leaves of $\widetilde{\lambda}$ which have x as an end point and which, when both oriented towards x, are images of one another by an element $\delta \in \pi_1(S)$. In particular, δ fixes the point x, and its axis d ■define? has x as an end point. Let g_1 be a half-geodesic in g converging to x, and let d_1 be a half-geodesic in d converging also to x. These two half-geodesics are asymptotic, and the projection of g_1 to S therefore spirals along the closed geodesic which is the projection of d_1. But this would imply that this closed geodesic is contained in λ, a contradiction. Therefore, the end points of g and g' are disjoint.

Pick a third leaf g'' of $\widetilde{\lambda}$ which separates g from g'. By the same argument as above, the end points of g'' are disjoint from those of g, and we can choose g'' close enough to g that they are also disjoint from those of g'. The set of those geodesics which cross g and g' is relatively compact in $G\big(\widetilde{S}\big)$. Passing to a subsequence if necessary (although a small argument actually shows that this is unnecessary), the leaf $\big(\widetilde{f}'\big)^{m_i}_*(l)$ of the pre-image of $f^{n_i}_*(\gamma)$ converges, as i tends to ∞, to some geodesic l_∞ which transversely crosses g''. On the other hand, by definition of the Hausdorff topology, l_∞ is a leaf of the pre-image of the geodesic lamination $\lambda' = \lim_{k\to\infty} f^{n_k}_*(\gamma)$. Since $\lambda \subset \lambda'$, the projections of g'' and l_∞ would provide two leaves of λ' which transversely meet each other, the final contradiction to our hypothesis that $h < 1$.

This concludes the proof that $h \geqslant 1$. □

Lemma. *The number h is strictly greater than* 1.

Proof. By Lemma 7.??, it suffices to how that $h \neq 1$. So, we suppose that $h = 1$, and we again look for a contradiction.

Let $\widetilde{f}'$ be as above, namely lifting some power f^p and such that the induced homeomorphism $\widetilde{f}'_*$ fixes a boundary leaf g of $\widetilde{\lambda}$. Let g' be a leaf of $\widetilde{\lambda}$ which is close to g and is not a boundary leaf and, as before, let X denote the space of leaves of $\widetilde{\lambda}$ that separate g from g'. Then, the argument we already used in the proof of Lemma ?? shows that, for g' close enough to g, either $\widetilde{f}'_*(X) \subset X$ or $X \subset \widetilde{f}'_*(X)$. The fact that g' is not a boundary leaf implies that, if $\widetilde{f}'_*(X) \neq X$, then $\widetilde{f}'_*(X) - X$ or $X - \widetilde{f}'_*(X)$ contains a non-empty open subset of $\widetilde{\lambda}$, and therefore must have non-zero μ–mass since the support of μ is exactly $\widetilde{\lambda}$. But our hypothesis that $h = 1$ implies that $\mu\big(\widetilde{f}'_*(X)\big) = \mu(X)$. Therefore, $\widetilde{f}'_*(X) = X$, from which it follows that $\widetilde{f}'_*(g') = g'$. Therefore, $\widetilde{f}'_*$ fixes every non-boundary leaf of $\widetilde{\lambda}$ that is sufficiently close to g. Since λ is minimal, non-boundary leaves are dense in $\widetilde{\lambda}$ and this proves that $\widetilde{f}'_*$ fixes a neighborhood of g in $\widetilde{\lambda}$.

Figure 7.2

Let us show that $\widetilde{f}'_*$ fixes all of $\widetilde{\lambda}$. For this, suppose that there is a leaf g' of $\widetilde{\lambda}$ with $\widetilde{f}'_*(g') \neq g'$. Among all leaves of $\widetilde{\lambda}$ that separate g from g' and are fixed by $\widetilde{f}'_*$, there is one g'' that is closest to g, by continuity of $\widetilde{f}'_*$. Since $\widetilde{f}'_*(g'') = g''$, the argument we just used for g shows that $\widetilde{f}'_*$ fixes a neighborhood of g'' in $\widetilde{\lambda}$ (note that we did not use the fact that g is a boundary leaf, and only the fact that $\widetilde{f}'_*$ respects the transverse orientation of g). Because g'' is closest to g' among all leaves fixed by $\widetilde{f}'_*$ and separating g from g', we conclude that g'' is isolated on the side facing g'. As a consequence, either there is a leaf g''' which is closest to g'' among all leaves of $\widetilde{\lambda}$ that separate g'' from g', or there is no leaf of $\widetilde{\lambda}$ between g'' and g', in which case we take $g''' = g'$. From the type of the components of $S - \lambda$, we see that every component of $\widetilde{S} - \widetilde{\lambda}$ is, either an open disk bounded by a finite chain of boundary leaves of $\widetilde{\lambda}$, each asymptotic to the next one, or an infinite strip bounded on one side by a component of the pre-image of ∂S and on the other side by an infinite chain of boundary leaves. We conclude that the leaf g''' belongs to the same chain of boundary leaves as g''. However, since $\widetilde{f}'_*$ fixes g'', it must fix any leaf of this chain, from which it follows that $\widetilde{f}'_*(g''') = g'''$, contradicting the definition of g''.

Therefore, $\widetilde{f}'_*$ fixes $\widetilde{\lambda}$. Looking at the action of $\pi_1(S)$ on the leaves of $\widetilde{\lambda}$, we conclude that the isomorphism of $\pi_1(S)$ determined by $\widetilde{f}'$ is the identity. As a consequence, $f' = f^p$ is homotopic to the identity. But this contradicts our hypothesis that f is homotopically non-periodic. □

PROPOSITION. *If $f : S \to S$ is a surface diffeomorphism which is irreducible and homotopically non-periodic, there exists two measured geodesic laminations μ_{s} and $\mu_{\mathrm{u}} \in \mathcal{ML}(S)$ and a number $h > 1$ such that $f_*(\mu_{\mathrm{s}}) = h\mu_{\mathrm{s}}$ and $f_*(\mu_{\mathrm{u}}) = h^{-1}\mu_{\mathrm{u}}$, where $f_* : \mathcal{ML}(S) \to \mathcal{ML}(S)$ is the piecewise linear isomorphism induced by f. In addition, each of the supports of μ_{s} and μ_{u} fills the surface, and they meet transversely.*

PROOF. We found a geodesic lamination λ filling the surface S, a transverse measure μ for λ, and a number $h > 1$ such that $f_*(\lambda) = \lambda$ and $f_*(\mu) = h\mu$. As a consequence, $f_*(\mu_{\mathrm{s}}) = h\mu_{\mathrm{s}}$ if $\mu_{\mathrm{s}} \in \mathcal{ML}(S)$ is the measured geodesic lamination defined by λ with the transverse emasure μ.

We can also apply the same process to f^{-1}. We then get a geodesic lamination λ', a transverse measure μ', and a number $h' > 1$ such that $f_*^{-1}(\lambda') = \lambda'$ and $f_*^{-1}(\mu') = \mu'$. If $\mu_{\mathrm{u}} \in \mathcal{ML}(S)$ denotes the measured geodesic lamination defined by λ' with the transverse measure μ', we now have $f_*^{-1}(\mu_{\mathrm{u}}) = h'\mu_{\mathrm{u}}$, namely $f_*(\mu_{\mathrm{u}}) = (h')^{-1}\mu_{\mathrm{u}}$.

Therefore, we only have two things to prove: that λ' transversely meets λ and that $h' = h$.

Because λ and λ' are minimal and fill the surface, they can only meet transversely unless they are equal. It immediately follows from Lemma 7.?? that λ and λ' cannot be equal. Indeed, Lemma 7.?? shows that, if $f_*(\mu_1) = h_1\mu_1$ for some transverse measure μ_1 for λ, then necessarily $h_1 \geqslant 1$. Since $f_*(\mu') = (h')^{-1}\mu'$ and $(h')^{-1} < 1$, we conclude that μ' canot be a transverse measure for λ, and that $\lambda' \neq \lambda$. Consequently, λ and λ' meet transversely.

To prove that $h' = h$, we use the fact that f_* respects the intersection form i, which is immediate from the definition of this form ■make sure immediate. Then,

$$i(\mu_{\mathrm{s}}, \mu_{\mathrm{u}}) = i(f_*(\mu_{\mathrm{s}}), f_*(\mu_{\mathrm{s}})) = i\left(h\mu_{\mathrm{s}}, (h')^{-1}\mu_{\mathrm{u}}\right) = h(h')^{-1} i(\mu_{\mathrm{s}}, \mu_{\mathrm{u}}).$$

Since $i(\mu_{\mathrm{s}}, \mu_{\mathrm{u}}) \neq 0$ because λ and λ' have a non-empty transverse intersection, it follows that $h' = h$. □

We will see in §7.?? that the measured geodesic laminations μ_{s} and $\mu_{\mathrm{u}} \in \mathcal{ML}(S)$ are unique up to multiplication by real numbers.

7.3. The Nielsen decomposition

We saw in Lemma 7.?? that f is homotopically periodic if and only if there exists an $n > 0$ such that $f^n(\gamma)$ is isotopic to γ for every essential simple closed curve in S. The extreme opposite occurs when $f^n(\gamma)$ is never isotopic to γ for every $n > 0$ and every essential simple closed curve γ which is not parallel to a component of ∂S. In this case, we will say that $f : S \to S$ is ***homotopically aperiodic***. Passing to the closed geodesics homotopic to the simple closed curves, we see that f is homotopically aperiodic if and only if $f_*^n(\gamma) \neq \gamma$ for every $n > 0$ and every simple closed geodesic $\gamma \in \mathcal{L}(S)$ which is not a component of ∂S.

When $f : S \to S$ is reducible, there exists an essential multicurve γ such that $f(\gamma)$ is isotopic to γ. By an isotopy, we can arrange that $f(\gamma)$ is actually *equal* to γ. Then, f induces a diffeomorphism f_γ of the surface S_γ obtained by splitting S open along γ. Formally, S_γ is the metric completion of $S - \gamma$ for the path metric induced by an arbitrary Riemannian metric on S; it is also the union of $S - \gamma$ and of the unit normal bundle of γ in S, with the obvious topology.

The Nielsen Decomposition Theorem states:

THEOREM 7.1. *If $f : S \to S$ is a diffeomorphism of the surface S, then f can be isotoped so that it respects a (possibly empty) essential multicurve γ in such a way that, if $f_\gamma : S_\gamma \to S_\gamma$ denotes the diffeomorphism induced by f, the restriction of f_γ to each component of S_γ is either homotopically periodic or homotopically aperiodic. If, in addition, we require that γ has a minimum number of components among all essential multicurves satisfying the above property, then γ is unique up to isotopy.*

PROOF OF EXISTENCE. The existence of such a decomposition is a relatively simple consequence of Proposition 7.??.

To prove it, we first consider the case where f is irreducible. In this case, γ will be empty. We want to show that, either f is homotopically periodic, or f is homotopically aperiodic.

If f is not homotopically periodic, Proposition 7.?? provides a measured geodesic lamination $\mu_{\mathrm{s}} \in \mathcal{ML}(S)$ and a number $h > 1$ such that $f_*(\mu_{\mathrm{s}}) = h\mu_{\mathrm{s}}$ and such that the support of μ_{s} fills the surface. If $\gamma' \in \mathcal{L}(S)$ is a simple closed geodesic which is not a component of ∂S, then γ' must transversely cross the support of μ_{s} since it fills the surface. Let us endow γ' with the Dirac transverse measure of weight 1, and identify γ' to the corresponding measured geodesic lamination. Then, $i(\gamma', \mu_{\mathrm{s}}) \neq 0$ since γ' crosses the support of μ_{s}. Now, for every $n > 0$,

$$i(f_*^n(\gamma'), \mu_{\mathrm{s}}) = i\left(\gamma', f_*^{-n}(\mu_{\mathrm{s}})\right) = h^{-n} i(\gamma', \mu_{\mathrm{s}}) \neq i(\gamma', \mu_{\mathrm{s}})$$

since $h > 1$. It follows that $f_*^n(\gamma') \neq \gamma'$. Since this holds for every simple closed geodesic γ' which is disjoint from the boundary, this proves that f is homotopically periodic. This concludes the proof of Theorem 7.?? under the additional hypothesis that f is irreducible.

Now, consider the general case. The number of components of an essential multicurve γ is uniformly bounded by $\frac{1}{2}(3|\chi(S)| - b)$, where b is the number of components of ∂S. Among all essential multicurves γ such that $f(\gamma)$ is iotopic to γ, we can therefore pick one which has a maximum number of components. For such a γ, isotop f so that $f(\gamma) = \gamma$, and consider the diffeomorphism $f_\gamma : S_\gamma \to S_\gamma$ induced by f on the surface S_γ obtained by splitting S along γ. By maximality of the number of components of γ, the restriction of f_γ to each component of S_γ is necessarily ireducible. The above analysis of the irreducible case then shows that the restriction of f_γ to each component of S_γ is, either homotopically periodic, or homotopically aperiodic.

This proves the existence of the decomposition of $f : S \to S$ into homotopically periodic and homotopically aperiodic pieces. □

PROOF OF UNIQUENESS. To prove the uniqueness property in Theorem ??, suppose that we have two essential multicurves γ and γ' which satisfy the condition of the theorem and which have a minimum number of components among all essential multicurves satisfying these conditions. We can isotop γ and γ' so that they consist of closed m_0–geodesics. The advantage of this is that γ and γ' are in taut position with respect to each other, and that no two components of γ and γ' can be isotopic to each other unless they coincide. Then, by Lemma A.??, we can isotop f so that it respects both γ and γ'. As in the proof of Lemma ??, we can also arrange that f is periodic on the graph $\gamma \cup \gamma'$, and even on a neighborhood of $\gamma \cup \gamma'$. Consider the diffeomorphisms $f_\gamma : S_\gamma \to S_\gamma$ and $f_{\gamma'} : S_{\gamma'} \to S_{\gamma'}$ induced by f.

LEMMA. *The diffeomorphism f_γ is homotopically periodic on any component S_1 of S_γ whose interior meets γ'.*

PROOF. It is probably worth specifying the meaning of the statement, since f_γ does not necessarily respect S_1: We are trying to prove that there is an $n \neq 0$ such that f_γ^n respects S_1 and the restriction of f_γ^n to S_1 is homotopic to the identity.

Since we arranged f to be periodic on a neighborhood of $\gamma \cup \gamma'$, let $n > 0$ be such that f^n is the identity on a neighborhood $\gamma \cup \gamma'$. In particular, f_γ respects S_1. Let γ'_1 be $S_1 \cap \gamma'$, namely the closure in S_1 of the intersection of γ' with the component of $S - \lambda$ corresponding to the interior of S_1. By contruction, f_γ^n is the identity on a neighborhood U_1 of the union of γ'_1 and of the components of ∂S_1 that meet it. We can choose U_1 regular enough that its frontier ∂U_1 is a 1–dimensional submanifold of S_1.

If some component of ∂U_1 is an essential simple closed curve which is not parallel to a boundary component of S_1, this curve shows that f_γ is not homotopically aperiodic on S_1. Therefore, the restriction of f_γ has to be homotopically periodic.

Otherwise■clear?, all components of $S - U_1$ are open disks or semi-open annuli containing a component of ∂S_1. Since f_γ is periodic on U_1, it follows that f_γ is homotopically periodic on S_1. □

LEMMA. *The multicurves γ and γ' have no transverse intersection point.*

PROOF. Suppose otherwise. Let γ_0 be the union of those components of γ which have a transverse non-empty intersection with γ', and let γ'_0 be the union of those components of γ' which have a transverse non-empty intersection with γ, and suppose that γ'_0 is non-empty. Note that f respects γ'', and therefore induces a homeomorphism $f_{\gamma-\gamma_0}$ of $S_{\gamma-\gamma_0}$.

Choose an $n \neq 0$ such that f_γ^n is homotopic to the identity on each component of S_γ where f_γ is homotopically periodic, and such that fn respects each oriented component of γ'. In particular, by Lemma 7.??, f_γ^n is homotopic to the identity on each component of S_γ that meets γ_0'. It follows that $f^n_{\gamma-\gamma_0}$ is homotopic to a homeomorphism g of $S_{\gamma-\gamma_0}$ which, on each component S_1 of $S_{\gamma-\gamma_0}$ that meets γ_0, is the identity outside of a small tubular neighborhood U of $\gamma_0 \cap S_1$. As a consequence, the restriction of $f^n_{\gamma-\gamma_0}$ to such a component S_1 is homotopic to a product of Dehn twists along the 2–sided components of $\gamma_0 \cap S_1$ ■Lemma in Appendix for Möbius strips?. By definition of γ_0 and γ_0', every component γ_1 of γ_0 transversely meets a component γ_1' of γ_0', which is contained in the interior of $S_{\gamma-\gamma_0}$; from the property that $f^n_{\gamma-\gamma_0}(\gamma_1') = \gamma_1'$, we conclude that the exponent of the Dehn twist along the component γ_1 must be 0 (see Lemma in Appendix). In other words, the restriction of $f^n_{\gamma-\gamma_0}$ to each component S_1 of $S_{\gamma-\gamma_0}$ that meets γ_0 is homotopic to the identity. On the other hand, each component of $S_{\gamma-\gamma_0}$ that does not meet γ_0 corresponds to a component of S_γ, and $f_{\gamma-\gamma_0}$ is therefore homotopically periodic or homotopically aperiodic there. It follows that $f_{\gamma-\gamma_0}$ is homotopically periodic or homotopically aperiodic on each component of $S_{\gamma-\gamma_0}$.

But this contradicts the hypothesis that γ is minimal for the conclusions of Theorem 7.??, and therefore proves Lemma 7.??. □

As a consequence, the intersection $\gamma \cap \gamma'$ consists of components of γ adn γ'.

LEMMA. *No component of γ is disjoint from γ'.*

PROOF. Choose an $n \neq 0$ such that f^n_γ is homotopic to the identity on each component of S_γ where f_γ is homotopically periodic, and such that $f^n_{\gamma'}$ is homotopic to the identity on each component of $S_{\gamma'}$ where $f_{\gamma'}$ is homotopically periodic. Set $\gamma_0 = \gamma - \gamma'$.

As in the proof of Lemma 7.??, for each component S_1 of $S_{\gamma-\gamma_0}$ that meets γ_0, the restriction $f^n_{\gamma-\gamma_0}$ is homotopic to a product of Dehn twists along the 2-sided components of $\gamma_0 \cap S_1$. Each 2-sided component $\gamma 1$ of $\gamma_0 \cap S_1$ is contained in a component S_1' of $S_{\gamma'}$, where $f_{\gamma'}$ is homotopically periodic by Lemma 7.?? applied to γ' instead of γ. Pick a simple closed geodesic γ_1'' in $S_{\gamma\cup\gamma'-\gamma_1}$ which has a non-empty transverse intersection with γ_1. We can consider γ_1'' as contained in S_1 and in S_1' and, by choice of n, $f^n_{\gamma-\gamma_0}(\gamma_1'') = f^n_{\gamma'}(\gamma_1'')$ is homotopic to γ_1''. Again this implies that the exponent of the Dehn twist along γ_1 must be 0. Since this holds for every such component γ_1, this proves that $f^n_{\gamma-\gamma_0}$ is homotopic to the indentity on each component S_1 of $S_{\gamma-\gamma_0}$ that meets γ_0. On the other hand, each component of $S_{\gamma-\gamma_0}$ that does not meet γ_0 corresponds to a component of S_γ, and $f_{\gamma-\gamma_0}$ is therefore homotopically periodic or homotopically aperiodic there. It follows that $f_{\gamma-\gamma_0}$ is homotopically periodic or homotopically aperiodic on each component of $S_{\gamma-\gamma_0}$.

Again, this contradicts the hypothesis that γ is minimal for the conclusions of Theorem 7.??, and therefore proves Lemma 7.??. □

Lemmas 7.?? and 7.?? show that γ is contained in γ'. By symmetry, γ' is contained in γ, and we conclude that $\gamma = \gamma'$. This completes the proof of Theorem 7.??. □

Theorem 7.?? was proved by Nielsen in [**Nie1**]. He later improved its conclusion by showing in [**Nie2**] that every homotopically periodic diffeomorphism is isotopic to a periodic diffeomorphism. We will prove this second result in Chapter ??, using hyperbolic geometry. This will enable us to arrange that f_γ is actually periodic on the components of S_γ where f_γ is homotopically periodic, except possibly on a small collar neighborhood of γ.

The conclusion of Theorem 7.?? for the homotopically aperiodic parts was greatly improved by Thurston (although many of these properties were already anticipated by Nielsen; see [**Gil**][**Mil**]). Thurston showed that every homotopically aperiodic isotopy class of surface homeomorphisms contains a representative of a certain type, called pseudo-Anosov, and has important dynamical and uniqueness properites. We now study these pseudo-Anosov homeomorphisms.

7.4. Pseudo-Anosov homeomorphisms

The Nielsen Decomposition Theorem 7.?? was proved by Nielsen in [**Nie1**]. He later improved its conclusion by showing in [**Nie2**] that every homotopically periodic diffeomorphism is isotopic to a periodic diffeomorphism. We will prove this second result in Chapter ??, using hyperbolic geometry. This will enable us to arrange that f_γ is actually periodic on the components of S_γ where f_γ is homotopically periodic, except possibly on a small collar neighborhood of γ.

The conclusion of Theorem 7.?? for the homotopically aperiodic parts was greatly improved by Thurston (although many of these properties were already anticipated by Nielsen; see [**Gil**][**Mil**]). Thurston showed that every homotopically aperiodic isotopy class of surface homeomorphisms contains a representative of a certain type, called pseudo-Anosov, and has important dynamical and uniqueness properites. We now study these pseudo-Anosov homeomorphisms.

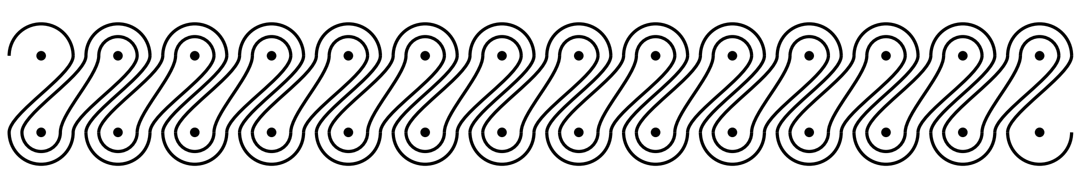

CHAPTER VIII

TRANSVERSE HÖLDER DISTRIBUTIONS FOR GEODESIC LAMINATIONS

8.1. An example revisited

In the example of §5.2, we encountered a geodesic lamination λ with a new type of transverse structure, namely a distribution defined on each arc k transverse to λ. We would like to say that this distribution is invariant under homotopy respecting λ. However, it turns out that this statement does not really make sense at this level of generality. Indeed, a distribution on the transverse arc k is a continuous linear functional on the space of differentiable functions on k. If a homotopy respecting λ sends k to another transverse arc k' and if we want to use this homotopy to identify distributions on k to distributions on k', we need the homotopy to send differentiable functions on k to differentiable functions on k', and therefore we need the homotopy to be differentiable. However, a homotopy respecting λ is very seldom differentiable.

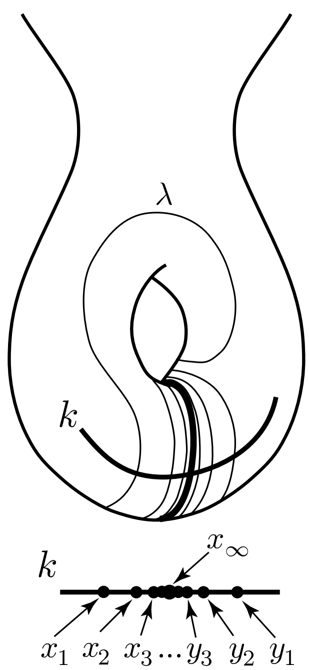

FIGURE 8.1

To understand the problem, let us examine our example in more detail. The geodesic lamination λ (called λ_0^+ in §5.2) consists of a closed geodesic g_0 and of an infinite geodesic g_1 whose ends both spiral around g_0, on opposite sides of g_0 and in opposite directions. Note that the complement $S-\lambda$ admits two spikes, spiralling on opposite sides of g_0.

Let k be an arc transverse to λ and cutting the closed leaf g_0 in exactly one point x_∞. The rest of the intersection $k\cup\lambda$ consists of isolated points $x_1, \ldots, x_i, \ldots$ and $y_1, \ldots, y_j, \ldots$, where the x_i are in the same component of $k-x_\infty$ and monotonically converge to x_∞, and where the y_j are in the other component of $k-x_\infty$ and also monotonically converge to x_∞.

If k_i is the component of $k-\lambda$ separating x_i from x_{i+1}, the k_i monotonically move towards the end of a spike in $S-\lambda$ and, consequently, its length decreases exponentially with i. To quantify this more precisely, suppose that λ is geodesic for a metric m of constant curvature -1, to simplify. Let L be the length of the closed m–geodesic g_0. As i tends to infinity, the distance from k_{i-1} to k_i in $S-\lambda$ converges to L. By an estimate of hyperbolic geometry in the corresponding spike of $S-\lambda$, we conclude that the ratio from the length of k_i to the length of k_{i-1} tends to e^{-L} as i tends to infinity. Therefore, the distance from x_i to x_{i+1} is asymptotically a constant times e^{-iL} as i tends to infinity. The same estimate holds for the distance from y_j to y_{j+1}.

The end of g_1 that contains the x_i spirals along g_0; orient g_0 in the corresponding direction. Move the arc k to a transverse arc k' by a homotopy respecting λ so that the intersection point $x_\infty = k\cap g_0$ moves to $x'_\infty = k'\cap g_0$ by an oriented distance of $D\neq 0$ along g_0. The homotopy sends each intersection point $x_i\in k\cap\lambda$ to a point $x'_i\in k'\cap\lambda$, abd each y_j to a point $y'_j\in k'\cap\lambda$. The same hyperbolic estimate as above shows that, as i tends to ∞, the ratio $d\left(x'_i, x'_{i+1}\right)/d\left(x_i, x_{i+1}\right)$ converges to $e^{-D}\sin\theta'/\sin\theta$, where θ (resp. θ') denotes the angle between k (resp. k') and g_0 at x_∞ (resp. x'_∞). On the other hand, because the end of g_1 containing the y_j spirals in the opposite direction, $d\left(y'_j, y'_{j+1}\right)/d\left(y_j, y_{j+1}\right)$ converges to $e^{+D}\sin\theta'/\sin\theta$. Since $e^{-D}\neq e^{+D}$, this shows that the homotopy from k to k' cannot be differentiable at x_∞.

In this specific example, we can make the homotopy from k to k' differentiable everywhere except at x_∞. However, for a 'generic' geodesic lamination λ, every point of intersection of λ with a transverse arc k can usually be approximated by components d, d' of $k-\lambda$ such that the two leaves of λ passing through the end points of d are asymptotic on one side of k, and the two leaves passing through the end points of d' are asymptotic on the other side; by the same argument as above, any homotopy of k respecting λ which moves x by a non-trivial amount cannot be differentiable at x.

At least, the estimates used should convince the reader that, if a homotopy respecting the geodesic lamination λ sends the transverse arch k to a transverse arc k, this homotopy can be chosen so that the resulting map $h: k\to k'$ is ***Lipschitz bicontinuous***, namely such that there exists constants $0<A<B$ such that $Ad\left(x,y\right)\leqslant d\left(h\left(x\right), h\left(y\right)\right)\leqslant Bd\left(x,y\right)$ for every x, $y\in k$. In other words, the m–geodesic lamination λ admits a Lipschitz transverse structure.

However, this Lipschitz transverse structure depends on the metric m, which is very unfortunate since we are trying to understand the properties of the space $\mathcal{ML}\left(S\right)$ in a metric independent way. To illustrate this, let us return to the example of the beginning of this section, with an m–geodesic lamination $\lambda = g_0 \cup g_1$ and with a transverse arc k meeting the closed leaf g_0 in one point x_∞. If m' is another hyperbolic metric on S, let λ' be the m'–geodesic lamination corresponding to λ. It is possible to find a homeomorphism φ of S which sends λ to λ',and sends k to an arc k' transverse to λ'. If we consider the points $x_i \in g_1 \cap \lambda$ as before, we saw that the distances $d\left(x_i, x_{i+1}\right)$ are asymptotic to a geometric progression of ration e^{-D}, where D is the m–length of the closed leaf of λ. Similarly, for the points $x_i' = \varphi\left(x_i\right)$, the distances $d\left(x_i', x_{i+1}'\right)$ are asymptotic to a geometric progression of ration $e^{-D'}$, where D is the m'–length of the closed leaf of λ'. If we chose m' so that D' is different from D, we see that φ cannot even be Lipschitz bicontinuous on k.

These considerations show that we have to be careful when defining transverse distributions for geodesic laminations, in particular if we want to do this in a metric independent way. The case of transverse measures, where we did not encounter these difficulties, points to a method of overcoming these problems. Indeed, a (Radon) measure on a manifold can be considered as a distribution, namely a functional on the space of differentiable functions, with the additional regularity property that it continuously extends to the space of all continuous functions. Similarly, we can restrict attention to distributions, defined on arcs transverse to the geodesic laminations, which continuously extend to a family of functions which is invariant under homotopy respecting λ and under change of the metric. We will see that the boundary at infinity $\partial_\infty \widetilde{S}$ admits a natural Hölder structure, so that the family of Hölder continuous functions is well suited to our purposes. This leads us to the introduction of Hölder distributions.

8.2. Hölder distributions

Given a metric space (X, d), recall that a function $\varphi : X \to \mathbb{R}$ is ***Hölder continuous*** if there exists two constants $A \geqslant 0$ and ν with $0 < \nu \leqslant 1$ such that $|\varphi(x) - \varphi(y)| \leqslant A d(x, y)^\nu$ for every x, $y \in X$. The number ν is a ***Hölder exponent*** for φ, and the ***Hölder norm of exponent*** ν of φ is

$$\|\varphi\|_\nu = \sup_x |\varphi(x)| + \sup_{x \neq y} |\varphi(x) - \varphi(y)|\, d(x, y)^{-\nu}.$$

Let $H(X)$ denote the space of Hölder continuous functions $\varphi : X \to \mathbb{R}$ with compact support. For every ν and every compact subset K of X, let $H_\nu(X; K)$ denote the space of those $\varphi \in H(X)$ which have Hölder exponent ν and whose support is contained in K, endowed with the topology defined by the norm $\| \ \|_\nu$. Note that there is a continuous inclusion map $H_\nu(X; K) \to H_{\nu'}(X; K')$ for every $\nu' \leqslant \nu$ and every compact subset K' containing K. Then, $H(X)$ is the union of all $H_\nu(X; K)$ as K ranges over all compact subsets of X and ν ranges over all positive numbers. By definition, a ***Hölder distribution*** on X is a linear functional $H(X) \to \mathbb{R}$ whose restriction to each $H_\nu(X; K)$ is continuous. When X is a differentiable manifold, a Hölder distribution is a distribution in the usual sense with some additional regularity properties.

A Radon measure on X, namely a Borelian signed measure which assigns finite mass to each compact subset, provides an example of such a Hölder distribution. But, there are many Hölder distributions which are not signed measures. Here is a simple example, reminiscent of the example of §5.2: For each Hölder continuous function $\varphi : [0,1] \to \mathbb{R}$, we can consider the sum $\alpha(\varphi) = \sum_{n=0}^{\infty} (\varphi(2^{-n}) - \varphi(0))$, which converges since it is bounded by a geometric series; this defines a Hölder distribution α on $[0,1]$, which is not a measure since explicit constructions show that $|\alpha(\varphi)|$ is not bounded by a constant times the L^∞ norm of φ.

Hölder distributions share an essential feature with signed measures, namely they are completely determined by their restriction to their support, as proved by the following elementary lemma. This property is well known to be false for many distributions, for instance the distribution $\alpha(\varphi) = \frac{d\varphi}{dt}(0)$ on $\mathbb{R}$. The fact that it holds for Hölder distributions will play a crucial rôle in our classification of transverse Hölder distributions for a given geodesic lamination.

For a Hölder distribution α on a metric space X, its ***support*** is the smallest closed subset K_α of X such that $\alpha(\varphi) = 0$ for every Hölder continuous function $\varphi : X \to \mathbb{R}$ whose support is compact and disjoint from K_α.

Proposition (Support Lemma). *Let K be a compact subset of a metric space X, and let α be a continuous linear form on the space $H_\nu(X,K)$ of Hölder continuous functions of exponent ν on X with support contained in K. If, for $\nu' > \nu$, the function $\varphi \in H_{\nu'}(X,K) \subset H_\nu(X,K)$ is identically 0 on the support K_α of α, then $\alpha(\varphi) = 0$.*

Proof. Let $\xi : \mathbb{R} \to \mathbb{R}$ be the piecewise linear function which is identically 1 on $\left]-\infty, \frac{1}{2}\right]$, which is identically 0 on $[1, \infty[$, and which is linear on $\left[\frac{1}{2}, 1\right]$. Note that $|\xi(u) - \xi(v)| \leqslant 2|u - v|$ for every u, v.

For every $\varepsilon > 0$, consider the function $\varphi_\varepsilon : X \to \mathbb{R}$ defined by $\varphi_\varepsilon(x) = \varphi(x)\,\xi(d(x, K_\alpha)/\varepsilon)$, where $d(\ ,\ )$ denotes the metric of X. Observe that φ_ε coincides with φ on a neighborhood of K_α, and therefore that $\alpha(\varphi_\varepsilon) = \alpha(\varphi)$.

Also, φ_ε is identically 0 outside of the ε–neighborhood of K_α. Since φ is identically 0 on K_α and since $|\xi| \leqslant 1$, it follows that $|\varphi_\varepsilon(x)| \leqslant \varepsilon^{\nu'} \|\varphi\|_{\nu'}$ for every $x \in X$.

Let us estimate the norm $\|\varphi_\varepsilon\|_\nu$. For this, we have to estimate the supremum of $|\varphi_\varepsilon(x) - \varphi_\varepsilon(y)|\, d(x,y)^{-\nu}$.

For x, y with $d(x,y) \geqslant \varepsilon$,

$$|\varphi_\varepsilon(x) - \varphi_\varepsilon(y)|\, d(x,y)^{-\nu} \leqslant 2\varepsilon^{\nu'} \|\varphi\|_{\nu'}\, d(x,y)^{-\nu} \leqslant 2\varepsilon^{\nu'-\nu} \|\varphi\|_{\nu'}.$$

For x, y with $d(x,y) < \varepsilon$, we may assume without loss of generality that at least one of them, say x, is in the ε–neighborhood of K_α, so that $|\varphi(x)| \leqslant$

$\varepsilon^{\nu'} \|\varphi\|_{\nu'}$. Then,

$$\begin{aligned} |\varphi_\varepsilon(x) - \varphi_\varepsilon(y)|\, d(x,y)^{-\nu} &= \left|\varphi(x)\,\xi\left(\varepsilon^{-1} d(x, K_\alpha)\right) - \varphi(y)\,\xi\left(\varepsilon^{-1} d(y, K_\alpha)\right)\right| d(x,y)^{-\nu} \\ &\leqslant |\varphi(x)| \left|\xi\left(\varepsilon^{-1} d(x, K_\alpha)\right) - \xi\left(\varepsilon^{-1} d(y, K_\alpha)\right)\right| d(x,y)^{-\nu} \\ &\qquad + |\varphi(x) - \varphi(y)| \left|\xi\left(\varepsilon^{-1} d(y, K_\alpha)\right)\right| d(x,y)^{-\nu} \\ &\leqslant \varepsilon^{\nu'} \|\varphi\|_{\nu'}\, 2\varepsilon^{-1} d(x,y)\, d(x,y)^{-\nu} \\ &\qquad + \|\varphi\|_{\nu'}\, d(x,y)^{\nu'}\, d(x,y)^{-\nu} \\ &\leqslant 2\varepsilon^{\nu'} \|\varphi\|_{\nu'}\, \varepsilon^{-1}\varepsilon^{1-\nu} + \|\varphi\|_{\nu'}\, \varepsilon^{\nu'-\nu} \\ &\leqslant 3\varepsilon^{\nu'-\nu} \|\varphi\|_{\nu'}\,. \end{aligned}$$

since $\nu \leqslant 1$ and $\nu' > \nu$.

Combining both cases, we see that the supremum of $|\varphi_\varepsilon(x) - \varphi_\varepsilon(y)|\, d(x,y)^{-\nu}$ is bounded by $3\varepsilon^{\nu'-\nu} \|\varphi\|_{\nu'}$. It follows that $\|\varphi_\varepsilon\|_\nu \leqslant \left(\varepsilon^{\nu'} + 3\varepsilon^{\nu'-\nu}\right) \|\varphi\|_{\nu'}$. In particular, the function φ_ε tends to 0 in the space $H_\nu(X;K)$ as ε tends to 0. As a consequence, $\alpha(\varphi_\varepsilon)$ tends to 0 as ε tends to 0, by continuity of α. Since we observed that $\alpha(\varphi_\varepsilon) = \alpha(\varphi)$ for every ε, this proves that $\alpha(\varphi) = 0$. □

For every closed subset Y of a metric space X, every Hölder distribution β on Y extends to a Hölder distribution α on X by setting $\alpha(\varphi) = \beta\left(\varphi_{|Y}\right)$ for every Hölder continuous $\varphi : X \to \mathbb{R}$ with compact support.

LEMMA. *If Y is a closed subset of a metric space X, the above extension operator establishes a one-to-one correspondence between Hölder distributions on Y and Hölder distributions on X whose support in contained in Y.*

PROOF. Clearly, if the Hölder distribution α on X is obtained by extension of a distribution β of Y, its support is contained in Y. Conversely, let α be a Hölder distribution on X whose support is contained in Y. We want to associate to it a Hölder distribution β on Y. For this, let ψ be a Hölder continuous function with compact support and with Hölder exponent ν. We can extend ψ to a Hölder continuous function $\varphi' : X \to \mathbb{R}$, for instance by setting $\varphi'(x) = \inf_{y \in Y} \left(\psi(y) + 2\|\psi\|_\nu\, d(x,y)^\nu\right)$ where $d(\ ,\ )$ denotes the metric of X. Multiplying φ' by a Hölder continuous function with compact support which is equal to 1 on the support of ψ, we get a Hölder continuous function $\varphi : X \to \mathbb{R}$ with compact support with $\varphi_{|Y} = \psi$. Define $\beta(\psi) = \alpha(\varphi)$; by the Support Lemma 1, this $\beta(\psi)$ does not depend on the choice of the extension φ. In addition, by construction, $\beta(\psi)$ depends continuously on ψ, and β therefore defines a Hölder distribution on Y. Clearly, this restriction operator $\alpha \mapsto \beta$ is an inverse for the extension operator. □

8.3. Hölder geodesic currents and transverse Hölder distributions for geodesic laminations

Consider an m–geodesic lamination λ, where m is a negatively curved metric with geodesic boundary on S. If k and k' are two arcs transverse to λ are homotopic respecting λ, we first prove a regularity property for the map $h : k \to k'$ induced by this homotopy. Note that, when k or k' is not embedded, the notation $h : k \to k'$ is somewhat abusive, and h has to be undersood as a map between the domains of k and k'. Recall that a map h is ***Lipschitz continuous*** if it is Hölder continuous with Hölder exponent 1, namely if there is a constant a such that $d\left(h\left(x\right), h\left(y\right)\right) \leqslant Ad\left(x, y\right)$ for every x, y, and that h is ***Lipschitz bicontinuous*** if it is bijective and if both h and h^{-1} are Lipchitz continuous.

LEMMA. *Let k and k' be two arcs which are transverse to the m–geodesic lamination λ and which are homotopic by a homotopy respecting λ. Then, the homotopy can be chosen so that the map $h : k \to k'$ it induces is Lipschitz bicontinuous with respect to the metrics induced by m on k and k'.*

PROOF. The homotopy establishes a one-to-one correspondence between $k \cap \lambda$ and $k' \cap \lambda$, as well as between the components of $k - \lambda$ and the components of $k' - \lambda$. An easy geometric estimate shows that the length of a component of $k - \lambda$ is bounded above and below by constants times the length of the corresponding component of $k' - \lambda$. (The constants depend on the curvature of m, on the diameter of $k \cup k'$, and on the minimum angle between k, k', and λ). Since, by Theorem 1.??, $k \cap \lambda$ has Hausdorff dimension 0, the distance in k between two of its points is equal to the length of the components of $k - \lambda$ separating them. It follows that the correspondence $k \cap \lambda \to k' \cap \lambda$ is Lipschitz continuous, as well as its inverse. As a consequence, we can choose the homotopy so that the induced map $h : k \to k'$ is Lipschitz bicontinuous. □

In particular, the map $h : k \to k'$ provided by Lemma ?? is ***Hölder bicontinuous***, in the sense that both h and h^{-1} are Hölder continuous. An important consequence is that h enables us to transport any Hölder continuous function on k to a Hölder continuous function on k', and conversely, so that h identifies Hölder distributions on k and Hölder distributions on k'.

A ***transverse Hölder distribution*** α for the geodesic lamination λ is a Hölder distribution defined on each differentiable arc k transverse to λ, and such that every Hölder bicontinuous homotopy sending k to another arc k' while respecting λ sends the Hölder distribution defined on k to the Hölder distribution defined on k'.

The invariance property implies that the Hölder distribution deposited by α on a transverse arc k has support contained in $k \cap \lambda$. Indeed, this Hölder distribution has to be invariant under any Hölder bicontinuous homeomorphism of k fixing $k \cap \lambda$. Applying the following elementary lemma to the closure of any component d of $k - \lambda$, we conclude that $\alpha\left(\psi\right) = 0$ for every Hölder continuous function $\psi : k \to \mathbb{R}$ whose support is contained in d. Every Hölder continous function whose support is disjoint from $k \cap \lambda$ can be written as the sum of finitely many Hölder continuous functions ψ_i, where the support of ψ_i is contained in a component d_i of $k - \lambda$. We conclude that $\alpha\left(\psi\right) = 0$ for every Hölder continuous function $\psi : k \to \mathbb{R}$ whose support is contained in $k - \lambda$, and therefore that the support of the Hölder distribution deposited by α on k is contained in $k \cap \lambda$.

LEMMA. *Let α be a Hölder distribution on the interval $[0,1]$ which is invariant under all oriented Hölder bicontinuous homeomorphisms of $[0,1]$. Then, $\alpha(\psi) = 0$ for every Hölder continuous function $\psi : [0,1] \to \mathbb{R}$ with support contained in the interior of $[0,1]$.*

PROOF. Let $\nu < 1$ be a Hölder exponent for ψ. For every $k \geqslant 1$, let $\psi_k : [0,1] \to \mathbb{R}$ be the function defined by $\psi_k(x) = \psi(kx)$ if $kx \in [0,1]$ and $\psi_k(x) = 0$ otherwise. This function ψ_k is obtained by composing ψ with any Hölder bicontinuous homeomorphism of $[0,1]$ coinciding with $x \mapsto kx$ on a neighborhood of the preimage of the support of ψ under this map $x \mapsto kx$. As a consequence, $\alpha(\psi_k) = \alpha(\psi)$. Also, $\|\psi_k\|_\nu \leqslant k^\nu \|\psi\|_\nu$.

If the support of ψ is contained in the interior of an interval of length $l < 1$, the support of ψ_k is contained in the interior of an interval of length l/k. If n_k denotes the integer part of k/l, we can choose n_k disjoint intervals of length l/k in $[0,1]$. Let ψ'_k be the function defined by putting a translate of ψ_k in each of these intervals. Then, $\alpha(\psi'_k) = n_k \alpha(\psi_k) = n_k \alpha(\psi)$ and $\|\psi'_k\|_\nu \leqslant 2\|\psi_k\|_\nu \leqslant 2k^\nu \|\psi\|_\nu$.

If we consider the function $n_k^{-1}\psi'_k$, we conclude that $\left\|n_k^{-1}\psi'_k\right\|_\nu \leqslant 2n_k^{-1}k^\nu \|\psi\|_\nu$ tends to 0 as k tends to ∞ (since $\nu < 1$), while $\alpha\left(n_k^{-1}\psi'_k\right) = \alpha(\psi)$. By continuity of α, it follows that $\alpha(\psi) = 0$. □

Although this definition of transverse Hölder distributions for a geodesic lamination λ is relatively intuitive, it is not clear that it is independent of the choice of a negatively curved metric m on S. We now give an equivalent definition, based on the space $G(\widetilde{S})$ of geodesics of the universal covering $\widetilde{S}$.

The two spaces $\widetilde{S}_\infty$ and $G(\widetilde{S})$ have a well-defined ***Hölder structure***, namely a preferred metric defined up to the ***Hölder equivalence*** relation which identifies two metrics d_1 and d_2 when there are constants $\nu > 0$ and $K > 0$ such that $K^{-1}d_1(x,y)^{1/\nu} \leqslant d_2(x,y) \leqslant Kd(x,y)^\nu$ for every x, y. Indeed, the choice of a negatively curved metric with totally geodesic boundary on S and of a base point $\widetilde{x}_0 \in \widetilde{S}$ identifies $\widetilde{S}_\infty$ to a subset of the circle of directions at $\widetilde{x}_0$. This induces a metric on $\widetilde{S}_\infty$ by restriction of the angle metric of this circle of directions. It can be shown that, if we vary the metric or the base point, the Hölder equivalence class of this metric is unchanged (see Appendix A). Therefore, this defines a natural Hölder structure on $\widetilde{S}_\infty$, and consequently on $G(\widetilde{S}) \subset \left(\widetilde{S}_\infty \times \widetilde{S}_\infty\right)/\mathbb{Z}_2$. We will alwaysassume $G(\widetilde{S})$ endowed with a metric that is compatible with this Hölder structure.

Since the notion of Hölder distribution depends only on the Hölder equivalence class of the metric considered, we therefore have a natural space of Hölder distributions on $G(\widetilde{S})$.

Let a ***geodesic Hölder current*** on S be a Hölder distribution on $G(\widetilde{S})$ which is invariant under the action of $\pi_1(S)$.

PROPOSITION 5. *Given a geodesic lamination λ on S, there is a natural one-to-one correspondence between transverse Hölder distributions for λ and geodesic Hölder currents whose support in $G(\widetilde{S})$ is contained in the set $\widetilde{\lambda}$ consisting of the lifts of the leaves of λ.*

Proof. Consider a geodesic Hölder current α whose support is contained in $\widetilde{\lambda}$. By Lemma 2, this amounts to a $\pi_1(S)$–invariant Hölder distribution on $\widetilde{\lambda} \subset G(\widetilde{S})$. If k is a transverse arc for λ, lift it to an arc $\widetilde{k}$ in $\widetilde{S}$. Cutting k into smaller arcs if necessary, we can assume that $\widetilde{k}$ meets each geodesic of $\widetilde{\lambda}$ in at most one point. Then, by consideration of intersection points, we can identify $k \cap \lambda$ with the set $\widetilde{\lambda} \cap G(\widetilde{k})$ consisting of those geodesics of $\widetilde{\lambda}$ which cross $\widetilde{k}$. By an easy geometric estimate (see for instance Appendix A), the distance between the directions of two geodesics of $\widetilde{\lambda} \cap G(\widetilde{k})$ is bounded by a constant times the distance between their intersection points with $\widetilde{k}$. By definition of the Hölder structure of $G(\widetilde{S})$, it follows that the identification of $k \cap \lambda$ with $\widetilde{\lambda} \cap G(\widetilde{k})$ is Hölder bicontinuous. By Lemma 2, α induces a Hölder distribution on $\widetilde{\lambda} \cap G(\widetilde{k})$, and therefore on $k \cap \lambda$. By extension, we get a Hölder distribution on k whose support is contained in $k \cap \lambda$. The $\pi_1(S)$–invariance of α implies that this Hölder distribution does not depend on the choice of the lift $\widetilde{k}$. This defines a Hölder distribution on each arc k transverse to λ, which is easily seen to be invariant under Hölder bicontinuous homotopy of k respecting λ.

Conversely, let α be a transverse Hölder distribution for λ. For every $g \in \widetilde{\lambda}$, there is a neighborhood U of g in $\widetilde{\lambda}$ consisting of those geodesics which cross an arc $\widetilde{k}$ transverse to $\widetilde{\lambda}$. We can assume $\widetilde{k}$ small enough so that it meets each geodesic of $\widetilde{\lambda}$ in at most one point. If k denotes the projection of $\widetilde{k}$ to S, the transverse Hölder distribution α induces a Hölder distribution on k whose support is contained in $k \cap \lambda$, and therefore a Hölder distribution on $k \cap \lambda \cong U$ by Lemma 2. In addition, this Hölder distribution on U is independent of the choice of k by the invariance property of α. Therefore, α induces a Hölder distribution on a small neighborhood of each geodesic of $\widetilde{\lambda}$, which extends to a Hölder distribution on $\widetilde{\lambda}$ by using a suitable Hölder continuous partition of unity, which itself extends to a Hölder distribution on $G(\widetilde{S})$ with support contained in $\widetilde{\lambda}$. The fact that this Hölder distribution is invariant under the action of $\pi_1(S)$ is immediate.

These two transformations clearly form a one-to-one correspondence between the transverse distributions for λ and the geodesic Hölder currents whose support is contained in $\widetilde{\lambda}$. □

8.4. The basic geometric-analytic estimates

The first half of this monograph was in great part combinatorial in nature. The second half will involve many more analytic arguments. When going over the justifications of these analytic arguments, the reader will quickly realize that they are based on a very small number of basic estimates, which use global topological observations on geodesic laminations to estimate the growth of locally defined quantities. We have already used some of these estimates in previous chapters, but now is probably a good time to single them out before we use them more extensively.

Most of these estimates deal with the situation of a geodesic lamination λ

which is strongly carried by a fattened train track Φ, for some metric m_0 on S. Let k_0 be a generic tie of Φ. We are concerned with the ***gaps*** occurring on k_0 between the points of $k_0 \cap \lambda$, namely with the components of $k_0 - \lambda$.

If d is such a component of $k_0 - \lambda$, the ***divergence radius*** or ***depth*** $r(d)$ is defined as follows. First, to avoid any ambiguity, we lift the situation to the universal covering $\widetilde{S}$ of S. Consider the pre-images $\widetilde{\Phi}$ of Φ and $\widetilde{\lambda}$ of λ in $\widetilde{S}$ and lift k_0 to a tie $\widetilde{k}_0$ of $\widetilde{\Phi}$. By convention, $r(d) = 1$ when d contains one of the end points of k_0. Otherwise, there are two leaves g_d^+ and g_d^- of $\widetilde{\lambda}$ which pass through the end points of the component $\widetilde{d}$ of $\widetilde{k}_0 - \widetilde{\lambda}$ associated to d by the projection $\widetilde{k}_0 \to k_0$. Then, $r(d)$ is the largest integer $r \geqslant 1$ such that g_d^+ and g_d^- follow a common edge path $\langle e_{-r+1}, e_{-r+2}, \dots, e_0, \dots, e_{r-2}, e_{r-1} \rangle$ of $\widetilde{\Phi}$ that is centered on the edge e_0 containing $\widetilde{k}_0$. The fact that $r(d)$ is largest for this property implies that the bi-infinite edge paths respectively followed by g_d^+ and g_d^- diverge after e_{-r+1} or e_{r-1}. Intuitively, $r(d)$ measures how far in $S - \lambda$ the gap d is from the complement of Φ.

Proposition (Exponential decay for gap lengths). *Let the surface S be endowed with a metric m_0 of negative curvature and with geodesic boundary, and let k_0 be a generic tie of the fattened train track $\Phi \subset S$. Then there is a constant A such that, for every m_0–geodesic lamination λ which is strongly carried by Φ and for every component d of $k_0 - \lambda$ with divergence radius $r(d)$, the m_0–length of d is an $O\left(e^{-Ar(d)}\right)$, where the constant A and the constant hidden in the symbol $O(\ )$ depend only on the metric m_0 and on the fattened train track Φ.*

Proof. For convenience, we repeat the proof which we already used in ??. $\square$

Proposition (Bounded number of gaps of given depth). *There is a constant C, depending only on the topology of S, with the following property: For every m_0–geodesic lamination λ which is strongly carried by the fattened train track Φ and for every generic tie k_0 of Φ, the number of components d of $k_0 - \lambda$ with $r(d) = r$ is bounded by C for r sufficiently large.*

Proof. $\square$

Pick an orientation for the tie k_0. If $\alpha \in \mathcal{H}(\lambda; \mathbb{R})$ is an $\mathbb{R}$–valued transverse cocycle for the geodesic lamination λ, we already encountered in §6.2 the α–***height*** $h_\alpha(d) \in \mathbb{R}$ with respect to α of the component d of $k_0 - \lambda$. By definition, $h_\alpha(d) = \alpha(k_d)$ where k_d is a sub-arc of k_0 that joins the negative end point of k_0 to an arbitrary point in d.

Lemma (Linear growth for gap heights). *Let k_0 be an oriented generic tie of the fattened train track Φ, and let $\alpha \in \mathcal{H}(\lambda; \mathbb{R})$ be an $\mathbb{R}$–valued transverse cocycle for an m_0–geodesic lamination λ which is strongly carried by Φ. Then, for every component d of $k_0 - \lambda$, the α–height $h_\alpha(d)$ is an $O\left(\|\alpha\|_\Phi\, r(d)\right)$, where $\|\alpha\|_\Phi$ is the maximum of the absolute values $|\alpha(e)|$ as e ranges over the edges of Φ, where the constant hidden in the symbol $O(\ )$ depends only on the topology of Φ.*

Proof. In Lemma 6.??, we gave an explicit formula which enables us to compute $h_\alpha(d)$ in terms of the edge weights $\alpha(e)$. The estimate is an immediate corollary of this formula, the constant involving the maximum of the number of edges meeting at each switch of Φ. □

Proposition (Linear growth for edge path weights). *Let γ be an edge path of length $l(\gamma)$ of the train track Φ. Then, for every m_0–geodesic lamination λ which is carried by Φ and for every transverse cocycle $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ for λ, the weight $\alpha(\gamma)$ is an $O(\|\alpha\|_\Phi\, l(\gamma))$, where $\|\alpha\|_\Phi$ is the maximum of the absolute values $|\alpha(e)|$ as e ranges over the edges of Φ, and where the constant hidden in the symbol $O(\ )$ depends only on the topology of Φ.*

Proof.

□

A similar property occurs when we consider two measured m_0–geodesic laminations $\alpha_0,\ \alpha_1 \in \mathcal{ML}(S)$ which are strongly carried by the fattened train track Φ. As in ??, these measured geodesic laminations associate masses $\alpha_0(\gamma)$, $\alpha_1(\gamma) \geqslant 0$ to each edge path γ in Φ. They also associate weights $\alpha_0(e)$, $\alpha_1(e) \geqslant 0$ to the edges of Φ, and we can consider the supremum $\|\alpha_1 - \alpha_0\|_\Phi$ of the absolute values $|\alpha_1(e) - \alpha_0(e)|$ as e ranges over all edges of Φ.

Proposition (Linear growth for edge path masses). *Let γ be an edge path of length $l(\gamma)$ of the train track Φ. Then, for every measured geodesic laminations $\alpha_0,\ \alpha_1 \in \mathcal{ML}(S)$ which are (weakly) carried by Φ, the masses $\alpha_0(\gamma)$, $\alpha_1(\gamma)$ are such that $\alpha_1(\gamma) = \alpha_0(\gamma) + O(\|\alpha_1 - \alpha_0\|_\Phi\, l(\gamma))$, where the combinatorial norm $\|\alpha_1 - \alpha_0\|_\Phi$ is the maximum of the absolute values $|\alpha_1(e) - \alpha_0(e)|$ as e ranges over the edges of Φ, and where the constant hidden in the symbol $O(\ )$ depends only on the topology of Φ.*

Proof. □

Proposition (Polynomial growth for the number of simple edge paths). *In the fattened train track Φ, consider those edge paths which are realized by simple arcs carried by Φ. The number of all such edge paths of length r is bounded by a polynomial function of r.* □

8.5. The Gap Formula

In the next section, we will classify transverse Hölder distributions to a given geodesic lamination λ, in terms of transverse cocycles for λ. The key technical step in the proof of this classification is the following Gap Formula which, given a transverse Hölder distribution α for λ and a Hölder continuous function $\psi : k \to \mathbb{R}$ defined on an arc k transverse to λ, expresses the α–integral $\alpha(\psi)$ in terms of the jumps made by ψ at the component of $k - \lambda$ (=gaps).

More precisely, consider an *oriented* arc k transverse to λ, and a transverse Hölder distribution α for λ. For each component d of $k - \lambda$, let the α–***height*** $h_\alpha(d) \in \mathbb{R}$ of d with respect to α be the α–integral of a Hölder continuous function $\psi_d : k \to \mathbb{R}$ which coincides with the constant function 1 on the part of $k \cap \lambda$ that sits below d, for the orientation of k, and with the constant function 0 on the part of $k \cap \lambda$ that sits above d. By the Support Lemma ??, this α–height $h_\alpha(d)$ does not depend on the choice of the function ψ_d.

THEOREM ?? (GAP FORMULA). *Let α be a transverse Hölder distribution for the geodesic lamination λ. Let $\psi : k \to \mathbb{R}$ be a Hölder continuous function defined on an oriented arc k transverse to λ. Then,*

$$\alpha(\psi) = \alpha(k)\,\psi\left(x_k^+\right) + \sum_d h_\alpha(d)\left(\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right)$$

where the sum is over all components d of $k-\lambda$, where x_d^+ and x_d^- are the positive and negative end points of d for the orientation induced by the orientation of k, where x_k^+ is the positive end point of k, where $\alpha(k)$ denotes the α–integral of the constant function 1 on k, and where $h_\alpha(d)$ is the α–height of d with respect to α.

Note that the requirement that ψ is Hölder continuous is crucial, even when the left hand side $\alpha(\psi)$ happens to make sense. For instance, if α is a transverse measure which is non-trivial on k, consider the function $\psi : k \to \mathbb{R}$ defined by $\psi(x) = \alpha\left(\left[x, x_k^+\right]\right)$. If α has no atoms, for instance if λ has no closed leaf (compare ??), the function ψ is continuous, non-constant, and therefore has non-zero α–integral $\alpha(\psi)$; actually, one can easily see that $\alpha(\psi) = \frac{1}{2}\alpha(k)^2$. On the other hand, the value $\psi\left(x_k^+\right)$ and all the jumps $\psi\left(x_d^-\right) - \psi\left(x_d^+\right)$ are all equal to 0, so that the right hand side of the Gap Formula is 0 in this case. The reason for the discrepancy is of course that the function ψ is continuous but not Hölder continuous.

PROOF OF THEOREM ??. The proof will take a while.

The first step is the following lemma, which does not involve α.

LEMMA 3 (GAP LEMMA). *Let k be an oriented transverse arc transverse to the geodesic lamination λ. For every component d of $k - \lambda$, let x_d^+ and x_d^- be the positive and negative end points of d for the orientation induced by the orientation of k. Then, for every Hölder continuous function $\psi : k \to \mathbb{R}$ and every component d of $k - \lambda$,*

$$\psi\left(x_d^+\right) = \psi\left(x_k^+\right) + \sum_{d'>d}\left(\psi\left(x_{d'}^-\right) - \psi\left(x_{d'}^+\right)\right)$$

where the sum ranges over all those components d' of $k - \lambda$ which are above d for the orientation of k and where x_k^+ is the positive end point of k.

PROOF. The main point is that, by Theorem ??, $k \cap \lambda$ has Hausdorff dimension 0. Since ψ is Hölder continuous, it follows that the image $\psi(k \cap \lambda)$ also has Hausdorff dimension 0 in $\mathbb{R}$, and therefore has Lebesgue measure 0. This is sufficient to guarantee that the difference $\psi\left(x_d^+\right) - \psi\left(x_k^+\right)$ is the sum of the jumps $\psi\left(x_{d'}^-\right) - \psi\left(x_{d'}^+\right)$. We can now give a formal proof.

Assume that ψ has Hölder exponent $\nu > 0$.

For every $\varepsilon > 0$, the fact that $k \cap \lambda$ has Hausdorff dimension 0 implies that it is possible to cover it by finitely many disjoint sub-arcs $k_1, k_2, \ldots, k_n$ of k

whose lengths $l_1, l_2, \ldots, l_n$ are such that $\sum_{i=1}^{n} l_i^{\nu} < \varepsilon$. Choose the indexing so that $k_1, k_2, \ldots, k_n$ occur in this order for the orientation of k and, for $1 \leqslant i \leqslant n-1$, let d_i be the component of $k-\lambda$ that contains the part of k which is between k_i and k_{i+1}. Let d_n denote the component of k which contains x_k^+, and let d_0 be the one containing x_k^-.

If ε is smaller than the length of d to the power ν, then d must be equal to some d_p. Since $x_{d_n}^+ = b$,

$$\psi\left(x_d^+\right) = \psi(b) + \sum_{i=p+1}^{n} \left(\psi\left(x_{d_i}^-\right) - \psi\left(x_{d_i}^+\right)\right) + \sum_{i=p}^{n-1} \left(\psi(x_{d_i}^+) - \psi(x_{d_{i+1}}^-)\right).$$

The two points $x_{d_{i+1}}^-$ and $x_{d_i}^+$ are in the interval k_{i+1}, their distance is therefore bounded by the length l_{i+1}. It follows that:

$$\left|\psi\left(x_d^+\right) - \psi(b) - \sum_{i=p+1}^{n} \left(\psi\left(x_{d_i}^-\right) - \psi\left(x_{d_i}^+\right)\right)\right| \leqslant \sum_{i=p}^{n-1} \left|\psi(x_{d_i}^+) - \psi(x_{d_{i+1}}^-)\right|$$
$$\leqslant \|\psi\|_{\nu} \sum_{i=p}^{n-1} l_{i+1}^{\nu} \leqslant \|\psi\|_{\nu}\, \varepsilon.$$

Therefore, $\psi\left(x_d^+\right)$ can be arbitrarily approximated by the sum of $\psi(b)$ and of finitely many $\psi\left(x_{d'}^-\right) - \psi\left(x_{d'}^+\right)$ with $d' > d$. In addition, observe that the above family $\{d_i; p < i \leqslant n\}$ contains any arbitrary finite family of $d' > d$ if we choose ε so that $\varepsilon^{1/\nu}$ is less than the lengths of these d'. It follows that $\psi\left(x_d^+\right) = \psi(b) + \sum_{d'>d} \left(\psi\left(x_{d'}^-\right) - \psi\left(x_{d'}^+\right)\right)$ which concludes the proof of the lemma. □

For each component d of $k - \lambda$, let $\eta_d : k \to \mathbb{R}$ be the continuous function which is identically 1 on those points which are below d, is identically 0 on those points which are above d, and is linear with respect to arc length on d.

As in §7.2, let $H_\nu(X)$ denote the space of Hölder continuous functions $\varphi : X \to \mathbb{R}$ of Hölder exponent ν with compact support, endowed with the Hölder norm $\| \ \|_\nu$. The following lemma is a refinement of the Gap Lemma ??, and precises the convergence in that lemma.

Lemma ?? (Hölder Gap Lemma). *If $\psi : k \to \mathbb{R}$ is Hölder continuous with Hölder exponent ν, the series $\psi\left(x_k^+\right) + \sum_d \left(\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right)\eta_d$ converges to some function $\bar{\psi}$ in every Hölder space $H_{\nu'}(k)$ with $0 < \nu' < \nu$. In addition, $\bar{\psi}$ coincides with ψ on $k \cap \lambda$.*

Proof. In contrast to the proof of the Gap Lemma ??, the proof of Lemma ?? uses global properties of the geodesic lamination λ in a crucial way.

Since k is compact, the space $H_{\nu'}(k)$ is complete. It therefore suffices to prove the convergence of the series $\sum_d \left|\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right| \|\eta_d\|_{\nu'}$.

To estimate $\|\eta_d\|_{\nu'}$, let $\delta(\ ,\ )$ denote the arc length metric on k. Then, the maximum of the ratio $|\eta_d(x) - \eta_d(y)|\,\delta(x,y)^{-\nu'}$ is attained when x and y are the end points of the interval d. It follows that $\|\eta_d\|_{\nu'} = 1 + l(d)^{-\nu'}$ where $l(d)$ is the length of d. In particular, $\|\eta_d\|_{\nu'} \leqslant 2l(d)^{-\nu'}$ for all but finitely many d.

Also, $\left|\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right| \leqslant \|\psi\|_{\nu} l(d)^{\nu}$. It therefore suffices to prove the convergence of the series $\sum_d l(d)^{\nu-\nu'}$.

Consider the completion $\widehat{S-\lambda}$ of $S-\lambda$ with the path metric induced by the metric of S, as in §1.??. In §1.??, we saw that $\widehat{S-\lambda}$ is a surface with geodesic boundary, and is the union of a compact part and of finitely many spikes. Each component d of $k-\lambda$ gives an arc in $\widehat{S-\lambda}$ going from the boundary to the boundary.

By transversality, those components d of $k-\lambda$ which meet the compact part of $\widehat{S-\lambda}$ have length bounded away from 0. Therefore, there are only finitely many of them.

For each spike of $\widehat{S-\lambda}$, index the components d of $k-\lambda$ that are contained in this spike as $d_1, \ldots, d_n, \ldots$, in such a way that the d_n occur in this order as we move towards the end of the spike. Then an easy geometric estimate (see Appendix ??), using the negative curvature of the metric m of S, shows that the length of d_n is an $O\left(e^{-An}\right)$, where the constant $A > 0$ depends only on the curvature of m and on the infimum of the lengths of the components of $\lambda - k$.

It follows that the series $\sum_d l(d)^{\nu-\nu'}$ is bounded by the sum of finitely many series $\sum_{n=1}^{\infty} e^{-A\left(\nu-\nu'\right)n}$, one for each spike of $\widehat{S-\lambda}$. Since $B > 1$ and $\nu' < \nu$, these geometric series are convergent and this concludes the proof of the convergence of the series $\sum_d \left(\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right)\eta_d$ in the Hölder space $H_{\nu'}(k)$, as claimed. Consider the Hölder continuous function $\bar{\psi} = \psi\left(x_k^+\right) + \sum_d \left(\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right)\eta_d$.

For every component d of $k-\lambda$, $\eta_{d'}\left(x_d^+\right) = 0$ if $d' \leqslant d$ and $\eta_{d'}\left(x_d^+\right) = 1$ if $d' > d$. It follows that

$$\bar{\psi}\left(x_d^+\right) = \psi\left(x_k^+\right) + \sum_{d'>d}\left(\psi\left(x_{d'}^-\right) - \psi\left(x_{d'}^+\right)\right) = \psi\left(x_d^+\right)$$

by the Gap Lemma ??. Since $k \cap \lambda$ has Hausdorff dimension 0, its interior in k is empty and the x_d^+, where d ranges over all components of $k - \widetilde{\lambda}$, are dense in $k \cap \lambda$. By continuity, it follows that $\bar{\psi}$ coincides with ψ on $k \cap \lambda$. □

The Support Lemma ?? implies that $\alpha\left(\bar{\psi}\right) = \alpha(\psi)$ if $\bar{\psi}$ is the function provided by the Hölder Gap Lemma ??. On the other hand, by continuity of α in $H_{\nu'}(k)$,

$$\begin{aligned}\alpha\left(\bar{\psi}\right) &= \psi\left(x_k^+\right)\alpha(k) + \sum_d \left(\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right)\alpha(\eta_d) \\ &= \psi\left(x_k^+\right)\alpha(k) + \sum_d h_\alpha(d)\left(\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right).\end{aligned}$$

by definition of the α–height $h_\alpha(d)$. Since $\alpha(\psi) = \alpha\left(\bar{\psi}\right)$, this completes the proof of the Gap Formula, namely of Theorem ??. □

8.6. Transverse Hölder distributions and transverse cocycles

Transverse Hölder distributions are analytic objects, and $\mathbb{R}$–valued transverse cocycles are of a combinatorial nature. However, we now prove that these two notions are essentially equivalent. The correspondence is provided by the following elementary lemma.

LEMMA. *Let α be a transverse Hölder distribution for the geodesic lamination λ. Then, there is an $\mathbb{R}$–valued transverse cocycle which associates to each arc k transverse to λ the α–integral $\alpha(k) = \int_k 1\, d\alpha \in \mathbb{R}$ of the constant function 1 on k.*

PROOF. If the two transverse arc k, k' are homotopic respecting λ, the λ–invariance of the transverse distribution α immediately shows that $\alpha(k) = \alpha(k')$.

If we split the transverse arc k as the union of two transverse arcs k_1 and k_2 with disjoint interiors, consider a partition of unity defined by two Hölder continuous functions $\varphi_1, \varphi_2 : k \to \mathbb{R}$ such that $\varphi_1 + \varphi_2 = 1$ and such that each φ_i coincides with the constant function 1 on $\lambda \cap k_i$; then, $\alpha(k) = \alpha(\varphi_1) + \alpha(\varphi_2) = \alpha(k_1) + \alpha(k_2)$ by the Support Lemma ??. If follows that the number $\alpha(k) \in \mathbb{R}$ is an additive function of the transverse arc k.

Therefore, the rule $k \mapsto \alpha(k) \in \mathbb{R}$ is a transverse cocycle for λ. □

THEOREM. *If we associate, as in Lemma ?? above, an $\mathbb{R}$–valued transverse cocycle to each transverse Hölder distribution for the geodesic lamination λ, this defines a one-to-one correspondence between the space of transverse Hölder distributions for λ and the space $\mathcal{H}(\lambda; \mathbb{R})$ of $\mathbb{R}$–valued transverse cocycles for λ.*

PROOF. Let k be an oriented arc which is transverse to k. For each component d of $k - \lambda$, pick a sub-arc k_d of k that joins the negative end point of k to some point of d. By definition of the α–length and by the Support Lemma ??, $\alpha(k_d) = h_\alpha(d)$ for every transverse distribution α for λ. Then, we can reprase the Gap Formula by saying that

$$\alpha(\psi) = \alpha(k)\,\psi\left(x_k^+\right) + \sum_d \alpha(k_d)\left(\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right)$$

for every Hölder continuous function $\psi : k \to \mathbb{R}$. In particular, the Hölder distribution defined by α on k is completely determined by the transverse cocycle associated to α.

Conversely, let α be an $\mathbb{R}$–valued transverse cocycle for λ. If k is an arc transverse to λ, we define a Hölder distribution on k by the Gap Formula. Namely, choose an orientation for k, and associate to each Hölder continuous function $\psi : k \to \mathbb{R}$ the quantity

$$\alpha(\psi) = \alpha(k)\,\psi\left(x_k^+\right) + \sum_d \alpha(k_d)\left(\psi\left(x_d^-\right) - \psi\left(x_d^+\right)\right)$$

with the same notation as above. The fact that the series converges and is a continuous function of ψ is proved by the arguments of the proof of Lemma ?? need $\alpha(k_d) = O(r(d))$, borrow from Chap. 6.

If the oriented arc k' is homotopic to k respecting λ, the homotopy associates to each component d of $k - \lambda$ a component d' of $k' - \lambda'$, and to each function $\psi : k \to \mathbb{R}$ a function $\psi' : k' \to \mathbb{R}$. For the obvious notation, $\psi'\left(x_{d'}^{\pm}\right) = \psi\left(x_d^{\pm}\right)$ by construction, and $\alpha\left(k'_{d'}\right) = \alpha(k_d)$ by λ–invariance of the transverse cocycle, so that $\alpha(\psi') = \alpha(\psi)$. The Hölder distribution α is therefore invariant under homotopy respecting λ.

However, it is less obvious that the formula for $\alpha(\psi)$ is unchanged if we reverse the orientation of k. Indeed, reversing the orientation of k replaces $x_d^{\pm}$ by $x_d^{\mp}$, x_k^{+} by the negative en point x_k^{-}, and the arc k_d by the closure k'_d of $k - k_d$. Note that $\alpha(k_d) + \alpha(k'_d) = \alpha(k)$ by additivity of the transverse cocycle α. Therefore, the formula associated to the opposite orientation of k gives

$$
\begin{aligned}
&\alpha(k)\,\psi\left(x_k^{-}\right) + \sum_d \alpha\left(k'_d\right)\left(\psi\left(x_d^{+}\right) - \psi\left(x_d^{-}\right)\right)\\
&= \alpha(k)\,\psi\left(x_k^{-}\right) + \sum_d \left(\alpha(k) - \alpha(k_d)\right)\left(\psi\left(x_d^{+}\right) - \psi\left(x_d^{-}\right)\right)\\
&= \alpha(k)\,\psi\left(x_k^{+}\right) + \sum_d \alpha(k_d)\left(\psi\left(x_d^{-}\right) - \psi\left(x_d^{+}\right)\right)\\
&\qquad + \alpha(k)\left(\psi\left(x_k^{-}\right) - \psi\left(x_k^{+}\right) - \sum_d \left(\psi\left(x_d^{-}\right) - \psi\left(x_d^{+}\right)\right)\right)\\
&= \alpha(k)\,\psi\left(x_k^{+}\right) + \sum_d \alpha(k_d)\left(\psi\left(x_d^{-}\right) - \psi\left(x_d^{+}\right)\right)\\
&= \alpha(k)
\end{aligned}
$$

by application of the Gap Lemma ?? to the component of $k - \lambda$ that contains x_k^{-}. As a consequence, $\alpha(\psi)$ is independent of the choice of the orientation of k.

This enables us to associate a transverse Hölder distribution to each transverse cocycle. Applying the formula to the constant function 1, we see that this operation is really the inverse of the operation associating a transverse cocycle to each transverse Hölder distribution. Therefore, this operation defines a one-to-one correspondence between transverse Hölder distributions and $\mathbb{R}$–valued transverse cocycles for λ. □

An immediate corollary of Theorem ?? and of Theorem 6.?? is that the transverse Hölder distributions for the geodesic lamination λ form a finite dimensional vector space, of dimension $|\chi(\lambda)| + n_0(\lambda)$.

Incidentally, Theorem ?? should explain the use of the letter $\mathcal{H}$ in the notation $\mathcal{H}(\lambda; \mathbb{R})$ for the space of transverse cocycles for λ. Indeed, transverse Hölder distributions were discovered first, and the equivalent notion of transverse cocycle was introduced only after the combinatorial nature of transverse Hölder distributions had emerged.

8.7. Transverse Hölder distributions as derivatives of transverse measures

Let $\alpha_t \in \mathcal{ML}(S)$, $t \in [0, \varepsilon[$, be a 1–parameter family of measured geodesic laminations, which has a tangent vector $\dot{\alpha}_0$ at $t = 0$. In §§6.?? and 8.??, we saw how to associate to $\dot{\alpha}_0$ a geodesic lamination with a transverse Hölder distribution. For applications, we need the stronger property that the transverse Hölder distribution $\dot{\alpha}_0$ is really the derivative of the transverse measures α_t. Namely, interpreting α_t and $\dot{\alpha}_0$ as Hölder geodesic currents, we want to prove that $\dot{\alpha}_0(\varphi)$ is equal to the right derivative of $t \mapsto \alpha_t(\varphi)$ at $t = 0$ for every Hölder continuous function $\varphi : G(\widetilde{S}) \to \mathbb{R}$ with compact support. This section is devoted to the proof of this property, which will require analytic arguments which are somewhat more elaborate than those of used earlier in this chapter.

We consider a 1–parameter family α_t, $t \in [0, t_0]$, of measured geodesic laminations and we assume that this path in $\mathcal{ML}(S)$ has a tangent vector $\dot{\alpha}_0$ at α_0. We want to show that the transverse Hölder distribution associated to $\dot{\alpha}_0$ by Propositions 6.?? and 8.?? is really the derivative of the transverse measures α_t. Namely, interpreting α_t and $\dot{\alpha}_0$ as Hölder geodesic currents, we want to prove that, for every Hölder continuous function $\varphi : G(\widetilde{S}) \to \mathbb{R}$ with compact support, the number $\dot{\alpha}_0(\varphi)$ associated to φ by $\dot{\alpha}_0$ is equal to the right derivative $\frac{\partial}{\partial t^+}\alpha_t(\varphi)_{|t=0}$. Incidentally, a side benefit of this will be to remove a current ambiguity of our 'dot' notation, which does not distinguish between $\frac{\partial}{\partial t^+}(\alpha_t)_{|t=0}(\varphi)$ and $\frac{\partial}{\partial t^+}\alpha_t(\varphi)_{|t=0}$.

Theorem. *Let $\alpha_t \in \mathcal{ML}(S)$, $t \in [0, \varepsilon[$, be a 1–parameter family of measured geodesic lamination which has a right derivative $\dot{\alpha}_0$ at $t = 0$. This derivative $\dot{\alpha}_0$ is a tangent vetor to $\mathcal{ML}(S)$ and, by Theorems 6.?? and 8.??, can be identified to a Hölder geodesic current. Then, for every Hölder continuous function $\varphi : G\left(\widetilde{S}\right) \to \mathbb{R}$, the derivative $\frac{\partial}{\partial t^+}\alpha_t(\varphi)_{|t=0}$ exists and is equal to the number $\dot{\alpha}_0(\varphi)$ associated to φ by this Hölder geodesic current $\dot{\alpha}_0$.*

Proof. The proof will take a while. Because of the above mentioned ambiguity in the notation, we will reserve the symbol $\dot{\alpha}_0$ to denote the tangent vector, its associated transverse cocycle, its associated transverse Hölder distribution or its associated Hölder geodesic current. In particular, $\dot{\alpha}_0(\varphi)$ will here denote the $\dot{\alpha}_0$–integral of φ, and not the derivative $\frac{\partial}{\partial t^+}\alpha_t(\varphi)_{|t=0}$.

By Proposition 3.??, we can endow S with a negatively curved metric m with geodesic boundary and we can find a fattened train track Φ that strongly carries the m_0–geodesic laminations λ_t underlying the α_t, after a possible restriction of the domain $[0, \varepsilon[$ of t. In addition, by the proof of Proposition 3.?? (see in particular ??), we can choose Φ so that it strongly carries any m_0–geodesic lamination containing λ_0.

We will henceforth use m_0 as our auxiliary metric, so that 'geodesic' will always mean 'm_0–geodesic'.

Considering the α_t as measure geodesic currents, all geodesics of the support $\widetilde{\lambda}_t \subset G(\widetilde{S})$ of α_t are strongly carried by the pre-image $\widetilde{\Phi}$ of Φ in $\widetilde{S}$. In the space of those geodesics of $\widetilde{S}$ which are strongly carrried by Φ, those geodesics which meet a given tie of $\widetilde{\Phi}$ form an open subset of this space. Therefore, using a suitable partition of unity, we can restrict attention to the case where there is a generic tie k_0 of $\widetilde{\Phi}$ which transversely meets every geodesic of the support of φ.

Consider the essential support λ_{0+} of the α_t as t tends to 0^+, as defined in §5.??. Since λ_{0+} contains λ_0, it is strongly carried by Φ by choice of this fattened train track. Let $\widetilde{\lambda}_{0+}$ denote the pre-image of λ_{0+} in $\widetilde{S}$.

Arbitrarily choose an orientation for k_0. If d is a component of $k_0 - \widetilde{\lambda}_{0+}$, recall that the ***height*** of d for the transverse cocycle $\dot{\alpha}_0$ is the number $h_{\dot{\alpha}_0}(d) = \dot{\alpha}_0(k_d)$, where k_d is an arbitrary sub-arc of k_0 joining the negative end point of k_0 to a point in d. Also, let g_d^+ (resp. g_d^-) be the leaf of λ_{0+} passing through the positive (resp. negative) end point of d; note that $g_d^\pm$ will be undefined if the corresponding end point of d is also an end point of k_0. Then, because every geodesic of the support of φ transversely crosses k_0, the Gap Formula of Theorem 8.?? and the definition of the Hölder geodesic current associated to a transverse Hölder distribution (Proposition 5.??) show that the $\dot{\alpha}_0$–integral of φ is equal to

$$\dot{\alpha}_0(\varphi) = \sum_d h_{\dot{\alpha}_0}(d)\left(\varphi\left(g_d^-\right) - \varphi\left(g_d^+\right)\right)$$

where the sum is over all components d of $k_0 - \widetilde{\lambda}_{0+}$ and where, by convention, $\varphi\left(g_d^\pm\right) = 0$ for the two terms where the geodesic $g_d^\pm$ is undefined.

We need to estimate the rate of convergence of the above sum, in terms of the depths $r(d)$ of the gaps d.

LEMMA. *There is a constant $A > 0$ such that, for every integer $r \geqslant 1$,*

$$\dot{\alpha}_0(\varphi) = \sum_{r(d)<r} h_{\dot{\alpha}_0}(d)\left(\varphi\left(g_d^-\right) - \varphi\left(g_d^+\right)\right) + \|\dot{\alpha}_0\|_\Phi \|\varphi\|_\nu O\left(e^{-\nu A r}\right)$$

where $\|\dot{\alpha}_0\|_\Phi$ denotes the maximum of the edge weights $\dot{\alpha}_0(e)$ as e ranges over the edges of Φ, where $\|\varphi\|_\nu$ is the Hölder norm of exponent ν, and where the constant A and the constant hidden in the symbol $O(\ )$ is independent of α.

PROOF. This is a consequence of the proof of the convergence of the Gap Formula, in ??. Indeed,

$$\begin{aligned}\varphi\left(g_d^-\right) - \varphi\left(g_d^+\right) &= \|\varphi\|_\nu O\left(d\left(g_d^+, g_d^-\right)^\nu\right) \\ &= \|\varphi\|_\nu O\left(d\left(k_0 \cap g_d^+, k_0 \cap g_d^-\right)^\nu\right) \\ &= \|\varphi\|_\nu O\left(e^{-\nu A' r}\right)\end{aligned}$$

for some constant $A' > 0$, where the first equality comes from the definition of $\|\varphi\|_\nu$, the second one from Proposition ??, and the last one from Proposition ??. Also, $|h_{\dot{\alpha}_0}(d)| = \|\dot{\alpha}_0\|_\Phi O(r(d))$ by Proposition ??. Since the number of gaps d with $r(d) = r$ is bounded independent of r (Proposition ??), it follows that for every r

$$\sum_{r(d)=r} h_{\dot{\alpha}_0}(d)\left(\varphi\left(g_d^-\right) - \varphi\left(g_d^+\right)\right) = \|\dot{\alpha}_0\|_\Phi \|\varphi\|_\nu O\left(re^{-\nu A' r}\right) = \|\dot{\alpha}_0\|_\Phi \|\varphi\|_\nu O\left(e^{-\nu A r}\right)$$

for any $A < A'$, where the sum is over those d with $r(d) = r$. Writing

$$\dot{\alpha}_0(\varphi) = \sum_{r=1}^{\infty} \sum_{r(d)=r} h_{\dot{\alpha}_0}(d) \left(\varphi\left(g_d^-\right) - \varphi\left(g_d^+\right)\right),$$

the estimate then follows from the formula for the difference between a geometric series and its partial sums. □

We now investigate the right derivative

$$\frac{\partial}{\partial t^+} \alpha_t(\varphi)_{|t=0} = \lim_{t \to 0^+} \left(\alpha_t(\varphi) - \alpha_0(\varphi)\right)/t,$$

using edge paths.

Given $r \geqslant 1$, let Γ_r denote the set of all edge paths $\gamma = \langle e_{-r}, e_{-r+1}, \ldots, e_0, \ldots, e_{r-1}, e_r \rangle$ of $\widetilde{\Phi}$ which are realized by some geodesic g_γ, where e_0 is the edge containing k_0 and where we identify two edge paths when they differ by reversal of orientation. For every $\gamma \in \Gamma_r$, pick a geodesic g_γ realizing γ. ■weakly, strongly?

For each edge path $\gamma \in \Gamma_r$, let $G_\gamma \subset G\left(\widetilde{S}\right)$ be the set of those m_0–geodesics which realize γ. The G_γ with $\gamma \in \Gamma_r$ are pairwise disjoint, and their diameter decrease exponentially with r by ??. We can therefore use them to approximate by Riemann sums the integral of φ for the signed measure $(\alpha_t - \alpha_0)/t$. This gives

$$\frac{\alpha_t(\varphi) - \alpha_0(\varphi)}{t} = \lim_{r \to \infty} \sum_{\gamma \in \Gamma_r} \frac{\alpha_t(\gamma) - \alpha_0(\gamma)}{t} \varphi(g_\gamma)$$

for every $t > 0$.

We want to estimate the rate of convergence of this limit. For this, we first show that the number of terms contributing to the above sum has polynomial growth in r.

For $\gamma \in \Gamma_r$, let $\sigma(\gamma) \in \Gamma_{r-1}$ be the edge path obtained from γ by removing its two end edges. Note that, for every $\gamma \in \Gamma_r$, G_γ is the disjoint union of those $G_{\gamma'}$ with $\gamma' \in \Gamma_{r+1}$ and $\sigma(\gamma') = \gamma$. It follows that $\alpha(\gamma) = \sum_{\sigma(\gamma')=\gamma} \alpha(\gamma')$ for every $\gamma \in \Gamma_r$ and every measured geodesic lamination α which is strongly carried by Φ. Also, since φ is Hölder continuous of exponent ν, $\varphi(g_\gamma) - \varphi\left(g_{\sigma(\gamma)}\right) = \|\varphi\|_\nu O\left(e^{-\nu A r}\right)$ for some $A > 0$. In addition, by Proposition ??, $\alpha_t(\gamma) - \alpha_0(\gamma) = \|\alpha_t - \alpha_0\|_\Phi O(r)$ if $\gamma \in \Gamma_r$. ■details? Make it a lemma? Finally, Proposition ?? shows that the number of those $\gamma \in \Gamma_r$ with $\alpha_t(\gamma) \neq 0$ for some t is an $O(r^n)$ for some n, since such an edge path is realized by a geodesic of $\widetilde{S}$ which projects to a simple geodesic of S. Combining these facts and regrouping terms, it follows that the difference

$$\sum_{\gamma' \in \Gamma_{r+1}} \frac{\alpha_t(\gamma') - \alpha_0(\gamma')}{t} \varphi(g_{\gamma'}) - \sum_{\gamma \in \Gamma_r} \frac{\alpha_t(\gamma) - \alpha_0(\gamma)}{t} \varphi(g_\gamma)$$

is equal to

$$\sum_{\gamma' \in \Gamma_{r+1}} \frac{\alpha_t(\gamma') - \alpha_0(\gamma')}{t} \left(\varphi(g_{\gamma'}) - \varphi\left(g_{\sigma(\gamma')}\right)\right) = \frac{\|\alpha_t - \alpha_0\|_\Phi}{t} \|\varphi\|_\Phi O\left(r^n e^{-\nu A r}\right).$$

Summing these differences from r to infinity, we conclude that

$$\begin{aligned}\frac{\alpha_t(\varphi)-\alpha_0(\varphi)}{t} &= \lim_{r\to\infty}\sum_{\gamma\in\Gamma_r}\frac{\alpha_t(\gamma)-\alpha_0(\gamma)}{t}\varphi(g_\gamma)\\ &= \sum_{\gamma\in\Gamma_r}\frac{\alpha_t(\gamma)-\alpha_0(\gamma)}{t}\varphi(g_\gamma)+\frac{\|\alpha_t-\alpha_0\|_\Phi}{t}\|\varphi\|_\nu\, O\left(r^n e^{-\nu A r}\right).\end{aligned}$$

We can then let t tend to 0^+. By finiteness of Γ_r, we conclude:

LEMMA 11. *For every $r \geqslant 1$, the supremum limit and the infimum limit of $(\alpha_t(\varphi)-\alpha_0(\varphi))/t$ as t tends to 0^+ are both of the form*

$$\sum_{\gamma\in\Gamma_r}\dot\alpha_0(\gamma)\varphi(g_\gamma)+\|\dot\alpha_0\|_\Phi\|\varphi\|_\nu\, O\left(r^n e^{-\nu A r}\right).$$

□

Let us connect the estimates of Lemmas ?? and ??.

LEMMA 12.

$$\sum_{\gamma\in\Gamma_r}\dot\alpha_0(\gamma)\varphi(g_\gamma)=\sum_{r(d)<r}h_{\dot\alpha_0}(d)\left(\varphi\left(g_d^-\right)-\varphi\left(g_d^+\right)\right)+\|\dot\alpha_0\|_\Phi\|\varphi\|_\nu\, O\left(re^{-\nu A r}\right).$$

PROOF. If $\dot\alpha_0(\gamma)\neq 0$ note that, by the Tracking Lemma 1, γ is realized by a geodesic of $\widetilde\lambda_{0^+}$. Therefore, if Γ_r' consists of those $\gamma\in\Gamma_r$ which are realized by geodesics of $\widetilde\lambda_{0^+}$,

$$\sum_{\gamma\in\Gamma_r}\dot\alpha_0(\gamma)\varphi(g_\gamma)=\sum_{\gamma\in\Gamma_r'}\dot\alpha_0(\gamma)\varphi(g_\gamma).$$

For every $\gamma\in\Gamma_r'$, the intersection points of k_0 with the geodesics of $\widetilde\lambda_{0^+}$ realizing γ form a closed interval in $k_0\cap\widetilde\lambda_{0^+}$, for the ordering of $k_0\cap\widetilde\lambda_{0^+}$ induced by the orientation of k_0. It follows that there are two components d_γ^+ and d_γ^- of $k_0-\widetilde\lambda_{0^+}$ such that the geodesics of $\widetilde\lambda_{0^+}$ realizing γ are exactly those which hit k_0 above d_γ^- and below d_γ^+.

Note that, for $d=d_\gamma^+$, the geodesic g_d^- realizes γ and g_d^+ does not; it follows that $r\left(d_\gamma^+\right)<r$. Similarly, $r\left(d_\gamma^-\right)<r$. Conversely, every component d of $k_0-\widetilde\lambda_{0^+}$ with $r(d)<r$ is equal to $d=d^-_{\gamma_d^+}=d^+_{\gamma_d^-}$ where γ_d^+ and γ_d^- are the edge paths in Γ_r respectively realized by the geodesics g_d^+ and g_d^- (if defined).

By definition of $h_{\dot\alpha_0}$, the difference $h_{\dot\alpha_0}\left(d_\gamma^+\right)-h_{\dot\alpha_0}\left(d_\gamma^-\right)$ is equal to $\dot\alpha_0(k_\gamma)$ where k_γ is a sub-arc of k_0 joining a point of d_γ^- to a point of d_γ^+. However, by definition of d_γ^- and d_γ^+, those leaves of $\widetilde{\lambda}_{O^+}$ which cut k_γ are exactly those which realize the edge path γ. It follows that $h_{\dot\alpha_0}\left(d_\gamma^+\right) - h_{\dot\alpha_0}\left(d_\gamma^-\right) = \dot\alpha_0\left(\gamma\right)$.

As a consequence,

$$\begin{aligned}\sum_{\gamma\in\Gamma_r} \dot\alpha_0\left(\gamma\right)\varphi\left(g_\gamma\right) &= \sum_{\gamma\in\Gamma_r} h_{\dot\alpha_0}\left(d_\gamma^+\right)\varphi\left(g_\gamma\right) - \sum_{\gamma\in\Gamma_r} h_{\dot\alpha_0}\left(d_\gamma^-\right)\varphi\left(g_\gamma\right)\\ &= \sum_{r(d)<r} h_{\dot\alpha_0}\left(d\right)\varphi\big(g_{\gamma_d^-}\big) - \sum_{r(d)<r} h_{\dot\alpha_0}\left(d\right)\varphi\big(g_{\gamma_d^+}\big)\\ &= \sum_{r(d)<r} h_{\dot\alpha_0}\left(d\right)\left(\varphi\big(g_{\gamma_d^-}\big) - \varphi\big(g_{\gamma_d^+}\big)\right)\\ &= \sum_{r(d)<r} h_{\dot\alpha_0}\left(d\right)\left(\varphi\left(g_d^-\right) - \varphi\left(g_d^+\right)\right) + \left\|\varphi\right\|_\nu O\Big(e^{-\nu A r}\sum_{r(d)<r}\big|\dot h_0\left(d\right)\big|\Big)\end{aligned}$$

since g_d^- realizes the same path $\gamma_d^- \in \Gamma_r$ as $g_{\gamma_d^-}$ and is therefore at distance $O\left(e^{-\nu A r}\right)$ from $g_{\gamma_d^-}$, and since the same property holds for g_d^+ and $g_{\gamma_d^+}$. (And using the convention that $\varphi\left(g_d^\pm\right) = 0$ when the geodesic $g_d^\pm$ is not defined.)

By Propositions ?? and ??, $h_{\dot\alpha_0}\left(d\right) = \left\|\dot\alpha_0\right\|_\Phi O\left(r\left(d\right)\right)$ and the number of gaps d with $r\left(d\right) = r$ is uniformly bounded. It follows that $\sum_{r(d)<r}\left|h_{\dot\alpha_0}\left(d\right)\right| = \left\|\dot\alpha_0\right\|_\Phi O\left(r\right)$, which completes the proof of Lemma 12. □

Combining Lemmas 11, 12 and Complement 9, we conclude that the supremum limit and the infimum limit of $\left(\alpha_t\left(\varphi\right) - \alpha_0\left(\varphi\right)\right)/t$ as t tends to 0^+ are both of the form

$$\sum_d \dot h_0\left(d\right)\left(\varphi\left(g_d^-\right) - \varphi\left(g_d^+\right)\right) + \left\|\dot\alpha_0\right\|_\Phi \left\|\varphi\right\|_\nu O\left(r^n e^{-\nu A r}\right).$$

Since this is true for every $r \geqslant 1$, it follows by letting r tend to ∞ that these supremum limit and infimum limit are both equal to the sum of the above series, namely that

$$\begin{aligned}\frac{\partial}{\partial t^+}\alpha_t\left(\varphi\right)_{|t=0} &= \lim_{t\to 0^+}\left(\alpha_t\left(\varphi\right) - \alpha_0\left(\varphi\right)\right)/t\\ &= \sum_d \dot h_0\left(d\right)\left(\varphi\left(g_d^-\right) - \varphi\left(g_d^+\right)\right)\\ &= \dot\alpha_0\left(\varphi\right),\end{aligned}$$

which concludes the proof of Theorem 7. □

8.8. A first application: The tangent map of the length function

Consider a metric m of negative curvature on the surface S for which the boundary ∂S is totally geodesic. In §4.??, we associated to m a ***length function*** $l_m : \mathcal{C}\left(S\right) \to \mathbb{R}^+$. In particular, we can restrict l_m to a length function $l_m : \mathcal{ML}\left(S\right) \to \mathbb{R}^+$. The machinery developped in this chapter provides an immediate computation of the tangent map $T_{\alpha_0} l_m : T_{\alpha_0}\mathcal{ML}(S) \to \mathbb{R}$ of this length function at each $\alpha_0 \in \mathcal{ML}(S)$.

Indeed, by the definition of tangent maps (see §??), if $\dot\alpha_0$ is a vector tangent to $\mathcal{ML}(S)$ at α_0, its image $T_{\alpha_0} l_m(\dot\alpha_0)$ under the tangent map is the right derivative $\frac{\partial}{\partial t^+} l_m(\alpha_t)_{|t=0}$ where $t \mapsto \alpha_t$ is an arbitrary curve in $\mathcal{ML}(S)$ that passes through α_0 with right derivative $\dot\alpha_0$ at $t = 0$.

, let us recall how the length $l_m(\alpha)$ of a geodesic current $\alpha \in \mathcal{C}(S)$ is defined: Consider the projective tangent bundle $PT(S)$ of S, consisting of pairs (x, d) where $x \in S$ and d is a line passing through the origin in the tangent plane of S at x. If we endow S with the metric m, every m–geodesic can be lifted to $PT(S)$ by considering the tangent line at each of its points. Let $PT_0(S) \subset PT(S)$ be the union of the lifts of those geodesics of S which do not transversely hit the boundary ($PT_0(S) = PT(S)$ when $\partial S = \emptyset$). These lifts of geodesics foliate $PT_0(S)$, in the sense that every point of $PT_0(S)$ has a neighborhood homeomorphic to some space $D \times [0, 1]$ where each arc $* \times [0, 1]$ is contained in a lift of geodesic. This foliation $\mathcal{F}$ is the ***geodesic foliation*** of $PT_0(S)$. Locally, a point of $PT_0(S)$ is characterized by the leaf of $\mathcal{F}$ containing it, and by where it sits on that leaf. Consequently, given $\alpha \in \mathcal{C}(S)$, we can consider on $PT_0(S)$ the measure which is locally the product of α and of the length measure deposited by m on the leaves of $\mathcal{F}$. Then, $l_m(\alpha)$ is defined as the total mass of this measure.

More precisely, choose a continuous partition of unity $\xi_i : PT(S) \to \mathbb{R}$, $i = 1$, … , n, such that $\sum_{i=1}^n \xi_i = 1$ and such that the support of each ξ_i is contained in the interior of a flow box B_i for $\mathcal{F}$ in $PT_0(S)$; namely, there is for each i a homeomorphism $\sigma_i : D_i \times [0, 1] \to B_i$ for some space D_i, such that each $\sigma_i(g \times [0, 1])$ is contained in a leaf of $\mathcal{F}$. Lifting the situation to the universal covering $\widetilde{S}$ and assuming the B_i sufficiently small, we can identify D_i to a subset of $G(\widetilde{S})$, well-defined modulo the action of $\pi_1(S)$ on $G(\widetilde{S})$. Since $\alpha \in \mathcal{C}(S)$ is invariant under the action of $\pi_1(S)$, it follows that α induces a measure α on D_i. Then,

$$l_m(\alpha) = \sum_{i=1}^n \int_{D_i} \int_0^1 \xi_i(\sigma_i(g, t))\, dm(t)\, d\alpha(g) = \sum_{i=1}^n \alpha(\varphi_i)$$

where $\varphi_i : G(\widetilde{S}) \to \mathbb{R}$ is the continuous map with compact support defined by $\varphi_i(g) = \int_0^1 \xi_i(\sigma_i(g, t))\, dm(t)$ when $g \in D_i \subset G(\widetilde{S})$ and $\varphi_i(g) = 0$ otherwise.

It turns out that the last term of this formula makes sense when α is only a Hölder distribution, provided that we choose the σ_i and ξ_i Hölder continuous (which we can always assume). Indeed, under this regularity hypothesis, it follows from the definition of the Hölder structure of $G(\widetilde{S})$ that the maps φ_i are Hölder continuous. Therefore, given $\alpha \in \mathcal{H}(S)$, we can define

$$l_m(\alpha) = \sum_{i=1}^n \alpha(\varphi_i).$$

By linearity of the formula and by invariance of α under the action of $\pi_1(S)$, this is clearly independent of the choice of the ξ_i, σ_i and lifts of B_i to $PT(\widetilde{S})$. Also, this $l_m(\alpha)$ is a continuous function of α by definition of the topology of $\mathcal{H}(S)$. This proves:

THEOREM 24. *Given a metric m of negative curvature and with totally geodesic boundary on S, there is a continuous linear function $l_m : \mathcal{H}(S) \to \mathbb{R}$ such that, when $\alpha \in \mathcal{C}(S) \subset \mathcal{H}(S)$ corresponds to a closed geodesic λ endowed with a weight a, $l_m(\alpha)$ is equal to a times the length of this m–geodesic.* □

In particular, by continuity and linearity of l_m:

COROLLARY 25. *If $t \mapsto \alpha_t$ is a 1–parameter family of measured geodesic laminations which admits a right derivative $\dot{\alpha}_0 \in \mathcal{H}(S)$ at $t = 0$, then*

$$\frac{\partial}{\partial t^+} l_m(\alpha_t)_{|t=0} = l_m(\dot{\alpha}_0).$$

In particular, the length function $l_m : \mathcal{ML}(S) \to \mathbb{R}^+$ has a differential at each point of $\mathcal{ML}(S)$, and this differential is linear on each face of the piecewise linear structure of $\mathcal{ML}(S)$. □

In [**Bo5**], we give an explicit formula, based on the Thurston symplectic form on $\mathcal{H}(\lambda)$, which expresses this variation $l_m(\dot{\alpha}_0)$ in terms of the shear coordinates associated to m and to any maximal geodesic lamination λ containing the support of $\dot{\alpha}_0$.

CHAPTER IX

TEICHMÜLLER SPACE

9.1. Hyperbolic metrics

9.2. The Fricke-Klein coordinates

9.3. The shear coordinates associated to a finite ideal triangulation

9.4. The shearing transverse cocycle

On the surface S, consider a hyperbolic metric m and a maximal geodesic lamination λ. Lift the situation to the universal covering $\widetilde{S}$, where λ has preimage $\widetilde{\lambda}$. Recall that a ***plaque*** of $\widetilde{S}-\widetilde{\lambda}$ is the closure in $\widetilde{S}$ of a component of $\widetilde{S}-\widetilde{\lambda}$. Since λ is maximal, each plaque of $\widetilde{S}-\widetilde{\lambda}$ is an ideal triangle, namely a hyperbolic triangle with its vertices at infinity.

Given two leaves g and h of $\widetilde{\lambda}$, the geodesic lamination $\widetilde{\lambda}$ gives a preferred isometry $\theta^{gh}: g \to h$ defined as follows.

Indeed, consider the closure Σ of the component of $\widetilde{S}-g\cup h$ that is adjacent to both g and h. The leaves of $\widetilde{\lambda}$ that separate g from h provide a partial foliation of the strip Σ, which can be uniquely extended to a global foliation $\mathcal{G}$ of Σ by geodesics as follows: Since λ is maximal, every component of the complement of these leaves of $\widetilde{\lambda}$ in Σ is a hyperbolic wedge, bounded by two asymptotic geodesics; and such a wedge admits a unique foliation by geodesics, all asymptotic on one side. An estimate in hyperbolic plane geometry shows that two disjoint geodesics which pass through two nearby points do so with directions differing by at most a constant times the distance between these two points (see for instance [**CEG**, §5.2.6]). It follows that the normals to the leaves of $\mathcal{G}$ form a Lipschitz vector field on Σ. We can therefore integrate this vector field, to get a foliation $\mathcal{H}$ of Σ which is everywhere orthogonal to $\mathcal{G}$. Each leaf of $\mathcal{H}$ goes from g to h, and this defines a map $\theta^{gh}: g \to h$. Also, $\mathcal{H}$ respects distances on the leaves of $\mathcal{G}$ by the formula for the first variation, and it follows that θ^{gh} is an isometry. Note that $\theta^{hg} = \left(\theta^{gh}\right)^{-1}$.

Now, consider two plaques P and Q of $\widetilde{S}-\widetilde{\lambda}$. Let g be the leaf of $\widetilde{\lambda}$ in the boundary of P which is closest to Q, and let h be the leaf in the boundary of Q which is closest to P. Orient h as part of the boundary of Q, with the orientation induced by the orientation of S. The plaque Q also determines a preferred base point on h, namely the projection to h of the third vertex of the ideal triangle Q. Similarly, the plaque P determines an orientation and a base point on the geodesic g. For the oriented isometric parametrization of h by $\mathbb{R}$ which sends 0 to the base point, let $\sigma(P,Q) \in \mathbb{R}$ be the coordinate of the image of the base point of g under $\theta^{gh} : g \to h$. In other words, for the isometric parametrization of g and h defined by the choices of orientation and base point, the isometry $\theta^{gh} : g \to h$ corresponds to the map $t \mapsto \sigma(P,Q) - t$.

Since $\theta^{hg} = \left(\theta^{gh}\right)^{-1}$, $\sigma(Q,P)$ is equal to $\sigma(P,Q)$.

Also, consider three plaques P, Q and R of $\widetilde{S} - \widetilde{\lambda}$ such that Q separates P from R. Let g be the leaf of $P \cap \widetilde{\lambda}$ that is closest to Q, h the leaf of $Q \cap \widetilde{\lambda}$ closest to P, k the leaf of $Q \cap \widetilde{\lambda}$ closest to R and l the leaf of $R \cap \widetilde{\lambda}$ closest to Q. The map θ^{gl} decomposes as $\theta^{gl} = \theta^{kl} \circ \theta^{hk} \circ \theta^{gh}$. Since Q admits an isometry exchanging h and k, the orientation-reversing map θ^{hk} sends the base point of h to the base point of k. It immediately follows that $\sigma(P,R) = \sigma(P,Q) + \sigma(Q,R)$.

Therefore, the rule $(P,Q) \mapsto \sigma(P,Q)$ defines an $\mathbb{R}$–valued transverse cocycle σ for λ, in the sense of §1. This transverse cocycle is the ***shearing cocycle*** of the hyperbolic metric m.

If we change the metric m to a metric m' by an isotopy $\varphi : S \to S$, then φ sends λ to the corresponding m'–geodesic lamination λ'. It immediately follows that m and m' have the same shearing cocycles. Therefore, the shearing cocycle σ depends only on the class of m in $\mathcal{T}(S)$.

We can give another description of the number $\sigma(P,Q)$, which will be convenient later on.

Let $\widetilde{\lambda}_{PQ}$ be the set of those leaves of $\widetilde{\lambda}$ which separate P from Q, and orient these leaves to the left as seen from P. Let k be an oriented arc transverse to $\widetilde{\lambda}_{PQ}$ joining P to Q.

For each component d of $k - \widetilde{\lambda}$, let x_d^+ and x_d^- be its positive and negative end points, respectively. If d is not one of the components d^+ and d^- containing the positive and negative end points of d, respectively, then $x_d^\pm$ is contained in a leaf $g_d^\pm$ of $\widetilde{\lambda}_{PQ}$ which is adjacent to a component of $\widetilde{S} - \widetilde{\lambda}$. As before, the component of $\widetilde{S} - \widetilde{\lambda}$ containing d determines a base point on $g_d^\pm$, namely the projection of the third vertex. Let $f : g_d^\pm \to \mathbb{R}$ be the unique oriented isometry sending this base point to 0. This associates two numbers $f\left(x_d^+\right)$ and $f\left(x_d^-\right)$ to each component d of $k - \widetilde{\lambda}_{PQ}$ which is different from the end components d^+ and d^-. When $d = d^+$ or d^-, we can similarly define $f\left(x_{d^+}^-\right)$ and $f\left(x_{d^-}^+\right)$.

Lemma 7. *With the above data,*

$$\sigma(P,Q) = f\left(x_{d^-}^+\right) - f\left(x_{d^+}^-\right) + \sum_{d \neq d^+, d^-} \left(f\left(x_d^+\right) - f\left(x_d^-\right) \right)$$

where the sum is taken over all components d of $\widetilde{S} - \widetilde{\lambda}_{PQ}$ which are different from the end components d^+, d^-.

Proof. We can parametrize the component Σ of $\widetilde{S} - P \cup Q$ that separates P from Q by a strip $\mathbb{R} \times [a,b]$ so that the leaves of $\mathcal{G}$ correspond to $y =$ constant and the leaves of $\mathcal{H}$ correspond to $x =$ constant. In addition, since $\mathcal{H}$ respects the length along the leaves of $\mathcal{G}$, we can assume that this length along $\mathcal{G}$ is given by $|dx|$. Finally, having oriented the leaves of λ_{PQ} from right to left as seen from P, we can require that this orientation corresponds to the orientation by increasing values of x on the lines $y =$ constant.

By definition of $\sigma(P, Q)$, it is immediate that $\sigma(P, Q) = \int_{x_{d^-}^+}^{x_{d^+}^-} dx + f(x_{d^-}^+) - f(x_{d^+}^-)$.

The subarc $[x_{d^-}^+, x_{d^+}^-]$ of k is the union of $k \cap \widetilde{\lambda}_{PQ}$ and of the subarcs $[x_d^-, x_d^+]$, with d ranging over all components of $k - \widetilde{\lambda}_{QP}$ different from the end components d^+, d^-. Since $k \cap \widetilde{\lambda}$ has measure 0 on k, the integral term can therefore be decomposed as $\int_{x_{d^-}^+}^{x_{d^+}^-} dx = \sum_{d \neq d^+, d^-} \int_{x_d^-}^{x_d^+} dx$. Consequently, it suffices to prove that $\int_{x_d^-}^{x_d^+} dx = f(x_d^+) - f(x_d^-)$ for every d.

Given a component $d \neq d^+$, d^- of $k - \widetilde{\lambda}_{PQ}$, the component Σ_d of $\widetilde{S} - \widetilde{\lambda}_{PQ}$ that contains it is a wedge separated by the two geodesics g_d^+ and g_d^-. This wedge admits an isometry exchanging g_d^+ and g_d^-. This isometry respects $\mathcal{G} \cap \Sigma_d$, and therefore respects each leaf of $\mathcal{H} \cap \Sigma_d$. In particular, the base points of g_d^+ and g_d^- are located on the same leaf of $\mathcal{H}$. It immediately follows that $\int_{x_d^-}^{x_d^+} dx = f(x_d^+) - f(x_d^-)$. □

An immediate corollary of Lemma 7 is the following.

LEMMA 8. *With the data of Lemma 7,*

$$\left|\sigma(P, Q) - f(x_{d^-}^+) + f(x_{d^+}^-)\right| \leqslant l_m(k - P \cup Q)$$

where l_m denotes the length with respect to the metric m.

PROOF. By Lemma 7, it suffices to show that each term $\left|f(x_d^+) - f(x_d^-)\right|$ is bounded by the length of d. But we just saw that $f(x_d^+) - f(x_d^-)$ is equal to $\int_{x_d^-}^{x_d^+} dx$ which, up to sign, is equal to the length of the projection of d to any leaf of $\mathcal{G}$ parallel to $\mathcal{H}$. Since this projection is distance non-increasing, the result follows. □

9.5. The shearing transverse cocycle and the length function

THEOREM 9. *Given a maximal geodesic lamination λ, let σ_m be the shearing cocycle of the hyperbolic metric m. Then, for every transverse cocycle $\alpha \in \mathcal{H}(\lambda; \mathbb{R})$ for the geodesic lamination λ, its length $l_m(\alpha)$ is equal to $\tau(\alpha, \sigma_m)$.*

PROOF. As before, let $\widehat{\lambda} \to \lambda$ be the orientation covering of λ. Choose a neighborhood U of λ such that each component of $U - \lambda$ is an open annulus;

for instance, we could take U to consist of those points which are at distance less than ε from λ, for ε small enough. Extend $\widehat{\lambda} \to \lambda$ to a covering $\widehat{U} \to U$.

Along the leaves of the oriented lamination $\widehat{\lambda}$, the length measure induced by m defines a differential 1–form w_m. The direction of the leaf of $\widehat{\lambda}$ at a point $x \in \widehat{\lambda}$ is a Lipschitz function of x. Therefore, w_m can be extended to a closed Lipschitz differential 1–form $w_m \in \Omega^1_{\mathrm{Lip}}(\widehat{U})$. Since w_m is closed, it defines a cohomology class $[w_m] \in H^1(\widehat{U}; \mathbb{R})$.

By definition of the length function,

$$l_m(\alpha) = \tfrac{1}{2} l_m(\widehat{\alpha}) = \tfrac{1}{2} \iint_{\widehat{\lambda}} w_m \, d\widehat{\alpha} = \tfrac{1}{2} \langle [w_m], [\widehat{\alpha}] \rangle$$

where the last term denotes the evaluation of the cohomology class $[w_m] \in H^1(\widehat{U}; \mathbb{R})$ on the homology class $[\alpha] \in H_1(\widehat{U}; \mathbb{R})$, and where the last equality comes from the realization of $[\widehat{\alpha}]$ by a geometric current supported by $\widehat{\lambda}$.

By definition of τ, the proof of Theorem 9 will therefore be completed by the following lemma.

Lemma 10. *For every homology class $c \in H_1(\widehat{U}; \mathbb{R})$, the evaluation $\langle [w_m], c \rangle$ is equal to the intersection number $c \cdot [\widehat{\sigma}_m]$.*

Proof. Let W be a component of $U - \lambda$. By hypothesis on U, W is an open annulus bounded on one side by 3 leaves of λ. Consequently, its preimage $\widehat{W}$ in $\widehat{U}$ is an open annulus bounded on one side by 6 leaves of $\widehat{\lambda}$, with alternating orientations. Recall that each leaf of λ in the boundary of W has a preferred base point, coming from the projection of the third cusp of the component of $S - \lambda$ adjacent to that leaf. This determines a preferred base point O_g on each leaf g of $\widehat{\lambda}$ in the boundary of $\widehat{W}$.

Consider two consecutive leaves g, h of $\widehat{\lambda}$ in the boundary of $\widehat{W}$, and integrate w_m along an arc k joining O_g to O_h in $\widehat{W} \cup g \cup h$ which is made up of three pieces: first an arc in g joining O_g to a point x_g very close to the spike of $\widehat{W}$ separating g from h; then a small jump from x_g to its projection point x_h on h; and finally an arc in h joining x_h to O_h. By definition of w_m, the contribution of the first and last arcs to this integral are, in absolute value, respectively equal to the distances between O_g and x_g and between O_h and x_h. If x_g is far enough near the cusp, these two distances are approximately the same because O_g and O_h are at the same horocyclic distance from the spike (which comes from the fact that, as seen in §1, the same property holds for their projections in S). Also, because of alternating orientations, the integral of w_m along the first and last arc have opposite signs; their sum is therefore very small. It follows that the integral of w_m along k is arbitrarily small if we choose x_g close enough to the cusp. On the other hand, different choices for x_g give homotopic arcs, along which the integral of w_m is unchanged since w_m is closed. Therefore, the integral of w_m along k is actually 0.

Since $\widehat{W}$ is an annulus, it follows that the integral of w_m along every closed curve in $\widehat{W}$ is equal to 0. We can therefore define a function f_m on $\widehat{W}$ by the property that $f_m(x)$ is the integral of w_m along any arc in $\widehat{W}$ joining some base point O_g to x. This defines a function f_m on $\widehat{U}-\widehat{\lambda}$ such that $df_m = w_m$. There is of course no way to extend f_m to a global antiderivative of w_m over $\widehat{U}$. (The geodesic lamination λ carries at least one transverse measure, whose length has to be positive).

Let k be an oriented arc in $\widehat{U}$ that is transverse to $\widehat{\lambda}$. For each component d of $k-\widehat{\lambda}$, let x_d^+ and x_d^- be the positive and negative end points of d, respectively. From what precedes, and because $k\cap\lambda$ has Lebesgue measure 0,

$$\int_k w_m = \sum_d \int_d w_m = \sum_d \left(f_m\left(x_d^+\right) - f_m\left(x_d^-\right)\right)$$

where the sum is over all components d of $k-\widehat{\lambda}$, and where $f_m\left(x_d^\pm\right)$ is defined by continuous extension of the restriction of f_m to d. (In particular, if d is adjacent to d' so that $x_d^+ = x_{d'}^-$, it may very well happen that $f_m\left(x_d^+\right) \neq f_m\left(x_{d'}^-\right)$.)

We can compare this formula to that of Lemma 7. Note that, if g is a leaf of $\widehat{\lambda}$ in the boundary of a component $\widehat{W}$ of $\widehat{U}-\widehat{\lambda}$, and if we continuously extend f_m to $\widehat{W}\cup g$, the extension of f_m to g is just the oriented isometry from g to $\mathbb{R}$ which sends the base point to 0. Consequently, if k is an arc transverse to $\widehat{\lambda}$ which is small enough so that the leaves of $\widehat{\lambda}$ cross k in the same direction,

$$\int_k w_m - f_m\left(x_k^+\right) + f_m\left(x_k^-\right) = \varepsilon\sigma_m\left(k'\right) = \varepsilon\int_{k'}\sigma_m = \varepsilon\int_k \widehat{\sigma}_m$$

where k' is the projection of k to U, where $\varepsilon = +1$ if the leaves of $\widehat{\lambda}$ cross k from right to left, where $\varepsilon = -1$ if they cross from left to right, and where x_k^+ and x_k^- are the positive and negative end points of k, respectively. By interpreting $\widehat{\sigma}_m$ as a geometric current, we can incorporate the ε in an intersection number, and the above equality becomes

$$k\cdot\widehat{\sigma}_m = \int_k w_m - f_m\left(x_k^+\right) + f_m\left(x_k^-\right).$$

By additivity, this equation actually holds for every arc k transverse to $\widehat{\lambda}$, without any restriction on the direction in which the leaves of $\widehat{\lambda}$ cross k.

Now, consider a class $c\in H_1\left(\widehat{U};\mathbb{R}\right)$. This class can be represented by a cycle $\sum_{i=1}^n a_i k_i$, with $a_i\in\mathbb{R}$ and with the arcs k_i transverse to $\widehat{\lambda}$. Then,

$$\begin{aligned}\langle [w_m], c\rangle &= \sum_{i=1}^n a_i \int_{k_i} w_m = \sum_{i=1}^n a_i\left(k_i\cdot\widehat{\sigma}_m + f_m\left(x_{k_i}^+\right) - f_m\left(x_{k_i}^-\right)\right)\\ &= \sum_{i=1}^n a_i k_i\cdot\widehat{\sigma}_m = c\cdot\left[\widehat{\sigma}_m\right]\end{aligned}$$

where the f_m terms cancel out because $\sum_{i=1}^n a_i k_i$ has boundary 0. This concludes the proof of Lemma 10, and therefore of Theorem 9. □

9.6. The shearing transverse cocycle determines the metric

We want to show that, if two hyperbolic metrics have the same shearing transverse cocycle, then they represent the same class in $\mathcal{T}(S)$.

Consider two hyperbolic metrics m_1 and m_2 and a maximal geodesic lamination λ. As indicated in §1, λ can be represented by an m_1–geodesic lamination λ_1 and an m_2–geodesic lamination λ_2. Lift the situation to the universal covering $\widetilde{S}$, where λ, λ_1 and λ_2 have respective preimages $\widetilde{\lambda}$, $\widetilde{\lambda}_1$, $\widetilde{\lambda}_2$.

Since λ_i represents λ, there is a leaf of $\widetilde{\lambda}_i$ which is naturally associated to each leaf of $\widetilde{\lambda}$. Therefore, for each plaque P of $\widetilde{S}-\widetilde{\lambda}$, there is a plaque P_i of $\widetilde{S}-\widetilde{\lambda}_i$ which is naturally associated to P, as well as a homeomorphism $P \to P_i$ well defined up to isotopy. By composition, we get a preferred isotopy class of homeomorphisms $P_1 \to P_2$. Since any two ideal triangles are isometric, this isotopy class is represented by a unique isometry $\varphi_P : P_1 \to P_2$, called the ***plaque map***.

Define the ***shear map*** $\varphi : \widetilde{S}-\widetilde{\lambda}_1 \to \widetilde{S}-\widetilde{\lambda}_2$ by the property that, on each component of $\widetilde{S}-\widetilde{\lambda}_1$, the map φ coincides with the corresponding plaque map $\varphi_P : P_1 \to P_2$. Note that φ is an isometry from the metric m_1 to the metric m_2.

Lemma 11. *If the two metrics m_1 and m_2 have the same shearing cocycle, the shear map φ continuously extends to an isometry $(\widetilde{S}, m_1) \to (\widetilde{S}, m_2)$.*

Proof. We first show that φ admits a continuous extension which is locally Lipschitz.

Consider two points $x_1, y_1 \in \widetilde{S}-\widetilde{\lambda}_1$, respectively contained in the plaques P_1 and Q_1 of $\widetilde{S}-\widetilde{\lambda}_1$. Let Σ_1 be the component of $\widetilde{S}-\operatorname{int}(P_1 \cup Q_1)$ which separates P_1 from Q_1. As in §2, there is a unique foliation $\mathcal{G}_1$ of Σ_1 by m_1–geodesics such that every leaf of $\widetilde{\lambda}_1$ separating x_1 from y_1 is a leaf of $\mathcal{G}_1$. Again as in §2, let $\mathcal{H}_1$ be the foliation of Σ_1 orthogonal to $\mathcal{G}_1$.

Consider the m_1–geodesic arc α_1 joining x_1 to y_1, and let $u_1 = \alpha_1 \cap \Sigma_1 \cap P_1$ and $v_1 = \alpha_1 \cap \Sigma_1 \cap Q_1$ be the two end points of $\alpha_1 \cap \Sigma_1$. In Σ_1, u_1 and v_1 can be also be connected by the union of a leaf γ_1 of $\mathcal{H}_1$ and of an arc δ_1 contained in the leaf $\Sigma_1 \cap Q_1$ of $\mathcal{G}_1$. Let β_1 be the arc obtained from α_1 by replacing $\alpha_1 \cap \Sigma_1$ by $\gamma_1 \cup \delta_1$.

The projection of Σ_1 onto $\Sigma_1 \cap Q_1$ along the leaves of $\mathcal{H}_1$ is distance nonincreasing. It follows that the length of δ_1 is bounded by the length of $\alpha_1 \cap \Sigma_1$, and therefore by the distance $d(x_1, y_1)$. By the Jacobi equation, the projection from Σ_1 to γ_1 along the leaves of $\mathcal{G}_1$ is locally Lipschitz, where the local Lipschitz constant can be taken to be the exponential of the projection distance. As in the case of the length of δ_1, this projection distance is bounded by $d(x_1, y_1)$. It follows that the length of γ_1 is bounded by $e^{d(x_1,y_1)}$ times the m_1–length of $\alpha_1 \cap \Sigma_1$, and therefore by $e^{d(x_1,y_1)} d(x_1, y_1)$. Altogether, we conclude that the length of β_1 is bounded by $\left(2 + e^{d(x_1,y_1)}\right) d(x_1, y_1)$.

Now, consider $x_2 = \varphi(x_1)$, $y_2 = \varphi(y_1)$. Let P_2 and Q_2 be the plaques of $\widetilde{S}-\widetilde{\lambda}_2$ respectively containing x_2 and y_2, and let Σ_2 be the closure of the component of $\widetilde{S}-P_2 \cup Q_2$ that separates P_2 from Q_2. As before, let $\mathcal{G}_2$ be the foliation of Σ_2 by m_2–geodesics such that every leaf of $\widetilde{\lambda}_2$ separating x_2 from y_2 is a leaf of $\mathcal{G}_2$, and let $\mathcal{H}_2$ be the orthogonal foliation. The point $u_2 = \varphi_P(u_1) \in P_2 \cap \Sigma_2$ can be joined to the point $v_2 = \varphi_Q(v_1) \in Q_2 \cap \Sigma_2$ by the union of a leaf γ_2 of $\mathcal{H}_2$ and of an arc δ_2 contained in the leaf $\Sigma_2 \cap Q_2$ of $\mathcal{G}_2$. Let β_2 be the arc connecting x_2 to y_2 which is the union of $\varphi(\alpha_1 \cap P_1)$, γ_2, δ_2 and $\varphi(\alpha_1 \cap Q_1)$.

Because m_1 and m_2 have the same shearing cocycle, the end point $\gamma_2 \cap Q_2$ of γ_2 is the image of $\gamma_1 \cap Q_1$ under the plaque map $Q_1 \to Q_2$. It follows that δ_2 is the image of δ_1 under the same plaque map; in particular, the m_1–length of δ_1 is equal to the m_2–length of δ_2.

In Σ_1, consider a wedge W_1 delimited by two asymptotic leaves of $\widetilde{\lambda}_1$ separating x_1 from y_1, such that the interior of W_1 does not meet $\widetilde{\lambda}_1$. Let $R_1 \subset W_1$ be the plaque of $\widetilde{S} - \widetilde{\lambda}_1$ that is adjacent to the same two leaves of $\widetilde{\lambda}_1$, and let W_2 and R_2 be the wedge and plaque in Σ_2 respectively corresponding to W_1 and R_1. The fact that m_1 and m_2 have the same shearing cocycle implies that the plaque map $R_1 \to R_2$ sends the end point of $\gamma_1 \cap R_1$ that is closest to u_1 to the end point of $\gamma_2 \cap R_2$ that is closest to u_2 . As a consequence, the isometric extension of this plaque map to $W_1 \to W_2$ sends $\gamma_1 \cap W_1$ to $\gamma_2 \cap W_2$. In particular, the m_1–length of $\gamma_1 \cap W_1$ is equal to the m_2–length of $\gamma_2 \cap W_2$.

Since $\gamma_1 \cap \widetilde{\lambda}_1$ has 1–dimensional Lebesgue measure 0, the length of γ_1 is equal to the (infinite) sum of the lengths of the $\gamma_1 \cap W_1$, where W_1 ranges over all wedges in Σ_1 as above. Since the same property holds for γ_2, it follows that the m_1–length of γ_1 is equal to the m_2–length of γ_2.

This proves that each of the four pieces forming β_1 has the same length as the corresponding piece of β_2. As a consequence, the m_1–length of β_1 is equal to the m_2–length of β_2.

Therefore, $d(x_2, y_2) \leqslant l_{m_2}(\beta_2) = l_{m_1}(\beta_1) \leqslant \left(2 + e^{d(x_1,y_1)}\right) d(x_1, y_1)$. Since this holds for every x_1, $y_1 \in \widetilde{S} - \widetilde{\lambda}_1$, it follows that φ admits a continuous extension $\varphi : \left(\widetilde{S}, m_1\right) \to \left(\widetilde{S}, m_2\right)$ which is locally Lipschitz.

We now prove that φ is distance non-increasing. For this, consider two points x_1 and $y_1 \in \widetilde{S}$ which are not on the same leaf of $\widetilde{\lambda}_1$, and let α_1 be the m_1–geodesic arc joining x_1 to y_1. Since $\alpha_1 \cap \widetilde{\lambda}_1$ has 1–dimensional Lebesgue measure 0 and since φ is locally Lipschitz, the image $\varphi\big(\alpha_1 \cap \widetilde{\lambda}_1\big)$ also has 1–dimensional Lebesgue measure 0 (for the metric m_2). Also, because φ is isometric on $\widetilde{S} - \widetilde{\lambda}_1$, the m_2–length of $\varphi\big(\alpha_1 - \widetilde{\lambda}_1\big)$ is equal to the m_1–length of $\alpha_1 - \widetilde{\lambda}_1$. Therefore, the m_2–length of $\varphi(\alpha_1)$ is equal to the m_1–length of α_1, and

$$d_{m_2}(\varphi(x_1), \varphi(y_1)) \leqslant l_{m_2}(\varphi(\alpha_1)) = l_{m_1}(\alpha_1) = d_{m_1}(x_1, y_1).$$

By density, this inequality $d_{m_2}(\varphi(x_1), \varphi(y_1)) \leqslant d_{m_1}(x_1, y_1)$ holds for every x_1, $y_1 \in \widetilde{S}$, namely even if the two points are on the same leaf of $\widetilde{\lambda}_1$. In other words, $\varphi : \left(\widetilde{S}, m_1\right) \to \left(\widetilde{S}, m_2\right)$ is distance non-increasing.

By symmetry, the shear map $\varphi^{-1} : \widetilde{S} - \widetilde{\lambda}_2 \to \widetilde{S} - \widetilde{\lambda}_1$ extends to a distance non-increasing map $\left(\widetilde{S}, m_2\right) \to \left(\widetilde{S}, m_1\right)$. It follows that $\varphi : \left(\widetilde{S}, m_1\right) \to \left(\widetilde{S}, m_2\right)$ is an isometry. □

THEOREM 12. *Two hyperbolic metrics m_1 and m_2 have the same shearing transverse distribution if and only if $m_1 = m_2$ in $\mathcal{T}(S)$.*

PROOF. If m_1 and m_2 have the same shearing transverse distribution, let $\varphi : \big(\widetilde{S}, m_1\big) \to \big(\widetilde{S}, m_2\big)$ be the isometry provided by Lemma 11. Since the shear map $\varphi : \widetilde{S} - \widetilde{\lambda}_1 \to \widetilde{S} - \widetilde{\lambda}_2$ commutes with the action of $\pi_1(S)$, φ induces an isometry $\psi : \big(S, m_1\big) \to \big(S, m_2\big)$ which is homotopic to the identity. In particular, m_1 and m_2 represent the same element of $\mathcal{T}(S)$. □

9.7. The local realization of shearing cocycles

In this section we show that, given a maximal geodesic lamination λ, the map $\mathcal{T}(S) \to \mathcal{H}(\lambda, \mathbb{R})$ which associates its shearing cocycle to a hyperbolic metric is open. By Theorem 12, this implies that this map is a homeomorphism onto its image. Its precise image will be determined in §6.

PROPOSITION 13. *Let m_0 be a hyperbolic metric with associated shearing cocycle σ_0 for the maximal geodesic lamination λ. Then, every $\sigma \in \mathcal{H}(\lambda; \mathbb{R})$ that is sufficiently close to σ_0 is the shearing cocycle of some hyperbolic metric m.*

PROOF. The proof will require several steps.

Set $\alpha = \sigma - \sigma_0 \in \mathcal{H}(\lambda; \mathbb{R})$.

Represent λ by an m_0–geodesic lamination which we will also denote by λ, and let $\widetilde{\lambda}$ be the preimage of λ in the universal covering $\widetilde{S}$. Consider two plaques P and Q of $\widetilde{S} - \widetilde{\lambda}$.

For every plaque R separating P from Q, let g_R^P and g_R^Q be the geodesics in the boundary of R which are closest to P and Q, respectively. Orient these geodesics to the left, as seen from P. Also, given an oriented geodesic of g of $\widetilde{S}$ and a number $u \in \mathbb{R}$, let $T_g^u : \widetilde{S} \to \widetilde{S}$ denote the m_0–isometry which respects g and act by translation of oriented amplitude u on g.

Let $\mathcal{P}_{PQ}$ be the set of all plaques of $\widetilde{S} - \widetilde{\lambda}$ that separate P from Q. Given a finite subset $\mathcal{P}$ of $\mathcal{P}_{PQ}$, index its elements as $P_1, P_2, \ldots, P_n$ so that the index i of P_i increases as one goes from P to Q, and consider

$$\varphi_{\mathcal{P}} = T_{g_1^P}^{\alpha(P,P_1)} T_{g_1^Q}^{-\alpha(P,P_1)} T_{g_2^P}^{\alpha(P,P_2)} T_{g_2^Q}^{-\alpha(P,P_2)} \ldots T_{g_n^Q}^{-\alpha(P,P_n)} T_{g_Q^P}^{\alpha(P,Q)}$$

where $g_i^P = g_{P_i}^P$, $g_i^Q = g_{P_i}^Q$, and g_Q^P is the geodesic in the boundary of Q that is closest to P. This formula is perhaps easier to read and understand if we notice that each P_i contributes a term $T_{g_i^P}^{\alpha(P,P_i)} T_{g_i^Q}^{-\alpha(P,P_i)}$.

Now, we let the finite subset $\mathcal{P}$ converge to $\mathcal{P}_{PQ}$ and we consider the limit

$$\varphi_{PQ} = \lim_{\mathcal{P} \to \mathcal{P}_{PQ}} \varphi_{\mathcal{P}}.$$

By convention, we decide that φ_{PP} is the identity. Of course, we first have to prove that the above limit exists.

LEMMA 14. *Let k be the lift to $\widetilde{S}$ of a simple geodesic arc in S transverse to λ. Then, if $\alpha \in \mathcal{H}(\lambda)$ is sufficiently small (depending on k) and if the two plaques P and Q meet k, the map $\varphi_{\mathcal{P}}$ converges to an m_0–isometry φ_{PQ} as $\mathcal{P}$ tends to $\mathcal{P}_{PQ}$.*

PROOF. For notational convenience, set

$$\psi_{\mathcal{P}} = \varphi_{\mathcal{P}} T_{g_Q^P}^{-\alpha(P,Q)} = T_{g_1^P}^{\alpha(P,P_1)} T_{g_1^Q}^{-\alpha(P,P_1)} T_{g_2^P}^{\alpha(P,P_2)} T_{g_2^Q}^{-\alpha(P,P_2)} \dots T_{g_n^Q}^{-\alpha(P,P_n)}.$$

Identify the m_0–isometry group of $\widetilde{S}$ to the matrix group $\mathrm{SO}(2,1) \subset \mathrm{GL}_3(\mathbb{R})$, and endow it with the norm $\|A\| = \max_{x\in\mathbb{R}^3} \|Ax\| / \|x\|$. The main property we want is that this norm satisfies $\|AB\| \leqslant \|A\| \, \|B\|$.

We first show that the norm $\|\psi_{\mathcal{P}}\|$ is uniformly bounded, if α is small enough.

For every i, the distance between the geodesics g_i^P and g_i^Q is bounded by a constant times the length of $k \cap P_i$. By Lemma 5, this distance is therefore an $O\left(e^{-Ar(k\cap P_i)}\right)$ for some constant $A > 0$, where we identify k to its projection to S and $k \cap P_i$ to the corresponding component of $k - \lambda$. By an easy hyperbolic estimate, it follows that

$$T_{g_i^P}^{\alpha(P,P_i)} T_{g_i^Q}^{-\alpha(P,P_i)} = \mathrm{Id} + O\left(e^{|\alpha(P,P_i)|} e^{-Ar(k\cap P_i)}\right).$$

As a consequence,

$$\|\psi_{\mathcal{P}}\| \leqslant \prod_{i=1}^{n} \left(1 + O\left(e^{|\alpha(P,P_i)|} e^{-Ar(k\cap P_i)}\right)\right).$$

By Lemmas 4 and 6, the series $\sum_{R\in\mathcal{P}_{PQ}} e^{|\alpha(P,R)|} e^{-Ar(k\cap R)}$ is bounded by the sum of finitely many geometric series $\sum_{r=0}^{\infty} e^{C\|\alpha\|(r+1)} e^{-Ar}$. It is therefore convergent if $\|\alpha\| < A/C$.

It follows that, if the transverse distribution α is small enough, the norm $\|\psi_{\mathcal{P}}\|$ is uniformly bounded.

Let $\mathcal{P}_n$, $n \in \mathbb{N}$, be an increasing sequence of finite subsets converging to $\mathcal{P}_{PQ}$, with the cardinal of $\mathcal{P}_n$ equal to n. The map $\psi_{\mathcal{P}_{n+1}}$ is obtained from $\psi_{\mathcal{P}_n}$ by inserting a term $T_{g_R^P}^{\alpha(P,R)} T_{g_R^Q}^{-\alpha(P,R)}$ in its expression. More precisely, there are subsets $\mathcal{P}$ and $\mathcal{P}'$ of $\mathcal{P}_{PQ}$ such that $\psi_{\mathcal{P}_n} = \psi_{\mathcal{P}} \psi_{\mathcal{P}'}$ and $\psi_{\mathcal{P}_{n+1}} = \psi_{\mathcal{P}} T_{g_R^P}^{\alpha(P,R)} T_{g_R^Q}^{-\alpha(P,R)} \psi_{\mathcal{P}'}$.

Then

$$\begin{aligned}\left\|\psi_{\mathcal{P}_{n+1}} - \psi_{\mathcal{P}_n}\right\| &\leqslant \|\psi_{\mathcal{P}}\| \left\| T_{g_R^P}^{\alpha(P,R)} T_{g_R^Q}^{-\alpha(P,R)} - \mathrm{Id} \right\| \|\psi_{\mathcal{P}'}\| \\ &= O\left(e^{|\alpha(P,R)|} e^{-Ar(k\cap R)}\right) = O\left(e^{C\|\alpha\|(r(k\cap R)+1)} e^{-Ar(k\cap R)}\right)\end{aligned}$$

by Lemmas 5 and 6, and because we just proved that $\|\psi_{\mathcal{P}}\|$ and $\|\psi_{\mathcal{P}'}\|$ are uniformly bounded.

By Lemma 4, it follows that the sequence $\psi_{\mathcal{P}_n}$ is Cauchy, and therefore convergent, if $\|\alpha\| < A/C$.

This proves that $\psi_{\mathcal{P}}$ has a limit ψ_{PQ} as $\mathcal{P}$ tends to $\mathcal{P}_{PQ}$, provided that $\alpha \in \mathcal{H}(\lambda)$ is small enough. The same clearly holds for $\varphi_{\mathcal{P}} = \psi_{\mathcal{P}} T_{g_Q^P}^{\alpha(P,Q)}$. □

For future reference, we note the following estimate.

Lemma 15. *Under the hypotheses of Lemma 14, there is a constant $B > 0$, depending on k and α, such that φ_{PQ} can be decomposed as $\varphi_{PQ} = \psi_{PQ} T_{g_Q^P}^{\alpha(P,Q)}$ with $\|\psi_{PQ} - \mathrm{Id}\| = O\left(\sum_{R\in\mathcal{P}_{PQ}} e^{-Br(k\cap R)}\right)$.*

Proof. $\|\psi_{PQ} - \mathrm{Id}\| = \lim_{n\to\infty} \|\psi_{\mathcal{P}_n} - \psi_{\mathcal{P}_0}\| \leqslant \sum_{n=0}^{\infty} \left\|\psi_{\mathcal{P}_{n+1}} - \psi_{\mathcal{P}_n}\right\|$. □

Having proved the convergence in Lemma 14, we can now give another description of φ_{PQ} which is perhaps more intuitive. We are still assuming that P and Q meet the lift k of a simple geodesic arc transverse to λ in S.

Given an integer $r > 0$, let $\mathcal{P}_{PQ}^r$ consist of the finitely many $R \in \mathcal{P}_{PQ}$ such that $r(k\cap R) \leqslant r$. Index the elements of $\mathcal{P}_{PQ}^r$ as $P_1, P_2, \ldots, P_n$ so that the index i of P_i increases as one goes from P to Q. For notational convenience, set $P_0 = P$ and $P_{n+1} = Q$. For every i, choose a geodesic h_i which separates the interior of P_i from the interior of P_{i+1}, and orient h_i to the left as seen from P. Then, set

$$\varphi_{PQ}^r = T_{h_0}^{\alpha(P_0,P_1)} T_{h_1}^{\alpha(P_1,P_2)} \ldots T_{h_n}^{\alpha(P_n,P_{n+1})}.$$

(Compare [**EpM**, §3].)

Lemma 16. *Under the hypotheses of Lemma 14, as r tends to infinity, φ_{PQ}^r tends to φ_{PQ} if $\alpha \in \mathcal{H}(\lambda)$ is small enough.*

Proof. We will estimate the difference between φ_{PQ}^r and

$$\varphi_{\mathcal{P}_{PQ}^r} = T_{g_1^P}^{\alpha(P_0,P_1)} T_{g_1^Q}^{-\alpha(P_0,P_1)} T_{g_2^P}^{\alpha(P_0,P_2)} T_{g_2^Q}^{-\alpha(P_0,P_2)} \ldots T_{g_n^Q}^{-\alpha(P_0,P_n)} T_{g_{n+1}^P}^{\alpha(P_0,P_{n+1})}.$$

For this, it will be more convenient to rewrite φ_{PQ}^r as

$$\varphi_{PQ}^r = T_{h_0}^{\alpha(P_0,P_1)} T_{h_1}^{-\alpha(P_0,P_1)} T_{h_1}^{\alpha(P_0,P_2)} T_{h_2}^{-\alpha(P_0,P_2)} \ldots T_{h_n}^{-\alpha(P_0,P_n)} T_{h_n}^{\alpha(P_0,P_{n+1})},$$

noting that $\alpha(P_i, P_{i+1}) = \alpha(P_0, P_{i+1}) - \alpha(P_0, P_i)$, and to consider

$$\psi_{PQ}^r = \varphi_{PQ}^r T_{h_n}^{-\alpha(P_0,P_{n+1})} = T_{h_0}^{\alpha(P_0,P_1)} T_{h_1}^{-\alpha(P_0,P_1)} T_{h_1}^{\alpha(P_0,P_2)} T_{h_2}^{-\alpha(P_0,P_2)} \ldots T_{h_n}^{-\alpha(P_0,P_n)}$$

and

$$\psi_{\mathcal{P}_{PQ}^r} = \varphi_{\mathcal{P}_{PQ}^r} T_{g_{n+1}^P}^{-\alpha(P_0,P_{n+1})} = T_{g_1^P}^{\alpha(P_0,P_1)} T_{g_1^Q}^{-\alpha(P_0,P_1)} T_{g_2^P}^{\alpha(P_0,P_2)} T_{g_2^Q}^{-\alpha(P_0,P_2)} \ldots T_{g_n^Q}^{-\alpha(P_0,P_n)}.$$

The isometry ψ^r_{PQ} is obtained from $\psi_{\mathcal{P}^r_{PQ}}$ by replacing each term $T^{\alpha(P_0,P_i)}_{g_i^P} T^{-\alpha(P_0,P_i)}_{g_i^Q}$ by $T^{\alpha(P_0,P_i)}_{h_{i-1}} T^{-\alpha(P_0,P_i)}_{h_i}$.

For every i, the two geodesics g_i^Q and g_{i+1}^P follow the same edge path of length $2r$ in the train track associated to the arc k; otherwise, there would be another $R \in \mathcal{R}^r_{PQ}$ between P_i and P_{i+1}. Since h_i is between these two geodesics, it also follows the same edge path. Therefore, the distance between any two of these three geodesics is bounded by a constant times e^{-Ar}, for the constant $A > 0$ of Lemmas 3 and 5.

In particular, the distance from g_i^P to h_{i-1} and the distance from g_i^Q to h_i are both $O\left(e^{-Ar}\right)$. Also, the distance between g_i^P and g_i^Q is an $O\left(e^{-Ar(k\cap P_i)}\right)$ by Lemma 5. Since $r\left(k \cap P_i\right) \leqslant r$, it follows that the distance from h_{i-1} to h_i is also an $O\left(e^{-Ar(k\cap P_i)}\right)$.

From the second statement, it follows that

$$\begin{aligned} T^{\alpha(P_0,P_i)}_{h_{i-1}} T^{-\alpha(P_0,P_i)}_{h_i} &= \mathrm{Id} + O\left(e^{|\alpha(P_0,P_i)|} e^{-Ar(k\cap P_i)}\right) \\ &= \mathrm{Id} + O\left(e^{C\|\alpha\|(r(k\cap P_i)+1)} e^{-Ar(k\cap P_i)}\right) \end{aligned}$$

by Lemma 6. If ψ is any isometry obtained from $\psi_{\mathcal{P}^r_{PQ}}$ by replacing some of the n terms $T^{\alpha(P_0,P_i)}_{g_i^P} T^{-\alpha(P_0,P_i)}_{g_i^Q}$ by $T^{\alpha(P_0,P_i)}_{h_{i-1}} T^{-\alpha(P_0,P_i)}_{h_i}$ or by the identity, it follows as in the proof of Lemma 14 that

$$\begin{aligned} \log \|\psi\| &= O\left(\sum_{i=1}^{n} e^{C\|\alpha\|(r(k\cap P_i)+1)} e^{-Ar(k\cap P_i)}\right) \\ &= O\left(\sum_{R\in\mathcal{P}_{PQ}} e^{C\|\alpha\|(r(k\cap R)+1)} e^{-Ar(k\cap R)}\right) \end{aligned}$$

As a consequence, if $\|\alpha\| < A/C$, the norm of such a ψ is uniformly bounded.

Let ψ_i be obtained from $\psi_{\mathcal{P}^r_{PQ}}$ by replacing each $T^{\alpha(P_0,P_j)}_{g_j^P} T^{-\alpha(P_0,P_j)}_{g_j^Q}$ with $j \leqslant i$ by $T^{\alpha(P_0,P_j)}_{h_{j-1}} T^{-\alpha(P_0,P_j)}_{h_j}$, so that $\psi_0 = \psi_{\mathcal{P}^r_{PQ}}$ and $\psi_n = \psi^r_{PQ}$. To estimate the difference between ψ_{i-1} and ψ_i, note that we can write these as $\psi_{i-1} = \psi T^{\alpha(P_0,P_i)}_{g_i^P} T^{-\alpha(P_0,P_i)}_{g_i^Q} \psi'$ and $\psi_i = \psi T^{\alpha(P_0,P_i)}_{h_{i-1}} T^{-\alpha(P_0,P_i)}_{h_i} \psi'$, where ψ and ψ' are obtained from $\psi_{\mathcal{P}^r_{PQ}}$ by replacing some $T^{\alpha(P_0,P_j)}_{g_j^P} T^{-\alpha(P_0,P_j)}_{g_j^Q}$ by $T^{\alpha(P_0,P_j)}_{h_{j-1}} T^{-\alpha(P_0,P_j)}_{h_j}$ or by the identity. By the above observation $\|\psi\|$ and $\|\psi'\|$ are uniformly bounded. Also, we noted that the distance from g_i^P to h_{i-1} and the distance from g_i^Q to h_i are both $O\left(e^{-Ar}\right)$. It follows that

$$\begin{aligned} \|\psi_{i-1} - \psi_i\| &\leqslant \|\psi\| \, \|\psi'\| \left\| T^{\alpha(P_0,P_j)}_{g_j^P} T^{-\alpha(P_0,P_j)}_{g_j^Q} - T^{\alpha(P_0,P_j)}_{h_{j-1}} T^{-\alpha(P_0,P_j)}_{h_j} \right\| \\ &= O\left(e^{2|\alpha(P_0,P_j)|} e^{-Ar}\right) = O\left(e^{2C\|\alpha\|(r+1)} e^{-Ar}\right) \end{aligned}$$

and therefore that

$$\left\| \psi^r_{PQ} - \psi_{\mathcal{P}^r_{PQ}} \right\| = \|\psi_n - \psi_0\| \leqslant nO\left(e^{2C\|\alpha\|(r+1)}e^{-Ar}\right) = O\left(re^{2C\|\alpha\|(r+1)}e^{-Ar}\right)$$

since $n = O(r)$ by Lemma 4.

It follows that ψ^r_{PQ} and $\psi_{\mathcal{P}^r_{PQ}}$ have the same limit as r tends to infinity, provided that $\|\alpha\| < A/2C$. On the other hand, h_n converges to g^P_{n+1} as r tends to infinity. Therefore, the limit of $\varphi^r_{PQ} = \psi^r_{PQ} T^{\alpha(P_0,P_{n+1})}_{h_n}$ is the same as the limit of $\varphi_{\mathcal{P}^r_{PQ}} = \psi_{\mathcal{P}^r_{PQ}} T^{\alpha(P_0,P_{n+1})}_{g^P_{n+1}}$, namely is equal to φ_{PQ}, if α is small enough. □

There is a natural generalization of Lemma 16, closer to the one used in [**EpM**] for the construction of earthquake, which would lead to an even more intuitive definition of φ_{PQ}. As before, index the elements of a finite subset $\mathcal{P}$ of $\mathcal{P}_{PQ}$ as $P_1, P_2, \dots, P_n$ so that the index i of P_i increases as one goes from P to Q. Then, we could expect φ_{PQ} to be the limit as $\mathcal{P}$ tends to $\mathcal{P}_{PQ}$ of $T^{\alpha(P_0,P_1)}_{h_0} T^{\alpha(P_1,P_2)}_{h_1} \dots T^{\alpha(P_n,P_{n+1})}_{h_n}$, where the geodesic h_i separates P_i from P_{i+1} and is oriented to the left as seen from P. This approach would certainly lead to a more intuitive definition of φ_{PQ}, but is unfortunately too naive. Indeed, it is not hard to see that the above limit does not necessarily exist if the transverse distribution α associated to the transverse cocycle α is not a measure. So, only the restricted limit of Lemma 16 makes sense.

From Lemma 16, we get the following properties, which were not obvious from the definition of φ_{PQ}. (Note that even the second one is non-trivial if Q does not separate P from R.)

Lemma 17. *If $\alpha \in \mathcal{H}(\lambda)$ is small enough for the conclusions of Lemmas 14 and 16 to hold then, for every plaques P, Q, R of $\widetilde{S} - \widetilde{\lambda}$ meeting k, $\varphi_{QP} = \varphi^{-1}_{PQ}$ and $\varphi_{PR} = \varphi_{PQ}\varphi_{QR}$.* □

We can now get rid of the assumption that P and Q both meet a suitable arc k.

Lemma 18. *If $\alpha \in \mathcal{H}(\lambda)$ is sufficiently small then, for every plaques P, Q of $\widetilde{S} - \widetilde{\lambda}$, the map $\varphi_{\mathcal{P}}$ converges to an m_0–isometry φ_{PQ} as $\mathcal{P}$ tends to $\mathcal{P}_{PQ}$. In addition, $\varphi_{QP} = \varphi^{-1}_{PQ}$ and $\varphi_{PR} = \varphi_{PQ}\varphi_{QR}$ for every plaques P, Q, R.*

Proof. Select in S finitely many simple geodesic arcs $k_1, \dots, k_n$ transverse to λ, such that each component of $S - \lambda$ meets at least one of the k_i.

For each pair of plaques P, Q of $\widetilde{S} - \widetilde{\lambda}$, we can find a finite sequence of plaques $P = P_0, P_1, \dots, P_n, P_{n+1} = Q$ such that each P_j separates P_{j-1} from P_{j+1} and such that P_j and P_{j+1} meet the same lift $\widetilde{k}_{i_j}$ of some k_{i_j}. Then, for $\|\alpha\|$ sufficiently small (depending on the k_i), Lemma 14 proves the existence of a limit $\varphi_{P_jP_{j+1}}$ for every j. This guarantees the existence of the limit $\varphi_{PQ} = \lim_{\mathcal{P}\to\mathcal{P}_{PQ}} \varphi_{\mathcal{P}} = \varphi_{P_0P_1}\varphi_{P_1P_2}\dots\varphi_{P_nP_{n+1}}$.

The second statement easily follows from Lemma 17. (Hint for the case where none of the three plaques P, Q, R separates the two other ones: Consider the unique plaque N which separates any two of these plaques.) □

Now, consider the action of the fundamental group $\pi_1(S)$ on $\widetilde{S}$. By invariance of α under this action, we have that $\varphi_{(\gamma P)(\gamma Q)} = \gamma \varphi_{PQ} \gamma^{-1}$ for every $\gamma \in \pi_1(S)$ and every plaques P and Q.

Fix a base plaque P_0 of $\widetilde{S} - \widetilde{\lambda}$, and define $\rho(\gamma) = \varphi_{P_0(\gamma P_0)}\gamma$. Then, it immediately follows from Lemma 18 and the above property that ρ defines a group homomorphism from $\pi_1(S)$ to the group of m_0–isometries of $\widetilde{S}$.

By definition of φ_{PQ}, the interiors of $\varphi_{PQ}(Q)$ and P are always disjoint. In particular, for every $\gamma \in \pi_1(S)$ which is not the identity, $\rho(\gamma) P_0 = \varphi_{P_0(\gamma P_0)}(\gamma P_0)$ is disjoint from the interior of P_0. It follows that $\rho(\gamma)$ cannot be very close to the identity, and therefore that the representation ρ is a discrete homomorphism from $\pi_1(S)$ into the m_0–isometry group of $\widetilde{S}$.

Consider the surface $S' = \widetilde{S}/\rho$. The metric m_0 on $\widetilde{S}$ induces a hyperbolic metric m' on S'. Since ρ defines an isomorphism between $\pi_1(S)$ and $\pi_1(S')$, we have a preferred isotopy class of diffeomorphisms $\psi : S \to S'$. Let m be the hyperbolic metric on S obtained by pulling back the metric m' under ψ. Note that the class of m in $\mathcal{T}(S)$ does not depend on the choice of ψ.

The proof of Proposition 13 will then be completed by the following statement.

Lemma 19. *The shearing cocycle of the metric m is equal to $\sigma = \sigma_0 + \alpha$.*

Proof. To understand the shearing cocycle σ_m of m, we first have to understand the m–geodesic lamination λ_m corresponding to λ.

Define a map $\widetilde{\varphi} : \widetilde{S} - \widetilde{\lambda} \to \widetilde{S}$ by the property that $\widetilde{\varphi}$ coincides with $\varphi_{P_0 P}$ on each plaque P of $\widetilde{S} - \widetilde{\lambda}$. This $\widetilde{\varphi}$ will more or less correspond to the shear map in the sense of §4, modulo composition with a suitable lift of ψ.

Note that $\widetilde{\varphi}\gamma = \rho(\gamma)\widetilde{\varphi}$ for every $\gamma \in \pi_1(S)$, and that $\widetilde{\varphi}$ is m_0–isometric. Therefore, $\widetilde{\varphi}$ induces an isometric map $\varphi : (S - \lambda, m_0) \to (S', m')$.

If P and Q are two distinct plaques of $\widetilde{S} - \widetilde{\lambda}$, it follows from the fact that the interior of $\varphi_{PQ}(Q)$ is disjoint from P that $\widetilde{\varphi}(P)$ and $\widetilde{\varphi}(Q)$ have disjoint interiors. As a consequence, φ is injective. Since $(S - \lambda, m_0)$ and (S', m') have the same area, we conclude that the image of φ is dense in S.

Therefore, the image of $\widetilde{\varphi}$ in dense in $\widetilde{S}$. In particular, every point in the complement $\widetilde{S} - \widetilde{\varphi}\big(\widetilde{S} - \widetilde{\lambda}\big)$ is in the geodesic limit of a sequence $\widetilde{\varphi}(g_i)$, where each g_i is in the boundary of a plaque of $\widetilde{S} - \widetilde{\lambda}$. It follows that $\widetilde{S} - \widetilde{\varphi}\big(\widetilde{S} - \widetilde{\lambda}\big)$ is a $\rho(\pi_1(S))$–invariant geodesic lamination $\widetilde{\lambda}'$ of $\widetilde{S}$, which projects to an m'–geodesic lamination λ' of S'.

The m–geodesic lamination $\lambda_m = \psi^{-1}(\lambda')$ of S is the m–geodesic lamination corresponding to λ. Indeed, a leaf of $\widetilde{\lambda}$ (resp. $\widetilde{\lambda}_m$) is completely determined by the way it separates the plaques of $\widetilde{S} - \widetilde{\lambda}$ (resp. $\widetilde{S} - \widetilde{\lambda}_m$), and $\widetilde{\psi}^{-1} \circ \widetilde{\varphi} : \widetilde{S} - \widetilde{\lambda} \to \widetilde{S} - \widetilde{\lambda}_m$ respects the combinatorics of these plaques while commuting with the action of $\pi_1(S)$.

We will prefer to work with λ' rather than λ_m, to avoid the interference of too many ψ with the arguments. Let $\sigma_{m'}$ be the shearing cocycle of the metric m' with respect to λ'. Then the isometry $\psi : (S, m) \to (S', m')$ sends σ_m to $\sigma_{m'}$. Namely, $\sigma_m(P, Q) = \sigma_{m'}(\widetilde{\varphi}(P), \widetilde{\varphi}(Q))$ if we note that $\widetilde{\varphi}(P)$ and $\widetilde{\varphi}(Q)$

are plaques of $\widetilde{S}-\widetilde{\lambda}'$ and are sent by the appropriate lift of ψ to the plaques of $\widetilde{S}-\widetilde{\lambda}_m$ corresponding to P and Q, respectively.

We want to show that $\sigma_m(P,Q)=\sigma_0(P,Q)+\alpha(P,Q)$ for every plaques P, Q of $\widetilde{S}-\widetilde{\lambda}$. By additivity of transverse cocycles, we can restrict attention to the case where P and Q both meet a transverse geodesic arc k whose projection to S is simple.

As usual, let $\mathcal{P}_{PQ}$ be the set of all plaques of $\widetilde{S}-\widetilde{\lambda}$ that separate P from Q. Choose a finite subset $\mathcal{P}$ of $\mathcal{P}_{PQ}$, index its elements as $P_1, P_2, \dots, P_n$ so that the index i of P_i increases as one goes from P to Q, and set $P_0=P$ and $P_{n+1}=Q$.

We want to compare $\sigma_m(P_i,P_{i+1})=\sigma_{m'}(\widetilde{\varphi}(P_i),\widetilde{\varphi}(P_{i+1}))$ to $\sigma_0(P_i,P_{i+1})$. Note that, to compute $\sigma_{m'}(\widetilde{\varphi}(P_i),\widetilde{\varphi}(P_{i+1}))$, we do not need the whole lamination $\widetilde{\lambda}'$. We only need to know the metric m_0 on $\widetilde{S}$, the two triangles $\widetilde{\varphi}(P_i)$ and $\widetilde{\varphi}(P_{i+1})$, and the family $\widetilde{\lambda}'_{PQ}$ of those leaves of $\widetilde{\lambda}'$ that separate $\widetilde{\varphi}(P_i)$ and $\widetilde{\varphi}(P_{i+1})$ (the only requirement here being that two geodesics of $\widetilde{\lambda}'_{PQ}$ which are adjacent to the same component of $\widetilde{S}-\widetilde{\lambda}'_{PQ}$ are asymptotic). Let $s\big(\widetilde{\varphi}(P_i),\widetilde{\varphi}(P_{i+1});\widetilde{\lambda}'_{PQ}\big)$ be the number defined by this procedure. Of course, in this case, $s\big(\widetilde{\varphi}(P_i),\widetilde{\varphi}(P_{i+1});\widetilde{\lambda}'_{PQ}\big)=\sigma_{m'}(\widetilde{\varphi}(P_i),\widetilde{\varphi}(P_{i+1}))$ and, similarly, $s\big(P_i,P_{i+1};\widetilde{\lambda}_{PQ}\big)=\sigma_0(P_i,P_{i+1})$.

Now, remember that $\widetilde{\varphi}(P_i)=\varphi_{P_0P_i}(P_i)$ and $\widetilde{\varphi}(P_{i+1})=\varphi_{P_0P_{i+1}}(P_{i+1})$. Since $\varphi_{P_0P_i}$ is an isometry, we conclude that $s\big(\widetilde{\varphi}(P_i),\widetilde{\varphi}(P_{i+1});\widetilde{\lambda}'_{PQ}\big)=s\big(P_i,\varphi_{P_iP_{i+1}}(P_{i+1});\varphi$

From Lemma 15, there is a constant $B>0$ such that $\varphi_{P_iP_{i+1}}$ can decomposed as $\varphi_{P_iP_{i+1}}=\psi_{P_iP_{i+1}}T_g^{\alpha(P_i,P_{i+1})}$ with $\left\|\psi_{P_iP_{i+1}}-\mathrm{Id}\right\|=O\left(\sum_{R\in\mathcal{P}_{P_iP_{i+1}}}e^{-Br(k\cap R)}\right)$, where g is the geodesic in the boundary of P_{i+1} that is closest to P_i, oriented to the left as seen from P_i.

By definition of s, $s\big(P_i,T_g^{\alpha(P_i,P_{i+1})}(P_{i+1});\widetilde{\lambda}_{PQ}\big)=s\big(P_i,P_{i+1};\widetilde{\lambda}_{PQ}\big)+\alpha(P_i,P_{i+1})$.

Let k_i be the subarc of k joining P_i to P_{i+1}. Note that k_i also joins P_i to $T_g^{\alpha(P_i,P_{i+1})}(P_{i+1})$. Within an error of $l_{m_0}(k_i)$, $s\big(P_i,T_g^{\alpha(P_i,P_{i+1})}(P_{i+1});\widetilde{\lambda}_{PQ}\big)$ does not depend on the lamination $\widetilde{\lambda}_{PQ}$ by Lemma 8. An additional application of Lemma 8 gives that

$$
\begin{aligned}
\sigma_m(P_i,P_{i+1}) &= s\big(\widetilde{\varphi}(P_i),\widetilde{\varphi}(P_{i+1});\widetilde{\lambda}'_{PQ}\big)\\
&= s\big(P_i,\psi_{P_iP_{i+1}}T_g^{\alpha(P_i,P_{i+1})}(P_{i+1});\varphi^{-1}_{P_0P_i}\big(\widetilde{\lambda}'_{PQ}\big)\big)\\
&= s\big(P_i,T_g^{\alpha(P_i,P_{i+1})}(P_{i+1});\widetilde{\lambda}_{PQ}\big)+O\Bigg(\sum_{R\in\mathcal{P}_{P_iP_{i+1}}}e^{-Br(k\cap R)}\Bigg)+O\left(l_{m_0}(k_i)\right)\\
&= s\big(P_i,P_{i+1};\widetilde{\lambda}_{PQ}\big)+\alpha(P_i,P_{i+1})+O\Bigg(\sum_{R\in\mathcal{P}_{P_iP_{i+1}}}e^{-Br(k\cap R)}\Bigg)+O\left(l_{m_0}(k_i)\right)\\
&= \sigma_0(P_i,P_{i+1})+\alpha(P_i,P_{i+1})+O\Bigg(\sum_{R\in\mathcal{P}_{P_iP_{i+1}}}e^{-Br(k\cap R)}\Bigg)+O\left(l_{m_0}(k_i)\right).
\end{aligned}
$$

Summing over all i, we obtain

$$\begin{aligned}\sigma_m(P,Q) &= \sigma_0(P,Q) + \alpha(P,Q) + \sum_{i=0}^{n} O\Big(\sum_{R\in\mathcal{P}_{P_iP_{i+1}}} e^{-Br(k\cap R)}\Big) + \sum_{i=0}^{n} O\left(l_{m_0}(k_i)\right) \\ &= \sigma_0(P,Q) + \alpha(P,Q) + O\Big(\sum_{R\in\mathcal{P}_{PQ}-\mathcal{P}} e^{-Br(k\cap R)}\Big) + O\Big(l_{m_0}\Big(k - P\cup Q\cup \bigcup_{R\in\mathcal{P}} R\Big)\Big).\end{aligned}$$

By Lemma 4, the series $\sum_{R\in\mathcal{P}_{PQ}} e^{-Br(k\cap R)}$ is convergent. Letting $\mathcal{P}$ tend to $\mathcal{P}_{PQ}$, this enables us to conclude that $\sigma_m(P,Q) = \sigma_0(P,Q) + \alpha(P,Q)$.

This completes the proof of Lemma 19, and therefore of Proposition 13. □

9.8. The global realization of shearing cocycles

In this section, we determine which transverse cocycles for λ can occur as shearing cocycles of hyperbolic metrics.

There is an obvious necessary condition for a given transverse cocycle α for λ to be the shearing cocycle σ_m of a hyperbolic metric m. Indeed, by Theorem 9, the m–length $l_m(\mu)$ of another $\mu \in \mathcal{H}(\lambda)$ is equal to $\tau(\mu,\sigma_m) = \tau(\mu,\alpha)$. If, in addition, this μ is a non-zero transverse measure for λ, then it follows from the definition of $l_m(\mu)$ that this length is positive. Consequently, for α to be the shearing distribution of some hyperbolic metric, it is necessary that $\tau(\mu,\alpha) > 0$ for every transverse non-zero measure μ for λ. Quite remarkably, this condition turns out to be sufficient.

THEOREM 20. *A transverse Hölder distribution α for λ is the shearing distribution of some hyperbolic metric if and only if $\tau(\mu,\alpha) > 0$ for every non-zero transverse measure μ for λ.*

PROOF. As in §5, let us endow $\mathcal{H}(\lambda;\mathbb{R})$ with a norm $\|\ \|$. For a hyperbolic metric m_0, Proposition 13 provides a ball $B(\sigma_{m_0},\varepsilon_0) \subset \mathcal{H}(\lambda;\mathbb{R})$ around σ_{m_0} such that every transverse cocycle in this ball is also the shearing distribution of some hyperbolic metric. Let us examine the proof of Proposition 13 in detail, to see what determines ε_0.

We start with a topological data (independent of the metric m_0) consisting of simple arcs $\bar{k}_1, \dots, \bar{k}_n$, transverse to λ, such that every component of $S-\lambda$ meets some $\bar{k}_i$. We also require that, for any hyperbolic metric m_0 and after making λ m_0–geodesic by a first isotopy, each $\bar{k}_i$ can be isotoped respecting λ to a simple m_0–geodesic arc k_i. An easy way to achieve this is to choose each $\bar{k}_i$ contained in some non-backtracking simple closed curve transverse to λ, which we can always do.

Given a hyperbolic metric m_0 and geodesic arcs k_i isotopic to the $\bar{k}_i$ as above, Lemmas 3 and 6 associate constants A_i, N_i and C_i to each k_i. Note that A_i

depends on k_i and on the metric m_0, but that C_i does not and depends only on the topology of $\bar{k}_i$ and λ.

Then, if we examine the proof of Proposition 13, and in particular the proof of Lemmas 14 and 16, we see that we can take $\varepsilon_0 = \min_i A_i/2C_i$.

Now, let us change the perspective of the problem. Consider a hyperbolic metric m_0 whose shearing cocycle σ_0 is within ε of the complement of the image of the map $\Sigma : \mathcal{T}(S) \to \mathcal{H}(\lambda;\mathbb{R})$ which associates its shearing cocycle σ_m to each metric m.

This means that $\varepsilon_0 \leqslant \varepsilon$, and therefore that there is a k_i for which it is impossible to find constants A_i, N_i which satisfy the conclusion of Lemma 3 and such that $A_i/2C_i > \varepsilon$. In other words, there is a k_i such that, for every N, there is an arc b_N contained in a leaf of λ which cuts $n_N \geqslant N$ times the arc k_i and whose length $l_{m_0}(b_N)$ is such that $l_{m_0}(b_N) \leqslant 2\varepsilon C_i (n_N - 1)$.

Let n'_N denote the number of times the arc b_N crosses the union of the k_j. On each arc k transverse to λ, consider the Dirac measure μ_k^N of weight $1/n'_N$ based at the finite set $k \cap b_N$. Since every leaf of λ meets some k_j, the total mass of μ_k^N is uniformly bounded in N. It follows that we can extract a subsequence $(N_p)_{p\in\mathbb{N}}$ such that, for every transverse arc k, the measure $\mu_k^{N_p}$ weakly converges to some measure μ_k as p tends to ∞. Since $n'_{N_p} \geqslant n_{N_p} \geqslant N_p$ tends to ∞, these measures μ_k are invariant under homotopy of k respecting λ, and therefore define a transverse measure μ for λ.

By construction, the length $l_{m_0}(\mu)$ is equal to the limit of the

$$l_{m_0}\left(b_{N_p}\right)/n'_{N_p} \leqslant 2\varepsilon C_i \left(n_{N_p} - 1\right)/n'_{N_p} \leqslant 2\varepsilon C_i.$$

It follows that $l_{m_0}(\mu) \leqslant 2\varepsilon C_i$.

On the other hand, $\sum_j \mu(k_j)$ is the limit of the $\sum_j \mu_{k_j}^{N_p}(k_j) = 1$, and is therefore equal to 1. Since $\mu(k_j) \leqslant C_j \|\mu\|$ by definition of C_j in Lemma 6, we conclude that $\|\mu\| \geqslant \left(\sum_j C_j\right)^{-1}$.

As a conclusion, if the shearing cocycle σ_0 of the metric m_0 is within ε of the complement of the image of the map $\mathcal{T}(S) \to \mathcal{H}(\lambda;\mathbb{R})$, there is a transverse measure μ for λ such that

$$\tau(\mu,\sigma_0)/\|\mu\| = l_{m_0}(\mu)/\|\mu\| \leqslant 2\varepsilon C_i \left(\sum_j C_j\right)^{-1} \leqslant 2\varepsilon.$$

By weak compactness of the space of transverse measures μ with $\|\mu\| = 1$ and by continuity of τ it follows that, if $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ is in the boundary of the image of $\mathcal{T}(S) \to \mathcal{H}(\lambda;\mathbb{R})$, then there exists a transverse measure μ with $\tau(\mu,\alpha) = 0$.

Therefore, the image of $\mathcal{T}(S) \to \mathcal{H}(\lambda;\mathbb{R})$ is closed in the set $\mathcal{C}(\lambda)$ of those $\alpha \in \mathcal{H}(\lambda;\mathbb{R})$ such that $\tau(\mu,\alpha) > 0$ for every transverse measure μ. This image is also open in $\mathcal{C}(\lambda)$ by Proposition 13. Since $\mathcal{C}(\lambda)$ is defined by linear inequalities, it is connected. It follows that $\mathcal{C}(\lambda)$ is exactly equal to the image of $\mathcal{T}(S) \to \mathcal{H}(\lambda;\mathbb{R})$. □

Corollary 21. *Consider the map* $\Sigma : \mathcal{T}(S) \to \mathcal{H}(\lambda; \mathbb{R})$ *which associates its shearing cocycle* σ_m *to each hyperbolic metric* m*. The image of* Σ *is an open convex cone in* $\mathcal{H}(\lambda; \mathbb{R})$ *bounded by finitely many faces.*

Proof. By [**Ka**] (compare [**Pa1**][**PeH**][**Bo4**, §4]), the geodesic lamination admits only finitely many ergodic transverse measures $\mu_1, \dots, \mu_n$, and every transverse measure is a linear combination with non-negative coefficients of these μ_i. □

9.9. Hyperbolic surfaces with cusps

9.10. Earthquakes

9.11. The Nielsen realization problem

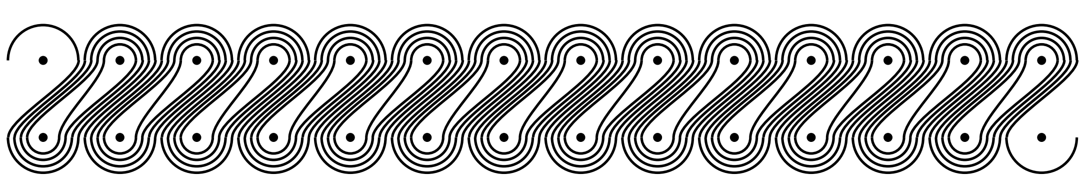

APPENDIX A

BASIC TOPOLOGICAL PROPERTIES OF SURFACES

APPENDIX B

SURFACES OF NEGATIVE CURVATURE

every surface admits a hyperbolic metric with totally geodesic boundary and/or finite area
Every closed curve is homootpic to a unique closed geodesic, every simple closed curve is homotopic to a simple closed geodesic, or to twice an orientation reversing simple closed geodesic.
boundary at infinity,
quasi-geodesics
quasi-isometries, Milnor-Svarc, boundary at infinity in terms of π_1
the Hölder structure of the bdry at infinity
case with finite area, Floyd

B.1. Riemannian metrics

We are gathering in this section some of the basic facts in riemannian geometry which we need in our monograph. We refer the reader to standard textbooks on this topic, such as ??, for a proof of these facts.

We first consider the case of manifolds without boundary. A ***riemannian metric*** on the manifold M without boundary is the data of a positive definite bilinear form, denoted by $\langle\ ,\ \rangle_x$ or just $\langle\ ,\ \rangle$, on each tangent space T_xM, in such a way that the bilinear form $\langle\ ,\ \rangle_x$ depends differentiably on the point $x \in M$. The ***norm*** of a vector $v \in T_xM$ is $\|v\| = \langle v, v\rangle^{\frac{1}{2}}$.

If we are given a riemannian metric m on M, the m–***length*** of a parametrized curve $\alpha : [a, b] \to M$ is equal to the integral

$$\int_a^b \left\| \frac{d\alpha}{dt} \right\| dt.$$

As in calculus, this length is invariant under change of parametrization of α. The curve α is ***parametrized by arc length*** if $\left\|\frac{d\alpha}{dt}\right\| = 1$ everywhere, namely if the length of each restriction $\alpha_{|[u,v]}$ is equal to $|v - u|$.

A ***unit speed geodesic*** is a parametrized curve $\gamma : I \to M$ defined on an interval $I \subset \mathbb{R}$, parametrized by arc length, and which is locally length minimizing in the sense that, every t in the interior of I has a neighborhood $]0, \varepsilon[$ such that, for every u, v, the restriction $\gamma_{|[u,v]}$ has minimum length among all the curves going from $\gamma(u)$ to $\gamma(v)$. A classical fact is that this length minimizing property can be expressed in local coordinates by a second order ordinary differential equation, so that the following holds: For every $x \in M$ and every $v \in T_xM$

with $\|v\| = 1$, there is a unit speed geodesic $\gamma :]-\varepsilon, \varepsilon[\rightarrow M$, defined on a neighborhood $]-\varepsilon, \varepsilon[$ of 0, with $\gamma(0) = x$ and $\frac{d}{dt}\gamma(t)_{|t=0} = v$; in addition, two such geodesics coincide where they are both defined.

We will identify two unit speed geodesics $\gamma_1 : I_1 \rightarrow M$ and $\gamma_2 : I_2 \rightarrow M$ when they differ by composition by a map $t \mapsto \pm t + \text{constant}$ sending I_1 to I_2. A ***geodesic*** will be an equivalence class for this relation.

A ***closed geodesic*** is a pair consisting of a geodesic $\gamma :]-\infty, +\infty[\rightarrow M$, considered modulo reparametrization by a map $t \mapsto \pm t + \text{c}$, and of a number $L > 0$ such that $\gamma(s + L) = \gamma(s)$ for every s. We emphasize that the number L is part of the data of a closed geodesic. In particular, if (γ, L) is a closed geodesic, then (γ, nL) is also a closed geodesic for every positive integer n, which is different from (γ, L) if $n \geqslant 2$. We will say that (γ, nL) is the ***closed geodesic which wraps*** n times around (γ, L).

A geodesic is ***bi-infinite*** if it is represented by a unit speed geodesic $\gamma :]-\infty, +\infty[\rightarrow M$ defined over the whole real line $]-\infty, +\infty[$. In particular, a closed geodesic defines a bi-infinite geodesic.

Let us now consider the case of manifolds with boundary. Given a manifold M with boundary, we can slightly enlarge it to a manifold M^+ without boundary, of the same dimension. Indeed, any locally finite atlas for M provides such an extension. In addition, the 'germ' of M^+ is unique in the sense that, if two such extensions M_1^+ and M_2^+ are respectively associated to the atlases $\mathcal{A}_1$ and $\mathcal{A}_2$, the changes of charts between charts in $\mathcal{A}_1$ and charts in $\mathcal{A}_2$ provide a natural identification between a neighborhood of M in M_1^+ and a neighborhood of M in M_2^+.

A ***riemannian metric*** on M is the restriction to M of a riemannian metric on M^+, as defined above. A ***unit speed geodesic*** or a ***geodesic***in M is, respectively, a unit speed geodesic or a geodesic of M^+ which is contained in M. The description of geodesics by an ordinary differential equaiton shows that this depends only on the riemannian metric of M, and not on its extension to a riemannian metric on M^+.

The manifold M has ***locally convex boundary*** for the riemannian metric m if every $x \in \partial M$ admits arbitrarily small neighborhoods U in the extension manifold M^+ such that $M \cap U$ is convex, in the sense that, for every $y, z \in M \cap U$, every shortest curve from y to z in M^+ is actually contained in $M \cap U$. Again, this property can be translated in local coordinates into a property of the derivatives of the coefficients of the metric, which shows that its depends only on the riemannian metric of M, and not on its extension to a riemannian metric on M^+.

A riemannian metric m on a connected manifold M induces a distance function d on M, defined by the property that the distance $d(x, y)$ is the infimum of the m–lengths of all the curves going from x to y. The only non-trivial thing to check here is that $d_m(x, y) = 0$ only when $x = y$. In addition, it can easily be shown that the topology induced on M by this distance function coincides with the original topology of M.

We will say that the riemannian metric m is ***complete*** if (M, d_m) is a complete metric space. We will mostly deal with two situations where the set-up automatically guarantees that the riemannian metric considered is complete. The first case is when the manifold M is compact. The second case is when M is a covering of a compact manifold N, and where the riemannian metric of M is the lift of a riemannian metric on N.

The celebrated Hopf-Rinow Theorem states:

LEMMA. *Let M be a connected riemannian manifold with locally convex boundary. If, in addition, M is complete then:*

a) Every unit speed geodesic $\gamma : I \to M$ admits a unique maximal extension $\gamma : J \to M$ where the interval $J \supset I$ is a closed subset of $\mathbb{R}$ and where, if the endpoint a of J is different from $\pm\infty$, the extension $\gamma : J \to M$ transversely hits ∂M at the point $\alpha(a)$.

b) For any two $x, y \in M$, there is a geodesic which goes from x to y and has minimum length among all the curves going from x to y. □

A geodesic $\gamma : I \to M$ is ***complete***if it is equal to its maximal extension, as in Lemmma ??a). Note that every bi-infinite geodesic is complete, and Lemma ?? shows that the converse is true if M is complete without boundary.

We now discuss curvature properties of riemannian metrics. Traditionally, the sectional curvature of the metric m along a plane $\Pi \subset T_xM$ is defined by a long expression involving connections and/or various derivatives of the coefficients of the metric in local coordinates. However, for all practical purposes, we will only need to understand how it impacts the variation of geodesics.

We will restrict ourselves to the case where M is a surface S, so that we do not have to worry about picking a plane $\Pi \subset T_xM = T_xS$. Some formulas are also easier in this case. The sectional curvature along the plane $\Pi = T_xS$ is then called the ***Gauss curvature*** of the metric m at the point $x \in S$.

In the surface S, consider a 1–***parameter family of unit speed geodesics***, namely a differentiable map $(s,t) \mapsto \alpha_t(s) \in S$ defined on the product $I \times J$ of two intervals and such that, for every $t \in J$, the map $s \mapsto \alpha_t(s)$ is a unit speed geodesic. We should here think of s as the arc length parameter along the geodesic α_t, while t is the 'time' parameter as the geodesic α_t moves in the deformation. We can decompose the variation vector $V(s,t) = \frac{\partial}{\partial t}\alpha_t(s)$ into the sum of a tangent component $V^T(s,t)$ parallel to the tangent vector $\frac{\partial}{\partial s}\alpha_t(s)$ of α_t, and of a normal component $V^\perp(s,t)$ orthogonal to $\frac{\partial}{\partial s}\alpha_t(s)$.

LEMMA. *On the surface S, endowed with a riemannian metric m, let $(s,t) \mapsto \alpha_t(s)$, $(s,t) \in I \times J$, be a 1–parameter family of unit speed geodesics, and consider the variation vector $V(s,t) = \frac{\partial}{\partial t}\alpha_t(s)$. Then, for a fixed t, the norm $\left\|V^T(s,t)\right\|$ of the tangent component of V is a constant function of s. The norm $\left\|V^\perp(s,t)\right\|$ of the normal component satisfies the following differential equation*

$$\frac{\partial^2}{\partial s^2}\left\|V^\perp(s,t)\right\| + K(\alpha_t(s))\left\|V^\perp(s,t)\right\| = 0$$

where $K(x) \in \mathbb{R}$ is the Gauss curvature of the metric m at $x \in S$. □

The equation of Lemma ?? is another way to express the famous ***Jacobi equation*** for the variation vector V of a 1–parameter family of geodesics.

PROPOSITION. *Let $\widetilde{M}$ be a simply connected manifold endowed with a complete riemannian metric of non-positive sectional curvature, for which the boundary $\partial\widetilde{M}$ is locally convex. For every $x,\, y \in \widetilde{M}$ there is a unique geodesic γ_{xy} going from x to y. In addition, the length of γ_{xy} is equal to the distance $d\,(x,y)$, and this geodesic depends differentiably on x and y in the following sense: If we consider γ_{xy} as a unit speed parametrized geodesic $\gamma_{xy} : [0, d\,(x,y)] \to M$ with $\gamma_{xy}\,(0) = x$ and $\gamma_{xy}\,(d\,(x,y)) = y$, then $d\,(x,y)$ and $\gamma_{xy}\,(t)$ depend differentiably on x, y and t.* □

PROPOSITION. *Let $\widetilde{M}$ be a simply connected manifold endowed with a complete riemannian metric of non-positive sectional curvature, for which the boundary $\partial\widetilde{M}$ is locally convex. Let $\alpha : I \to \widetilde{M}$ and $\beta : J \to \widetilde{M}$ be two unit speed parametrized geodesics. Then the function $f\,(s,t) = d\,(\alpha\,(s)\,, \beta\,(t))$ is convex on the product of the two intervals I and J.* □

B.2. Surfaces of negative curvature

Let us apply this to the special case where the curvature of S is negative everywhere, and where the geodesics α_t pass through a single point, or are orthogonal to a given geodesic. Note that, if the surface S is compact and if $K\,(x)$ at each $x \in K$, then there exists constants a, b such that $-a^2 \leqslant K\,(x) \leqslant -b^2$, by continuity.

LEMMA. *Suppose that, at each point x of the surface S, the Gauss curvature $K\,(x)$ of the metric m is such that $-a^2 \leqslant K\,(x) \leqslant -b^2$ for every $x \in S$, for two constants a, $b > 0$. Let $(s,t) \mapsto \alpha_t\,(s)$, $(s,t) \in I \times J$, be a 1–parameter family of unit speed geodesics such that $\alpha_t\,(0) = x_0$ for every $t \in J$. Consider the initial tangent vector $\frac{\partial}{\partial s}\alpha_t\,(s)_{|s=0} \in T_{x_0}S$ and its time derivative $\frac{\partial}{\partial t}\frac{\partial}{\partial s}\alpha_t\,(s)_{|s=0} \in T_{x_0}S$. Then, $V\,(s,t) = \frac{\partial}{\partial t}\alpha_t\,(s)$ is orthogonal to the tangent vector $\frac{\partial}{\partial s}\alpha_t\,(s)$, and*

$$b^{-1}\sinh(bs)\left\|\frac{\partial}{\partial t}\frac{\partial}{\partial s}\alpha_t\,(s)_{|s=0}\right\| \leqslant \|V\,(s,t)\| \leqslant a^{-1}\sinh(as)\left\|\frac{\partial}{\partial t}\frac{\partial}{\partial s}\alpha_t\,(s)_{|s=0}\right\|$$

SKETCH OF PROOF. This easily follows from Lemma ??. Indeed, since $\left\|V^T\,(s,t)\right\|$ is independent of s and $V\,(0,t) = 0$, it follows that $V^T\,(s,t) = 0$ everywhere, and therefore that $V\,(s,t)$ is orthogonal to $\frac{\partial}{\partial s}\alpha_t\,(s)$. In particular, $V^T\,(s,t) = V\,(s,t)$. Also, it is not too hard to see that $\frac{\partial}{\partial s}\|V\,(s,t)\|_{|s=0} = \left\|\frac{\partial}{\partial t}\frac{\partial}{\partial s}\alpha_t\,(s)_{|s=0}\right\|$; this can be seen by analogy with the euclidean case if we intuitively believe that the curvature only impacts derivatives of $V\,(s,t)$ of order 2 and higher, and can be rigorously proved using the connection defined by the metric. Since $\|V\,(0,t)\| = 0$, the inequality then follows by comparing the solutions of the order 2 differential equation of Lemma ?? to the solutions of the equations $y'' + ay = 0$ and $y'' + by = 0$ with the same initial data. □

LEMMA. *Suppose that, at each point x of the surface S, the Gauss curvature $K(x)$ of the metric m is such that $-a^2 \leqslant K(x) \leqslant -b^2$ for every $x \in S$, for two constants a, $b > 0$. Let $(s,t) \mapsto \alpha_t(s)$, $(s,t) \in I \times J$, be a 1–parameter family of unit speed geodesics such that the curve $t \mapsto \alpha_t(0)$ is a unit speed geodesic which is orthogonal to the α_t; in particular, $\frac{\partial}{\partial t}\alpha_t(0)$ is orthogonal to $\frac{\partial}{\partial t}\alpha_t(s)_{|s=0}$. Then, $V(s,t) = \frac{\partial}{\partial t}\alpha_t(s)$ is orthogonal to the tangent vector $\frac{\partial}{\partial s}\alpha_t(s)$, and*

$$\cosh(bt) \leqslant \|V(s,t)\| \leqslant \cosh(at).$$

SKETCH OF PROOF. Again, the fact that $V(s,t)$ is orthogonal to $\frac{\partial}{\partial s}\alpha_t(s)$ follows from the fact that $\left\|V^T(s,t)\right\|$ is independent of s and is equal to 0 when $s = 0$. The fact that $t \mapsto \alpha_t(0)$ is a unit speed geodesics implies that $\|V(0,t)\| = 1$ and that $\frac{\partial}{\partial s}\|V(s,t)\|_{|s=0} = 0$ (there does not seem to be any easy way to avoid connections this time). The inequality of Lemma ?? again follows by comparing the solutions of the order 2 differential equation of Lemma ?? to the solutions of the equations $y'' + ay = 0$ and $y'' + by = 0$ with the same initial data. □

LEMMA. *Let $\widetilde{S}$ be a simply connected surface endowed with a complete metric whose curvature is everywhere less than or equal to $-a^2$, for some $a > 0$, and such that the boundary $\partial\widetilde{S}$ is locally convex. Let γ be a geodesic going from x to y, and let γ' be a geodesic going from x' to y'. Suppose that the distance from x to x', and the distance from y to y' are both at most $A > 0$. Let the point $z \in \gamma$ be at distance at least $B > 0$ from x and y in γ. Then, z is at distance at most $c(A,a)\,e^{-aB}$ from γ', where the constant $c(A,a)$ depends only on A and a.*

In particular, if the end points of the geodesic arcs γ and γ' are not too far from each other and if these arcs are very long, then most of γ is very close to γ'.

PROOF OF LEMMA ??. Consider the geodesic arc γ'' going from x to y'. Let L, L', L'' be the respective lengths of γ, γ', γ'', and let L_0 be the minimum of $\{L, L', L''\}$. Note that, by the triangle inequality, L, L' and L'' are all in the interval $[L_0, L_0 + 2A]$. Also, the existence of z guarantees that $B \leqslant L/2 \leqslant L_0/2 + A$.

Let $\alpha : [0,D] \to \widetilde{S}$ be the unit speed geodesic going from $y = \alpha(0)$ to $y' = \alpha(D)$ and, for each $t \in [0,D]$, consider the unit speed geodesic $\beta_t : [0, d(t)] \to \widetilde{S}$ going from $x = \beta_t(0)$ to $\alpha(t) = \beta_t(d(t))$ where $d(t)$ is the distance from x to $\alpha(t)$. In particular, $\beta_0 = \gamma$ and $\beta_D = \gamma''$. Also, D is the distance from y to y', and therefore $D \leqslant A$ by hypothesis.

Let $s_0 \in [0,D]$ be such that $\beta_0(s_0) = z$. We want to estimate the length of the curve $t \mapsto \beta_t(s_0)$, going from $z = \beta_0(s_0) = \gamma(s_0)$ to the point $z'' = \beta_D(s_0) = \gamma''(s_0)$.

Consider the variation vector $V(s,t) = \frac{\partial}{\partial t}\beta_t(s)$. Since $V(0,t) = 0$, it follows from Lemma ?? that its tangent component $V^T(s,t)$ is always 0. From the differential equation of Lemma ??, we also conclude that the norm of $V(s,t) =$

$V^{\perp}(s,t)$ is bounded above by the solution of the differential equation $\frac{d^2}{ds^2}u = a^2u$ that is such that $u(0)=0$ and $u(d(t)) = \|V(d(t),t)\|$. Namely,

$$\|V(s,t)\| \leqslant \|V(d(t),t)\| \sinh(as)/\sinh(ad(t)).$$

Also, note that

$$\begin{aligned}\frac{d}{dt}\alpha(t)_{|t=t_0} &= \frac{d}{dt}\beta_t(d(t))_{|t=t_0}\\ &= \frac{d}{dt}\beta_t(d(t_0))_{|t=t_0} + \frac{d}{ds}\beta_{t_0}(s)_{|s=d(t_0)}\frac{d}{dt}d(t)_{|t=t_0}\\ &= V(d(t_0),t_0) + \frac{d}{ds}\beta_{t_0}(s)_{|s=d(t_0)}\frac{d}{dt}d(t)_{|t=t_0}.\end{aligned}$$

Because $V(s,t_0)$ is orthogonal to the vector $\frac{d}{ds}\beta_{t_0}(s)$, and because $\left\|\frac{d}{dt}\alpha(t)_{|t=t_0}\right\| = 1$ since α is parametrized by arc length, we conclude that $\|V(d(t_0),t_0)\| \leqslant 1$ for every t_0.

It follows that the distance from $z = \beta_0(s_0)$ to $z'' = \beta_D(s_0)$ is at most

$$\begin{aligned}\int_0^D \left\|\frac{d}{dt}\beta_t(s_0)\right\| dt &= \int_0^D \|V(s_0,t)\|\,dt\\ &\leqslant \int_0^D \|V(d(t),t)\| \sinh(as_0)/\sinh(ad(t))\,dt\\ &\leqslant A\sinh(a(L_0+2A-B))/\sinh(a(L_0-A/2))\\ &\leqslant 2Ae^{5aA/2}e^{-aB}\end{aligned}$$

since $s_0 \leqslant L - B \leqslant L_0 + 2A - B$ by hypothesis on z, and $d(t) \geqslant L_0 - A/2$ by the triangle inequality.

We can now apply the same process to the family of geodesics joining y' to the points on the geodesic arc going from x to x'. This provides a curve joining z'' to a point z' on γ', whose length is again bounded by $2Ae^{5aA/2}e^{-aB}$. (Note that s_0 is now replaced by the distance $L'' - s_0$ from z'' to y', which satisfies the same estimate $L'' - s_0 \leqslant L'' - B \leqslant L_0 + 2A - B$ as s_0). As a consquence, the distance from z to z' is at most $4Ae^{5aA/2}e^{-aB}$.

Therefore, for every point z in γ which is at distance $\geqslant B$ from the end points x and y, we were able to find a point z' in γ' with $d(z,z') \leqslant c(A,a)e^{-aB}$, with $c(A,a) = 4Ae^{5aA/2}$. This concludes the proof of Lemma ??. □

Lemma. *Let $\widetilde{S}$ be a simply connected surface endowed with a complete metric with locally convex boundary and whose curvature is everywhere bounded between two constants $-a^2$ and $-b^2$, for $0 < b < a \leqslant \infty$. Let T be a triangle with totally geodesic sides in $\widetilde{S}$, with vertices x_1, x_2, x_3. Let $\theta_i \in [0,\pi]$ denote the internal angle of T at x_i. Then:*

a) $\pi - a^2\,\mathrm{area}(T) \leqslant \theta_1+\theta_2+\theta_3 \leqslant \pi - b^2\,\mathrm{area}(T)$;

b) $\operatorname{area}(T) \leqslant d(x_1, x_2)/b$;
c) $\theta_3 \leqslant 2b\,d(x_1, x_2)\,e^{b\,d(x_1,x_2)}e^{-b\max\{d(x_3,x_1),d(x_3,x_2)\}}$.

PROOF. The first property is a consequence of the Gauss-Bonnet formula, which in this case states that

$$\iint_T K(x)\,dx = \chi(T) - \sum_{i=1}^{3}(\pi - \theta_i).$$

Since the Euler characteristic $\chi(T)$ is equal to 2π and since the curvature $K(x)$ is everywhere bounded between $-a^2$ and $-b^2$, Property a) immediately follows.

To prove Properties b) and c), consider the unit speed geodesic $\alpha : [0, d(x_1, x_2)] \to \widetilde{S}$ going from $x_1 = \alpha(0)$ to $x_2 = \alpha(d(x_1, x_2))$, and let $\beta_t : [0, d(t)] \to \widetilde{S}$ be the unit speed geodesic going from $x_3 = \beta_t(0)$ to $\alpha(t) = \beta_t(d(t))$, where $d(t)$ is the distance from x_3 to $\alpha(t)$.

Consider the variation vector $V(s,t) = \frac{\partial}{\partial t}\beta_t(s)$. As in the proof of Lemma ??,

$$\frac{d}{dt}\alpha(t) = \frac{d}{dt}\beta_t(d(t)) = V(d(t), t) + \frac{\partial}{\partial s}\beta_t(s)_{|s=d(t)}\,\frac{d}{dt}d(t),$$

from which we conclude that $\|V(d(t), t)\| \leqslant 1$, because $\left\|\frac{d}{dt}\alpha(t)\right\| = 1$ and $V(s,t)$ is orthogonal to $\frac{\partial}{\partial s}\beta_t(s)$. Comparing $\|V(s,t)\|$ to the solution of the equation $y'' = b^2 y$ that is such that $y(0) = 0$ and $y(d(t)) = 1$ (namely to the function $y(s) = \sinh(bs)/\sinh(b\,d(t))$), we obtain that

$$\|V(s,t)\| \leqslant \sinh(bs)/\sinh(bd(t))$$

since $\frac{d^2}{ds^2}\|V(s,t)\| \geqslant b^2\|V(s,t)\|$ by Proposition ??.

As a consequence,

$$\begin{aligned}
\operatorname{area}(T) &= \int_0^{d(x_1,x_2)}\int_0^{d(t)}\|V(s,t)\|\,ds\ \ dt \\
&\leqslant \int_0^{d(x_1,x_2)}\int_0^{d(t)}\sinh(bs)/\sinh(b\,d(t))\ \ ds\,dt \\
&\leqslant \int_0^{d(x_1,x_2)}\frac{\cosh(b\,d(t)) - 1}{a\sinh(b\,d(t))}dt \\
&\leqslant d(x_1, x_2)/b.
\end{aligned}$$

This proves Property b).

Finally, let $\theta(t)$ be the angle from the unit vector $\frac{d}{ds}\beta_0(s)_{|s=0}$ to $\frac{d}{ds}\beta_t(s)_{|s=0}$. Note that $\theta(0) = 0$ and $\theta(d(x_1, x_2)) = \theta_3$. By Lemma ??,

$$\begin{aligned}
\left|\frac{d}{dt}\theta(t)\right| &\leqslant b\,\|V(d(t), t)\|/\sinh(b\,d(t)) \leqslant a/\sinh(b\,d(t)) \\
&\leqslant 2b e^{-b\,d(t)} \leqslant 2b e^{b\,d(x_1,x_2)}e^{-b\max\{d(x_3,x_1),d(x_3,x_2)\}}
\end{aligned}$$

since $d(t) \geqslant \max\{d(x_3,x_1), d(x_3,x_2)\} - d(x_1,x_2)$ by the triangle inequality. Property c) immediately follows by integrating in t. □

We often make use of the ***projective tangent bundle*** $PT(S)$**,** consisting of all pairs (x,l) where $x \in S$ is a point of S and where l is a line (through the origin) in the tangent vector space T_xS of S at x. If we endow the surface S with a riemannian metric, it is also technically convenient to consider the ***unit tangent bundle*** $T^1(S)$ of the riemannian surface S, consisting of all pairs (x,v) where $x \in S$ and where $v \in T_xS$ is such that $\|v\| = 1$. The relationship between these two spaces is that $PT(S)$ is the quotient of $T^1(S)$ under the free action of the group $\mathbb{Z}/2$ generated by the map $(x,v) \mapsto (x,-v)$.

The natural projections $PT(S) \to S$ and $T^1(S) \to S$ are locally trivial bundle maps, with fibers homeomorphic to the circle. In particular, $PT(S)$ and $T^1(S)$ are 3–dimensional manifolds, with boundary corresponding to the pre-image of ∂S under the projection map. Note that these 3–manifolds are orientable, and even canonically oriented, even when S is non-orientable.

The tangent space of $T^1(S)$ at (x,v) naturally splits as the direct sum of a ***vertical line*** L and of a ***horizontal plane*** P defined as follows. The line L is the tangent space $T_vp^{-1}(x)$ of the fiber $p^{-1}(x)$ at the vector v. The horizontal plane P is defined as the image of a left inverse $\sigma : T_xS \to T_{(x,v)}T^1(S)$ of the tangent map $T_{(x,v)}p : T_{(x,v)}T^1(S) \to T_xS$, constructed by parallel transport along curves passing through x. This map σ is a little easier to describe in dimension 2: If $w \in T_xS$ and is $g_w :]-\varepsilon, +\varepsilon[\to S$ is the unit speed geodesic with $g_w(0) = x$ and $\frac{d}{ds} g_w(s)_{|s=0} = w/\|w\|$, then $\sigma(w) = \|w\| \frac{d}{ds}(g_W(s), v_s) \in T_{(x,v)}T^1(S)$ where $v_t \in T_{g_w(s)}S$ denotes the unit vector which makes the same angle with $\frac{d}{ds} g_w(s)$ at $g_w(s)$ as $v = v_0$ with $w/\|w\| = \frac{d}{ds} g_w(s)_{|s=0}$ at $x = g_w(0)$. If we extend this definition by $\sigma(0) = 0$, it can be shown that the map $\sigma : T_xS \to T_{(x,v)}T^1(S)$ so defined is linear. Clearly, $T_xp \circ \sigma = \mathrm{Id}$ and, since L is contained in the kernel of T_xp, it immediately follows that the tangent space $T_{(x,v)}T^1(S)$ is the direct sum of L and of the plane P image of σ.

This gives us a way to introduce a canonical metric on $T^1(S)$. The quadratic form on T_xS defined by the metric m retricts to a riemannian metric on the circle $p^{-1}(x) \subset T_xS$, which induces a quadratic form on $L = T_vp^{-1}(x)$. Similarly, m induces a quadratic form on $P \cong T_xS$ if we identify the plane P to T_xS by the restriction of the tangent map T_xp (or by its inverse σ). We can then consider on $T_{(x,v)}T^1(S) = P \oplus L$ the positive definite quadratic form which restricts to the above forms on P and L and which makes these two spaces orthogonal. This defines on $T^1(S)$ a riemannian metric, which we will also denote by m if there is no ambiguity.

The above metric on $T^1(S)$ is invariant under the map $(x,v) \mapsto (x,-v)$, and therefore induces a riemannian metric on $PT(S) = T^1(S)/(\mathbb{Z}/2)$.

LEMMA. *Let the surface S be endowed with a complete riemannian metric with convex boundary and whose curvature is everywhere bounded between two constants $-a^2$ and $-b^2$, for $0 < b < a \leqslant \infty$. Let g be a bi-infinite geodesic in S which passes through the point $x \in S$ and is at that point tangent to the line $l \subset T_xS$, and let g' be another bi-infinite geodesic passing through a point*

$x' \in S$ and tangent to the line $l' \subset T_{x'}S$ there. If the two geodesics g and g' are disjoint, then the distance from (x,l) to (x',l') in $PT(S)$ is bounded by $c\,d(x,x')$, where the constant c depends only on a and b, and where $d(x,x')$ is the distance from x to x' in S.

In particular, disjoint geodesics which pass through nearby points in S must do so with directions which are close to each other.

PROOF. Lifting the situation to the universal cover $\widetilde{S}$ of S if necessary, we can assume without loss of generality that S is simply connected. Pick an orientation for S. Let α be an oriented geodesic going from x to x', orient g and g' to the right of α so that, if θ (resp. θ') denotes the angle from the oriented geodesic α to the oriented geodesic g (resp. g') at x (resp. x'), these two angles are in the interval $]0,\pi[$. By definition of the metric of $PT(S)$ and $T^1(S)$, it clearly suffices to show that $|\theta-\theta'| \leqslant c_1 d(x,x')$ for some constant c_1 depending only on a and b.

Pick a point $y \in g$ very far to the left of α, and consider the triangle T with vertices x, x' and y. Let $\theta_{x'}$ and θ_y be the internal angles of T at x' and y, respectively. Note that the internal angle of T at x is equal to θ. Also, because $T-\alpha$ is completely contained in the component of $S - g \cup g'$ tant contains y (by convexity of the frontier of the components of $S-g$ and $S-g'$ that contain y), $\theta_{x'} \leqslant \pi - \theta'$.

Then,

$$\theta' - \theta \leqslant \pi - \theta_{x'} - \theta \leqslant \theta_y + a^2 \operatorname{area}(T) \leqslant \theta_y + a^2 d(x,x')/b$$

by Lemma ??(a) and (b). Lemma ??(c) shows that θ_y is arbitrary small if y is chosen sufficiently far from x on g. We conclude that $\theta'-\theta \leqslant a^2 d(x,x')/b$.

Applying the same reasoning to a point $z \in g$ which is very far to the right of α, we similarly obtain that $\theta - \theta' \leqslant a^2 d(x,x')/b$, and therefore that $|\theta'-\theta| \leqslant a^2 d(x,x')/b$, which concludes the proof. □

B.3. Existence of metrics of negative curvature

The typical exemple of a surface with a riemannian metric of negative curvature is the ***hyperbolic plane***$\mathbb{H}^2$ defined as follows. As a set, $\mathbb{H}^2$ is the open unit disk $\left\{x \in \mathbb{R}^2; d_{\mathrm{e}}(x,O) < 1\right\}$ in the euclidean plane, where $d_{\mathrm{e}}(\ ,\ )$ denote the euclidean distance and where O is the origin of $\mathbb{R}^2$. The hyperbolic metric of $\mathbb{H}^2$ is the riemannian metric which at x is $2/\left(1-d_{\mathrm{e}}(O,x)^2\right)$ times[2] the euclidean metric.

The ***circle at infinity*** of $\mathbb{H}^2 \subset \mathbb{R}^2$ is the circle $\partial_\infty \mathbb{H}^2 = \left\{x \in \mathbb{R}^2; d_{\mathrm{e}}(x,O) < 1\right\}$ delimiting $\mathbb{H}^2$ in $\mathbb{R}^2$. Note that $\partial_\infty\mathbb{H}^2$ is disjoint from $\mathbb{H}^2$, and at this point seems to depend on our specific description of $\mathbb{H}^2$ as a subset of $\mathbb{R}^2$. However, we will see in Section B.5 that $\partial_\infty\mathbb{H}^2$ admits a more intrinsic description, in terms of the geometry of the riemaniann manifold $\mathbb{H}^2$.

---

[2]We are here using the topologist's convention to rescale metrics: When we multiple a riemannian metric by $\lambda > 0$, we mean that the distances are multiplied by $\lambda > 0$. In the same situation, a differential geometer would say that the metric is multiplied by λ^2.

Lemma. *The hyperbolic plane $\mathbb{H}^2$ has the following properties:*

(i) the riemannian metric of $\mathbb{H}^2$ has constant curvature -1 and is complete;

(ii) the bi-infinite geodesics are exactly the intersections of $\mathbb{H}^2$ with euclidean circles of $\mathbb{R}^2$ (including straight lines) that are orthogonal to the circle at infinity $\partial_\infty \mathbb{H}^2$;

(iii) for every x, $y \in \mathbb{H}^2$ and every linear map $\varphi : T_x\mathbb{H}^2 \to T_y\mathbb{H}^2$ which sends the quadratic form defined by the hyperbolic riemannian metric on the tangent space $T_x\mathbb{H}^2$ to that defined on $T_y\mathbb{H}^2$, there is a unique hyperbolic isometry $\Phi : \mathbb{H}^2 \to \mathbb{H}^2$ whose tangent map $T_x\Phi : T_x\mathbb{H}^2 \to T_y\mathbb{H}^2$ is equal to φ. □

Exercise. Show that the intersection with $\mathbb{H}^2$ of any euclidean line passing through the origin is a geodesic of $\mathbb{H}^2$. Then use Proposition ?? to check that the curvature of $\mathbb{H}^2$ at each x is really equal to -1. Also, show that every closed hyperbolic ball centered at O is compact, and conclude that $\mathbb{H}^2$ is a complete metric space.

Lemma. *For every a, b, $c > 0$, there is a right-angled hexagon H in $\mathbb{H}^2$, with vertices A, A', B, B', C, C' occuring in this order as one goes around ∂H, such that the hyperbolic distances $d\,(A, A')$, $d\,(B, B')$ and $d\,(C, C')$ are respectively equal to a, b and c. In addition, the hexagon H is unique up to isometry of $\mathbb{H}^2$.*

Proof. Set A to be the origin O, and let A' be the point on the euclidean segment $[0, 1[\times \{0\} \subset \mathbb{H}^2 \subset \mathbb{R}^2$ which is at hyperbolic distance a from $A = O$ (the reader can check that A' has euclidean coordinates $(\tanh(a/2), 0)$ in $\mathbb{R}^2$). Pick a number $x \geqslant 0$, and let C'_x be the point on the euclidean segment $\{0\} \times [0, 1[\subset \mathbb{H}^2$ which is at distance x from O. Let g be the bi-infinite geodesic of $\mathbb{H}^2$ which passes through A' and is orthogonal to the geodesic $\mathbb{H}^2 \cap \mathbb{R} \times \{0\}$, and let h_x be the bi-infinite geodesic which passes through C'_x and is orthogonal to the geodesic $\mathbb{H}^2 \cap \{0\} \times \mathbb{R}$. Finally, consider the point $C_x \in h_x$ which is at distance c from C'_x and has positive first coordinate in $\mathbb{R}^2$, and let k_x be the bi-infinite geodesic which passes through h_x and is there orthogonal to h_x.

Figure ??. Constructing right-angles hexagons.

We claim that, if x is large enough, the geodesics h_x and k_x are contained in an arbitrarily small neighborhood of the point $(0, 1)$ in $\mathbb{R}^2$. Indeed, when $d = 0$, we can find a geodesic $g1$ which is orthogonal to the geodesic $\mathbb{H}^2 \cap \{0\} \times \mathbb{R}$, such that h_0 and k_0 are both contained in a component P of $\mathbb{H}^2 - g_1$. Note that P is also the component of $\mathbb{H}^2 - g_1$ that contains $(0,1)$ in its euclidean closure in $\mathbb{R}^2$. By Lemma ??(iii), there is an isometry $\Phi_x : \mathbb{H}^2 \to \mathbb{H}^2$ which respects the oriented geodesic $\mathbb{H}^2 \cap \{0\} \times \mathbb{R}$ and which sends the point $O = A$ to the point C'_x. By definition, $h_x = \Phi_x(h_0)$ and $k_x = \Phi_x(k_0)$. If x is very large, the geodesic $\Phi_x(g_1)$ is contained in an arbitrarily small neighborhood of $(0,1)$, and so is the component $\Phi_x(P)$ of $\mathbb{H}^2 - \Phi_x(g_1)$. Since h_x and k_x are contained in $\Phi_x(P)$, we conclude that they are containes in an arbitrarily small neighborhood of $(0,1)$ provided that x is small enough.

In particular, for x large enough, the geodesics g and k_x are disjoint and their end points on the circle at infinity $\partial_\infty \mathbb{H}^2$ are distinct. The line of $\mathbb{R}^2$ passing through the end points of g intersects the line passing through the end points of k_x at some point D_x which is outside the closure of $\mathbb{H}^2$ in $\mathbb{R}^2$. Let l_x be the bi-infinite geodesic of $\mathbb{H}^2$ which is contained in a euclidean circle centered at D_x. Geometrically, l_x is the geodesic whose end points are the two points where the tangent to the circle at infinity $\partial_\infty \mathbb{H}^2$ passes through D_x. From this description, it is clear that l_x meets the geodesic g at a point B_x, and the geodesic k_x at a point B'_x.

We claim that lx is orthogonal to k_x. Indeed, consider any one of the end points $Ex \in \partial_\infty \mathbb{H}^2$ of l_x, and let F_x and G_x be the end points of k_x. Among the lines passing through D_x, one of them is tangent to $\partial_\infty \mathbb{H}^2$ at E_x, and another one meets $\partial_\infty \mathbb{H}^2$ at F_x and G_x. In elementary euclidean geometry, the power of the point D_x with respect to the circle $\partial_\infty \mathbb{H}^2$ is equal to $d_e(D_x, E_x)^2 = d_e(D_x, F_x)\, d_e(D_x, G_x)$. Note that $d_e(D_x, B'_x) = d_e(D_x, E_x)$ since these two quantities are equal to the radius of the circle containing l_x. If we now consider the circle containing k_x, which also meets a line passing through D_x at F_x and G_x and which contains the point B'_x, the equation $d_e(D_x, B'_x)^2 = d_e(D_x, F_x)\, d_e(D_x, G_x)$ implies that the line passing through D_x and B'_x is tangent to k_x. It follows that l_x, which is contained in a circle centered at D_x, is orthogonal to k_x at B'_x.

Similarly, l_x is orthogonal to g, and we conclude that the internal angles of the hexagon H_x with vertices A, A', B_x, B'_x, C_x, C'_x are all equal to $\pi/2$.

Let b_x be the hyperbolic distance between B_x and B'_x. By construction, b_x is a continuous function of x. When x is very large, k_x is in a very small neighborhood of the point $(0,1)$, and it easily follows from the definition of the hyperbolic metric that bx tends to ∞ as x tends to ∞. Conversely, as we decrease x, we eventually reach a value x_0 where one of the end points of the geodesic k_{x_0} coincides with one of the end points of g, in which case b_x tends to 0 as x tends to x_0 on the right. By the Intermediate Value Theorem, there consequently exists a number x for which $b_x = b$. The corresponding right-angles hexagon Hx_0 then satisfies the conditions requested.

Conversely, to prove the uniqueness, consider a right-angled hexagon H, with vertices A, A', B, B', C, C', such that the hyperbolic distances $d(A, A')$, $d(B, B')$ and $d(C, C')$ are respectively equal to a, b and c. Using Lemma ??(iii), we can modify H by an isometry of $\mathbb{H}^2$ so that the vertex A coincides with the origin O, and so that the vertices A' and C' are respectively located on $]0,1[\times\{0\}$

and $\{0\} \times]0,1[$ in $\mathbb{H}^2 \subset \mathbb{R}^2$. In other words, we can arrange that $H = H_x$ for some $x > 0$ with $b_x = b$. The uniqueness of H will then follow if we can prove that the map $x \mapsto b_x$ is injective.

□

B.4. Closed geodesics

B.5. The boundary at infinity

B.6. Quasi-isometric geometry

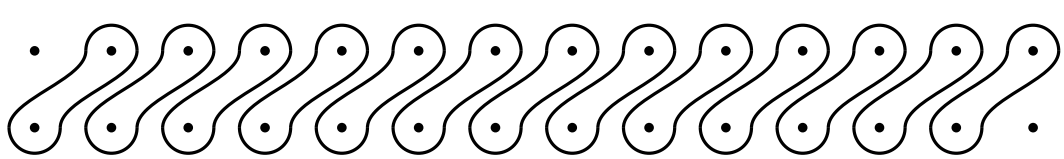

REFERENCES

[ArF] P. Arnoux, A. Fathi, *Un exemple de difféomorphisme peudo-Anosov*, C. R. Acad. Sci. Paris Sér. I Math. **312** (1991), 241–244.

[ArY] P. Arnoux, J.-C. Yoccoz, *Construction de difféomorphismes peudo-Anosov*, C. R. Acad. Sci. Paris Sér. I Math. **292** (1981), 75–78.

[Bau1] M. Bauer, *Examples of psudo-Anosov homeomorphisms*, Trans. Amer. Math. Soc. **330** (1992), 333–359.

[Bau2] M. Bauer, *An upper bound for the least dilation*, Trans. Amer. Math. Soc. **330** (1992), 361–370.

[Ber] H. Bergström, *Weak convergence of measures*, Academic press, New York, 1982.

[Bers] L. Bers, *An extremal problem for quasi-conformal mappings and a theorem by Thurston*, Acta Math. **141** (1978), 73–98.

[BeH] M. Bestvina, M. Handel, *Train-tracks for surface homeomorphisms*, Topology **34** (1995), 109–140.

[BeH] M. Bestvina, M. Handel, *Train-tracks and automorphisms of free groups*, Ann. Math. **135** (1992), 1–51.

[BLM] J.S. Birman, A. Lubotzky, J. McCarthy, *Abelian and solvable subgroups of the mapping class group*, Duke Math. J. **50** (1983), 1107–1120.

[BiS] J.S. Birman, C. Series, *Geodesics with bounded intersection numbers on surfaces are sparsely distributed*, Topology **24** (1985), 217-225.

[Bir] J.S. Birman, *Braids, links and mapping class groups*, Ann. of Math. Studies vol. 84, Princeton Univ. Press, 1974.

[Boa] R.P. Boas, *Möbius shorts*, Math. Mag. **68** (1995), 127.

[Bon] F. Bonahon, *Cobordism of automorphisms of surfaces*, Ann. Sci. École Norm. Sup. **16** (1983), 237–270.

[Bon] F. Bonahon, *Ribbon fibred knots, cobordism of surface diffeomorphisms and pseudo-Anosov diffeomorphisms*, Math. Proc. Cambridge Philos. Soc. **94** (1983), 235–251.

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

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

[Bon] F. Bonahon, *Structures géométriques sur les variétés de dimension* 3 *et applications*, Thèse d'état, Université d'Orsay, 1985.

[Bon] F. Bonahon, *Earthquakes on Riemann surfaces and on measured laminations*, Trans. Amer. Math. Soc. **330** (1992), 69–95.

[Bon] F. Bonahon, *Geodesic currents on negatively curved groups*, Arboreal group theory (R. Alperin, ed.), MSRI Publications vol. 19, Springer-Verlag, Berlin Heidelberg New York, 1991, pp. 143–168.

[Bon] F. Bonahon, *Ensembles limites et applications*, (45 minute address), Proceedings of the International Congress of Mathematicians, Kyoto, Japan, 1990, Springer-Verlag, Berlin Heidelberg New York, 1991, pp. 599–608.

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

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

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

[Bon] F. Bonahon, *The boundary geometry of 3–dimensional hyperbolic convex cores*, preprint.

[Bon] F. Bonahon, *A Schläfli-type formula for the volume of hyperbolic convex cores*, preprint.

[Bou] N. Bourbaki, *Éléments de mathématiques*, livre VI (Intégration), Hermann, Paris, 1965.

[CEG] R.D. Canary, D.B.A. Epstein, P. Green, *Notes on notes of Thurston*, Analytical and Geometrical aspects of Hyperbolic space (D.B.A. Epstein, ed.), L.M.S. Lecture Notes Series vol. 111, Cambridge Univ. Press, 1987, pp. 3–92.

[Cas] A. Casson, *Automorphisms of surfaces after Thurston and Nielsen*, (Lecture Notes by S. Bleiler), University of Texas, Austin, 1983.

[CaB] A. Casson, S.A. Bleiler, *Automorphisms of surfaces after Thurston and Nielsen*, Cambridge University Press, 1988.

[DaN] C. Danthony, A. Nogueira, *Involutions linéaires et feuilletages mesurés*, C.R. Acad. Sci. Paris Sér. I Math. **307** (1988), 409–412.

[DaN] C. Danthony, A. Nogueira, *Measured foliations on nonorientable surfaces*, Ann. Sci. Ec. Norm. Sup. **23** (1990), 469–494.

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

[EpM] D.B.A. Epstein, A. Marden, *Convex hulls in hyperbolic spaces, a theorem of Sullivan, and measured pleated surfaces*, Analytical and geometric aspects of hyperbolic space (D.B.A. Epstein, ed.), L.M.S. Lecture Note Series vol. 111, Cambridge University Press, 1986, pp. 113–253.

[Fat] A. Fathi, *The Poisson bracket on the space of measured foliations on a surface*, Duke Math. J. **55** (1987), 693–697.

[Fat] A. Fathi, *Dehn twists and pseudo-Anosov diffeomorphisms*, Invent. Math. **87** (1987), 129–151.

[Fat] A. Fathi, *Démonstration d'un théorème de Penner sur la décomposition des twists de Dehn*, Bull. Soc. Math. France **120** (1992), 467–484.

[FLP] A. Fathi, F. Laudenbach, V. Poenaru, *Travaux de Thurston sur les surfaces*, Astérisque 66–67, Société Mathématique de France, 1979.

[Flo] W.J. Floyd, *Group completions and limit sets of Kleinian groups*, Invent. Math. **57** (1980), 205–218.

[Flo] W.J. Floyd, *Incompressible surfaces in 3–manifolds*, Low dimensional topolofy and Kleinian groups (D.B.A. Epstein, ed.), London Math. Soc. Lecture Note Ser. vol. 112, Cambridge University Press, Cambridge New York, 1986, pp. 131–143.

[Ger] M. Gerber, *Conditional stability and real analytic pseudo-Anosov maps*, Mem. Amer. Math. Soc. vol. 54, American Mathematical Society, 1988.

[GeK] M. Gerber, A. Katok, *Smooth models of Thurson's pseudo-Anosov maps*, Ann. Sci. École Norm. Sup. **15** (1982), 173–204.

[Ghy] E. Ghys, *Le cercle à l'infini des surfaces à courbure négative*, (45 minute address), Proceedings of the International Congress of Mathematicians, Kyoto, Japan, 1990, Springer-Verlag, Berlin Heidelberg New York, 1991, pp. 501–509.

[GhH] E. Ghys, P. de la Harpe et al., *Sur les groupes hyperboliques d'après Gromov*, (with contributions of W. Ballman, A. Haefliger, E. Salem, R. Strebel, M. Troyanov), manuscript, Ecole normale Supérieure de Lyon and Université de Genève ■■published!, 1989.

[Gol] W. Goldman, *The symplectic nature of fundamental groups of surfaces*, Adv. in Math. **54** (1984), 200–225.

[Gro] M. Gromov, *Hyperbolic groups*, Essays in group theory (S.M. Gersten, ed.), M.S.R.I. Publications vol. 8, Springer-Verlag, Berlin Heidelberg New York, 1987, pp. 75–263.

[GuP] V. Guillemin, Pollack, *Differential topology*, Prentice Hall, Englewood Cliffs, 1974.

[HaC] H. Hamidi-Tehrani, Z.-H. Chen, *Surface diffeomorphisms via train tracks*, Topology Appl. **73** (1996), 141–167.

[Han] M. Handel, *Entropy and semi-conjugacy in dimension two*, Ergodic Theory Dynamical Systems **8** (1988), 585–596.

[Han] M. Handel, W.P. Thurston, *New proofs of some results of Nielsen*, Adv. in Math. **56** (1985), 173–191.

[Hat] A.E. Hatcher, *Measured lamination spaces for surfaces, from the topological viewpoint*, Topology Appl. **30** (1988), 63–88.

[HaO] A. Hatcher, U. Oertel, *Affine lamination spaces for surfaces*, Pacific J. Math. **154** (1992), 87–101.

[Hir] K. Hiraide, *Expansive homeomorphisms of compact surfaces are pseudo-Anosov*, Osaka J. Math. **27** (1990), 117–162.

[Iva] N.V. Ivanov, *Coefficients of expansion of pseudo-Anosov homeomorphisms*, Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI) **167** (1988), 111-116. (Russian)

[Kat] A. Katok, *Invariant measures of flows on oriented surfaces*, Soviet Math. Dokl. **14** (1973), 1104–1108.

[Kea] M. Keane, *Interval exchange transformations*, Math. Zeit. **141** (1975), 25–31.

[Kea] M. Keane, *Non-ergodic interval exchange transformations*, Israel J. Math. **2** (1977), 188–196.

[KeS] L. Keen, C. Series, *Pleating coordinates for the Teichmüller space of a punctured torus*, Bull. Amer. Math. Soc. **26** (1992), 141–146.

[KeS] L. Keen, C. Series, *Pleating coordinates for the Maskit embedding of theTeichmüller space of punctured tori*, Topology **32** (1993), 719–749.

[Ken] R. Kenyon, *A group of paths in* $\mathbb{R}^2$, Trans Amer. Math. Soc. **348** (1996), 3155–3172.

[Ker] S.P. Kerckhoff, *The Nielsen realization problem*, Ann. of Math. **117** (1983), 235–265.

[Ker] S.P. Kerckhoff, *The measure of the limit set of the handlebody group*, Topology **29** (1990), 27–40.

[Ker] S.P. Kerckhoff, *Simplicial systems for interval exchange maps and measured foliations*, Erg. Th. Dyn. Syst. **5** (1985), 257-271.

[KeN] H.B. Keynes, D. Newton, *A minimal non-ergodic interval exchange transformation*, Math. Z. **148** (1976), 101–105.

[Kou] C. Kourouniotis, *Bending in the space of quasi-Fuchsian structures*, Glasgow Math. J. **33** (1991), 41–49.

[Lev] G. Levitt, *Pantalons et feuilletages des surfaces*, Topology **21** (1982), 9–33.

[Lev] G. Levitt, *Foliations and laminations on hyperbolic surfaces*, Topology **22** (1983), 119–135.

[Lic] W.B.R. Lickorish, *A finite set of generators for the homeotopy group*, Proc. Camb. Phil. Soc. **60** (1964), 769-778.

[Lon] D.D. Long, *Constructing pseudo-Anosov maps*, Knot theory and manifolds (Vancouver, 1983), Lecture Notes in Math. vol. 1144, Springer-Verlag, 1983, pp. 108–114.

[Los] J.E. Los, *Pseudo-Anosov maps and invariant train tracks in the disc: a finite algorithm*, Proc. Londin Math. Soc. **66** (1993), 400–430.

[MaP] I. Maniscalco, A. Portolano, *The space of closed simple curves and measured foliations on noncompact orientable surfaces*, Rend. Circ. Mat. Palermo **36** (1988), 401–406.

[Mas] H. Masur, *Interval exchange transformations and measured foliations*, Ann. Math. **115** (1982), 169-200.

[Mas] H. Masur, *Measured foliations and handlebodies*, Erg. Theory and Dyn. Syst. **6** (1986), 99-116.

[Mas] H. Masur, *Hausdorff dimension of the set of nonergodic foliations of a quadratic differential*, Duke Math. J. **66** (1992), 387–442.

[MaS] H. Masur, J. Smillie, *Hausdorff dimension of sets of nonergodic measured foliations*, Ann. of Math. **134** (1991), 455–543.

[MCP] J. McCarthy, A. Papadopoulos, *Dynamics on Thurston's sphere of projective measured foliations*, Comment. Math. Helv. **64** (1989), 133–166.

[Mil] R.T. Miller, *Geodesic laminations from Nielsen's viewpoint*, Adv. in Math. **45** (1982), 189–212.

[Mil] J.W. Milnor, *Topology from the differentiable viewpoint*, The University Press of Virginia, Charlotteville, 1965.

[Mil] J.W. Milnor, *A note on curvature and fundamental group*, J. Differential Geometry **2** (1968), 1–7.

[Min] Y. Minsky, *Harmonic maps into hyperbolic 3–manifolds*, Trans. Amer. Math. Soc. **332** (1992), 607–632.

[Min] Y. Minsky, *On rigidity, limit sets, and end invariants of hyperbolic 3–manifolds*, J. Amer. Math. Soc. **7** (1994), 539–588.

[Mos] L. Mosher, *The classification of pseudo-Anosovs*, Low dimensional topoloy and Kleinian groups (D.B.A. Epstein, ed.), London Math. Soc. Lecture Note Ser. vol. 112, Cambridge University Press, Cambridge New York, 1986, pp. 13–75.

[Mos] L. Mosher, *Tiling the projective foliation space of a punctured surface*, Trans. Amer. Math. Soc. **306** (1988), 1–70.

[Nie] J. Nielsen, *Untersuchungen zur Topologie des geschlossenen zweitigfen Flächen* I, Acta Math. **50** (1927), 189–358II, Acta Math. **53** (1929), 1–76III, Acta Math. **58** (1931), 87–167.

[Nie] J. Nielsen, *Abbildungsklassen endlicher Ordnung*, Acta Math. **75** (1943), 23–115.

[Nie] J. Nielsen, *Surface transformations classes of algebraically finite type*, Danske Vid. Selsk. Mat.-Fys. Medd. **21** (1944), 1–89.

[MSRI] J. Alonso, T. Brady, D. Cooper, T. Delzant, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, H. Short, *Notes on negatively curved groups*, manuscript, Mathematical Sciences Research Institute, Berkeley, 1989.

[Ohs] K. Ohshika, *Ending laminations and boundaries for deformation spaces for Kleinian groups*, J. London Math. Soc. **42** (1990), 111–121.

[Ota] J.-P. Otal, *Le théorème d'hyperbolisation pour les variétés fibrées de dimension* 3, Astérisque vol. 235, Société Mathématique de France, 1996.

[Pap] A. Papadopoulos, *Difféomorphisms pseudo-Anosov et automorphismes symplectiques de l'homologie*, Ann. Sci. Éc. Norm. Sup. **15** (1982), 543–546.

[Pap] A. Papadopoulos, *L'extension du flot de Fenchel-Nielsen au bord de Thurston de l'espace de Teichmüller*, Comptes Rendus Acad. Sci. Paris **302** (1986), 325–327.

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

[Pap] A. Papadopoulos, *Geometric intersection functions and Hamiltonian flows on the space of measured foliations of a surface*, Pacific J. Math. **124** (1986), 375–402.

[Pap] A. Papadopoulos, *Sur le bord de Thurston de l'espace de Teichmüller d'une surface non compacte*, Math. Ann. **282** (1988), 353–359.

[Pap] A. Papadopoulos, *On Thurston's boundary of Teichmüller space and the extension of earthquakes*, Top. Appl. **41** (1991), 147–177.

[PaP] A. Papadopoulos, R.C. Penner, *A characterization of pseudo-Anosov foliations*, Pacific J. Math. **130** (1987), 359–377.

[PaP] A. Papadopoulos, R.C. Penner, *Enumerating pseudo-Anosov foliations*, Pacific. J. Math. **142** (1990), 159–173.

[PaP] A. Papadopoulos, R.C. Penner, *The Weil-Petersson symplectic structure at Thurston's boundary*, Trans. Amer. Math. Soc. **335** (1993), 891–904.

[Pen] R.C. Penner, *A construction of pseudo-Anosov homeomorphisms*, Trans. Amer. Math. Soc. **310** (1988), 179–197.

[Pen] R.C. Penner, *Bounds on least dilatations*, Proc. Amer. Math. Soc. **113** (1991), 443–450.

[PeH] R.C. Penner, J.L. Harer, *Combinatorics of train tracks*, Ann. Math. Studies vol. 125, Princeton Univ. Press, 1992.

[Pla] J. Plante, *Foliations with measure preserving holonomy*, Ann. Math. **127** (1975), 327–362.

[PoM] A. Portolano, I. Maniscalco, *Curves as measured foliations on noncompact surfaces*, Rend. Circ. Mat. Palermo **42** (1993), 161–180.

[Rau] G. Rauzy, *Echanges d'intervalles et transformations induites*, Acta Arith. **34** (1979), 315–328.

[Ree] M. Rees, *An alternative approach to the ergodic theory of measured foliations on surfaces*, Erg. Th. and Dynam. Syst. **1** (1981), 461–488.

[RuS] D. Ruelle, D. Sullivan, *Cycles for the dynamical study of foliated spaces and complex manifolds*, Invent. Math. **36** (1976), 225–255.

[Sig] K. Sigmund, *Generic properties of invariant measures for Axiom A diffeomorphisms*, Invent. Math. **11** (1970), 99–109.

[Sig] K. Sigmund, *On dynamical systems with the Specification Property*, Trans. Amer. Math. Soc. **190** (1974), 285–299.

[Sko] R.K. Skora, *Splitting of surfaces*, Bull. Amer. Math. Soc. **23** (1990), 85–90.

[Sko] R.K. Skora, *Geometric actions of surface groups on* Λ*–trees*, Comment. Math. Helv. **65** (1990), 519–533.

[Slu] L. Slutkin, *Computations of the transverse measured foliations associated with a pseudo-Anosov automorphism*, J. Differential Geom. **39** (1994), 359–378.

[Sul] D. Sullivan, *Cycles for the dynamical study of foliated spaces and complex manifolds*, Invent. Math. **36** (1976), 225–255.

[Šva] A.S. Švarc,, *A volume invariant of coverings*, Dokl. Akad. Nauk SSSR **105** (1955), 32–34.. (Russian)

[Thu] W.P. Thurston, *The topology and geometry of* $3-$*manifolds*, Lecture notes, Princeton University, 1976–79.

[Thu] W.P. Thurston, *Hyperbolic geometry and* $3-$*manifolds*, Low-dimensional Topology (R. Brown and T. Thickstun, eds.), London Math. Soc. Lecture Notes Series 48, Cambridge University Press, 1982, pp. 9–25.

[Thu] W.P. Thurston, *Three-dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. **6** (1982), 357–381.

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

[Thu] W.P. Thurston, *Earthquakes in two-dimensional hyperbolic geometry*, Low-dimensional Topology and Kleinian groups (D.B.A. Epstein, ed.), L.M.S. Lecture Notes Series vol. 112, Cambridge University press , 1986, pp. 91–112.

[Thu] W.P. Thurston, *Minimal stretch maps between hyperbolic surfaces*, preprint.

[Vee] W. Veech, *Gauss measures for transformations on the space of interval exchange maps*, Ann. Math. **115** (1982), 201-242.

[Wei] H.N. Weiss, *Non-smooth geodesic flows and the earthquake flow on Teichmüller space*, Ergodic Theory Dynamical Systems **9** (1989), 571–586.

[Wei] H.N. Weiss, *The geometry of measured geodesic laminations and measured train tracks*, Ergodic Theory Dynamical Systems **9** (1989), 587–604.

[Wol] S. Wolpert, *The Fenchel-Nielsen deformation*, Ann. of Math. **115** (1982), 501–528.

[Wol] S. Wolpert, *On the symplectic geometry of deformations of a hyperbolic surface*, Ann. of Math. **117** (1983), 207–234.

[Wol] S. Wolpert, *Thurston's Riemannian metric for Teichmüller space*, J. Diff. Geom. **23** (1986), 143–174.

[Wu] Y.-Q. Wu, *Canonical reducing curves of surface homeomorphisms*, Acta Math. Sinica **3** (1987), 305–313.

[Zeg] A. Zeghib, *Laminations et hypersurfaces géodésiques des variétés hyperboliques*, Ann. Sci. Ec. Norm. Sup. **24** (1991), 171–188.

[Zhi] A.Y. Zhirov, *On the minimum dilation of pseudo-Anosov diffeomorphisms of a double torus*, Uspekhi Mat. Nauk **50** (1995), 197–198. (Russian)

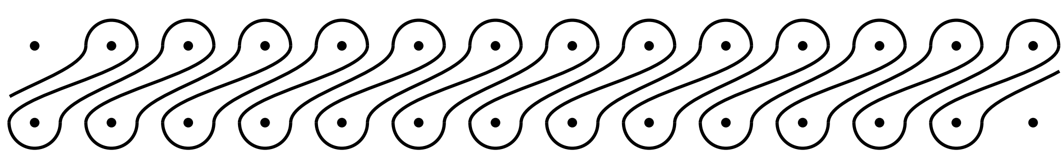

INDEX