

MAPPING CLASS GROUP ORBIT CLOSURES FOR NON-ORIENTABLE SURFACES

VIVEKA ERLANDSSON, MATTHIEU GENDULPHE, IRENE PASQUINELLI
AND JUAN SOUTO



Abstract. Let S be a connected non-orientable surface with negative Euler characteristic and of finite type. We describe the possible closures in \mathcal{ML} and \mathcal{PML} of the mapping class group orbits of measured laminations, projective measured laminations and points in Teichmüller space. In particular we obtain a characterization of the closure in \mathcal{ML} of the set of weighted two-sided curves.

1 Introduction

In this paper we study the closures of the orbits of the action of the mapping class group of a complete connected hyperbolic surface on the space of measured laminations (Theorem 1.4), projective measured laminations (Theorems 1.1 and 1.3) and Teichmüller space (Theorem 1.5). When reading this the reader might well be surprised. They might be thinking that all of this is already known; that it is a classical result that the action of the mapping class group on \mathcal{PML} is minimal in the sense that all orbits are dense [FLP], that \mathcal{PML} is the limit set of the action of the mapping class group on Teichmüller space, and that the orbit closures of the action of Map on \mathcal{ML} were already described by Mirzakhani and Lindenstrauss [LM08]. And the reader would be correct if they restricted themselves to orientable surfaces. Things are actually quite different in the non-orientable world.

For starters, it is due to Scharlemann [Sch82] and Danthony-Nogueira [DN90] that the set of (projective) measured laminations which have a one-sided closed leaf is open and has full measure, where a simple closed essential curve is *one-sided* if it has a regular neighborhood homeomorphic to a Möbius band—otherwise it is *two-sided*. Since the set of measured laminations without one-sided component is also mapping class group invariant, we get in particular that the action $\text{Map}(S) \curvearrowright \mathcal{PML}(S)$ is not minimal whenever S is non-orientable: the set

$$\mathcal{PML}^+(S) = \{\lambda \in \mathcal{PML}(S) \text{ without closed one-sided components}\}$$

The first and third authors gratefully acknowledge support from EPSRC grant EP/T015926/1. This research was also supported by the UiT Aurora project MASCOT. No data was used in this research.

is a non-empty, closed, invariant proper subset. Our first result is that the action of the mapping class group on $\mathcal{PML}^+(S)$ is minimal, or rather that this set is the unique closed minimal subset of $\mathcal{PML}(S)$:

Theorem 1.1. *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. The set $\mathcal{PML}^+(S)$ is the unique non-empty closed subset of $\mathcal{PML}(S)$ which is invariant and minimal under the action of $\text{Map}(S)$.*

In Theorem 1.1, as in the remaining of the paper, we say that a hyperbolic surface S is *exceptional* if it is either a pair of pants or non-orientable with $\chi(S) = -1$. Otherwise the surface is *non-exceptional*. The non-orientable exceptional surfaces are the two-holed projective plane, the one-holed Klein bottle, and the connected sum of three projective planes. The mapping class group, the spaces of measured laminations and projective measured laminations are well understood if the surface is exceptional, and the interested reader will have no difficulty to clarify matters on their own for those cases, but we discuss briefly exceptional surfaces in Section 2.5 below.

The closure of the mapping class group orbit of a two-sided curve γ is a closed invariant subset of \mathcal{PML}^+ . Theorem 1.1 implies thus that $\mathcal{PML}^+(S) = \overline{\text{Map}(S) \cdot \gamma}$. It follows a fortiori that the set of projective classes of two-sided curves is dense in $\mathcal{PML}^+(S)$. This answers a question which seems to have been around for some time [Bes]:

Theorem 1.2. *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. The set of two-sided curves is dense in $\mathcal{PML}^+(S)$.*

If Theorem 1.2 follows from Theorem 1.1 because \mathcal{PML}^+ is minimal, the fact that it is actually the unique non-empty closed minimal subset of $\mathcal{PML}(S)$ implies that $\mathcal{PML}^+(S)$ is contained in the closure of every orbit $\lambda \in \mathcal{PML}(S)$. To describe the actual closure of $\text{Map}(S) \cdot \lambda$ for an arbitrary λ we need to introduce some notation.

Following Lindenstrauss-Mirzakhani [LM08] we consider the decomposition $\lambda = \gamma_\lambda + \lambda'$ where γ_λ is the atomic part of λ , and let R_λ be the union of those connected components of $S \setminus \gamma_\lambda$ which contain a non-compact leaf of λ . The pair $(R_\lambda, \gamma_\lambda)$ is, in the terminology of [LM08], a *complete pair*. We associate to the pair $(R_\lambda, \gamma_\lambda)$ first the set $\gamma_\lambda + \mathcal{ML}^+(R_\lambda)$ of measured laminations of the form $\gamma_\lambda + \mu$ where $\mu \in \mathcal{ML}^+(R_\lambda)$ is a measured lamination supported by R_λ and without one-sided leaves, and then the orbit of this set under the mapping class group action:

$$\mathcal{G}_\lambda = \cup_{\phi \in \text{Map}(S)} \phi(\gamma_\lambda + \mathcal{ML}^+(R_\lambda)) \subset \mathcal{ML}(S).$$

Unsurprisingly we denote by $\mathcal{PG}_\lambda \subset \mathcal{PML}(S)$ the image of \mathcal{G}_λ in the space of projective measured laminations.

With this notation in place we can give a precise description for the closure of $\text{Map}(S) \cdot \lambda$ in \mathcal{PML} for an arbitrary λ .

Theorem 1.3. *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. We have $\overline{\text{Map}(S)} \cdot \lambda = \mathcal{PG}_\lambda \cup \mathcal{PML}^+(S)$ for any projective measured lamination $\lambda \in \mathcal{PML}(S)$.*

The reader might be wondering what happens if we consider the action of the mapping class group on $\mathcal{ML}(S)$ instead. In fact, even if this question had not crossed their mind it would have to be considered: understanding orbit closures for the action $\text{Map}(S) \curvearrowright \mathcal{ML}(S)$ is an integral part of the proof of Theorem 1.3. We prove:

Theorem 1.4. *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. We have $\overline{\text{Map}(S)} \cdot \lambda = \mathcal{G}_\lambda$ for any measured lamination $\lambda \in \mathcal{ML}(S)$.*

In the orientable case, Theorem 1.4 is due to Lindenstrauss-Mirzakhani (see Theorem 1.2 in [LM08]). They obtain it as a consequence of their main theorem, the classification of mapping class group invariant measures on $\mathcal{ML}(S)$. It seems that, for the time being, the latter kind of result is out of scope in the non-orientable setting—for example, everything that Lindenstrauss and Mirzakhani do, relies on the fact that the moduli space of S has finite volume, and this is false if the surface is non-orientable. The proof of Theorem 1.4 that we give below is actually pretty elementary and can also be enjoyed by those readers who only care about orientable surfaces.

As we just pointed out, the moduli space of a non-orientable surface has infinite volume. Indeed, the usual analogy between the mapping class group and an arithmetic lattice breaks in the non-orientable case. In fact, in the non-orientable case it might well be closer to the mark to compare the mapping class group with a Kleinian group of the second kind, or rather with some higher rank version of this. Recall that a Kleinian group is of the second kind if its discontinuity domain is not empty, or equivalently if its limit set is a proper subset of the boundary of hyperbolic space. Thinking of the Thurston boundary $\partial\mathcal{T}(S) = \mathcal{PML}(S)$ of Teichmüller space as being the analogue to the boundary at infinity of hyperbolic space, we interpret the following result as describing the limit set of the action of the mapping class group on Teichmüller space:

Theorem 1.5. *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. We have $\overline{\text{Map}(S)} \cdot X \cap \partial\mathcal{T}(S) = \mathcal{PML}^+(S)$ for any point X in Teichmüller space $\mathcal{T}(S)$.*

Having presented the main results of the paper, we comment briefly on their proofs and the outline of the paper. First, we give a very short argument explaining that filling and uniquely ergodic measured laminations are limits of two-sided curves.

An argument. Let us suppose that λ is a filling and uniquely ergodic measured lamination. As every projective measured lamination, λ can be approximated by a sequence (α_i) of weighted multicurves. Some of these might well be one-sided but,

invoking for example the classification theorem of surfaces, we can get for all i a two-sided simple curve β_i with $\iota(\beta_i, \alpha_i) = 0$. Passing to a subsequence we can assume that the sequence (β_i) converges projectively to some $\mu \in \mathcal{PML}(S)$. Continuity of the intersection number implies then that $\iota(\mu, \lambda) = 0$. Now, since λ is filling we get that both μ and λ have the same support. Since it is also uniquely ergodic we get that λ and μ differ by a factor, meaning that they are equal in \mathcal{PML} . This proves that the filling uniquely ergodic lamination $\lambda \in \mathcal{PML}(S)$ is the limit of the sequence of two-sided curves (β_i) .

At this point the reader might be ready to point out we also proved that every element in the closure of the set of filling uniquely ergodic measured laminations is also a limit of two sided curves, thinking that the set of filling uniquely ergodic measured laminations must be dense. While it is true that they are dense in say \mathcal{PML} if the surface is orientable, this is not true for non-orientable surfaces as we already pointed out earlier. In the non-orientable case we get from say Theorem 1.1 that the set of filling uniquely ergodic laminations is dense in \mathcal{PML}^+ . If the reader knew how to prove this without using the results of this paper, then everything we do here could be done in 5 to 10 pages. Indeed, the bulk of the work is to deal with measured laminations which are not ergodic.

Section-by-section summary. Once this introduction is concluded, we have a section on preliminaries. We recall a few well-known facts about topology of surfaces, mapping class groups, laminations, and hyperbolic surfaces. There is nothing new here but the reader who is not used to thinking about non-orientable surfaces might still want to have a look.

In Section 3 we recall what train tracks are and introduce what we call *uniform*, or rather (C, λ) -*uniform* train tracks. In a nutshell, these are train tracks carrying a given lamination λ and where, from the point of view of λ , all edges have comparable lengths. Existence of such train tracks, or rather the fact that every train track can be refined to such a train track, will be proved in Appendix A. What we do in Section 3 is to prove two technical results—needed later on—about such uniform train tracks. We think that uniform train tracks might turn out to be useful in other settings as well.

In any case, uniform train tracks play a key role in the proof of the following theorem in Section 4:

Theorem 4.1 (*Informal version*). *Let $\lambda \in \mathcal{L}(S)$ be a lamination and let μ_1, μ_2 be distinct unit length ergodic measured laminations with support λ . If τ is a sufficiently nice uniform train track carrying λ then there are disjoint non-empty sub-train tracks τ_1 and τ_2 such that any unit length measured lamination carried by τ_i is near μ_i .*

The actual statement of Theorem 4.1 is pretty technical—the informal version given here just captures the gist of it.

REMARK. Theorem 4.1 is motivated by a result of Lenzhen-Masur [LM10, Proposition 1]. It is not completely clear to us if, as stated, we could just have used what they prove instead of working with Theorem 4.1. But in any case we would have had to comment on the proof because they work in the world of quadratic holomorphic differentials and hence, at least formally, deal only with orientable surfaces. We found however much simpler to just give a direct proof of Theorem 4.1. Maybe a reader who is better versed than ourselves in taking geometric limits and exploiting compactness properties of spaces of manifolds would be able to see clearly enough through their argument to modify the proof so that it also works in the non-orientable case.

Theorem 4.1 is maybe the main technical result of this paper and will play a key role in the proof of Theorem 1.2. It has however applications that might be of independent interest. For example we obtain the following as a corollary:

COROLLARY 5.3 *Let $c(S)$ and $c^+(S)$ be, respectively, the maximal number of components of a multicurve and of a two-sided multicurve in S . Every lamination $\lambda \in \mathcal{L}(S)$ supports at most $c(S)$ mutually singular ergodic transverse measures. Moreover, if λ has no one-sided leaves then it supports at most $c^+(S)$ mutually singular ergodic transverse measures.*

Note that if S has non-orientable genus k and r boundary components then we have

$$c(S) = 2k - 3 + r \text{ and } c^+(S) = \left\{ \begin{array}{ll} \frac{1}{2}(3k - 7 + 2r) & \text{if } k \text{ is odd} \\ \frac{1}{2}(3k - 8 + 2r) & \text{if } k \text{ is even} \end{array} \right\},$$

meaning that $c(S)$ is about $k/2$ larger than $c^+(S)$.

Also note that the corollary says that a maximal multicurve carries the maximal number of mutually singular ergodic transverse measure. In particular, the bound is optimal. However, it is an interesting question—which we do not know the answer to—whether this bound can be achieved by a connected recurrent lamination, or what the best bound in that case is.

REMARK. Corollary 5.3 is due to Levitt [Lev83] in the orientable case. He uses a beautiful idea due to Katok [Kat73] and we really encourage the reader to have a look at Levitt’s argument—it is very nice mathematics.

Let us continue now with the summary of the paper. In Section 5 we prove Theorem 1.2. Indeed, although we might have given the impression that Theorem 1.2 was a corollary of Theorem 1.1, it is rather the other way around: all the theorems mentioned earlier build on Theorem 1.2. The main idea of its proof is as follows: Let $\mu \in \mathcal{ML}$ be a representative of the element in \mathcal{PML}^+ that we want to approximate and consider its ergodic decomposition. For the sake of concreteness say that $\mu = \mu_1 + \mu_2$ with μ_i ergodic. Let then τ be a “sufficiently nice” train track carrying the support of μ and let $\tau_1, \tau_2 \subset \tau$ be the sub-train tracks provided by Theorem 4.1. We can then approximate μ by weighted (simple) multicurves $c_1 \cdot \gamma_1 + c_2 \cdot \gamma_2$ with γ_i carried by τ_i . But we must make sure that we can choose the γ_i to be two-sided.

To do that we will need to first characterize which train tracks do not carry any two-sided curves and use then a (rather minimal) quantification of Scharlemann's result on the openness of the set of measured laminations which have a one-sided component.

Armed with Theorem 1.2 we attack the other results stated in the introduction. In Section 6 we prove Theorem 1.4. The proof only relies on the density of two-sided curves, that is Theorem 1.2, and of elementary facts about measured laminations. In Section 7 we deduce Theorems 1.3 and 1.1 from Theorem 1.4. Finally, in Section 8 we prove Theorem 1.5.

As we mentioned earlier, we conclude with an appendix in which we prove the existence of uniform train tracks.

On the genesis of this paper. This paper grew out of the work of the second of us. In fact, one of the inclusions in the orbit closure results mentioned above were proved by the second author in the unpublished preprint [Gen17], and equality was established for surfaces of genus one. The genus one assumption guarantees that the complement of a one-sided curve is orientable and hence that one can use all the standard theory to deal with measured laminations disjoint from such a curve. The novel ingredients here are Theorem 1.2 and the fact that we give a proof of Theorem 1.4 which does not rely on the Lindenstrauss-Mirzakhani classification of mapping class group invariant measures on the space of measured laminations.

It should be said that the orbit closure theorems were only a fraction of the content of [Gen17]—we hope to revisit other parts of the said paper at a later point.

Another paper. At the moment of posting this paper to the arXiv we noticed that just a couple of days earlier Khan [Kha21] had posted a very nice paper on this topic. Some of his results are of completely different nature than what we are doing here: for example, in Theorem 5.2 he shows that certain subsets of Teichmüller space are non quasi-convex, answering in the negative a question asked in [Gen17]. Other results are however somewhat weaker results than the ones we present here. For example, his Theorem 3.3, although stated with different words, asserts that certain kinds of measured laminations, for example orientable ergodic laminations, are limits of two-sided curves. It goes without saying that prior to the completion of our respective papers, neither him nor ourselves were aware of each other's work.

We really encourage the reader to have a look at Khan's paper because the arguments he gives when dealing with particular cases of laminations are much prettier and simpler than what awaits them if they continue reading.

2 Preliminaries

We assume that the reader feels at home in the world of surfaces, that they know what the mapping class group is, that they are familiar with laminations, measured laminations and Teichmüller spaces. We breeze through these concepts below, but should the reader wish more details we refer to the beautiful books [FM12, CB88,

PH92]. However, even if the reader knows all these things inside out, it might well be that they have only encountered them in the orientable setting. This is why we start by recalling some basic facts about the topology of non-orientable surfaces.

2.1 Non-orientable surfaces. We should probably say once what a surface is: it is nothing other than a 2-dimensional smooth manifold, possibly with non-empty boundary. Moreover, unless we say explicitly otherwise, we will assume that our surfaces S are connected and have *finite topological type*, by what we understand that S is obtained from a compact surface Σ by deleting some of its boundary components, possibly none and possibly all. We will indeed be mostly interested in the case that the surface S has empty boundary, meaning that $S = \Sigma \setminus \partial\Sigma$ where Σ is a compact surface.

As in the orientable case, the classification theorem is the foundation on which one builds any work on surfaces:

Classification theorem. *Two compact connected surfaces Σ and Σ' are homeomorphic if and only if the following two hold:*

- Σ and Σ' have the same number of boundary components and the same Euler characteristic.
- Σ and Σ' are either both orientable or both non-orientable.

We refer for example to [Sti93] for a proof of this theorem (under the, always satisfied, assumption that the involved surfaces are triangulable).

A simple curve in a surface is *essential* if it neither bounds a disk, nor an annulus, nor a Möbius band. Equivalently, a simple curve is essential if and only if it is homotopically non-trivial, non-peripheral, and primitive. Note that, as long as S has negative Euler characteristic, every orientation reversing simple curve is essential. Note also that the boundary of a regular neighborhood of such an orientation reversing simple curve, although it is simple and homotopically non-trivial, is not essential. In the sequel we will say that an essential simple curve is *one-sided* if it reverses orientation and *two-sided* if it preserves orientation. A *topological multicurve* is the union of pairwise non-isotopic disjoint simple curves—it is essential, (resp. two-sided, resp. one-sided) if all its components are.

To conclude with these topological preliminaries, recall that we say that a surface S with negative Euler characteristic is *exceptional* if it is either a pair of pants or non-orientable with $\chi(S) = -1$. For clarification, here is a list of the possible exceptional surfaces:

- S *orientable*: Pair of pants
- S *non-orientable*: Two-holed projective plane, one-holed Klein bottle, connected sum of three projective planes.

The main reason to ignore exceptional surfaces is that, in all cases except the connected sum of three projective planes, they support no measured laminations—at least none containing non-compact leaves. Also the case of connected sum of three

projective planes is a bit special, and for technical reasons we exclude it since sometimes we need to have space enough to find curves with certain properties. For example we will want to know the following:

LEMMA 2.1. *If S is a non-exceptional connected surface then there are two-sided simple curves $\alpha, \beta \subset S$ which fill S .*

Recall that a collection $\{\alpha_1, \dots, \alpha_r\}$ of curves *fills* S if $\sum_i \iota(\alpha_i, \gamma) > 0$ for every essential curve γ —here $\iota(\cdot, \cdot)$ is the geometric intersection number.

We are sure that Lemma 2.1 is well-known and that is why we just sketch the proof in the case of non-orientable surfaces. Via the classification theorem of surfaces we get that there is $\eta \subset S$ consisting of either a single one-sided simple curve (if S has odd genus) or two disjoint one-sided simple curves (if the genus is even) such that $S \setminus \eta$ is connected and orientable. For the sake of concreteness we assume that we are in the former case, leaving the other case to the reader.

Let Σ be the metric completion of $S \setminus \eta$ and note that a connected component of $\partial\Sigma$ maps in a two-to-one way onto η —the remaining components are mapped homeomorphically onto ∂S . Choose two distinct points $x, y \in \eta$ and let x', x'' and y', y'' be the corresponding points in $\partial\Sigma$. Since $\chi(\Sigma) = \chi(S \setminus \eta) = \chi(S) \leq -2$ we can find two disjoint, non-boundary parallel simple arcs $\kappa', \kappa'' \subset \Sigma$ with endpoints $\partial\kappa' = \{x', y'\}$ and $\partial\kappa'' = \{x'', y''\}$. When mapped into S , the two arcs κ' and κ'' match together and yield a two-sided curve $\alpha \subset S$.

Now let $\beta' \subset S \setminus \eta$ be any essential curve which fills with $\kappa' \cup \kappa''$. The image β of β' in S is the second curve we are after. This concludes the discussion of Lemma 2.1. \square

2.2 Mapping class groups. The mapping class group of a surface S is the group

$$\text{Map}(S) = \text{Homeo}(S) / \text{Homeo}_0(S)$$

of isotopy classes of homeomorphisms of S —the reader can replace homeomorphisms by diffeomorphisms if they so wish, and nothing will change.

Both in the orientable and in the non-orientable case, the prime example of an element in the mapping class group $\text{Map}(S)$ is a *Dehn-twist* along a two-sided curve γ . A Dehn-twist D_γ along γ is a mapping class represented by a homeomorphism obtained as follows: conjugate the homeomorphism

$$(\mathbb{R}/\mathbb{Z}) \times [0, 1] \rightarrow (\mathbb{R}/\mathbb{Z}) \times [0, 1], \quad (\theta, t) \mapsto (\theta + t, t)$$

via a homeomorphism between the standard annulus $(\mathbb{R}/\mathbb{Z}) \times [0, 1]$ and a regular neighborhood A of γ , and extend the so obtained homeomorphism $A \rightarrow A$ by the identity to a self-homeomorphism $S \rightarrow S$.

If γ and γ' are disjoint two-sided curves then they have disjoint regular neighborhoods. This means that Dehn-twists along γ and γ' commute.

Below we will need to understand what happens to certain curves when we apply Dehn-twists to them. The following lemma due to Ivanov [Iva, Lemma 4.2] is priceless:

LEMMA 2.2. (Ivanov). *Let α and β be simple essential curves and assume that $\gamma = \gamma_1 \cup \dots \cup \gamma_r$ is a two-sided simple topological multicurve. Let D_{γ_i} be a Dehn-twist along γ_i and for some $n_1, \dots, n_r \in \mathbb{Z}$ let $T = D_{\gamma_1}^{n_1} \circ \dots \circ D_{\gamma_r}^{n_r}$. Then we have*

$$\begin{aligned} \iota(T(\alpha), \beta) &\leq \left(\sum_i |n_i| \cdot \iota(\alpha, \gamma_i) \cdot \iota(\gamma_i, \beta) \right) + \iota(\alpha, \beta) \\ \iota(T(\alpha), \beta) &\geq \left(\sum_i (|n_i| - 2) \cdot \iota(\alpha, \gamma_i) \cdot \iota(\gamma_i, \beta) \right) - \iota(\alpha, \beta) \end{aligned}$$

Since Ivanov only considers the orientable case, let us comment briefly on the proof in the non-orientable case. First, note that the upper bound for $\iota(T(\alpha), \beta)$ is just the number that you get by applying the standard representative of T to α and counting intersection points without bothering to get rid of bigons. To get the lower bound lift everything to the orientation cover $S' \rightarrow S$ and apply Ivanov’s lemma there. □

Note now that, since measured laminations and a fortiori multicurves are determined by intersection numbers, we get from Ivanov’s lemma, and with the same notation as therein, that

$$\lim_{k \rightarrow \infty} \frac{1}{k} \cdot T^k(\alpha) = \sum_i (|n_i| \cdot \iota(\alpha, \gamma_i)) \cdot \gamma_i.$$

Since for every multicurve $\gamma = \gamma_1 \cup \dots \cup \gamma_r$ we can find a two-sided curve α with $\iota(\alpha, \gamma_i) \neq 0$ for all i we get then the following well-known fact:

LEMMA 2.3. *Every weighted two-sided multicurve is a limit of weighted two-sided curves.* □

Note that the lemma says that weighted two-sided curves are dense in weighted 2-sided multicurves, and hence can be seen as the baby case of Theorem 1.2.

Dehn-twists fix many curves: the curve one is twisting around, as well as all the curves disjoint from it. Pseudo-Anosov mapping classes are on the other extremum. Recall that, following Thurston, a mapping class is *pseudo-Anosov* if none of its positive powers fixes the homotopy class of an essential curve. As in the orientable setting, examples of pseudo-Anosov mapping classes can be constructed as follows: let D_α and D_β be Dehn-twists along two two-sided curves α and β which together fill the surface S and consider the composition $D_\alpha^n \circ D_\beta^n$. It follows (compare with [Man10]) for example from Lemma 2.2 that $D_\alpha^n \circ D_\beta^n$ is pseudo-Anosov for all large n . Alternatively, the reader can consider the lifts of D_α and D_β to the orientation cover and quote the result in the orientated case. Either way we get:

LEMMA 2.4. *If S is a connected non-essential hyperbolic surface then $\text{Map}(S)$ contains pseudo-Anosov elements.* \square

Non-orientable surfaces are confusing. This is why we add here some comments to put things into context, although none of these things will play a role in the sequel:

- (1) Note that if S is orientable then what we here called the mapping class group is often referred to as the *full mapping class group*, with the mapping class group itself being the subgroup consisting of elements represented by orientation preserving homeomorphisms. That last sentence makes no sense for non-orientable surfaces.
- (2) The reader might have been disturbed because we wrote “a” instead of “the” in the sentence “a Dehn-twist along γ ”. It is not a typo. The point is that, even up to isotopy, there is no uniqueness for the homeomorphism between $(\mathbb{R}/\mathbb{Z}) \times [0, 1]$ and the regular neighborhood A of our given two-sided curve γ , the homeomorphism we used to define the Dehn-twist. In the orientable world one can decide to choose it to be orientation preserving. This leads to the notion of *right Dehn-twist*. If one chooses it to be orientation reversing then one has the *left Dehn-twist*—these notions make no sense in the non-orientable world and hence “a” is the correct article.
- (3) Recall that the adjective “two-sided” only applies to essential curves. If we were to mimic the construction of a Dehn-twist but starting with a non-essential orientation preserving simple curve which is the boundary of a Möbius band then we would get the trivial mapping class. On the other hand we get for example from Lemma 2.2 that Dehn-twists along two-sided curves have infinite order in $\text{Map}(S)$.
- (4) There is nothing like a Dehn-twist along a one-sided curve. In fact, every homeomorphism of a surface which fixes the complement of a Möbius band is isotopic to the identity.

We refer to [FM12] for facts on the mapping class group, and to [Par14] and the references therein for facts on mapping class groups of specifically non-orientable surfaces.

2.3 Hyperbolic metrics. It follows from the classification theorem of surfaces, together with elementary constructions in hyperbolic geometry, that every connected surface of negative Euler characteristic admits a complete hyperbolic metric with totally geodesic boundary and finite volume. In fact, the Gauß-Bonnet theorem implies that the condition on the Euler-characteristic is not only sufficient but also necessary for the existence of such a metric. Under a hyperbolic surface, orientable or not, we will understand in the sequel a surface endowed with such a complete hyperbolic metric with (possibly empty) totally geodesic boundary and finite volume.

As in the orientable setting, we let the Teichmüller space $\mathcal{T}(S)$ of S be the set of all isotopy classes of hyperbolic metrics on S . The pull-back action of the group of

diffeomorphisms on the set of hyperbolic metrics induces an action $\text{Map}(S) \curvearrowright \mathcal{T}(S)$ of the mapping class group on Teichmüller space.

Both in the orientable and in the non-orientable settings one can come up with many reasonable, but equivalent, ways to endow $\mathcal{T}(S)$ with a topology. The obtained topology on $\mathcal{T}(S)$ is mapping class group invariant and $\mathcal{T}(S)$ becomes homeomorphic to a euclidean space. If for example one endows $\mathcal{T}(S)$ with the topology induced by the embedding via length functions into the space $\mathbb{R}_+^{\mathcal{S}(S)}$ where $\mathcal{S}(S)$ is the set of essential curves, then a concrete homeomorphism between $\mathcal{T}(S)$ and the euclidean space of appropriate dimension is obtained via Fenchel-Nielsen coordinates. See [PP16] for a discussion on the Teichmüller space and on Fenchel-Nielsen coordinates for non-orientable surfaces.

2.4 Laminations and measured laminations. Let S be a hyperbolic surface with empty boundary, orientable or not, and recall that we require it to be complete and of finite volume. A *lamination* is a compact subset of S which admits a decomposition into disjoint simple geodesics. Those geodesics are unique—they are the *leaves* of the lamination. A basic example of a lamination on S is the closed geodesic isotopic to a simple essential curve. Simple geodesic multicurves are a slight generalization of this example: a *simple geodesic multicurve* is a lamination obtained by taking the collection of simple closed geodesics corresponding to an essential simple topological multicurve.

The set $\mathcal{L}(S)$ of all laminations on S is compact when endowed with the topology induced by the Hausdorff distance. Also, if $\phi : S \rightarrow S'$ is a homeomorphism between hyperbolic surfaces then there is a homeomorphism $\phi_* : \mathcal{L}(S) \rightarrow \mathcal{L}(S')$ sending simple closed geodesics γ to the geodesic freely homotopic to $\phi(\gamma)$. Moreover, homotopic homeomorphisms $S \rightarrow S'$ induce the same homeomorphism $\mathcal{L}(S) \rightarrow \mathcal{L}(S')$. It follows that there is an action $\text{Map}(S) \curvearrowright \mathcal{L}(S)$ of the mapping class group on the space $\mathcal{L}(S)$ of laminations.

A *measured lamination* λ is a lamination $\lambda_{\mathcal{L}}$ endowed with a transverse measure. Recall that a *transverse measure* is an assignment of a Radon measure λ^I on every segment I transversal to the lamination in such a way that segments which are isotopic relative to the lamination have the same measure

$$\iota(\lambda, I) \stackrel{\text{def}}{=} \lambda^I(I).$$

The *support* $\text{supp}(\lambda)$ is the smallest sublamination of $\lambda_{\mathcal{L}}$ such that $\lambda^I = 0$ for every segment I with $I \cap \text{supp}(\lambda) = \emptyset$. The set of all measured laminations in S is denoted by $\mathcal{ML}(S)$.

Again, the most basic example of a measured lamination is given by an essential simple closed geodesics γ : the underlying lamination is the geodesic γ itself and the transverse measure is such that it assigns to the segment I the counting measure of the set $I \cap \gamma$. We will often identify the homotopy class of the essential curve γ , the geodesics γ in this class, and the corresponding element $\gamma \in \mathcal{ML}(S)$.

The union of non-intersecting laminations is a lamination and the sum of transverse measures on two disjoint laminations is a transverse measure on the union. It follows that the sum $\lambda_1 + \lambda_2$ of two measured laminations is a well-defined measured lamination as long as the supports of λ_1 and λ_2 do not meet transversally. For example, if $\gamma_1, \dots, \gamma_r$ are the components of a simple geodesic multicurve and if c_1, \dots, c_r are positive reals, then $\sum c_i \gamma_i$ is a perfectly sound measured lamination, a *weighted multicurve*. Along the same lines, any measured lamination $\lambda \in \mathcal{ML}(S)$ has a unique decomposition $\lambda = \gamma + \lambda'$ where γ is a weighted multicurve corresponding to the atomic part of λ and where λ' is a measured lamination whose support does not contain any closed leaves.

We say that a measured lamination λ is *ergodic* if whenever $X \subset \text{supp}(\lambda)$ is a saturated subset of its support (that is, a union of leaves) then for every segment I we have that either $\lambda^I(I \cap X) = 0$ or $\lambda^I(I \setminus X) = 0$. Two measured laminations $\lambda_1, \lambda_2 \in \mathcal{ML}(S)$ are *mutually singular* if for every segment I transversal to $\lambda_{\mathcal{L}} = \text{supp}(\lambda_1) \cup \text{supp}(\lambda_2) \in \mathcal{L}(S)$ there is $Z \subset I$ with $\lambda_1^I(Z) = 0$ and $\lambda_2^I(I \setminus Z) = 0$.

Up to scaling, every lamination supports only finitely many ergodic measured laminations (recall for that matter that we give a concrete bound in Corollary 5.3). Moreover, every measured lamination can be written as the sum of ergodic ones. To be more precise, for every $\lambda \in \mathcal{ML}(S)$ there are pairwise mutually singular ergodic measured laminations $\mu_1, \dots, \mu_r \in \mathcal{ML}(S)$ with $\text{supp}(\mu_i) \subset \text{supp}(\lambda)$ and with

$$\lambda = \mu_1 + \dots + \mu_r. \tag{2.1}$$

In fact, the collection $\{\mu_1, \dots, \mu_r\}$ is unique. The measured laminations μ_1, \dots, μ_r are the *ergodic components* of λ and (2.1) is its *ergodic decomposition*.

Measured laminations can be seen as currents (see [Bon86, ES] for details) and hence can be identified with measures on the unit tangent bundle T^1S which are invariant under both the geodesic flow and the geodesic flip. To describe how this identification goes, suppose that $\gamma \in \mathcal{ML}(S)$ is a closed geodesic and note that it corresponds to two orbits γ_+ and γ_- of period $\ell_S(\gamma)$ of the geodesic flow: traversing the geodesic in one direction and in the other direction. The flip and flow invariant measure $\hat{\gamma}$ associated to γ is the convex combination of the Lebesgue measure along those two orbits. In symbols, this means that for $f \in C^0(T^1S)$ we have

$$\int_{T^1S} f(v) d\hat{\gamma}(v) = \frac{1}{2} \left(\int_0^{\ell_S(\gamma)} f(\rho_t(v)) dt + \int_0^{\ell_S(\gamma)} f(\rho_t(-v)) dt \right)$$

where $v \in T^1S$ is any vector tangent to the geodesic γ . The factor $\frac{1}{2}$ is there to ensure that the total measure of $\hat{\gamma}$ agrees with the length of γ . In general we refer to the total measure

$$\ell_S(\mu) = \|\hat{\mu}\|$$

of $\hat{\mu}$ as the *length* of $\mu \in \mathcal{ML}(S)$.

REMARK. We mentioned above the *ergodic decomposition of a measured lamination*. To our surprise we found out that it is not easy to give a reference for its existence—surprising because it is definitively common knowledge. Let us thus briefly explain a way to get it. Given a measured lamination $\mu \in \mathcal{ML}(S)$, let $\hat{\mu}$ be the associated flip and flow invariant measure. In other words, μ is a $(\mathbb{R} \times \mathbb{Z}/2\mathbb{Z})$ -invariant measure on T^1S . Applying the usual ergodic decomposition theorem (or maybe noting that the argument used to prove it for iterates of an individual map applies word-by-word for actions of a group such as $\mathbb{R} \times \mathbb{Z}/2\mathbb{Z}$) we get that $\hat{\mu}$ can be written as a linear combination of ergodic measures, or rather as an integral over the space of ergodic measures. Now, via the identification between flip and flow invariant measures and currents, we get that each one of the measures we are integrating over arises from an ergodic measured lamination carried by the support of μ . This means that we can write μ as an integral over the space of ergodic measured laminations with support contained in $\text{supp}(\mu)$. As we pointed out earlier (compare with Corollary 5.3) the latter set is finite, meaning that instead of an integral we have a sum as in (2.1).

The space of measures on T^1S is naturally endowed with the weak-* topology. We get thus a topology on $\mathcal{ML}(S)$ via the embedding of the latter in the space of flip and flow invariant measures on T^1S . We list a few properties of the space $\mathcal{ML}(S)$:

- $\mathcal{ML}(S)$ is a metrizable space and the corresponding projective space $\mathcal{PML}(S)$ is compact. In fact, it turns out that $\mathcal{ML}(S)$ and $\mathcal{PML}(S)$ are homeomorphic to a Euclidean space and a sphere respectively, but we will not need these facts.
- The set of weighted closed geodesics $c \cdot \gamma$ with $c > 0$ is dense in $\mathcal{ML}(S)$.
- The geometric intersection number of two essential curves extends to a continuous bilinear symmetric map $\iota : \mathcal{ML}(S) \times \mathcal{ML}(S) \rightarrow \mathbb{R}_{\geq 0}$, the *intersection form*.
- A concrete distance on $\mathcal{ML}(S)$ can be obtained as follows: it is known (see for example [Ham03]) that there are finite collections \mathcal{A} of curves which *separate measured laminations*, by which we mean that for any two distinct $\mu, \mu' \in \mathcal{ML}(S)$ there is $\alpha \in \mathcal{A}$ with $\iota(\mu, \alpha) \neq \iota(\mu', \alpha)$. Such a finite collection \mathcal{A} separating measured laminations induces a distance

$$d_{\mathcal{A}}(\mu, \mu') = \max_{\alpha \in \mathcal{A}} |\iota(\mu, \alpha) - \iota(\mu', \alpha)| \quad (2.2)$$

on $\mathcal{ML}(S)$.

- The mapping class group acts on $\mathcal{ML}(S)$ by homeomorphisms, preserving the intersection form.
- It is the key part of Thurston's classification of mapping classes that every pseudo-Anosov mapping class has north–south dynamics on $\mathcal{PML}(S)$. Moreover, the two fixed points λ_{\pm} are such that $\iota(\lambda_{\pm}, \gamma) > 0$ for every simple curve γ .




Figure 1: Two two-holed projective plane. The curves γ and η are the core and dual to the core, while α is a boundary component (so a curve around one of the punctures).

Although there has been no sign of it in what we said so far, la raison d'être of this paper is that measured laminations behave really differently in the non-orientable compared to the orientable setting. The reason behind this is the following theorem of Scharlemann [Sch82]:

Theorem (Scharlemann). *Let S be a non-orientable hyperbolic surface and let $\lambda \in \mathcal{ML}(S)$ be a measured lamination in S . If the support of λ contains a one-sided closed geodesic γ , then γ is also contained in the support of every measured lamination in a neighborhood of λ .*

Since this result is of capital importance for our paper, let us sketch the proof. A first basic fact is that every two-holed projective plane P contains exactly two homotopy classes of essential curves: both of them are one-sided and they intersect each other once (see Figure 1). If we declare one of them to be *the core of P* , then the other one is *dual to the core*. Now, the basic observation is that any one-sided curve γ arises as the core of some two-holed projective plane, and if $P \subset S$ is such a two-holed projective plane with core γ and if $\eta \subset P$ is dual to γ , then a measured lamination $\lambda \in \mathcal{ML}(S)$ has γ as an atom if and only if

$$\max_{\alpha \subset \partial P} \iota(\lambda, \alpha) < \iota(\lambda, \eta) \quad (2.3)$$

This inequality is an open condition and this proves the theorem.

REMARK. Note also that if $c \cdot \gamma$ is a one-sided atom of λ and if P and η are as above, then γ has weight $c = 2 \cdot (\iota(\lambda, \eta) - \max_{\alpha \subset \partial P} \iota(\lambda, \alpha))$ and this quantity depends continuously on λ . This means that not only are one-sided atoms stable, but so is their weight.

Before moving on we should point out that Danthony-Nogueira [DN90] proved a much stronger version of Scharlemann's theorem: the set of measured laminations with a one-sided atom is not only open but also has full measure, and hence is dense.

2.5 Non-orientable exceptional surfaces. As mentioned already, there are exactly three non-orientable surfaces with $\chi(S) = -1$, namely the two-holed projective plane (i.e. the one-holed Möbius band), the one-holed connected sum of two projective planes (i.e. one-holed Klein bottle), and the (closed) connected sum of three projective planes (i.e. the connected sum of a torus and a Möbius band). We denote a surface of genus g and r punctures as $N_{g,r}$ and so since $\chi(S) = 2 - g - r$, these surfaces are $N_{1,2}$, $N_{2,1}$ and $N_{3,0}$ respectively. Below we describe their associated spaces of measured laminations and mapping class groups. We refer to [Gen05, GK14] for more details.

We already discussed $N_{1,2}$ above: there are only two simple closed curves γ and η , both of which are one-sided, and γ and η intersect. Moreover, any non-compact simple geodesic must spiral towards either a boundary component or one of these closed curves. It follows that every measured lamination is either a multiple of γ or a multiple of η . In particular, $\mathcal{ML}^+(N_{1,2}) = \emptyset$ and $\mathcal{PML}(N_{1,2}) = \mathcal{PML}^-(N_{1,2})$ consists of two points. The mapping class group $\text{Map}(N_{1,2})$ is the finite group $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$.

Next we consider $N_{2,1}$, the one-holed Klein bottle. There is exactly one two-sided simple closed curve α and any other simple closed curve intersects it. Letting γ_0 be a one-sided closed curve, all other one-sided closed curves are obtained from γ_0 through Dehn twists along α and hence the set of one-sided closed curves is given by the bi-infinite sequence (γ_n) where $\gamma_n = D_\alpha^n(\gamma_0)$ for each $n \in \mathbb{Z}$. Note that γ_n is disjoint from exactly two other closed curves: γ_{n-1} and γ_{n+1} . Moreover, any non-compact simple geodesic spirals towards one of the closed curves or around the boundary component. Consequently, any measured lamination is a simple multicurve having at most two components. In particular $\mathcal{PML}^+(N_{2,1})$ consists of one point and any $[\lambda] \in \mathcal{PML}^-(N_{2,1})$ is of the form $[t\gamma_n + (1-t)\gamma_{n+1}]$ for some n and $t \in [0, 1]$. Moreover, γ_n converges projectively to α as n tends to $\pm\infty$. So $\mathcal{PML}(N_{2,1})$ is a circle with a marked point corresponding to α . The mapping class group $\text{Map}(N_{2,1})$ is $D_\infty \times \mathbb{Z}/2\mathbb{Z}$, where D_∞ is the infinite dihedral group.

REMARK. We stress that neither $N_{1,2}$ nor $N_{2,1}$ carry measured laminations whose support has a non-compact leaf.

Finally we deal with the connected sum of three projective planes. In this surface there exists a unique one-sided closed curve γ disjoint from every two-sided closed curve and such that $N_{3,0} \setminus \{\gamma\}$ is orientable. Note that $N_{3,0} \setminus \gamma$ is a one-holed torus T . Hence we can embed $\mathcal{ML}(T)$ into $\mathcal{ML}(N_{3,0})$. Moreover, any two-sided simple closed curve is homotopic to one in T , and in fact, any $\lambda \in \mathcal{ML}^+(N_{3,0})$ is contained in T . Any other one-sided curve intersects γ in one point. We identify the circle $\mathcal{PML}(T)$ with $\mathcal{PML}^+(N_{3,0})$ in the 2-sphere $\mathcal{PML}(N_{3,0})$ and it divides $\mathcal{PML}^-(N_{3,0})$ into two disks: one corresponding to measured laminations having γ as a leaf and the other corresponding to all multicurves having a one-sided leaf other than γ as a component. Note that since every mapping class must fix γ we can identify $\text{Map}(N_{3,0})$ with $\text{Map}(T)$, the (full) mapping class group of the one-holed torus.

3 Uniform train tracks

Throughout this section we assume that S is a connected non-exceptional hyperbolic surface—we will write $\ell(\cdot) = \ell_S(\cdot)$ for lengths in S of curves, arcs, and measured laminations. We assume also that $\lambda \in \mathcal{L}(S)$ is a geodesic lamination without closed leaves. If the reader prefers, they can think of λ being recurrent (that is, the support of a measured lamination) and connected, and of S being non-orientable.

In the informal version of Theorem 4.1 in the introduction we said something about a sufficiently nice train track. Our goal now is to clarify what kind of train tracks we will consider and prove two technical lemmas needed in later sections. We start by recalling what train tracks are.

3.1 Train tracks. A *train track* τ in a hyperbolic surface S is a smoothly embedded graph such that every vertex is contained in the interior of a smoothly embedded arc $I \subset \tau$ and such that none of its complementary regions is a disk with at most 2 cusps, an annulus, a Möbius band, or a disk with at most one puncture. See [ES], or rather [Hat88, PH92], for details.

Except as an aid in a construction in the appendix, we will always assume that train tracks don't have superfluous vertices: connected components homeomorphic to a circle have a single vertex and all other connected components have vertices which are at least trivalent. This implies that the complexity of a train track τ is bounded by the topology of the surface S it lives in. We can for example bound the possible number of vertices by noting that each complementary region of τ contributes a multiple of $-\pi$ to the Gauß-Bonnet integrand and has at most 3 cusps:

$$|V(\tau)| = \text{number of vertices in } \tau \leq 6|\chi(S)|.$$

Since both the number of vertices and the number of edges are maximal if the train track is trivalent, we also get the following bound on the number of edges:

$$|E(\tau)| = \text{number of edges in } \tau \leq 9|\chi(S)|.$$

These bounds imply that, from a combinatorial point of view, there are only finitely many train tracks. Since we will need this fact here and there, we record it as a lemma:

LEMMA 3.1. *The surface S has only finitely many mapping class group orbits of train tracks.* □

A lamination λ is *carried* by a train track τ if there is a continuous map

$$\Psi : [0, 1] \times \lambda \rightarrow S, \quad (t, x) \mapsto \Psi_t(x) \tag{3.1}$$

such that Ψ_t is an embedding for all $t \in [0, 1)$, Ψ_0 is the identity and Ψ_1 is an immersion with image contained in τ . If the support of a measured lamination is carried by τ then we say that the measured lamination itself is carried by τ . The set of measured laminations carried by τ is denoted by $\mathcal{ML}(\tau)$. It is a closed subset of

$\mathcal{ML}(S)$. Also, the set of weighted multicurves is dense in $\mathcal{ML}(\tau)$ —recall that one of the main results of this paper, Theorem 1.2, is to understand what is the closure of the set of weighted *two-sided* multicurves.

Still with the same notation suppose that a measured lamination $\lambda \in \mathcal{ML}(\tau)$ is carried by the train track τ , let e be an edge of τ , and let I be a smooth segment meeting τ transversally in a single point of e . The *weight* $\omega_\lambda(e)$ of λ on e is then the limit

$$\omega_\lambda(e) \stackrel{\text{def}}{=} \lim_{t \rightarrow 1} \iota(\lambda, \Psi_t^{-1}(I)),$$

where Ψ_t is as in (3.1). It is well-known that $\lambda \in \mathcal{ML}(\tau)$ is uniquely determined by the vector $(\omega_\lambda(e))_{e \in E(\tau)} \in \mathbb{R}_{\geq 0}^{E(\tau)}$ whose entries are the weights of the edges of τ with respect to λ . We denote by

$$W(\tau) = \{(\omega_\lambda(e))_{e \in E(\tau)} \in \mathbb{R}_{\geq 0}^{E(\tau)} \text{ with } \lambda \in \mathcal{ML}(\tau)\}$$

the set of all the so-obtained vectors. It is well-known that the map $\mathcal{ML}(\tau) \rightarrow W(\tau)$ is a homeomorphism and that $W(\tau)$ is the intersection of a linear subspace of $\mathbb{R}^{E(\tau)}$, the set of solutions of the switch equations, with the positive quadrant $\mathbb{R}_{\geq 0}^{E(\tau)}$.

3.2 Almost geodesic train tracks. A train track τ in our hyperbolic surface S is ε_0 -geodesic if the edges have geodesic curvature less than ε_0 at every point. We will fix ε_0 such that the following holds:

- (*) If γ is any bi-infinite path in the hyperbolic plane with geodesic curvature less than ε_0 and if γ_* is the geodesic which stays at bounded distance of γ , then the closest point projection $\gamma_* \rightarrow \gamma$ is well-defined, moves points less than distance 1 away, and distorts length by at most a factor 2.

Fixing ε_0 once and forever so that all of this holds, we will refer to an ε_0 -geodesic train track as an *almost geodesic train track*.

The reason to consider ε_0 such that (*) holds is that if τ is an almost geodesic train track and if λ is a lamination carried by τ , then applying to each leaf the closest point projection to its image under Ψ_1 as in (3.1) we get a continuous map

$$\Phi = \Phi^{\lambda, \tau} : \lambda \rightarrow \tau \tag{3.2}$$

satisfying:

- $d_S(\Phi(x), x) \leq 1$ for all $x \in \lambda$, and
- the restriction of Φ to each individual leaf of λ is a smooth immersion satisfying $\frac{1}{2} \leq \left| \frac{\partial \Phi}{\partial t} \right| \leq 2$ at every point.

The map Φ as in (3.2) is the *carrying map* of λ in τ . With this notation we say that λ *fills* τ if Φ is surjective—if this is the case we say that τ carries λ *in a filling way*.




Figure 2: On the left a local picture of a lamination λ ; on the right a local picture of a train track τ carrying λ . The thick edge is a λ -loopy edge of τ .

REMARK. It is well-known that for any lamination λ there is an almost geodesic train track τ carrying λ , say in a filling way, and that if λ has no closed leaves we can find such a τ whose edges are as long as we want. There are different ways to construct such a train track, but for example one can follow Thurston [Thu80, Section 8.9] and take a neighborhood of λ and then collapse it along a foliation whose leaves are short segments transversal to λ . Another way is to create a one-vertex train track using the first return map to a short transversal of λ . We encourage the reader to think of yet other ways—there are many.

Before moving on note that if $\lambda \in \mathcal{ML}(\tau)$ is carried by an almost geodesic train track τ then we have

$$\ell(\lambda) \leq \sum_{e \in E(\tau)} \omega_\lambda(e) \cdot \ell(e) \leq 2 \cdot \ell(\lambda). \quad (3.3)$$

We will often want to replace the actual length of the lamination by the middle term in (3.3)—this is one of the reasons explaining why factors of 2 keep appearing in this paper.

3.3 Uniform train tracks. The chain of inequalities (3.3) probably makes apparent that almost geodesic train tracks are geometric objects—objects that reflect the geometry of S . It is often however very comfortable to work with train tracks as if they were purely combinatorial objects. In some sense, *uniform train tracks* are the kind of train tracks where these two points of view are married. In a nutshell, a uniform train track is an almost geodesic train track all of whose edges have comparable lengths. The actual definition is a bit more subtle because we are using λ to measure the length of edges. That is why we first need some terminology.

With notation as above, suppose that τ is an almost geodesic train track carrying our lamination λ , and let $\Phi : \lambda \rightarrow \tau$ be the carrying map. A pair $\{e, e'\}$ consisting of two half-edges of τ based at the same vertex is said to be a *turn* of τ . A turn is *λ -legal* if it is represented by $\Phi(I)$ for some smooth segment $I \subset \lambda$. An edge $e \in E(\tau)$ of τ is *λ -loopy* if both its half-edges are based at the same vertex and form a λ -legal turn, see Figure 2. Finally we will say that τ is *λ -generic* if it has no bivalent vertices, if all its vertices have at most valence 4, and if moreover all valence 4 vertices correspond to λ -loopy edges.




Figure 3: Left: A local picture (in a Möbius band) of the one-sided edge e (in black) with its vertex v and two adjacent half-edges (in gray). Right: The result after splitting. The labels ω_i denotes the λ -weight of the corresponding (half)edge.




Figure 4: Left: A local picture (in a Möbius band) of the one-sided edge e (in black) with its vertex v and two adjacent half-edges (in gray). Middle and right: The result after splitting. The labels ω_i denotes the λ -weight of the corresponding (half)edge.

REMARK. If the lamination λ is understood from the context then we might simply say that a turn is *legal* instead of λ -legal, that an edge is *loopy* instead of λ -loopy, and that the train track is question is *generic* instead of λ -generic and so on.

Recall now that we are assuming λ has no closed leaves. This will in particular imply that any λ -loopy edge must be two-sided (in the sense that a regular neighborhood of the edge is an annulus):

LEMMA 3.2. *If $\lambda \in \mathcal{ML}(\tau)$ has no closed leaves, then any λ -loopy edge of τ must be two-sided.*

Proof. Let $\lambda \in \mathcal{ML}(\tau)$ and e be a λ -loopy edge of τ with vertex v . Note that v has valence at least 4 since $\omega_\lambda(e) > 0$ and λ has no closed leaves. For concreteness we will assume v has valence exactly 4—the below argument can easily be generalized to higher valence, and in any case, we will mostly be concerned with generic train tracks. Now, suppose e were one-sided. Then in a regular neighborhood around e (which is a Möbius band) the train track τ would look like the left most figure in either Figure 3 or Figure 4. In the first case, we can split τ to get the righthand picture in Figure 3 and hence λ has an atom (since $\omega_1 = \omega_\lambda(e) > 0$). In the second case, since e is loopy, we can split τ as in the middle picture in Figure 4 and continue until we arrive in the left most picture in that figure. Note that, with notation as in the figure, $\omega_1 - \omega_2 > 0$ since e is λ -loopy. We are now in the situation of the first case and can split similarly as there. Hence, in either case, we get that λ must have an atom if e is one-sided. \square

Still with the same notation, we next define the λ -length of a closed edge $e \in E(\tau)$. Intuitively, we define it simply as the length of the edge if it is not λ -loopy, but when it is loopy, it is the length of the edge multiplied by the maximum number of times

a (connected) segment of λ wraps around it. More precisely, the λ -length of an edge e is

$$\ell^{\lambda, \tau}(e) = \max_{I \in \pi_0(\Phi^{-1}(e))} \ell(\Phi(I)), \quad (3.4)$$

where $\ell(\Phi(I))$ is the length of $\Phi(I)$ as a parametrized curve. Alternatively, it is the length of I with respect to the pull-back metric of τ to λ . Note that, indeed, the λ -length $\ell^{\lambda, \tau}(e)$ and the usual length $\ell(e)$ agree if e is not λ -loopy and that when e is loopy we have $\ell^{\lambda, \tau}(e) = k \cdot \ell(e)$ for some $k \geq 2$.

We will write $\ell^\lambda(e)$ instead of $\ell^{\lambda, \tau}(e)$ when the train track τ is understood from the context. Accordingly, there will be no τ in the superscript in our notation for the λ -length of τ itself and for the minimum of the λ -lengths of the edges of τ :

$$\ell^\lambda(\tau) = \sum_{e \in E(\tau)} \ell^{\lambda, \tau}(e) \quad \text{and} \quad m^\lambda(\tau) = \min_{e \in E(\tau)} \ell^{\lambda, \tau}(e). \quad (3.5)$$

With this notation in place, we are now ready to formally define what is a (C, λ) -uniform train track.

DEFINITION. *Suppose that a lamination λ in a hyperbolic surface S has no closed leaves. An almost geodesic train track τ is (C, λ) -uniform for some $C > 0$ if λ is carried by τ in a filling way and if $\ell^\lambda(\tau) \leq C \cdot m^\lambda(\tau)$.*

The existence of uniform train tracks, or more specifically the fact that any train track can be refined to a uniform one, will be proved in the appendix of this paper. We say τ can be *refined* to a train track τ' if there is a smooth map

$$\Psi : [0, 1] \times \tau' \rightarrow S$$

such that Ψ_t is an embedding for all $t \in [0, 1)$ and such that Ψ_1 is an immersion with image τ (compare to the definition of a lamination being carried by a train track in Section 3.1). With slight abuse of terminology we set $\Phi^{\tau', \tau} = \Psi_1 : \tau' \rightarrow \tau$ which we call a carrying map and refer to τ' as a *refinement* of τ . With this definition in place, we state the proposition we will prove in Appendix A:

PROPOSITION 3.3. *Let S be a hyperbolic surface. There is a constant $C > 0$ such that for any geodesic lamination $\lambda \subset S$ without closed leaves, any train track τ_0 carrying λ , and any L large enough, there is a refinement τ of τ_0 which is λ -generic, (C, λ) -uniform and satisfies $\ell^\lambda(\tau) \geq L$.*

It is maybe worth mentioning that the constant C is not obtained via a compactness theorem: one can actually give concrete estimates for it. We will for example see that one can take C to be $1000|\chi(S)|$ if we drop the requirement of τ being generic and, say, $2^{30|\chi(S)|}$ if we do not. Clearly, we are not trying to obtain optimal bounds.

We will assume Proposition 3.3 for now and prove it in Appendix A.

3.4 Two lemmas on uniform train tracks. As we mentioned earlier, uniform train tracks are geometric objects (they are almost geodesic) but they act as if they were combinatorial objects (all edges have basically the same length, meaning that up changing our measuring unit from nanometer to kyndemil we can act as if all edges had unit length). We exploit this in the proof of two technical lemmas needed below.

The first one basically asserts that if a measured lamination carried by a uniform train track spends little of its life in some edge, then one can find another measured lamination nearby which avoids that edge altogether. The second is a very weakly quantified version of Scharlemann’s theorem.

LEMMA 3.4. *For every geodesic lamination λ and all positive C and δ there is $\varepsilon < \delta$ such that for every (C, λ) -uniform, λ -generic train track τ the following holds: For every unit length $\mu \in \mathcal{ML}(\tau)$ there is a second unit length measured lamination $\nu \in \mathcal{ML}(\tau)$ with*

- (1) $\ell(e) \cdot |\omega_\mu(e) - \omega_\nu(e)| < \delta$ for all $e \in E(\tau)$, and
- (2) $\omega_\nu(e) = 0$ for every $e \in E(\tau)$ with $\ell(e) \cdot \omega_\mu(e) < \varepsilon$.

We encourage the reader to find counter examples to Lemma 3.4 if one drops the assumption that τ is uniform.

Proof. To begin with let τ^* be the train track obtained from τ by removing all open λ -loopy edges and let τ^{**} be the union of the closed λ loopy edges. In other words, τ^{**} is a disjoint union of circles, $\tau = \tau^* \cup \tau^{**}$ and τ^* and τ^{**} only meet at the vertices of τ^{**} . Note also that not only do we have $\mathbb{R}_{\geq 0}^{E(\tau)} = \mathbb{R}_{\geq 0}^{E(\tau^*)} \oplus \mathbb{R}_{\geq 0}^{E(\tau^{**})}$ but also

$$W(\tau) = W(\tau^*) \oplus W(\tau^{**}).$$

Said in plain terms, every solution of the switch equations of τ^* and every solution of the switch equations of τ^{**} give, when taken together, a solution to the switch equations of τ , and every solution arises in this way.

All of this means that, up to maybe solving the problem for some smaller δ , we can assume that τ is either the union of disjoint circles or a trivalent train track which has thus no loopy edges. For the sake of concreteness we will assume that we are in the second case—the case that τ is a union of disjoint circles can be dealt with using exactly the same argument, but we are sure that the reader can come up with a much simpler one.

Arguing by contradiction, suppose that the claim fails to be true for some λ , C and δ , meaning that there are $\varepsilon_i \rightarrow 0$ such that for each ε_i one can find some

- (C, λ) -uniform trivalent train track τ_i ,
- and $\mu_i \in \mathcal{ML}(\tau_i)$ with $\ell(\mu_i) = 1$

for which we cannot find any $\nu_i \in \mathcal{ML}(\tau_i)$ satisfying (1) and (2) in the statement for $\varepsilon = \varepsilon_i$.

We get from Lemma 3.1 that the surface S contains only finitely many homeomorphism types of train tracks. This means that, up to passing to a subsequence, we can assume that there is some train track τ and for each i a homeomorphism $\phi_i : S \rightarrow S$ with $\phi_i(\tau) = \tau_i$. The map ϕ_i induces a linear isomorphism—denoted by the same symbol—between the associated spaces of solutions of the switch equations:

$$\phi_i : W(\tau) \rightarrow W(\tau_i)$$

Our μ_i gives us a vector $\omega_{\mu_i} = (\omega_{\mu_i}(e))_{e \in E(\tau_i)} \in W(\tau_i)$. Consider the scaled vector

$$w_i = \ell(\tau_i) \cdot \omega_{\mu_i}$$

and note that

$$\sum_{e \in E(\tau_i)} w_i(e) = \sum_{e \in E(\tau_i)} \ell(\tau_i) \cdot \omega_{\mu_i}(e) \geq \sum_{e \in E(\tau_i)} \ell(e) \cdot \omega_{\mu_i}(e) \geq \ell(\mu_i) = 1$$

where the final inequality holds because of (3.3).

To get a bound in the other direction note that, since τ_i is not only (C, λ) -uniform but also trivalent and hence has no loopy edges, we have $\ell(\tau_i) = \ell^\lambda(\tau_i) \leq C \cdot \ell^\lambda(e) = C \cdot \ell(e)$ for every edge $e \in E(\tau_i)$. This implies that

$$\begin{aligned} \sum_{e \in E(\tau_i)} w_i(e) &= \sum_{e \in E(\tau_i)} \ell(\tau_i) \cdot \omega_{\mu_i}(e) \\ &\leq C \cdot \sum_{e \in E(\tau_i)} \ell(e) \cdot \omega_{\mu_i}(e) \\ &\leq C \cdot 2 \cdot \ell(\mu_i) = 2 \cdot C \end{aligned}$$

where again the last inequality holds because of (3.3). Now, the chain of inequalities

$$1 \leq \sum_{e \in E(\tau_i)} w_i(e) \leq 2 \cdot C$$

implies that the elements $(\phi_i)^{-1}(w_i) \in W(\tau)$, for all i , belong to a compact set and that, up to passing to a subsequence, the limit

$$w = \lim_i (\phi_i)^{-1}(w_i) \in W(\tau)$$

exists and is non-zero. Now, if we let $\nu_i \in \hat{\mathcal{ML}}(\tau_i)$ be the measured lamination corresponding to $\phi_i(w)$, then $\nu_i = \frac{1}{\ell(\hat{\nu}_i)} \hat{\nu}_i$ satisfies (1) and (2) for $\varepsilon = \frac{\delta}{2}$ and all sufficiently large i . This is a contradiction to our assumption—we are thus done with the proof of Lemma 3.4. □

We proceed now to our (rather weak) quantification of Scharlemann’s theorem:

LEMMA 3.5. *For all C and ε positive there is $\delta < \varepsilon$ such that if $\tau \subset S$ is a (C, λ) -uniform λ -generic train track, if $\mu, \nu \in \mathcal{ML}(\tau)$ are measured laminations carried by τ with*

$$\ell(e) \cdot |\omega_\mu(e) - \omega_\nu(e)| < \delta$$

for all $e \in E(\tau)$, and if ν is such that

- ν has unit length,
- ν has a one-sided atom with at least weight $\frac{\varepsilon}{\ell(\tau)}$, and
- $\omega_\nu(e) = 0$ for every loopy edge e ,

then μ has a one-sided atom.

The key point of this lemma is that the given δ does the trick for every train track τ satisfying the stated conditions.

Proof. Seeking a contradiction we assume that the lemma fails for some C and ε , meaning that there are

- a sequence (τ_n) of (C, λ) -uniform λ -generic train tracks,
- a sequence (ν_n) of unit length measured laminations carried by τ_n with $\omega_{\nu_n}(e) = 0$ for every loopy edge of τ_n ,
- a sequence $(c_n \cdot \gamma_n)$ of one-sided atoms of ν_n with weight $c_n \geq \frac{\varepsilon}{\ell(\tau_n)}$, and
- a sequence (μ_n) of measured laminations satisfying

$$\ell_S(e) \cdot |\omega_{\mu_n}(e) - \omega_{\nu_n}(e)| < \frac{1}{n}$$

for every edge $e \in E(\tau_n)$ and every n , and such that γ_n is not a one-sided atom of μ_n .

As in the proof of Lemma 3.4, let us denote by $\tau_n^{**} \subset \tau_n$ the union of the loopy edges of τ_n and note that we have

$$\ell(e) = \ell^{\lambda, \tau_n}(e) \geq \frac{1}{C} \ell^\lambda(\tau_n) \geq \frac{1}{C} \cdot \ell(\tau_n) \tag{3.6}$$

for every edge $e \in E(\tau_n) \setminus E(\tau_n^{**})$. Let K_n be the number of times that γ_n meets a vertex of τ_n . Since ν_n , and thus γ_n , do not travel through τ_n^{**} , we get that

$$1 = \ell(\nu_n) \geq \ell(c_n \cdot \gamma_n) \geq \frac{1}{2} \cdot \frac{\varepsilon}{\ell(\tau_n)} \cdot K_n \cdot \frac{1}{C} \cdot \ell(\tau_n)$$

where the last inequality arises as follows: the factor $\frac{1}{2}$ because we are measuring lengths in the almost geodesic train track τ_n , the next factor is the lower bound on the weight c_n , next is the number of edges that γ_n travels through in τ_n , and finally the lower bound for the length of each such edge. Cleaning up we get that

$$K_n \leq \frac{2C}{\varepsilon}$$

is uniformly bounded.

Now, by Lemma 3.1 there are finitely many homeomorphism types of train tracks. Since any train track carries only finitely many curves made out of a given number of edges, we get that, up to passing to a subsequence, we can assume that there is a train track τ , a sub-train track $\tau^{**} \subset \tau$, a one-sided curve γ carried by τ but disjoint of τ^{**} , and a sequence (ϕ_n) of homeomorphisms with

$$\tau_n = \phi_n(\tau), \quad \tau_n^{**} = \phi_n(\tau^{**}) \text{ and } \gamma_n = \phi_n(\gamma).$$

We let $\hat{\nu}_n = \phi_n^{-1}(\nu_n)$ and note that $c_n \cdot \gamma$ is a one-sided atom of $\hat{\nu}_n$ for all n . This means, as (2.3) in the proof of Scharlemann’s theorem and the remark following it, that whenever $P \subset S$ is a two-holed projective plane with core curve γ and dual curve η we have

$$\iota(\hat{\nu}_n, \eta) = \max_{\alpha \subset \partial P} \iota(\hat{\nu}_n, \alpha) + \frac{1}{2} \cdot c_n \tag{3.7}$$

for all n . Note that τ^{**} can be seen as a two-sided multicurve because it is homeomorphic to τ_n^{**} , the union of loopy edges of τ_n . It follows that, since γ is disjoint of τ^{**} , we can also choose the two-holed projective plane to be contained in $S \setminus \tau^{**}$. It follows that ∂P and η only meet edges in $\tau \setminus \tau^{**}$.

We claim now that the measured laminations $\hat{\mu}_n = \phi_n^{-1}(\mu_n)$ also satisfy (2.3) for all sufficiently large n . Indeed, taking into account that the difference between the intersection numbers of $\hat{\mu}_n$ and $\hat{\nu}_n$ with ∂P is bounded by the product of the number of edges of $\hat{\tau}$ that ∂P meets times the maximal difference between the weights that $\hat{\mu}_n$ and $\hat{\nu}_n$ give to those edges, we get

$$\begin{aligned} |\iota(\hat{\mu}_n, \partial P) - \iota(\hat{\nu}_n, \partial P)| &\leq \text{const} \cdot \max_{e \in E(\tau) \setminus E(\tau^{**})} |\omega_{\hat{\mu}_n}(e) - \omega_{\hat{\nu}_n}(e)| \\ &\leq \text{const} \cdot \max_{e \in E(\tau_n) \setminus E(\tau_n^{**})} |\omega_{\mu_n}(e) - \omega_{\nu_n}(e)| \\ &\leq \text{const} \cdot \frac{1}{n} \cdot \max_{e \in E(\tau_n) \setminus E(\tau_n^{**})} \frac{1}{\ell(e)} \\ &\leq \text{const} \cdot \frac{1}{n} \cdot \frac{C}{\ell(\tau_n)}. \end{aligned}$$

Here the max is taken over $E(\tau) \setminus E(\tau^{**})$ because we have chosen P disjoint of τ^{**} . A similar computation yields that, up to possibly replacing the constant by another one,

$$|\iota(\hat{\mu}_n, \eta) - \iota(\hat{\nu}_n, \eta)| \leq \frac{\text{const}}{n \cdot \ell(\tau_n)}$$

Taking these two bounds with (3.7), and changing once again our constant, we get that

$$\iota(\hat{\mu}_n, \eta) - \max_{\alpha \subset \partial P} \iota(\hat{\mu}_n, \alpha) \geq \frac{1}{2} \cdot c_n - \frac{\text{const}}{n \cdot \ell(\tau_n)}$$

for all n . Since we have the bound $c_n \geq \frac{\varepsilon}{\ell(\tau_n)}$ we deduce that

$$\iota(\hat{\mu}_n, \eta) > \max_{\alpha \subset \partial P} \iota(\hat{\mu}_n, \alpha)$$

for all large n . This means that (2.3) holds and thus that $\hat{\mu}_n$ has γ as an atom. Since the pairs $(\hat{\mu}_n, \gamma)$ and (μ_n, γ_n) just differ by the homeomorphism ϕ_n we get that for all large n the measured lamination μ_n has the one-sided curve γ_n as an atom. This is a contradiction, and we are done. \square

4 Separating ergodic components

Continuing with the same notation as in the previous section we assume that S is a connected non-exceptional hyperbolic surface and that $\lambda \in \mathcal{L}(S)$ is a connected geodesic lamination without closed leaves. We will also assume that it is recurrent. Since λ is connected, this last assumption just means that $\lambda = \text{supp}(\mu)$ for every non-trivial measured lamination $\mu \in \mathcal{ML}(\lambda)$ supported by λ .

Our goal now is to prove Theorem 4.1:

Theorem 4.1. *Let S be a connected and possibly non-orientable hyperbolic surface, let $\lambda \in \mathcal{L}(S)$ be a connected recurrent lamination, and let τ_0 be a train track carrying λ . Let $\mu_1, \dots, \mu_r \in \mathcal{ML}(\lambda)$ be distinct unit length ergodic measured laminations with support λ , and let $\mathcal{A} \subset S$ be a collection of simple curves separating measured laminations.*

There is C such that for every $\varepsilon > 0$ there are

- *a (C, λ) -uniform λ -generic refinement τ^ε of τ_0 , and*
- *unit length measured laminations $\nu_1^\varepsilon, \dots, \nu_r^\varepsilon \in \mathcal{ML}(\tau^\varepsilon)$*

satisfying the following:

- (1) $\ell_S(e) \cdot |\omega_{\mu_i}(e) - \omega_{\nu_i^\varepsilon}(e)| \leq \delta(\varepsilon)$ for every $i = 1, \dots, r$ and $e \in E(\tau^\varepsilon)$.
- (2) If $\tau_j^\varepsilon \subset \tau^\varepsilon$ is the sub-train track filled by ν_j^ε then we have $\tau_i^\varepsilon \cap \tau_j^\varepsilon = \emptyset$ for all $i \neq j \in \{1, \dots, r\}$.
- (3) We have $d_{\mathcal{A}}(\mu_i^\varepsilon, \mu_i) \leq \delta(\varepsilon)$ for all i, ε , and unit length $\mu_i^\varepsilon \in \mathcal{ML}(\tau_i^\varepsilon)$.

Here $\delta(\cdot)$ is a function with $\lim_{\varepsilon \rightarrow 0} \delta(\varepsilon) = 0$ and $d_{\mathcal{A}}$ is as in (2.2).

As pointed out by a referee, in the above statement it might be more natural to start by giving δ , drop ε altogether, and just label everything in terms of δ . However, in the course of the proof (e.g. see Proposition 4.3) τ^ε and ν_i^ε are defined in terms of a parameter ε which we think of as tending to 0. We thus find it more convenient to think of δ and ε as different parameters.

Given that the statement of Theorem 4.1 is rather technical, the reader might be wondering what it actually means. It might help to re-visit the informal version that we stated in the introduction and compare with the one here. Also to help clarify the statement of the theorem, and because it will be useful when proving Corollary 5.3, we prove the following result:

COROLLARY 4.2. *Let $\lambda \in \mathcal{L}(S)$ be a lamination and $\mu_1, \dots, \mu_r \in \mathcal{ML}(\lambda)$ be distinct unit length ergodic measured laminations with support λ . There is a sequence of weighted simple multicurves (γ^i) with $\gamma^i = c_1^i \cdot \gamma_1^i + \dots + c_r^i \cdot \gamma_r^i$ such that $\mu_j = \lim_{i \rightarrow \infty} c_j^i \cdot \gamma_j^i$ for all $j = 1, \dots, r$.*

Corollary 4.2 is due to Lenzhen-Masur [LM10, Theorem C] in the orientable case.

In some sense, the proof of Theorem 1.2 from the introduction boils down to proving that if λ in Corollary 4.2 has no one sided leaves, then the multicurves γ_i can be chosen to be two-sided; compare with the proof of Theorem 1.2 in Section 5.2.

Proof. We will just prove the claim under the assumption that λ is connected and recurrent, leaving to the reader to make the necessary changes to deal with the disconnected case. These assumptions imply that $\lambda = \text{supp}(\mu_j)$ for all j . Let also \mathcal{A} be as in the statement of Theorem 4.1 and, for $\varepsilon \rightarrow 0$ let τ^ε and $\tau_1^\varepsilon, \dots, \tau_r^\varepsilon$ be the produced train tracks. Let then $c_j^\varepsilon \cdot \gamma_j^\varepsilon \in \mathcal{ML}(\tau_j^\varepsilon)$ be a unit length weighted curve carried by τ_j^ε —as we mention earlier the set of weighted curves is dense, and in particular non-empty, in the set of measured laminations carried by a train track. Now, we have

$$\mu_j = \lim_{\varepsilon \rightarrow 0} c_j^\varepsilon \cdot \gamma_j^\varepsilon$$

because of (3) in the theorem. Moreover, $c_1^\varepsilon \cdot \gamma_1^\varepsilon + \dots + c_r^\varepsilon \cdot \gamma_r^\varepsilon$ is a simple multicurve because, by (2) in the theorem, we have $\tau_\varepsilon^i \cap \tau_\varepsilon^j = \emptyset$ for all $i \neq j$. We are done. \square

4.1 Proof of Theorem 4.1. We turn now to the proof of Theorem 4.1. We will complete it first modulo the following result that we deal with below. It basically says that, as long as the shortest edge in the train track has sufficiently long λ -length, then the measured laminations μ_1, \dots, μ_r spend most of their time in different parts of the train track. More precisely, if we write

$$E(\tau, \mu, \varepsilon) = \{e \in E(\tau) \text{ with } \ell(e) \cdot \omega_\mu(e) > \varepsilon\} \tag{4.1}$$

then the proposition says that for all ε the sets $E(\tau, \mu_i, \varepsilon)$ and $E(\tau, \mu_j, \varepsilon)$ are disjoint for suitably chosen τ and $\mu_i \neq \mu_j$:

PROPOSITION 4.3. *Suppose that $\lambda, \mu_1, \dots, \mu_r$ and \mathcal{A} are as in Theorem 4.1. For every ε there is L_ε such that for every almost geodesic train track τ carrying λ and with $m^\lambda(\tau) \geq L_\varepsilon$ the following holds:*

- We have $E(\tau, \mu_i, \varepsilon) \cap E(\tau, \mu_j, \varepsilon) = \emptyset$ for all $i \neq j$.
- Moreover, any unit length $\nu_i \in \mathcal{ML}(\tau)$ with

$$\{e \in E(\tau) \text{ with } \omega_{\nu_i}(e) \neq 0\} \subset E(\tau, \mu_i, \varepsilon)$$

satisfies $d_{\mathcal{A}}(\nu_i, \mu_i) < \delta(\varepsilon)$ with $\lim_{\varepsilon \rightarrow 0} \delta(\varepsilon) = 0$.

Proposition 4.3 is basically a property of mutually singular ergodic measures invariant under some dynamical system. The main difference between Theorem 4.1 and Proposition 4.3 is that the latter does not assert anything about the topology of the sets of edges $E(\tau, \mu_i, \varepsilon)$ —for all we know, this set could consist of a single edge, meaning that there is no non-trivial measured lamination satisfying the condition in the second statement of the proposition. We work with uniform train tracks exactly to get some control of these sets. This is where Lemma 3.4 will come in handy.

Assuming Proposition 4.3 for the time being, we prove Theorem 4.1:

Proof of Theorem 4.1. First note that if λ has a closed leaf then we actually have that λ is itself a circle and hence that the theorem is trivially true. We suppose from now on that λ has no closed leaves.

Let C be the constant provided by Proposition 3.3 and let L_ε be the constant provided by Proposition 4.3 for our ε . By Proposition 3.3 there is a (C, λ) -uniform λ -generic refinement $\tau = \tau^\varepsilon$ of τ_0 with $\ell^\lambda(\tau) \geq C \cdot L_\varepsilon$. Uniformity implies that $m^\lambda(\tau) \geq L_\varepsilon$, and hence that Proposition 4.3 applies.

We start however by applying Lemma 3.4 to our ε and to each μ_i , letting $\nu_i = \nu_i^\varepsilon \in \mathcal{ML}(\tau)$ be the obtained measured lamination. We are going to prove that these measured laminations have the properties claimed in the statement of the theorem.

To begin with, note that (1) in the theorem holds because ν_i satisfies (1) in Lemma 3.4.

To see that (2) is satisfied, note that, by Lemma 3.4, the smallest sub-train track $\tau_i = \tau_i^\varepsilon \subset \tau$ carrying ν_i is such that $E(\tau_i) \subset E(\tau, \mu_i, \varepsilon)$. Now Proposition 4.3 implies that $E(\tau_i) \cap E(\tau_j) = \emptyset$ for all $i \neq j$. This implies that $\tau_i \cap \tau_j = \emptyset$ because they are both sub-train tracks of the λ -generic train track τ . We have proved that (2) in the theorem is also satisfied.

To conclude, note that the final statement in Proposition 4.3 implies that also (3) in the theorem is satisfied. □

The remaining of this section is devoted to prove Proposition 4.3.

4.2 Approximating segments. Given a point $x \in \lambda$ and some $L > 0$ we denote by

$$I_{x,L}^\lambda : [-L, L] \rightarrow \lambda$$

the path parametrized by arc length and mapping 0 to x . We will not distinguish between the path $t \mapsto I_{x,L}^\lambda(t)$ and its opposite $t \mapsto I_{x,L}^\lambda(-t)$. In fact, abusing notation we often identify $I_{x,L}^\lambda$ with its image, acting as if $I_{x,L}^\lambda = I_{x,L}^\lambda([-L, L])$. We care about these segments because, for large L , they do a good job approximating λ . It will however be important to only work with segments $I_{x,L}^\lambda$ which are basically determined by the point $\Phi(x)$, where Φ is the carrying map of λ in τ . Note however first that $\Phi(I_{x,L}^\lambda)$ and $\Phi(I_{x',L}^\lambda)$ could actually be very different even if $\Phi(x) = \Phi(x')$: they might be close for quite a while, but then, upon passing a vertex, they might

diverge into different edges of τ . This is why we consider, for each closed edge $e \in E(\tau)$, the set

$$\mathbb{I}_L^\lambda(e) = \{x \in \lambda \text{ with } \Phi(I_{x,L}^\lambda) \subset e\}$$

of those points x such that $I_{x,L}^\lambda$ projects into e . Denote by

$$\mathbb{I}_L^\lambda(\tau) = \cup_{e \in E(\tau)} \mathbb{I}_L^\lambda(e)$$

the union of all those sets. We stress that to define $\mathbb{I}_L^\lambda(e)$ and thus $\mathbb{I}_L^\lambda(\tau)$ we considered closed edges. This only matters, but then it matters a lot, if e is a loopy edge. The following lemma asserts that for points in $\mathbb{I}_L^\lambda(\tau)$ the segment $I_{x,L}^\lambda$ is to all effects determined by x .

LEMMA 4.4. *For any finite set of closed geodesics \mathcal{A} there is a constant $c(\mathcal{A})$ such that if τ is an almost geodesic train track and if λ is a lamination carried by τ , then we have*

$$\left| |\alpha \cap I_{x,L}^\lambda| - |\alpha \cap I_{x',L}^\lambda| \right| \leq c(\mathcal{A})$$

for any $\alpha \in \mathcal{A}$ and any two $x, x' \in \mathbb{I}_L^\lambda(\tau)$ with $\Phi(x) = \Phi(x')$.

Proof. There is d_0 such that whenever we have $\Phi(x) = \Phi(x')$ for $x, x' \in \mathbb{I}_L^\lambda(\tau)$, then (up to reversing one of the segments if necessary) there is a homotopy from $I_{x,L}^\lambda$ to $I_{x',L}^\lambda$ whose tracks have length at most d_0 —for example $d_0 = 6$ does the trick. Now, for any given closed geodesic α there is a bound $c(\alpha)$ for the maximal number of times that it can meet a geodesic segment of length d_0 in general position. This quantity, or rather its double, bounds then the difference between the number of times that α can meet $I_{x,L}^\lambda$ and $I_{x',L}^\lambda$ if $x, x' \in \mathbb{I}_L^\lambda(\tau)$ are such that $\Phi(x) = \Phi(x')$. The claim follows when we take $c(\mathcal{A}) = 2 \cdot \max_{\alpha \in \mathcal{A}} c(\alpha)$. \square

Recall now that any measured lamination, and thus any unit length $\mu \in \mathcal{ML}(\lambda)$ can be interpreted as a flip and flow invariant probability measure $\hat{\mu}$ on T^1S . Let then $\bar{\mu}$ be the push-forward of $\hat{\mu}$ under the projection map $T^1S \rightarrow S$. We refer to $\bar{\mu}$ as *the probability measure on S associated to the unit length lamination μ* .

We will want to make sure that we arrange things in such a way that $\mathbb{I}_L^\lambda(\tau)$ has large $\bar{\mu}$ -measure. Indeed, if one fixes the train track τ first and chooses L second it might well be that $\mathbb{I}_L^\lambda(\tau)$ is empty—it makes more sense to do things the other way around:

LEMMA 4.5. *For all L_1 and ε there is L_2 such that the following holds: If $\mu \in \mathcal{ML}(\lambda)$ is a unit length measured lamination supported by λ and if τ is an almost geodesic train track with $m^\lambda(\tau) > L_2$ then*

$$\bar{\mu}(\mathbb{I}_{L_1}^\lambda(\tau)) \geq 1 - \varepsilon$$

where $\bar{\mu}$ is the probability measure on S associated to μ .

Proof. With some L_2 to be determined below, suppose that τ is almost geodesic with $m^\lambda(\tau) > L_2$ and let $\Phi : \lambda \rightarrow \tau$ be the corresponding carrying map. Note that to get the desired bound it suffices to prove that

$$\bar{\mu}(\Phi^{-1}(e) \setminus \mathbb{I}_{L_1}^\lambda(e)) < \varepsilon \cdot \bar{\mu}(\Phi^{-1}(e)) \tag{4.2}$$

for all $e \in E(\tau)$.

The set $\Phi^{-1}(e)$ is the disjoint union of segments. If e is not loopy then Φ maps each such segment s homeomorphically onto e and we thus have $\ell(\Phi(s)) = \ell^\lambda(e)$ and hence that $\ell(s) \geq \frac{1}{2} \cdot \ell(e) > \frac{1}{2} \cdot L_2$ because Φ distorts distances by at most a factor of 2. If e is loopy then we can either have $\ell(\Phi(s)) = \ell^\lambda(e)$ or $\ell(\Phi(s)) = \ell^\lambda(e) - \ell(e)$ because some of the arcs may turn around e once less than others. In any case, we get from the definition of λ -loopy that $\ell^\lambda(e) \geq 2 \cdot \ell(e)$ for every λ -loopy edge e . This means that we still have

$$\ell(\Phi(s)) \geq \frac{1}{4} \cdot L_2$$

for any component $s \in \pi_0(\Phi^{-1}(e))$, independently of e being loopy or not.

On the other hand we have for each such segment $s \in \pi_0(\Phi^{-1}(e))$ that

$$\ell(s \setminus \mathbb{I}_L^\lambda(e)) \leq 4 \cdot L_1 + 4$$

because s has two end-points, because Φ displaces points at most distance 1 and distorts distances by at most a factor 2.

This means that if we take $L_2 > (16 \cdot L_1 + 16) \cdot \varepsilon^{-1}$ we have that

$$\ell(s \setminus \mathbb{I}_L^\lambda(e)) \leq \varepsilon \cdot \ell(s).$$

Since this holds for every $s \in \pi_0(\Phi^{-1}(e))$ we get (4.2), as we wanted. □

To conclude with these generalities note that if $e \in E(\tau)$ is an edge in an almost geodesic train track carrying λ then we have $\bar{\mu}(\Phi^{-1}(e)) \geq \frac{1}{2} \cdot \ell(e) \cdot \omega_\mu(e)$ because each segment in $\Phi^{-1}(e)$ has at least length $\frac{1}{2} \cdot \ell(e)$ and there are $\omega_\mu(e)$ worth of segments. If e is loopy this is still true when we just cut e open at a point (equivalently, when we work with open edges). We record a minimal generalisation of this fact:

LEMMA 4.6. *If τ is an almost geodesic train track carrying λ and $\mu \in \mathcal{ML}(\lambda)$ is a unit length measured lamination supported by λ then we have*

$$\bar{\mu}(\Phi^{-1}(U)) \geq \frac{1}{2} \cdot \ell(U) \cdot \omega_\mu(e)$$

for every edge $e \in E(\tau)$ and every measurable subset $U \subset e$. Here $\bar{\mu}$ is the probability measure on S associated to μ and $\Phi : \lambda \rightarrow \tau$ is the carrying map. □

All of this concerns arbitrary measured laminations with support λ —things become much more interesting when we restrict ourselves to ergodic measured laminations.

4.3 The Birkhoff game. The segments $I_{x,L}^\lambda$ considered above will be important because they do not only approximate λ , but also, if x is well-chosen, they approximate concrete measured laminations supported by λ . Indeed, applying the Birkhoff ergodic theorem, we get that for every L^1 -function f the limit

$$\lim_{L \rightarrow \infty} \frac{1}{2L} \int_{-L}^L (f \circ I_{x,L}) dt$$

of time averages exists for $\bar{\mu}$ -almost every $x \in \text{supp}(\mu)$. Moreover, if μ is ergodic then we get that this limit of time averages actually agrees with the space average $\int f d\bar{\mu}$. The convergence in the ergodic theorem is only pointwise, but we get from Egorov’s theorem that it is uniform as long as we are willing to forgo a set of x ’s with small but positive measure. In this way we get:

LEMMA 4.7. *Let $\mu \in \mathcal{ML}(\lambda)$ be a unit length ergodic measured lamination supported by λ and let \mathcal{A} be a finite collection of closed geodesics. For every ε, δ there is a compact set $X_{\mu,\mathcal{A},\delta,\varepsilon} \subset \lambda$ with $\bar{\mu}(X_{\mu,\mathcal{A},\delta,\varepsilon}) \geq 1 - \varepsilon$ and some $L_0 \geq 0$ with*

$$\max_{\alpha \in \mathcal{A}} \left| \iota(\mu, \alpha) - \frac{1}{2L} |I_{x,L}^\lambda \cap \alpha| \right| \leq \delta$$

for all $x \in X_{\mu,\mathcal{A},\delta,\varepsilon}$ and $L \geq L_0$. □

With the same notation as above choose some $\varepsilon > 0$, let \mathcal{A} be a finite collection of closed geodesics and let $\mu, \mu' \in \mathcal{ML}(\lambda)$ be unit length ergodic measured laminations with support λ . Note that, as long as the intersection numbers $\iota(\mu, \alpha)$ and $\iota(\mu', \alpha)$ are distinct enough for some $\alpha \in \mathcal{A}$, we have that the sets $X_{\mu,\mathcal{A},\delta,\varepsilon}$ and $X_{\mu',\mathcal{A},\delta,\varepsilon}$ provided by Lemma 4.7 are disjoint. Indeed, we get from Lemmas 4.4 and 4.7 that for every $\delta > 0$ with

$$\delta < \frac{1}{2} \cdot d_{\mathcal{A}}(\mu, \mu'),$$

there is $L_1 > 0$ such that for every almost geodesic train track τ carrying λ we have

$$\Phi(X_{\mu,\mathcal{A},\delta,\varepsilon} \cap \mathbb{I}_{L_1}^\lambda(\tau)) \cap \Phi(X_{\mu',\mathcal{A},\delta,\varepsilon} \cap \mathbb{I}_{L_1}^\lambda(\tau)) = \emptyset. \tag{4.3}$$

Recall that $d_{\mathcal{A}}(\cdot, \cdot)$ is the distance induced by \mathcal{A} —see (2.2).

Let now L_2 be the constant provided by Lemma 4.5 for our L_1 and ε , and assume that $m^\lambda(\tau) > L_2$. Then we get from Lemmas 4.7 and 4.5

$$\begin{aligned} \bar{\mu}(X_{\mu,\mathcal{A},\delta,\varepsilon} \cap \mathbb{I}_{L_1}^\lambda(e)) &\geq \bar{\mu}(\Phi^{-1}(e)) - 2 \cdot \varepsilon \\ \bar{\mu}'(X_{\mu',\mathcal{A},\delta,\varepsilon} \cap \mathbb{I}_{L_1}^\lambda(e)) &\geq \bar{\mu}'(\Phi^{-1}(e)) - 2 \cdot \varepsilon \end{aligned}$$

for every edge $e \in E(\tau)$. Combining these two bounds with Lemma 4.6 we get that

$$\ell(e \setminus \Phi(X_{\mu,\mathcal{A},\delta,\varepsilon} \cap \mathbb{I}_{L_1}^\lambda(e))) \geq \frac{4 \cdot \varepsilon}{\omega_\mu(e)}$$

$$\ell(e \setminus \Phi(X_{\mu', \mathcal{A}, \delta, \varepsilon} \cap \mathbb{I}_{L_1}^\lambda(e))) \geq \frac{4 \cdot \varepsilon}{\omega_{\mu'}(e)}$$

From (4.3) we thus get that at most for one of μ and μ' we can have that the right side is larger than $\frac{1}{2} \cdot \ell(e)$, meaning that we have the implication

$$\ell(e) \cdot \omega_\mu(e) > 8 \cdot \varepsilon \Rightarrow \ell(e) \cdot \omega_{\mu'}(e) < 8 \cdot \varepsilon$$

and the same if we revers the roles of μ and μ' . With $E(\cdot, \cdot, \cdot)$ as in (4.1), this reads as

$$E(\tau, \mu, 8 \cdot \varepsilon) \cap E(\tau, \mu', 8 \cdot \varepsilon) = \emptyset.$$

We summarize what we got so far:

LEMMA 4.8. *Let $\lambda \in \mathcal{L}(S)$ be a lamination and \mathcal{A} a finite set of closed geodesics separating measured laminations. Let also μ_1 and $\mu_2 \in \mathcal{ML}(\lambda)$ be distinct, unit length, ergodic measured laminations with support λ . Then for every ε there is L_2 such that whenever τ is an almost geodesic train track carrying λ and with $m^\lambda(\tau) \geq L_2$ then $E(\tau, \mu_1, \varepsilon) \cap E(\tau, \mu_2, \varepsilon) = \emptyset$. \square*

We are now ready to conclude the proof of Proposition 4.3.

4.4 Proof of Proposition 4.3. From Lemma 4.8 we get that there is some L_ε such that for any almost geodesic train track τ carrying λ and with $m^\lambda(\tau) \geq L_\varepsilon$ for all $e \in E(\tau)$ we have

$$E(\tau, \mu_i, \varepsilon) \cap E(\tau, \mu_j, \varepsilon) = \emptyset$$

for all $i \neq j \in \{1, \dots, r\}$. This was the first claim of the proposition. Suppose now that we have $\nu_i \in \mathcal{ML}(\tau)$ with

$$\{e \in E(\tau) \text{ with } \omega_{\nu_i}(e) \neq 0\} \subset E(\tau, \mu_i, \varepsilon). \tag{4.4}$$

For $\alpha \in \mathcal{A}$ we can then compute

$$\iota(\nu_i, \alpha) = \frac{1}{2L_0} \int |I_{x, L_0}^{\nu_i} \cap \alpha| \, d\bar{\nu}_i(x),$$

where $L_0 = L(\mu, \mathcal{A}, \delta)$ is as in Lemma 4.7, and where we abused notation slightly by writing $I_{x, L_0}^{\nu_i}$ instead of $I_{x, L_0}^{\text{supp}(\nu_i)}$. Now, since $\text{supp}(\nu_i)$ is carried by the almost geodesics train track τ we get as in the argument leading to Lemma 4.4 that $|I_{x, L_0}^{\nu_i} \cap \alpha|$ is basically the same as $|I_{y, L_0}^\lambda \cap \alpha|$ for any $y \in \lambda \cap \mathbb{I}_{L_0}^\lambda(\tau)$ with $\Phi(y) = \Phi(x)$:

$$\frac{1}{2L_0} \left| |I_{x, L_0}^{\nu_i} \cap \alpha| - |I_{y, L_0}^\lambda \cap \alpha| \right| < \delta + \frac{c(\mathcal{A})}{L_0}$$

with $c(\mathcal{A})$ as in Lemma 4.4.

Combining assumption (4.4) and Lemma 4.7 we get that for most $x \in \text{supp}(\nu_i)$ there is $y_x \in X_{\mu_i, \mathcal{A}, \delta, \varepsilon}$ with $\Phi(x) = \Phi(y_x)$. Now, if the train track τ is such that $m^\lambda(\tau)$ is sufficiently large, we get from Lemma 4.5 that for most x the corresponding y_x belongs to $\mathbb{I}_{L_0}^\lambda(\tau)$, say for all x outside a set of ν_i -measure $\rho(L)$ where $\lim_{L \rightarrow 0} \rho(L) = 0$.

On the other hand there is a constant which just depends on \mathcal{A} and the geometry of the underlying surface, bounding the number of intersections that a segment of length L_0 can have with \mathcal{A} . We get altogether there is some $\text{const} > 0$ with

$$\begin{aligned} |\iota(\mu_i, \alpha) - \iota(\nu_i, \alpha)| &\leq \frac{1}{2L_0} \int ||I_{x, L_0}^{\nu_i} \cap \alpha| - |I_{y_x, L_0}^\lambda \cap \alpha|| \, d\bar{\nu}_i(x) \\ &\leq \left(\delta + \frac{c(\mathcal{A})}{L_0} \right) (1 - \rho(L)) + \text{const} \cdot \rho(L). \end{aligned}$$

This means that, if we wanted to guarantee that ν_i and μ_i have basically the same intersection number with α , first we should take $L_0 = L(\mu_i, \mathcal{A}, \delta)$ large as in Lemma 4.7 and then we should make L huge and assume that our train track τ satisfies $m^\lambda(\tau) \geq L$. We are done. □

5 Approximating by two-sided curves

In this section we will prove Theorem 1.2 from the introduction, the central result of this paper. The main tool is Theorem 4.1 from the previous section: we will show that, when λ has no one-sided leaves, each one of the train tracks τ_j^ε given by that theorem carries a two-sided curve. To do so we will show that if a train track does not carry a two-sided closed curve, then the only measured laminations it does carry are one-sided multicurves which are not too long. We then conclude the proof using Lemma 3.5 to argue that the measured laminations ν_i^ε given by Theorem 4.1 are not such multicurves. In any case, we start by investigating which train tracks carry two-sided curves.

5.1 Train tracks without two-sided curves. Our goal here is to describe which train tracks fail to carry any two-sided curves. For expedience, and since it will suffice for our purpose, we will only consider one-vertex train tracks.

PROPOSITION 5.1. *Let τ be a one-vertex train track. Either τ carries a two-sided simple closed curve or every measured lamination carried by τ is a one-sided multicurve, each of whose components passes through the vertex at most twice.*

Recall that we only apply the adjectives one-sided and two-sided to curves which are essential—this means in particular that orientation preserving curves bounding Möbius bands do not count as two-sided.

Starting with the discussion of the proposition, consider τ a one-vertex train track. Since it will help to visualize what we describe here, we will arbitrary choose to denote the outgoing half-edges at v as *top* half-edges and the incoming as *bottom* half-edges. We will have three types of edges:




Figure 5: On the left, a one-vertex train track as a graph. As a band complex in the middle. And on the right, the instructions manual indicating how the bands have to be plugged into the switchboard to build the band complex.

- (1) t–b edges, formed by a top half-edge and a bottom half-edge,
- (2) t–t edges, formed by two top half-edges,
- (3) b–b edges, formed by two bottom half-edges.

Since τ only has one vertex, every closed edge e is a loop and hence determines a homotopy class of a simple closed curve, which we call the *curve associated to e* —note that this curve is carried by e if and only if e is of form t–b. Moreover, the *curve associated to a pair* of edges of type t–t and b–b is the homotopy class of the simple closed curve obtained as the concatenation of the two edges. Finally, we say an edge is two-sided (resp. one-sided) if the curve associated to it is two-sided (resp. one-sided).

So far we have been thinking of train tracks as graphs, but it is well-known that one can think of them as being made out of bands (see for example [Ota96]). In our concrete setting we represent our single vertex by a relatively wide and low rectangle that we imagine as a switch board. Accordingly, we take a long and thin band for each edge, looking a bit like a flat cable, and we plug the two short sides of the band somewhere at the top and bottom sides of our switchboard, possibly both ends on the top, or bottom, or one end on the top and the other on the bottom. We do that in such a way that at the end of the day the 2-complex that we obtain can be embedded into the surface as a regular neighborhood of our train track—in the language of [ES] what we have gotten is a *thickening* of the train track. See the left and center pictures in Figure 5 for a schematic representation of the train track as a graph and as a band complex. The rightmost image in Figure 5 is the set of instructions that one needs to follow to recover the train track: plug in the bands into the switchboard as indicated in the instructions manual, the two ends of the band labeled “a” in the positions labeled by “a” with the orientation indicated by the arrows. Note in particular that for t–b edges, two arrows in the same direction tell us that the edge associated is two-sided, while two arrows going in different directions tell us that the edge associated is one-sided; for t–t and b–b edges it is the other way around.

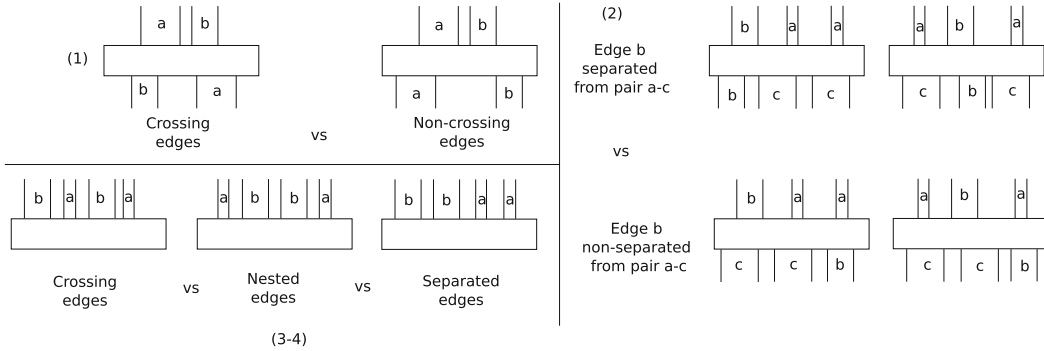


Figure 6: Some examples of crossing and non-crossing t - b edges (1), t - b edges separated and non-separated from a t - t , b - b pair (2), crossing, separated and nested t - t edges (3-4). Note that there are no arrows because the definitions are the same whether the edges are one- or two-sided.

Suggesting that the reader keeps Figure 6 in mind, we will now look at some possible mutual positions of edges.

DEFINITION 1. We say that

- (1) two t - b edges are crossing if the two corresponding curves γ and γ' have intersection number $i(\gamma, \gamma') = 1$,
- (2) a t - b edge e is separated from a pair of edges e_t of type t - t and e_b of type b - b if the curve γ corresponding to e and the curve γ' corresponding to the pair e_t, e_b have $i(\gamma, \gamma') = 0$,
- (3) two edges of type t - t (resp. b - b) are crossing if the corresponding curves γ and γ' have intersection number $i(\gamma, \gamma') = 1$,
- (4) two edges e and e' of type t - t (resp. b - b) which are not crossing, can be either nested or separated according to whether, for any fixed b - b (resp. t - t) edge e_0 , the curves γ and γ' associated to the pairs e, e_0 and e', e_0 have $i(\gamma, \gamma') = 0$ (then they are nested) or $i(\gamma, \gamma') = 2$ (then they are separated).

We are now ready to characterize those one-vertex train tracks that do not carry any two-sided simple closed curve.

LEMMA 5.2. Let τ be a one-vertex recurrent train track. If any of the following statements fails, then τ carries a two-sided simple closed curve.

- (1) All edges of type t - b are one-sided.
- (2) All edges of type t - t are two-sided (resp. one-sided), while all edges of type b - b are one-sided (resp. two-sided).
- (3) All pairs of t - b edges are non-crossing.
- (4) All t - b edges are separated from all pairs of t - t and b - b edges.
- (5) If there are two or more t - t edges, then they are crossing (resp. nested).
- (6) If there are two or more b - b edges, then they are nested (resp. crossing).




Figure 7: The two-sided simple closed curves carried by τ if (1), (2) or (3) in Lemma 5.2 fail.

(7) *There is at most one t-t edge or at most one b-b edge.*

Recall that a train track is recurrent if it carries a curve which gives positive weight to every edge.

Proof. First note that if (1) fails, meaning that there is a t-b edge which is two-sided, then the corresponding curve is a two-sided simple closed curve carried by τ . We can suppose from now on that (1) holds true.

Similarly, if there is a pair of edges of types t-t and b-b which are both one-sided or both two-sided, then the curve corresponding to the pair is a two-sided simple closed curve. This means that if (2) fails then τ carries a two-sided curve and we are done. We can assume from now on that (2) holds true. But not only that. We can assume, up to swapping the roles of top and bottom, that all the t-t edges are two-sided and all the b-b edges are one-sided.

Now assume that (3) fails, and hence that there is a pair of t-b edges which are crossing. Then one can construct a two-sided simple closed curve running along both of them as shown in Figure 7. Ergo, we can also assume that (3) holds.

If there is a t-b edge which is non-separated from a pair of t-t and b-b edges, then one can find a two-sided simple closed curve passing once through each edge. Figure 8 shows an example, and all other possible configurations can be obtained from the same figure by flipping the picture horizontally or vertically. We can thus assume that (4) holds true.

Assume now that there are two (two-sided) t-t edges that are either nested or separated. Recurrence implies that there must also be at least one (one-sided) b-b edge. As indicated by the top half of Figure 9, we get a two-sided simple closed curve which passes once along each t-t edge and twice along the b-b edge. This means that we can suppose that also (5) holds.

Similarly, if (6) fails we get a simple two-sided curve as in the bottom half of Figure 9. We can thus assume that (6) holds.




Figure 8: The two-sided simple closed curve carried by τ if (4) in Lemma 5.2 fails.




Figure 9: The two-sided curve carried by τ if either (5) or (6) in Lemma 5.2 fails.




Figure 10: The two-sided curve carried by τ if (7) in Lemma 5.2 fails.

Finally, assume that (7) fails, that is that we have two t-t edges and two b-b edges and note that by the above we can assume that the t-t are crossing and the b-b are nested. Then we can construct a two-sided simple closed curve passing once through each edge as in Figure 10.

As we wanted, we have found our two-sided curve if any one of the conditions (1)–(7) fails to be true. □

Our next goal is to figure out which measured laminations are carried by recurrent train tracks τ for which conditions (1)–(7) in Lemma 5.2 hold true.

Let us start by dissecting such a train track τ . From (2) we know that, up to switching the roles of top and bottom, we can assume that all t-t edges are two-sided

and that all b–b edges are one-sided. Moreover, from (7) we know that there can be either at most a single t–t edge or at most a single b–b edge. The discussions of both main cases (there is a single t–t edge or a single b–b edge) are very similar and the case that such edges do not exist is just a degenerate version of the main cases. To avoid repeating ourselves we will thus assume that there is a single b–b edge, leaving the other cases to the reader.

Summing up, we are assuming that there is one and only one b–b edge which is moreover one-sided (see Figure 11, part (a)).

Now recurrence implies that there is also at least one t–t edge: if no such t–t edge exists then all solutions of the switch equations give 0-value to our b–b edge. Now, every t–t edge is two-sided by (2), but we might have more than one such edge. However, by (5), each pair of t–t edges must be crossing, so their mutual positions need to be like in Figure 11, part (b).

Now we need to see where the possible t–b edges could be plugged in. First of all, if there are any t–b edges, then they are one-sided by (1). Now, by (4), they need to be separated from each t–t, b–b pair. This means that they can be in either of three possible positions with respect to the other edges, as in Figure 11, part (c). Finally, there could be more than one t–b edges in any of these three positions, but by (3) they need to be mutually non-crossing. See Figure 11 part (d) for the schematic description of our train track. (Figure 11 part (e) is a schematic description of the train track if there is instead a single t–t edge but maybe many b–b edges).

Once that we understand how a train track τ satisfying (1)–(7) in Lemma 5.2 looks like, let us figure out what are the measured laminations carried by τ . Given a solution $(\omega(e))_{e \in E(\tau)}$ to the switch equation at the single vertex v we can represent each edge e by a band of thickness $\omega(e)$. The band corresponding to the vertex then has thickness equal to the sum of the weights of all outgoing (top) half-edges, which is by definition the same as the sum of the weights of all the incoming (bottom) half-edges.

First note that the t–b edges give equal contributions to the incoming and outgoing components. Then, to satisfy the switch equation, one needs to have that the sum of the weights of the t–t edges must be equal to the weight of the single b–b edge.

This means that we can subdivide the rectangle representing the vertex into five (possibly degenerate) sub-rectangles (see Figure 12). In each rectangle the sum of the weights of the edges at the top is equal to the sum of the weights of the edges at the bottom. The first rectangle (from the left), A , has attached to it the top and bottom extremes of non-crossing t–b edges. The second rectangle B has attached to it one extreme of the single one-sided edge at the bottom and one extreme of each of the crossing two-sided edges at the top. The third rectangle C has the same non-crossing t–b edges attached to both the top and the bottom. The fourth rectangle D has the other extremes of the edges attached to B . The fifth rectangle E has again the same non-crossing t–b edges attached to both the top and the bottom.

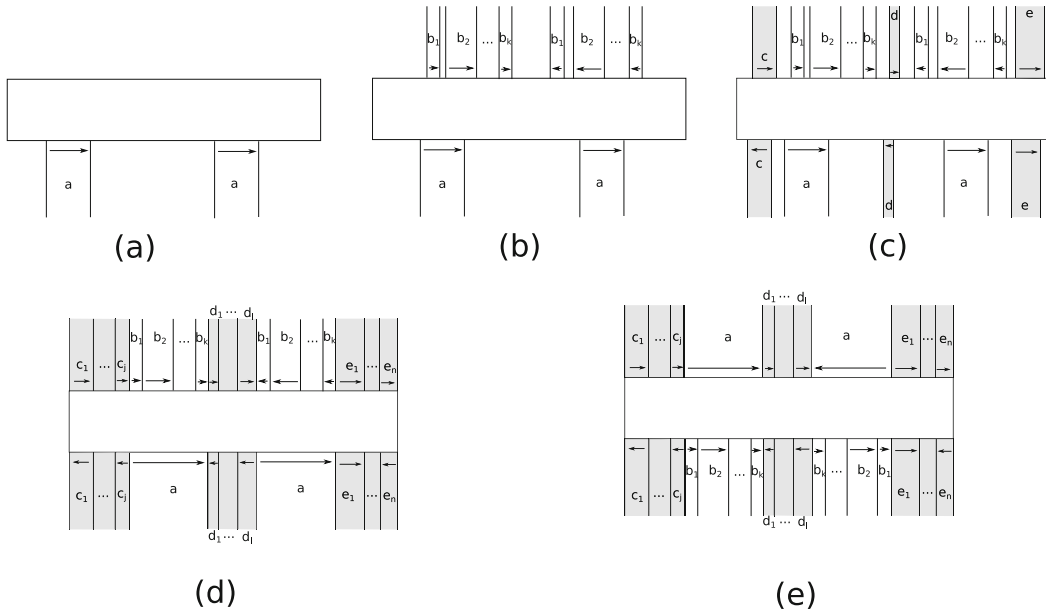


Figure 11: The instruction manual of a train track satisfying (1)–(7) in Lemma 5.2 in the “single b–b edge” model: first plug in the unique b–b edge (a), then plug in the t–t edges making sure that they are crossing (b), and finally start plugging in the t–b edges making sure that are separated from every t–t, b–b pair (c). The final product is represented in (d). Finally, (e) is a representation of the final product if what you have is the “single t–t edge” model instead.




Figure 12: Splitting the train track τ represented by Figure 11 part (d) into 4 disjoint train tracks.

In order to see which measured laminations are carried by τ , inside each sub-rectangle we need to match each interval at the top (resp. bottom) with an interval at the bottom (resp. top) of same length, without intersections. Inside A , the set of intervals at the top match exactly the ones at the bottom in length. This means that we can only connect a t–b band to itself. Then a measured lamination λ carried by τ can be written as $\lambda = c_1\gamma_1 + \dots + c_a\gamma_a + \lambda'$, where a is the number of t–b bands

attached to A , γ_i is the one-sided curve associated to each one of these bands and λ' is carried by the sub-train track obtained removing A . Note that each component γ_i passes through the vertex only once. The exact same thing can be done with E .

Inside B , the sum of the weights at the top is equal to the weight of the single band at the bottom. So we need to subdivide the bottom interval in b subintervals, where b is the number of crossing two-sided bands at the top of the same length of the ones at the top. Then we match the interval in the only possible way to not intersect. A similar thing needs to be done in D . Because of the way the bands attached to B and D are connected, the components of λ coming from these rectangles can only be one-sided weighted curves, each of which passes through the vertex exactly twice.

Finally, inside C , the intervals are again matching because there are only t - b edges. This means that the components given by C are again weighted one-sided curves passing through the vertex exactly once.

Summing up, we have proved that if a recurrent one-vertex train track τ satisfies conditions (1)–(7) in Lemma 5.2 then every measured lamination carried by τ is a one-sided weighted multicurve, each of whose components passes through the vertex at most twice. Since by Lemma 5.2 any recurrent one-vertex train track for which any one of the conditions (1)–(7) fails carries a two-sided curve, we have proved Proposition 5.1. \square

5.2 Density of two-sided curves in $\mathcal{ML}^+(S)$. We can now prove Theorem 1.2 from the introduction, stating that $\mathcal{PML}^+(S)$ is exactly the closure of the set of two-sided curves in $\mathcal{PML}(S)$:

Theorem 1.2 *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. The set of two-sided curves is dense in $\mathcal{PML}^+(S)$.*

Proposition 5.1 will play a key role in the proof of the above theorem. In particular, in the course of the proof we will have to pass from dealing with a given train track to a one-vertex train track. Let us explain how to do this. Suppose τ is a train track carrying a measured lamination ν and let e be an edge (since it is what we will do in the proof, the reader can take e to be an edge of maximal weight). Take a short transversal I to e , think of it as a transversal to ν , and consider the first return map. Identify arcs that are homotopic in (S, I) and then collapse I to a point v . The result is a refinement $\hat{\tau}$ of τ with a single vertex v .

Proof. Let $\mu \in \mathcal{ML}^+(S)$ be the measured lamination that we want to approximate by weighted simple two-sided curves. Denote by $\lambda = \text{supp}(\mu)$ the support of μ and let τ be a train track carrying λ in a filling way. We will approximate μ by two-sided multicurves carried by τ . This means in particular that we can assume without loss of generality that λ is connected. Noting that if λ is a closed curve then it is two-sided because we are assuming that $\mu \in \mathcal{ML}^+(S)$, we see that the only case we truly need to consider is that λ is connected but is not a closed curve—meaning in particular that λ has no compact leaves. This is the setting from now on.

Let then $\mu = \sum_{j=1}^n a_j \mu_j$ be the ergodic decomposition of μ , with $\mu_j \in \mathcal{ML}(\lambda)$ normalized in such a way that all of the μ_j have unit length, and fix also a finite collection \mathcal{A} of curves separating measured laminations. For ε very small—we will make it tend to 0—let τ^ε , ν_j^ε , and τ_j^ε , for $j = 1, \dots, n$, be given by Theorem 4.1. Forgetting for now the ν_j^ε 's, recall that there is some uniform C such that τ^ε is a (C, λ) -uniform refinement of τ , that τ_j^ε is a sub-train track of τ^ε for each j , that we have $\tau_i^\varepsilon \cap \tau_j^\varepsilon = \emptyset$ whenever $i \neq j$, and that we have

$$d_{\mathcal{A}}(\mu_j^\varepsilon, \mu_j) < \delta(\varepsilon)$$

for any unit length $\mu_j^\varepsilon \in \mathcal{ML}(\tau_j^\varepsilon)$. In the proof of Corollary 4.2 we argued that all of this implies that if we take any curve γ_j^ε carried by $\mathcal{ML}(\tau_j^\varepsilon)$ and if we chose weights a_j^ε wisely, then the weighted multicurves $\gamma^\varepsilon = \sum_{j=1}^n a_j^\varepsilon \cdot \gamma_j^\varepsilon$ converge to μ . What is different now is that we have to argue that we can choose the curves γ_j^ε to be two-sided.

This is where the ν_j^ε come into the game, but first note that any loopy edge is two-sided—this means that we can assume that our train tracks τ_j^ε do not have loopy edges. From Theorem 4.1 we get that ν_j^ε has unit length, is carried by τ_0^ε , and satisfies

$$\ell(e) \cdot |\omega_{\mu_j}(e) - \omega_{\nu_j^\varepsilon}(e)| < \delta(\varepsilon)$$

where $\delta(\varepsilon) \rightarrow 0$. The assumption that τ_j^ε has no loopy edges implies that $\omega_{\nu_j^\varepsilon}(e) = 0$ for every loopy $e \in E(\tau^\varepsilon)$, and hence that Lemma 3.5 applies. Since μ_j is supported by λ , we know for a fact that it has no atom, and a fortiori no one-sided atom. We get thus from Lemma 3.5 that any one-sided atom $c_j^\varepsilon \cdot \gamma_j^\varepsilon$ of ν_j^ε has to satisfy that

$$\lim_{\varepsilon \rightarrow 0} c_j^\varepsilon \cdot \ell(\tau^\varepsilon) = 0. \tag{5.1}$$

We are going to argue that this implies that τ_j^ε carries a two-sided curve, as we wanted to know.

Seeking a contradiction, suppose that this is not the case and let e be an edge of τ_j^ε with maximal possible weight $\omega_{\nu_j^\varepsilon}(e)$. Since ν_j^ε is unit length, since τ_j^ε is uniform without loopy edges, and since there is a universal bound on how many edges it can have, we have

$$\omega_{\nu_j^\varepsilon}(e) \geq \frac{\text{const}}{m^\lambda(\tau_j^\varepsilon)} \geq \frac{\text{const}}{m^\lambda(\tau^\varepsilon)} \tag{5.2}$$

where $m^\lambda(\cdot)$ is as in (3.5). Now refine τ_j^ε to a one-vertex train track $\hat{\tau}_j^\varepsilon$ as explained above, using a transversal to the edge e and the first return map of ν_j^ε . Noting that any two-sided curve in $\hat{\tau}_j^\varepsilon$ is also carried by τ_j^ε we get from our assumption and from Proposition 5.1 that every measured lamination carried by $\hat{\tau}_j^\varepsilon$ is a one-sided multicurve whose components pass through the vertex at most twice. This applies in particular to the largest sublamination $\sum_k c_k^{\varepsilon,j} \cdot \gamma_k^{\varepsilon,j}$ of ν_j^ε carried by $\hat{\tau}_j^\varepsilon$ (the whole

of ν_j^ε if τ_j^ε is connected). Noting that each one of the one-sided curves $\gamma_k^{\varepsilon,j}$ pass through the vertex of $\hat{\tau}_j$ at most twice, we get that $\sum_k c_k^{\varepsilon,j} \leq 2 \cdot \omega_{\nu_j^\varepsilon}(e)$. Since the number of components is bounded we have a constant such that there is always some k with $c_k^{\varepsilon,j} \geq \text{const} \cdot \omega_{\nu_j^\varepsilon}(e)$. Up to replacing the constant by another constant we get from (5.2) that for all ε and j the measured lamination ν_j^ε has an atom of at least weight $c_k^{\varepsilon,j} \geq \frac{\text{const}}{m^\lambda(\tau^\varepsilon)}$. Since τ^ε is (C, λ) -uniform we also get, up to again replacing the constant by yet another one, that ν_j^ε has an atom whose weight satisfies

$$c_k^{\varepsilon,j} \geq \frac{\text{const}}{\ell(\tau^\varepsilon)}.$$

This contradicts (5.1). □

Before moving on we point out that Theorem 1.2, or rather its proof, yields a bound for the number of mutually singular ergodic measured laminations with given support.

Indeed, suppose that λ is a recurrent lamination in S and let $\mu_1, \dots, \mu_n \in \mathcal{ML}(S)$ be mutually singular ergodic measured laminations with support λ . By Corollary 4.2 we can find a sequence of simple multicurves $c_i^1 \cdot \gamma_i^1 + \dots + c_i^n \cdot \gamma_i^n$ such that

$$\mu_j = \lim_{i \rightarrow \infty} c_i^j \cdot \gamma_i^j. \tag{5.3}$$

It follows directly that n is bounded by the maximal cardinality $c(S)$ of a multicurve in S . Now suppose that λ has no one-sided leaves. Then, as in the proof of Theorem 1.2 we get that we can assume that each of the curves γ_i^j in (5.3) can be chosen to be two-sided. This means that in this case n is bounded by the maximal cardinality $c^+(S)$ of a two-sided multicurve in S . We have proved the following:

COROLLARY 5.3. *Let $c(S)$ and $c^+(S)$ be respectively the maximal number of components of a multicurve and of a two-sided multicurve in S . Every lamination $\lambda \in \mathcal{L}(S)$ supports at most $c(S)$ mutually singular ergodic transverse measures. Moreover, if λ has no one-sided leaves then it supports at most $c^+(S)$ mutually singular ergodic transverse measures. □*

As we mentioned in the introduction, this result is due to Levitt [Lev83] in the orientable case.

6 Orbit closures in \mathcal{ML}

In this section we prove Theorem 1.4. We continue to assume that S is a non-exceptional hyperbolic surface. As in the introduction we associate to a measured lamination $\lambda \in \mathcal{ML}(S)$ a complete pair $(R_\lambda, \gamma_\lambda)$ as follows. Start by the decomposition $\lambda = \gamma_\lambda + \lambda'$ where γ_λ is the atomic part of λ and where λ' has no closed leaves. Let then R_λ be the possibly disconnected subsurface R_λ of S obtained by taking the

union of components of $S \setminus \gamma_\gamma$ which contain a non-compact leaf of λ . We note that none of the components of R_λ can be a pair of pants, a one-holed Möbius band, or a one-holed Klein bottle—in other words, no component of R_λ is exceptional.

As in the introduction we define

$$\gamma_\lambda + \mathcal{ML}^+(R_\lambda) = \{\gamma_\lambda + \mu \text{ with } \mu \in \mathcal{ML}^+(R_\lambda)\} \subset \mathcal{ML}(S)$$

and

$$\mathcal{G}_\lambda = \cup_{\phi \in \text{Map}(S)} \phi(\gamma_\lambda + \mathcal{ML}^+(R_\lambda)).$$

We can now recall the statement of Theorem 1.4:

Theorem 1.4 *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. We have $\overline{\text{Map}(S)} \cdot \lambda = \mathcal{G}_\lambda$ for any measured lamination $\lambda \in \mathcal{ML}(S)$.*

As we mentioned in the introduction, this theorem is due to Lindenstrauss-Mirzakhani [LM08] in the orientable case (see also Hamenstädt [Ham09]). They obtained it as a corollary of the classification of mapping class group invariant measures on $\mathcal{ML}(S)$, always for S orientable. Given that this measure classification does not exist (for good reasons) in the non-orientable case, we have to give a direct proof of the theorem. The argument below works both in the orientable and non-orientable case and uses only elementary facts about measured laminations and Theorem 1.2. Theorem 1.2 is well-known when S is orientable.

Let us start with a discussion of the theorem. The set \mathcal{G}_λ is, by its very definition, mapping class group invariant and contains λ . This means that the inclusion $\overline{\text{Map}(S)} \cdot \lambda \subset \mathcal{G}_\lambda$ follows once we know that the latter set is closed.

LEMMA 6.1. *For any $\lambda \in \mathcal{ML}(S)$ the set \mathcal{G}_λ is closed in $\mathcal{ML}(S)$.*

Proof. First note that we get from Scharlemann's theorem [Sch82] that the set $\mathcal{ML}^+(R_\lambda)$ is closed in $\mathcal{ML}(R_\lambda)$. It follows that also $\gamma_\lambda + \mathcal{ML}^+(R_\lambda)$ is closed in $\mathcal{ML}(S)$. This means that \mathcal{G}_λ is the union of closed sets. To prove that such a union is closed it suffices to show that any sequence $(\mu_n) \subset \mathcal{G}_\lambda$ which converges in the ambient space $\mathcal{ML}(S)$ is contained in finitely many of those closed sets, that is, in finitely many translates of $\gamma_\lambda + \mathcal{ML}^+(R_\lambda)$. Letting thus (μ_n) be such a sequence and noting that there is a sequence of mapping classes $(\phi_n) \in \text{Map}(S)$ with $\mu_n \in \phi_n(\gamma_\lambda + \mathcal{ML}^+(R_\lambda))$ for all n , we get that $\mu_n = \gamma_n + \lambda_n$ where $\gamma_n = \phi_n(\gamma_\lambda)$ and where $\lambda_n \in \mathcal{ML}^+(\phi_n(R_\lambda))$.

The assumption that μ_n converges as $n \rightarrow \infty$ implies in particular that its length $\ell(\mu_n)$ is bounded. It follows that the length of γ_n is also bounded, meaning that the multicurves γ_n belong to a finite set. Once we know γ_n there are only finitely many choices for $\phi_n(R_\lambda)$ because it is a union of components of $S \setminus \gamma_n$. It follows that the sequence (μ_n) belongs to the union of finitely many translates of $\gamma_\lambda + \mathcal{ML}^+(R_\lambda)$, as we needed to show. \square

So far we have established that $\overline{\text{Map}(S) \cdot \lambda} \subset \mathcal{G}_\lambda$. The main step towards proving the opposite inclusion is to prove that it holds true if λ has no atoms:

PROPOSITION 6.2. *If S is a connected non-exceptional hyperbolic surface then we have $\mathcal{ML}^+(S) \subset \overline{\text{Map}(S) \cdot \lambda}$ whenever $\lambda \in \mathcal{ML}(S)$ has no atoms.*

It is in the proof of Proposition 6.2 that Theorem 1.2 plays a key role.

Proof. Since the set of weighted two-sided geodesics is dense in $\mathcal{ML}^+(S)$ by Theorem 1.2, it is enough to prove that any $c \cdot \gamma_0$ with $c > 0$ and with γ_0 two-sided is a limit of translates of λ . Our first goal is to show that we can find translates of λ which are very close to the trivial measured lamination. To be able to make sense of words like “close” we fix a finite collection \mathcal{A} of curves separating measured laminations—we might assume without loss of generality that $\gamma_0 \in \mathcal{A}$ —and, as in (2.2), let $d_{\mathcal{A}}(\cdot, \cdot)$ be the induced distance on $\mathcal{ML}(S)$.

CLAIM. *For all $\varepsilon > 0$ there exists $\lambda_\varepsilon \in \text{Map}(S) \cdot \lambda$ with $\iota(\lambda_\varepsilon, \gamma_0) > 0$ and $d_{\mathcal{A}}(\lambda_\varepsilon, 0) < \varepsilon$.*

Proof. Note that there is some $\delta(\varepsilon) > 0$ such that if $\mu \in \mathcal{ML}(S)$ has length $\ell(\mu) < \delta(\varepsilon)$, then

$$d_{\mathcal{A}}(\mu, 0) = \max_{\eta \in \mathcal{A}} \iota(\mu, \eta) < \varepsilon$$

As we see, it suffices to guarantee that there is a $\lambda_\varepsilon \in \text{Map}(S) \cdot \lambda$ with $\iota(\lambda_\varepsilon, \gamma_0) > 0$ and $\ell(\lambda_\varepsilon) < \delta(\varepsilon)$. This is what we will do.

To begin with, note that, up to replacing λ by a translate by a pseudo-Anosov mapping class, we can assume that λ and γ_0 together fill the surface. Now, take a sequence $L_n \rightarrow \infty$ and for each n an almost geodesic train track τ_n carrying λ and whose edges have all at least length L_n —such a train track exists because λ has no closed leaves. Since λ and γ_0 fill, we can forget some of the members of our sequence (τ_n) and ensure that $\iota(\gamma_0, \mu) > 0$ for every $\mu \in \mathcal{ML}(\tau_n)$ and every n . Also, since there are only finitely many mapping class group orbits of train tracks by Lemma 3.1, we might pass to a subsequence such that for all n there is $\phi_n \in \text{Map}(S)$ with $\phi_n(\tau_n) = \tau_1$.

Set $\lambda_n = \phi_n(\lambda) \in \mathcal{ML}(\tau_1)$. Since τ_n is almost geodesic we have

$$L_n \cdot \omega_\lambda(e) \leq 2 \cdot \ell(\lambda)$$

for each edge $e \in E(\tau_n)$. This implies in turn that

$$\begin{aligned} \ell(\lambda_n) &\leq \sum_{e \in E(\tau_1)} \omega_{\lambda_n}(e) \cdot \ell(e) = \sum_{e \in E(\tau_1)} \omega_\lambda(\phi_n^{-1}(e)) \cdot \ell(e) \\ &\leq \left(\frac{2 \cdot \ell(\lambda)}{L_n} \right) \cdot \ell(\tau_1). \end{aligned}$$

We just have to choose n with $L_n = \delta(\varepsilon)^{-1} \cdot 2 \cdot \ell(\lambda) \cdot \ell(\tau_1)$. □

Now, let $D = D_{\gamma_0}$ be a Dehn twist along γ_0 . For each $\varepsilon > 0$ we will prove that there exists n_ε so that $D^{n_\varepsilon}(\lambda_\varepsilon)$ converges to $c \cdot \gamma_0$ as ε tends to 0, where $\lambda_\varepsilon \in \text{Map}(S) \cdot \lambda$ is provided by the claim. Since \mathcal{A} is a collection of curves separating measured laminations we get that in order to make sure that $D^{n_\varepsilon}(\lambda_\varepsilon)$ converges to $c \cdot \gamma_0$ we just have to make sure that

$$\lim_{\varepsilon \rightarrow 0} \iota(D^{n_\varepsilon} \lambda_\varepsilon, \eta) = \iota(c \cdot \gamma_0, \eta) \tag{6.1}$$

for all $\eta \in \mathcal{A}$. From Lemma 2.2 we get that for any such $\eta \in \mathcal{A}$

$$\begin{aligned} \iota(D^{n_\varepsilon} \lambda_\varepsilon, \eta) &\leq n_\varepsilon \cdot \iota(\lambda_\varepsilon, \gamma_0) \cdot \iota(\gamma_0, \eta) + \iota(\lambda_\varepsilon, \eta) \\ \iota(D^{n_\varepsilon} \lambda_\varepsilon, \eta) &\geq (n_\varepsilon - 2) \cdot \iota(\lambda_\varepsilon, \gamma_0) \cdot \iota(\gamma_0, \eta) - \iota(\lambda_\varepsilon, \eta) \end{aligned}$$

. Since $\iota(\lambda_\varepsilon, \gamma_0) \neq 0$ tends to 0 when ε goes to 0, we can choose $n_\varepsilon \rightarrow \infty$ with

$$\lim_{\varepsilon \rightarrow 0} n_\varepsilon \cdot \iota(\lambda_\varepsilon, \gamma_0) = 1.$$

Combining the two inequalities above, and taking into account that $\iota(\lambda_\varepsilon, \eta) < \varepsilon$, we get (6.1). This concludes the proof of the proposition. \square

We are now ready to conclude the proof of the theorem:

Proof of Theorem 1.4. Since $\text{Map}(S) \cdot \lambda \subset \mathcal{G}_\lambda$, and since the latter is closed by Lemma 6.1, we have $\overline{\text{Map}(S) \cdot \lambda} \subset \mathcal{G}_\lambda$.

For the reverse inclusion, let $\lambda = \gamma_\lambda + \lambda'$ be the unique decomposition of λ , where $\lambda' \in \mathcal{ML}(S)$ has no atoms. Let R_i be one of the connected components of R_λ , and let λ'_i be a component of λ' supported in R_i . As we pointed out earlier, R_i is non-exceptional. By Proposition 6.2 we have that $\mathcal{ML}^+(R_i) \subset \overline{\text{Map}(R_i) \cdot \lambda'_i}$. It follows thus that $\gamma_\lambda + \mu \in \overline{\text{Map}(S) \cdot \lambda}$ for any $\mu \in \mathcal{ML}^+(R_\lambda)$. In other words we have

$$\gamma_\lambda + \mathcal{ML}^+(R_\lambda) \subset \overline{\text{Map}(S) \cdot \lambda}.$$

Since the right side is mapping class group invariant it follows that

$$\mathcal{G}_\lambda = \text{Map}(S) \cdot (\gamma_\lambda + \mathcal{ML}^+(R_\lambda)) \subset \overline{\text{Map}(S) \cdot \lambda},$$

as we needed to prove. \square

7 Orbit closures in $\mathcal{PML}(S)$

We come now to the proofs of Theorem 1.1 and Theorem 1.3. We start by the latter: **Theorem 1.3** *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. We have $\overline{\text{Map}(S) \cdot \lambda} = \mathcal{PG}_\lambda \cup \mathcal{PML}^+(S)$ for any projective measured lamination $\lambda \in \mathcal{PML}(S)$.*

Theorem 1.3 will follow easily once we know that the set on the right is closed. This is what we prove next:

LEMMA 7.1. *The set $\mathcal{P}\mathcal{G}_\lambda \cup \mathcal{P}\mathcal{M}\mathcal{L}^+(S)$ is closed in $\mathcal{P}\mathcal{M}\mathcal{L}(S)$ for any $\lambda \in \mathcal{P}\mathcal{M}\mathcal{L}(S)$.*

Proof. Using for once square brackets to indicate projective classes, suppose that we have a sequence $([\mu_n])_n$ in $\mathcal{P}\mathcal{G}_\lambda \cup \mathcal{P}\mathcal{M}\mathcal{L}^+(S)$ which converges to some $[\mu] \in \mathcal{P}\mathcal{M}\mathcal{L}(S)$. We claim that $[\mu]$ also belongs to $\mathcal{P}\mathcal{G}_\lambda \cup \mathcal{P}\mathcal{M}\mathcal{L}^+(S)$. Since $\mathcal{P}\mathcal{M}\mathcal{L}^+(S)$ is closed, we might assume without loss of generality that our sequence is actually contained in $\mathcal{P}\mathcal{G}_\lambda$. This means that it is represented by a sequence $(\mu_n) \subset \mathcal{G}_\lambda \subset \mathcal{M}\mathcal{L}(S)$. If the sequence (μ_n) converges in $\mathcal{M}\mathcal{L}(S)$ then we are done because \mathcal{G}_λ is closed by Lemma 6.1. We might thus assume that μ_n only converges projectively to some representative $\mu \in \mathcal{M}\mathcal{L}(S)$ of the class $[\mu]$, meaning that there is a sequence $\varepsilon_n \rightarrow 0$ with

$$\mu = \lim_n \varepsilon_n \cdot \mu_n.$$

We claim that $\mu \in \mathcal{M}\mathcal{L}^+(S)$. Otherwise μ would have a one-sided atom $c \cdot \gamma$. We would then get from Scharlemann’s theorem that, for all large i , the geodesic γ is the support of an atom of $\varepsilon_n \cdot \mu_n$. In fact, the weight c_n in $\varepsilon_n \cdot \mu_n$ of this atom converges to c and hence stays bounded from below. Since there are $\phi_n \in \text{Map}(S)$ and $\nu_n \in \mathcal{M}\mathcal{L}^+(\phi_n(R_\lambda))$ with

$$\varepsilon_n \cdot \mu_n = \varepsilon_n \cdot \phi_n(\gamma_\lambda) + \nu_n$$

we get that our atom $c_n \cdot \gamma$, having relatively large weight, must be an atom of ν_n . This contradicts the assumption that ν_n has no one-sided leaves. We are done. \square

We are now ready to prove the theorem:

Proof of Theorem 1.3. Since $\text{Map}(S) \cdot \lambda$ is contained in $\mathcal{P}\mathcal{G}_\lambda$ we get from Lemma 7.1 the inclusion $\overline{\text{Map}(S) \cdot \lambda} \subset \mathcal{G}_\lambda \cup \mathcal{P}\mathcal{M}\mathcal{L}^+(S)$. To prove the opposite inclusion note that we get from Theorem 1.4 that $\mathcal{P}\mathcal{G}_\lambda \subset \overline{\text{Map}(S) \cdot \lambda}$. Also, the argument used in the proof of Proposition 6.2 implies that the projective class of every two-sided curve belongs to $\overline{\text{Map}(S) \cdot \lambda}$. Since two-sided curves are dense in $\mathcal{P}\mathcal{M}\mathcal{L}^+(S)$ by Theorem 1.2 we get that $\mathcal{P}\mathcal{G}_\lambda \cup \mathcal{P}\mathcal{M}\mathcal{L}^+(S) \subset \overline{\text{Map}(S) \cdot \lambda}$. We are done. \square

Note that Theorem 1.3 implies that $\mathcal{P}\mathcal{M}\mathcal{L}^+(S)$ is contained in the closure of any orbit of the action $\text{Map}(S) \curvearrowright \mathcal{P}\mathcal{M}\mathcal{L}(S)$. Theorem 1.1 follows immediately:

Theorem 1.1 *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. The set $\mathcal{P}\mathcal{M}\mathcal{L}^+(S)$ is the unique non-empty closed subset of $\mathcal{P}\mathcal{M}\mathcal{L}(S)$ which is invariant and minimal under the action of $\text{Map}(S)$.* \square

On the other hand, we note that $\mathcal{P}\mathcal{M}\mathcal{L}(S)$, for S non-orientable, has a dense orbit if and only if S is of genus 1. Recall that the *genus* of the surface is defined to be the maximum number of disjoint simple curves one can cut along without disconnecting the surface.

COROLLARY 7.2. *Suppose S is non-orientable. If S has genus $k = 1$, then $\mathcal{PML}(S)$ has a dense $\text{Map}(S)$ -orbit. If S has genus $k > 1$ then $\mathcal{PML}(S)$ is not the closure of any countable union of $\text{Map}(S)$ -orbits.*

Proof. For $k = 1$, let γ be a one-sided curve and set $\lambda = \gamma + \lambda'$ for some $\lambda' \in \mathcal{ML}(S \setminus \gamma)$. Note that every one-sided curve on S is in the $\text{Map}(S)$ -orbit of γ , that $S \setminus \gamma$ is orientable, and that $R_\lambda = S \setminus \gamma$. It follows that $\mathcal{PG}_\lambda = \mathcal{PML}(S) \setminus \mathcal{PML}^+(S)$ and hence that

$$\overline{\text{Map}(S) \cdot [\lambda]} = \mathcal{PG}^\lambda \cup \mathcal{PML}^+(S) = \mathcal{PML}(S).$$

For the second statement, note that when the genus is at least 2 then there are uncountably many projective classes of orbits of one-sided multicurves: for any 2 disjoint one-sided curves α, β , the classes $[\alpha + t\beta]$ are distinct and not in the same orbit for distinct $t > 0$. On the other hand, any countable union $\bigcup_{n \in \mathbb{N}} \mathcal{PG}^{\lambda_n}$ only contains countably many orbits of one-sided multicurves. \square

8 Orbit closures in Teichmüller space

There are many reasons why measured laminations and projective measured laminations play a key role when studying hyperbolic surfaces, but if one were to have to choose one then one would probably mention Thurston's compactification of Teichmüller space [Thu88]. Indeed, Thurston defined a mapping class group invariant topology on

$$\overline{\mathcal{T}}(S) = \mathcal{T}(S) \cup \mathcal{PML}(S)$$

with respect to which the inclusions $\mathcal{T}(S) \hookrightarrow \overline{\mathcal{T}}(S)$ and $\mathcal{PML}(S) \hookrightarrow \overline{\mathcal{T}}(S)$ are homeomorphisms onto their images, and where a sequence $(X_i) \subset \mathcal{T}(S)$ converges to a $[\lambda] \in \mathcal{PML}(S)$ if for some, and hence any representative $\lambda \in \mathcal{ML}(S)$ of the class $[\lambda]$, there is a sequence (ε_i) of positive real numbers with

$$\lim_{i \rightarrow \infty} \varepsilon_i \cdot \ell_{X_i}(\gamma) = \iota(\lambda, \gamma)$$

for every essential simple closed curve γ . The space $\overline{\mathcal{T}}(S)$ is compact when endowed with this topology—it is in fact homeomorphic to a ball, but we will not need that fact. We refer to [FLP] for a discussion of Thurston's compactification in the orientable case, to [PP16] for one in the non-orientable case, and to [Bon88] for one that works either way.

In this section we study where do orbits of the action $\text{Map}(S) \curvearrowright \mathcal{T}(S)$ accumulate in the Thurston boundary $\partial\mathcal{T}(S) = \mathcal{PML}(S)$.

Theorem 1.5 *Let S be a connected, possibly non-orientable, non-exceptional hyperbolic surface of finite topological type. We have $\overline{\text{Map}(S) \cdot X} \cap \partial\mathcal{T}(S) = \mathcal{PML}^+(S)$ for any point X in Teichmüller space $\mathcal{T}(S)$.*

In the course of the proof of the theorem we will need the following fact:

LEMMA 8.1. *Suppose $\lambda, \lambda' \in \mathcal{ML}(S)$ are such that $\iota(\lambda, \alpha) \leq \iota(\lambda', \alpha)$ for all simple closed curves α . Then $\text{supp}(\lambda) \subset \text{supp}(\lambda')$.*

Proof. First note that we might assume without loss of generality that λ has connected support, meaning that there are a priori three options: either $\iota(\lambda, \lambda') > 0$, or $\text{supp}(\lambda) \cap \text{supp}(\lambda') = \emptyset$, or $\text{supp}(\lambda) \subset \text{supp}(\lambda')$. To rule the first one out, let (γ_i) be a sequence of weighted curves which converges to λ' in $\mathcal{ML}(S)$. By the continuity of the intersection form we have that $\lim_i \iota(\lambda', \gamma_i) = 0$ while $\lim_i \iota(\lambda, \gamma_i) = \iota(\lambda, \lambda') > 0$. This contradicts our assumption.

The discussion of the second case is less clean. Suppose first that the support of λ is not a simple curve. This implies that every open neighborhood U of $\text{supp}(\lambda)$ contains a simple curve γ with $\iota(\lambda, \gamma) > 0$. If $\text{supp}(\lambda) \cap \text{supp}(\lambda') = \emptyset$ then we can choose the neighborhood U disjoint from λ' , meaning that $\iota(\lambda', \gamma) = 0$. This contradicts our assumption. It remains to consider the case when the support of λ is a simple curve—for the sake of concreteness we assume that the weight is one. Since it is the most interesting and hardest case we suppose that the curve λ is contained in a complementary region of $\text{supp}(\lambda')$ homeomorphic to an annulus with cusps, leaving the other cases to the reader. Let U be an open small neighborhood of λ' and \bar{U} its closure. Let also $[x, y] \subset S \setminus U$ be a simple arc which intersects $\text{supp}(\lambda)$ once, and with $x, y \in \bar{U}$. We can now find two disjoint simple curves $\gamma_x, \gamma_y \subset \bar{U}$ with $x \in \gamma_x$ and $y \in \gamma_y$, and with $\iota(\lambda, \gamma_x) < 1$ and $\iota(\lambda, \gamma_y) < 1$. Let γ be the simple curve obtained, up to homotopy, as the juxtaposition of 2 copies of $[x, y]$ and the curves γ_x and γ_y . We have $\iota(\gamma, \lambda) = 2$ while $\iota(\gamma, \lambda') < 2$. This contradicts once again our assumption.

Having ruled out the first and second possibilities we get that $\text{supp}(\lambda) \subset \text{supp}(\lambda')$, as we needed to prove. □

We are now ready to prove the theorem.

Proof of Theorem 1.5. Note that $\mathcal{PML}^+(S) \subset \overline{\text{Map}(S) \cdot X} \cap \partial\mathcal{T}(S)$ by minimality of $\mathcal{PML}^+(S)$, Theorem 1.1.

Suppose now that $[\lambda] \in \mathcal{PML}(S)$ is an accumulation point of $\text{Map}(S) \cdot X$, say $X_n = \phi_n(X)$ with ϕ_n pairwise distinct. Let α and β be two-sided simple closed curves which together fill S . Consider the function

$$I : \mathcal{ML}(S) \rightarrow \mathbb{R}$$

defined by

$$\lambda \mapsto \frac{\ell_X(\lambda)}{\iota(\alpha, \lambda) + \iota(\beta, \lambda)}.$$

Note that I is continuous, non-zero, and descends to a function on the compact space $\mathcal{PML}(S)$ and is hence bounded from above and below by positive constants. In particular, there exists $K > 1$ such that

$$\frac{1}{K} \leq \frac{\ell_X(\gamma)}{\iota(\alpha, \gamma) + \iota(\beta, \gamma)} \leq K$$

for all simple closed curves γ . It follows that

$$\frac{1}{K} \leq \frac{\ell_{\phi_n(X)}(\gamma)}{\iota(\phi_n(\alpha), \gamma) + \iota(\phi_n(\beta), \gamma)} \leq K \tag{8.1}$$

for all n and all simple closed curves γ .

Note now that, up to reversing the roles of α and β and passing to a subsequence, we can assume that $\ell_X(\phi_n(\alpha)) \geq \ell_X(\phi_n(\beta))$ for all n . This means that if we set $\varepsilon_n = \frac{1}{\ell_X(\phi_n(\alpha))}$ then we have that both the sequences $\varepsilon_n \cdot \phi_n(\alpha)$ and $\varepsilon_n \cdot \phi_n(\beta)$ are bounded in $\mathcal{ML}(S)$. This means that, up to possibly passing to a further subsequence, we can assume that the limits

$$\lim_{n \rightarrow \infty} \varepsilon_n \cdot \phi_n(\alpha) = \mu \text{ and } \lim_{n \rightarrow \infty} \varepsilon_n \cdot \phi_n(\beta) = \mu'$$

exist in $\mathcal{ML}(S)$. Note that μ has unit length and hence is not 0. Note also that $\mu, \mu' \in \mathcal{ML}^+(S)$ because α and β are two-sided and because $\mathcal{ML}^+(S)$ is closed by Scharlemann’s theorem.

Note that, since the mapping classes ϕ_n do not repeat, there is some γ_0 with $\ell_{\phi_n(X)}(\gamma_0) \rightarrow \infty$. Equation (8.1) implies then that also

$$\iota(\phi_n(\alpha), \gamma_0) + \iota(\phi_n(\beta), \gamma_0) \rightarrow \infty. \tag{8.2}$$

On the other hand

$$\lim_n \varepsilon_n \cdot (\iota(\phi_n(\alpha), \gamma_0) + \iota(\phi_n(\beta), \gamma_0)) = \iota(\mu, \gamma_0) + \iota(\mu', \gamma_0) < \infty. \tag{8.3}$$

Hence equations (8.2) and (8.3) together imply that $\varepsilon_n \rightarrow 0$.

Once we know that ε_n tends to 0 we have that

$$\iota(\mu, \mu') = \lim_n \iota(\varepsilon_n \cdot \phi_n(\alpha), \varepsilon_n \cdot \phi_n(\beta)) = \lim_n \varepsilon_n^2 \cdot \iota(\alpha, \beta) = 0$$

and hence that $\lambda' = \mu + \mu'$ is a perfectly sound measured lamination with

$$\lim_n \varepsilon_n \cdot (\iota(\phi_n(\alpha), \gamma) + \iota(\phi_n(\beta), \gamma)) = \iota(\lambda, \gamma) \tag{8.4}$$

for every curve γ . Moreover, since the support of λ' is the union of the supports of μ and μ' we have $\lambda' \in \mathcal{ML}^+(S)$.

Recall now that $[\lambda] \in \mathcal{PML}(S)$ is an accumulation point of the sequence $(\phi_n(X))$. From the choice of the sequence (ε_n) and from (8.1) we get that for all γ the sequence $(\varepsilon_n \cdot \ell_{\phi_n(X)}(\gamma))_n$ is bounded. This means that, once again up to passing to a subsequence, we might assume that the projective class $[\lambda] = \lim_n \phi_n(X) \in \overline{\mathcal{T}}(S)$ has a representative $\lambda \in \mathcal{ML}(S)$ with

$$\lim_n \varepsilon_n \cdot \ell_{\phi_n(X)}(\gamma) = \iota(\lambda, \gamma) \tag{8.5}$$

Combining (8.4) and (8.5) we get from (8.1) that

$$\iota(\lambda, \gamma) \leq \iota(K \cdot \lambda', \gamma)$$

for every curve γ . Lemma 8.1 implies thus that

$$\text{supp}(\lambda) \subset \text{supp}(\lambda').$$

Since $\lambda' \in \mathcal{ML}^+(S)$ we get that also $\lambda \in \mathcal{ML}^+(S)$, and hence that $[\lambda] = \lim_n \phi_n(X) \in \mathcal{PML}^+(S)$, as we needed to prove. We are done. \square

Acknowledgments

The last author would like to thank Anna Lenzhen for interesting conversation and pointing him to the paper [LM10]—it was an incredibly useful hint.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Appendix A. Existence of uniform train tracks

In this appendix we prove Proposition 3.3, which states that any train track carrying a lamination without closed leaves can be refined to be arbitrarily long while also uniform. We recall the precise statement:

Proposition 3.3 *Let S be a hyperbolic surface. There is a constant $C > 0$ such that for any geodesic lamination $\lambda \subset S$ without closed leaves, any train track τ_0 carrying λ , and any L large enough, there is a refinement τ of τ_0 which is λ -generic, (C, λ) -uniform and satisfies $\ell^\lambda(\tau) \geq L$.*

The reader is referred to Section 3 for the terminology and notation used in this appendix. In particular we remind the reader that τ is (C, λ) -uniform if it is almost geodesic, carries λ in a filling way, and the λ -lengths of all edges are comparable: $\ell^\lambda(\tau) \leq C \cdot m^\lambda(\tau)$.

Note that if τ_0 is any train track carrying the lamination λ , we can refine it to an almost geodesic (and as long as we want) train track τ'_0 , still carrying λ (using, for example, the Thurston's construction mentioned in the remark in Section 3.2). Since a refinement of a refinement of τ_0 is still a refinement of τ_0 , we can without loss of generality assume that τ_0 was already almost geodesic to begin with.




Figure 13: Splitting a cusp at v to remove the short edge e .

The idea of the proof of the proposition is to modify the train track τ_0 according to a series of procedures which increase the length of a short edge while controlling the total increase of length. Since the general idea is quite simple, but making it precise is tedious, we will first give an intuitive idea of the process. For simplicity assume τ_0 is trivalent and has a unique shortest edge e . We want to refine τ_0 such that this short edge disappears by getting concatenated with other edges, while only increasing the total length of the train track by a prescribed amount. There are two cases. If e is the only incoming edge at vertex v , then there is a cusp at v bounded by the two outgoing edges. We “split” this cusp by unzipping open e until its other vertex: see the top part of Figure 13. In the resulting train track the short edge has been absorbed into the adjacent ones and the total length has increased by $\ell^\lambda(e) = m^\lambda(\tau_0)$. In the other case e is one of two incoming edges at v . We again want to split the cusp at v but this now means unzipping the outgoing edge e' which could be very long. Hence, in order to control how much total length we add, if e' is very long we do not split the cusp all the way to the next vertex but to some predetermined distance (say twice the minimal edge length). Formally we will do this by adding a bivalent vertex v' on e' at that distance and split until v' . See bottom part of Figure 13. Note that in both situations, in the local picture around v , we have not only gotten rid of the short edge e , but the two new edges have lengths greater than $2\ell^\lambda(e)$. One can then repeat this process finitely many times obtaining a refinement where the shortest edge length has doubled, while the length has only increased by a comparatively small amount.

Now we will make all of this precise—the basic issue being that the process we described above works well if the involved train tracks are trivalent and that this property might get lost after even the first step. Recall from Section 3.3 the definition of a *refinement*; in this section, to give a refinement of a train track τ , we will give a smooth graph τ' and a smooth immersion $\tau' \looparrowright \tau$ which can be perturbed into an embedding $\tau' \hookrightarrow S$. To be fair, we are just going to give the graph τ' and a map $\Phi = \Phi^{\tau', \tau} : \tau' \rightarrow \tau$. It will be evident that a smooth structure exists on τ' which makes Φ a smooth immersion—note that if such a structure exists, it is unique. Finally, since we are thinking of τ' in terms of an immersion into τ , we identify the carrying maps of λ into τ and τ' respectively. Accordingly, we use lengths in τ to

measure the length and λ -length of edges of τ' . All of this will not pose a problem because all relevant inequalities here will be strict, meaning that if they are satisfied for length defined in this way, they also hold after perturbing Φ into an embedding. Along the same lines, the condition for being almost geodesic is open and this means that it makes sense to say that τ' is almost geodesic.

A final piece of terminology: we say that τ' is a *simplicial refinement* if the map Φ is a simplicial map, that is, maps vertices to vertices and edges to edges.

Now we define more rigorously what we mean by the process of splitting a cusp alluded to above, and in such a way that it also works for vertices of higher valence. Let e be a (half-)edge which is not the only incoming one at a vertex v . Then there is a (either one or two) complementary region F of $S \setminus \tau$ having a cusp at v and such that e is contained in a side of F ; we call each such cusp an *adjacent cusp at v* (with respect to e). Moreover, if e' is an outgoing half-edge at v and $\{e, e'\}$ is a λ -legal turn we say that e' is a λ -child of e . With this terminology and with the above intuitive idea in mind, we define what we mean by splitting a cusp:

Splitting a cusp. Let F be a complementary region of $S \setminus \tau$ with vertex v_0 a cusp of F and e_0, e_1 the two half-edges of τ incident to v_0 for which v_0 is an adjacent cusp. If e_0 and e_1 have a common λ -child denote it by \hat{e} . *Splitting the cusp v_0* is the process which produces the simplicial refinement τ' (possibly with bivalent vertices) of τ such that

$$|\Phi^{-1}(v)| = 1 \text{ for all } v \neq v_0 \text{ and } |\Phi^{-1}(v_0)| = 2$$

and

$$|\Phi^{-1}(e)| = 1 \text{ for all } e \neq \hat{e} \text{ and } |\Phi^{-1}(\hat{e})| = 2 \text{ (if } \hat{e} \text{ exists).}$$

Note that for every cusp of τ we split, we add the length of a common child to the total length of the resulting train track. That is, if τ' is a refinement resulting from splitting one cusp then we have

$$\ell^\lambda(\tau') = \ell^\lambda(\tau) + \ell^\lambda(\hat{e}).$$

See Figure 14 for some examples of splitting cusps.

Using this tool we define a few processes resulting in refinements. Recall that we say an edge e is λ -loopy if it has a single vertex and (e^+, e^-) is a λ -legal turn, where e^+, e^- are the two half-edges of e . If e is an edge with only one vertex and such that one half-edge is incoming and the other outgoing at that vertex but which is not λ -loopy (i.e. no leaf of λ takes that turn) we call it a *fake loopy* edge. A *non-loopy* edge is one that is neither loopy nor fake loopy (that is, a non-loopy edge is one which is not an embedded circle).

Combing a half-edge. Let e_0 be a half-edge which is part of a non-loopy edge, and let v_0 be the vertex adjacent to e_0 . We *comb* e_0 through the following steps:

- (1) Split any adjacent cusps and set $e'_0 = (\Phi^{\hat{\tau}, \tau})^{-1}(e_0)$ where $\hat{\tau}$ is the resulting simplicial refinement, and let $v'_0 \in (\Phi^{\hat{\tau}, \tau})^{-1}(v_0)$ be the vertex adjacent to e'_0




Figure 14: Examples of the process of splitting a cusp. The top row are local pictures of a train track with a cusp shaded, the corresponding two half-edges in bold, and the half-edges which are not a common child dashed. The bottom row is the result of splitting the cusp.




Figure 15: Examples of combing a half-edge. Left: local picture with the half-edge to be combed in bold and its children represented by the solid lines. Middle: result after splitting the adjacent cusps. Right: End result, after splitting the remaining cusps and removing bivalent vertices.

- (2) Split all cusps at v'_0 .
- (3) Remove any bivalent vertices.

Letting τ' be the resulting refinement of τ , note that $\Phi(e) = \Phi^{\tau',\tau}(e)$ is a concatenation of edges of τ for every edge e of τ' . In particular we have that any edge e of τ' such that $\Phi(e)$ intersects the interior of e_0 satisfies

$$\ell^\lambda(e) = \ell^\lambda(e_0) + \ell^\lambda(e')$$

for some child e' of e_0 . See Figure 15 for an example of combing an edge.

Unmasking a fake loopy edge. Suppose e_0 is a fake loopy edge with vertex v_0 in τ . There are two cases: either there is a λ -legal turn $\{e_1, e_2\}$ which crosses e_0 or there is not. In the second case, we simply replace v_0 with two vertices v'_0, v''_0 and replace e_0 with an edge connecting v'_0 and v''_0 such that the resulting train track still carries λ ; see top row in Figure 16. In the second case, we do the same but connect v'_0 and v''_0 with an additional edge e carrying segments that took the $\{e_1, e_2\}$ turn and then we comb e (and finally we delete any bivalent vertices). See the bottom row of Figure 16 for an example. Note that the length $\ell^\lambda(\tau')$ of resulting refinement τ' either is the same as that of τ or

$$\ell^\lambda(\tau') = \ell^\lambda(\tau) + \ell^\lambda(e_1) + \ell^\lambda(e_2)$$

where e_1, e_2 are two edges adjacent to v_0 .

Finally we consider the actual loopy edges. If e is a λ -loopy edge and e^+, e^- are its two half-edges, for every leaf l of λ for which $\Phi^{\lambda,\tau}(l)$ traverses e , $\Phi^{\lambda,\tau}(l)$ must




Figure 16: The two possibilities of unmasking a fake loopy edge.




Figure 17: The first step in unlooping a loopy edge.

take the turn $\{e^-, e^+\}$ a number of times. We define the λ -winding number of e to be $w = k - 1$ where k is the largest integer such that there is a leaf of λ which takes the turn $\{e^-, e^+\}$ k times. Note that if e has λ -winding number w , then every leaf l of λ such that $\Phi^{\lambda, \tau}(l)$ transverses e takes the turn $\{e^-, e^+\}$ either w or $w + 1$ times.

If the winding number w is not too large, it will be helpful below to replace a loopy edge e_0 with a concatenation of w edges by *unlooping* it according to the following process:

Unlooping a loopy edge. Suppose e_0 is a loopy edge with vertex v_0 . We will construct a refinement τ' of τ such that $\Phi^{-1}(e_0)$ consists of (one or two) non-loopy edges. Choose an incoming half-edge e_0^+ of e_0 , and denote the other half-edge by e_0^- . There is an edge e incoming at v_0 and adjacent to e_0^+ such that $\{e, e_0^-\}$ is a λ -legal turn. Split the cusp v_0 adjacent to e_0^+ and e . This results in a simplicial refinement τ' of τ with

$$\ell^{\lambda, \tau'}(e_0) = \ell^{\lambda, \tau}(e_0) - \ell(e_0).$$

Note that at this point we have not added any λ -length to the train track. Now, delete any bivalent vertices and continue inductively until we arrive at a refinement where the edge corresponding to e_0 is a fake loopy edge. We complete the process by unmasking it and deleting any bivalent vertices. The length of the refined track is related to the length of τ just like in the case of unmasking fake loopy edges.

Intuitively, the reader should picture the loopy edge e having been unwrapped into a non-loopy edge of length approximately $\ell^\tau(e)$ and generically also a second non-loopy edge obtained through the unmasking in the last step. See Figure 17.

We will apply the above processes to define our main procedure. As outlined in the beginning of this appendix, the idea is to modify a given train track trading any

short edges for longer ones while controlling the total length of the resulting train track.

Main Procedure. Let τ_0 be any (almost geodesic) train track carrying λ and recall that $\ell^\lambda(e) = \ell^{\lambda, \tau}(e)$, $\ell^\lambda(\tau)$ and $m^\lambda(\tau)$ were defined in (3.4) and (3.5).

We start by adding bivalent vertices to any non-loopy edge e with $\ell(e) > 4m^\lambda(\tau_0)$ such that it gets subdivided into a number of edges each having length $2m^\lambda(\tau_0)$ except possibly one which has length at least $2m^\lambda(\tau_0)$ but less than $4m^\lambda(\tau_0)$. In the resulting train track unmask all fake loopy edges, and if this results in an edge of length greater than $4m^\lambda(\tau_0)$, then add bivalent vertices to subdivide it as above. Finally, unloop any loopy edge e with $\ell^\lambda(e) \leq 4m^\lambda(\tau_0)$. Denote the resulting refinement by $\hat{\tau}_0$. Note that $m^\lambda(\hat{\tau}_0) = m^\lambda(\tau_0)$, that $\hat{\tau}_0$ has the same number of edges of length less than $2m^\lambda(\tau_0)$ as τ_0 does, and that for each (fake) loopy edge we have (unmasked or) unlooped we have increased the total length by at most $8m^\lambda(\tau_0)$ and hence $\ell^\lambda(\hat{\tau}_0) \leq \ell^\lambda(\tau_0) + 9|\chi(S)| \cdot 8m^\lambda(\tau_0)$ where the factor $9|\chi(S)|$ is a (very non-optimal) bound on how many loopy edges τ_0 can have.

Now let e_0 be an edge realizing $m^\lambda(\tau_0)$ in $\hat{\tau}_0$ and note that it is non-loopy and any adjacent edge has length at most $4m^\lambda(\tau_0)$ and at least $m^\lambda(\tau_0)$. Let e_0^+ be a half-edge of e_0 and v_0 its vertex. Combing e_0^+ (and then removing all remaining bivalent vertices) we obtain a refinement τ_0^1 of τ_0 carrying λ . Note that the number of edges e in τ_0^1 with $\ell^\lambda(e) = m^\lambda(\tau_0)$ is strictly less than the number of edges e in τ_0 with $\ell^\lambda(e) = m^\lambda(\tau_0)$, that the number of edges e in τ_0^1 with $\ell^\lambda(e) < 2m^\lambda(\tau_0)$ is strictly less than the number of edges e in τ_0 with $\ell^\lambda(e) < 2m^\lambda(\tau_0)$, and that $\ell^\lambda(\tau_0^1) \geq \ell^\lambda(\tau_0)$. Intuitively, the reader can think of any edge in τ_0^1 as corresponding to a concatenation of edges in $\hat{\tau}_0$ and any edge whose image contains e_0 is the concatenation of at least two and hence has length at least $2m^\lambda(\tau_0)$. Now, since combing an edge means splitting cusps and there are at most $6|\chi(S)|$ cusps in $S \setminus \hat{\tau}_0$ we have $\ell^\lambda(\tau_0^1) \leq \ell^\lambda(\hat{\tau}_0) + 6|\chi(S)| \cdot 4m^\lambda(\tau_0) \leq \ell^\lambda(\tau_0) + 96|\chi(S)| \cdot m^\lambda(\tau_0)$.

Now, if $m^\lambda(\tau_0^1) < 2m^\lambda(\tau_0)$ we repeat the above process, replacing τ_0 with τ_0^1 . We continue inductively as long as $m^\lambda(\tau_0^n) < 2m^\lambda(\tau_0)$ where τ_0^n denotes the refinement resulting from the n^{th} step. Note that in each step the number of edges of length less than $2m^\lambda(\tau_0)$ decreases and hence since there are finitely many edges, the procedure ends in finitely many steps. That is, there is a $k \leq 9|\chi(S)|$ such that $m^\lambda(\tau_0^k) \geq 2m^\lambda(\tau_0)$; we set $\tau_0^k := \tau_1$. Note that we have $\ell^\lambda(\tau_0) < \ell^\lambda(\tau_1) < \ell^\lambda(\tau_0) + 1000|\chi(S)| \cdot m^\lambda(\tau_0)$.

We are finally ready to prove the proposition:

Proof of Proposition 3.3. Let λ be a geodesic lamination without closed leaves and τ_0 any train track carrying it. Let $L \geq \ell^\lambda(\tau_0)$. As already mentioned after the statement of the proposition, up to replacing τ_0 by a refinement, we can assume that τ_0 is an almost geodesic train track and $\ell^\lambda(\tau_0) \geq L$.

Apply the Main Procedure to τ_0 resulting in the refinement τ_1 where $m^\lambda(\tau_1) \geq 2m^\lambda(\tau_0)$ and with $\ell^\lambda(\tau_1) < \ell^\lambda(\tau_0) + 1000|\chi(S)| \cdot m^\lambda(\tau_0)$. Continuing inductively, we

obtain the refinement t_k at the k^{th} step satisfying

$$m^\lambda(\tau_k) \geq 2^k m^\lambda(\tau_0)$$

and

$$\begin{aligned} L < \ell^\lambda(\tau_k) &< \ell^\lambda(\tau_0) + 1000|\chi(S)| \cdot \sum_{i=0}^{k-1} m^\lambda(\tau_i) \\ &< \ell^\lambda(\tau_0) + 1000|\chi(S)| \cdot m^\lambda(\tau_k) \cdot \sum_{i=1}^k 2^{-i} \\ &< \ell^\lambda(\tau_0) + 1000|\chi(S)| \cdot m^\lambda(\tau_k). \end{aligned}$$

It follows that for any $C > 1000|\chi(S)|$ there exists a k such that τ_k is C -uniform.

It remains to show that there is a *generic* such train track. Note that by choosing k large enough above, we can assume τ_k is uniform and all its edges having enormous λ -lengths since $m^\lambda(\tau_k)$ grows exponentially with k . The idea is then to add more vertices along these very long edges and spread out the edges so that every vertex is trivalent at the non-loopy edges, and to partly unwrap any loopy edge so that its vertex is 4-valent. We sketch this process and leave the precise details to the reader; in particular, for simplicity we assume there are no loopy edges. Note that there are at most $E = 18 \cdot |\chi(S)|$ half-edges in any train track and hence any vertex can be at most E -valent. Choose k large enough so that every edge in the C -uniform train track τ_k has length at least $2^E \cdot L$. Let v be a vertex which has maximal valence in τ_k , which we assume is greater than 3. Add a bivalent vertex at the midpoint of every edge incident at v , then split a (any) cusp at v , and finally delete any remaining bivalent vertex. The maximal valency of a vertex in the resulting refinement is either strictly less than in τ_k or it remains the same but has one less vertex of this valency. Moreover, every edge has length at least $2^{E-1} \cdot L$ and the total length has increased by at most $\ell^\lambda(\tau_k)/2$. We repeat the process until (in finitely many steps $n \leq E$) every vertex is trivalent. The resulting refinement τ satisfy $m^\lambda(\tau) \geq m^\lambda(\tau_k)/2^n$ and $\ell^\lambda(\tau) \leq \ell^\lambda(\tau_k) + \ell^\lambda(\tau_k) \cdot \sum_{i=1}^n 2^{-i} < 2\ell^\lambda(\tau_k)$. It follows that τ is generic and C' -uniform for $C' = 2^{E+1} \cdot C$. \square

References

- [Bes] M. BESTVINA. Questions in geometric group theory. Available at <https://www.math.utah.edu/~bestvina/eprints/questions-updated.pdf>.
- [Bon86] F. BONAHOON. Bouts des variétés hyperboliques de dimension 3. *Ann. of Math.* (2), (1)124 (1986), 71–158
- [Bon88] F. BONAHOON. The geometry of Teichmüller space via geodesic currents. *Invent. Math.*, (1)92 (1988), 139–162
- [CB88] A.J. CASSON and S.A. BLEILER. *Automorphisms of Surfaces After Nielsen and Thurston, Volume 9 of London Mathematical Society Student Texts*. Cambridge University Press, (1988).

- [DN90] C. DANTHONY and A. NOGUEIRA. Measured foliations on nonorientable surfaces. *Ann. Sci. École Norm. Sup. (4)*, (3)23 (1990), 469–494
- [ES] V. ERLANDSSON and J. SOUTO. *Mirzakhani’s curve counting and geodesic currents*. Progress in Mathematics, volume 345, Birkhäuser, (2022).
- [FM12] B. FARB and D. MARGALIT. *A Primer on Mapping Class Groups, Volume 49 of Princeton*. Princeton University Press, (2012).
- [FLP] A. FATHI, F. LAUDENBACH and V. POÉNARU. *Travaux de Thurston sur les Surfaces, Volume 66 of Astérisque*.
- [Gen05] M. GENDULPHE. Paysage systolique des surfaces hyperboliques compactes de caractéristique -1 (2005). [arXiv:math/0508036](https://arxiv.org/abs/math/0508036).
- [Gen17] M. GENDULPHE. What’s wrong with the growth of simple closed geodesics on nonorientable hyperbolic surfaces (2017). [arXiv:1706.08798](https://arxiv.org/abs/1706.08798).
- [GK14] M. GENDULPHE and Y. KOMORI. Polyhedral realization of a Thurston compactification. *Ann. Fac. Sci. Toulouse Math. (6)*, (1)23 (2014), 95–114
- [Ham03] U. Hamenstädt. Parametrizations of Teichmüller space and its Thurston boundary. In: *Geometric Analysis and Nonlinear Partial Differential Equations*. Springer (2003), pp. 81–88.
- [Ham09] U. HAMENSTÄDT. Invariant Radon measures on measured lamination space. *Invent. Math.*, (2)176 (2009), 223–273
- [Hat88] A.E. HATCHER. Measured lamination spaces for surfaces, from the topological viewpoint. *Topology Appl.*, (1)30 (1988), 63–88
- [Iva] N.V. IVANOV. *Subgroups of Teichmüller Modular Groups, Volume 115 of Translations of Mathematical Monographs*.
- [Kat73] A.B. KATOK. Invariant measures of flows on orientable surfaces. *Dokl. Akad. Nauk SSSR*, 211 (1973), 775–778
- [Kha21] S. KHAN. The limit set of non-orientable mapping class groups (2021). [arXiv:2110.00037](https://arxiv.org/abs/2110.00037).
- [LM10] A. LENZHEN and H. MASUR. Criteria for the divergence of pairs of Teichmüller geodesics. *Geom. Dedicata*, 144 (2010), 191–210
- [Lev83] G. LEVITT. Feuilletages des surfaces. *Thèse d’état soutenue à Paris 7*, (1983).
- [LM08] E. LINDENSTRAUSS and M. MIRZAKHANI. Ergodic theory of the space of measured laminations. *Int. Math. Res. Not. IMRN*, (4):Art. ID rnm126, 49, (2008).
- [Man10] J. MANGAHAS. Uniform exponential growth of subgroups of the mapping class group. *Geom. Funct. Anal.*, (5)19 (2010), 1468–1480
- [Ota96] J.-P. OTAL. *Le théorème d’hyperbolisation pour les variétés fibrées de dimension 3*. Number 235 in Astérisque. Société mathématique de France, (1996).
- [PP16] A. PAPADOPOULOS and R.C. PENNER. Hyperbolic metrics, measured foliations and pants decompositions for non-orientable surfaces. *Asian J. Math.*, (1)20 (2016), 157–182
- [Par14] L. PARIS. Mapping class groups of non-orientable surfaces for beginners. *Winter Braids Lect. Notes*, 1(Winter Braids IV (Dijon, 2014)):Exp. No. 3 (2014), 17
- [PH92] R.C. PENNER and J.L. HARER. *Combinatorics of Train Tracks, Volume 125 of Annals of Mathematics Studies*. Princeton University Press, (1992)
- [Sch82] M. SCHARLEMANN. The complex of curves on nonorientable surfaces. *J. London Math. Soc. (2)*, (1)25 (1982), 171–184
- [Sti93] J. STILLWELL. *Classical Topology and Combinatorial Group Theory, Volume 72 of Graduate Texts in Mathematics*, 2nd edn. Springer, (1993).

- [Thu80] W.P. THURSTON. The geometry and topology of three-manifolds. *Unpublished notes*, (1980). Available at <http://library.msri.org/books/gt3m/PDF/Thurston-gt3m.pdf>.
- [Thu88] W.P. THURSTON. On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. Amer. Math. Soc. (N.S.)*, (2)19 (1988), 417–431

V. Erlandsson
School of Mathematics, University of Bristol, Bristol BS8 1UG, UK.

and

Department of Mathematics and Statistics, UiT The Arctic University of Norway, Tromsø, Norway.

`v.erlandsson@bristol.ac.uk`

M. Gendulpe
Lycée La Folie Saint-James, Paris, France.

`matthieu@gendulpe.com`

I. Pasquinelli
School of Mathematics, University of Bristol, Bristol BS8 1UG, UK.

`irene.pasquinelli@bristol.ac.uk`

J. Souto
Université de Rennes, CNRS, IRMAR - UMR 6625, 35000 Rennes, France.

`jsoutoc@gmail.com`

Received: February 9, 2022

Revised: December 22, 2022

Accepted: January 23, 2023