PATH-CONNECTIVITY OF THE SET OF UNIQUELY ERGODIC AND COBOUNDED FOLIATIONS

JON CHAIKA AND SEBASTIAN HENSEL

Abstract. We show that if S is a closed surface of genus $g \geq 5$ or a surface of genus $g \geq 2$ with at least $p \geq 3$ marked points, then the set of uniquely ergodic foliations and the set of cobounded foliations is path-connected and locally path-connected.

1. Introduction

Projective measured foliations play a prominent role in Teichmüller theory, dynamics and the study of mapping class groups. In addition to the structure of individual foliations, the set $\mathcal{PMF}(S)$ of all foliations on a given finite type surface S has particular importance. $\mathcal{PMF}(S)$ carries a natural (weak-$*$) topology and is homeomorphic to a sphere of dimension $6g + 2p - 7$ if S has genus g and p punctures. One reason for its importance stems from the fact that $\mathcal{PMF}(S)$ can be identified with both the sphere of directions, and the boundary of infinity of Teichmüller space. One can also use $\mathcal{PMF}(S)$ to describe the Gromov boundary of the curve graph.

In this article we study global topological properties of two dynamically motivated subsets of $\mathcal{PMF}(S)$. The first is the set $\mathcal{UE}(S)$ of *uniquely ergodic* foliations, where a foliation F is called uniquely ergodic if it admits a unique transverse measure up to scale. The second is the set $\mathcal{COB}(S)$ of *cobounded foliations*, where F is called cobounded if a Teichmüller geodesic ray with vertical foliation F projects into a compact set of the moduli space of Riemann surfaces.

These sets have been intensely studied from a dynamical point of view, owing to their importance in Teichmüller theory. As a starting point, by a theorem of Masur [Mas], any cobounded foliation is uniquely ergodic, and we therefore have

$$\mathcal{COB}(S) \subset \mathcal{UE}(S) \subset \mathcal{PMF}(S)$$

Both $\mathcal{COB}(S)$ and $\mathcal{UE}(S)$ are dense in $\mathcal{PMF}(S)$ (but the same is also true for their complements). Masur and Veech [Mas, Vee] show that $\mathcal{UE}(S)$ has full measure in $\mathcal{PMF}(S)$. In contrast, the set $\mathcal{COB}(S)$ has measure zero.

It is known that there are embedded circles in $\mathcal{COB}(S)$ [LS2], but Masur and Smillie [MS] have shown that the complement $\mathcal{PMF}(S) \setminus \mathcal{UE}(S)$ has Hausdorff dimension stictly bigger than $\dim \mathcal{PMF}(S) - 1$, and hence one cannot expect to locally deform paths in order to avoid $\mathcal{PMF}(S) \setminus \mathcal{UE}(S)$.

Our main result shows that paths are nevertheless abundant in $\mathcal{COB}(S)$ and $\mathcal{UE}(S)$:

Theorem 1.1. *Let S be a closed surface of genus at least 5, or a surface of genus at least 2 with at least 3 punctures. Then the subsets $\mathcal{UE}(S), \mathcal{COB}(S)$ are path-connected, and locally path-connected.*

In fact, the proof shows something slightly stronger: any two points in $\mathcal{UE}(S)$ can be joined by a continuous path which is contained in $\mathcal{COB}(S)$ except possibly at its endpoints.

Our result can also be used to show that through any finite number of points in $\mathcal{UE}$ or $\mathcal{COB}$ there is an embedded circle in $\mathcal{UE}$ or $\mathcal{COB}$[1].

Proof Strategy and Structure of this Article. Our strategy in proving Theorem 1.1 has two main ingredients. On the one hand, we will develop in Section 3 a robust mechanism to construct paths of cobounded foliations in the sphere of projective measured foliations of a punctured surface. This construction was heavily inspired by the work in [LS1], and our main new contribution here is to use bad approximability of points under straight line flows on tori to certify coboundedness and to improve the paths built in [LS1] to consist of cobounded foliations. This will be done in Section 3.

The second ingredient consists in assembling these paths with a limiting procedure into paths that can reach any uniquely ergodic foliation. This is somewhat similar in spirit to our previous work [CH] on interval exchanges. However, the methods of this paper are softer and more flexible than the explicit construction in that work. Here, we use train track splitting sequences to define mapping class group sequences that exhibit contracting behaviour on $\mathcal{PMF}$. The main technical work to make this work happens in Section 2, and uses the hyperbolic geometry of curve graphs to show that these sequences act on $\mathcal{PMF}$ in a contracting way.

Section 4 then combines these two parts and shows the path-connectivity statement in Theorem 1.1 for punctured surfaces. This is also the prerequisite for Section 5, in which the path-connectivity statement of Theorem 1.1 is proved for closed surfaces.

Finally, in Section 6 we show how to leverage the constructions of paths to show local path-connectivity.

Further Questions. Finally, we want to highlight a few questions for further research suggested by Theorem 1.1 and its proof.

Question 1. Are $\mathcal{UE}(S)$ and $\mathcal{COB}(S)$ simply connected, if the genus of S is sufficiently large?

Question 2 (Gabai [Gab])**.** Is the set $\mathcal{AF}(S) \supset \mathcal{UE}(S)$ of arational foliations path-connected?

This question came up in Gabai's analysis of connectivity properties of the Gromov boundary of the curve graph (which is the quotient of $\mathcal{AF}(S)$ by the map which "forgets" the measure on the foliation). Gabai proves that this boundary is path-connected, but his methods does not apply to $\mathcal{AF}(S)$ directly. Leininger and Schleimer [LS1] proved that the set $\mathcal{AF}(S)$ of arational foliations is connected,

---

[1]To ensure that the circle is embedded, one has to use the proof of Theorem 1.1 rather than just the statement. We omit details, as the claim is not central to our discussion.

and contains a dense path-connected subset, but it is not clear that these paths can be extended to the closure. We suspect that our curve graph methods can recover Gabai's result that ending lamination space is path connected in the case of a surface of genus at least 5. Partly because such genus bounds would not be optimal, we have not investigated this thoroughly.

Our methods are at the moment also unable to deal with the case of arational foliations, mainly because the contraction properties in Section 2. This is due to the fact that in order to certify contraction we use the curve graph boundary, which is unable to distinguish different measures supported on a topological foliation.

Next, one could consider more restrictive subsets of $\mathcal{COB}(S)$. Namely, suppose we fix a constant $\epsilon > 0$. Call a foliation F ϵ-cobounded if a Teichmüller ray with vertical foliation F eventually stays in the ϵ-thick part of Teichmüller space.

Question 3. Is the set $\mathcal{COB}_\epsilon(S)$ of ϵ-cobounded foliations path-connected for any choice of ϵ?

Our methods do not yield this, since the paths (both in Section 4 and 5) need to degenerate very close to simple closed curves in order to apply the methods from Section 2. However, the basic paths from Section 3 can be guaranteed to have uniform thickness.

Finally, one motivating reason for studying paths of cobounded paths in the sphere of projective measured foliation stems from one of the central open questions in the study of mapping class groups and Teichmüller theory. Namely, Farb–Mosher [FM2] define *convex cocompact subgroups* in analogy to such Kleinian groups. At this time, all known examples of such groups are virtually free, and it is not clear if any other examples can exist. One touchstone question is therefore: is there a convex cocompact subgroup of the mapping class group, which is isomorphic to the fundamental group of a higher genus surface. Such a group G would give rise to a G–invariant circle in $\mathcal{COB}(S)$.

Question 4. Are there embedded circles in $\mathcal{COB}(S)$ which are invariant under groups that are not free?

Most likely, this question requires significant new tools. A weaker version of this question arises if we relax the invariance condition, e.g.

Question 5. Is there a finite subset $F \subset \mathrm{Mod}(S)$, and $P \subset F^2$ so that for x in Teichmüller space we have that the limit in $\mathcal{PMF}$ of

$$\{s_n ... s_1 x : (s_1 ... s_n) \in F^n \text{ and } (s_i, s_{i+1}) \in P \text{ for all } i < n\}_{n \in \mathbb{N}}$$

is a circle in $\mathcal{COB}(S)$. That is, is there a "convex cocompact shift of finite type" which has a circle limit set of cobounded foliations in $\mathcal{PMF}$?

One could also ask a similar question for semigroups.

2. Contractions on $\mathcal{PMF}$

We denote by $\mathcal{PMF}$ the sphere of projective measured foliations. Recall that a foliation is called *minimal*, if every regular leaf is dense. As mentioned in the introduction, we call a foliation F *uniquely ergodic*, if F admits a unique transverse measure up to scale. We call a foliation F *cobounded* if a Teichmüller ray with vertical foliation F is contained in some thick part of Teichmüller space. By Masur's criterion [Mas], cobounded foliations are uniquely ergodic, and it is well known that uniquely ergodic foliations are minimal.

Throughout this article, we will use the notion of measured foliations, although most literature on train tracks uses measured geodesic laminations instead. We refer the reader to [Lev] for an excellent dictionary between foliations and laminations on surfaces. Most of the time this will not be cause for confusion. We only want to emphasise that a minimal foliation in our sense corresponds to a minimal *and filling* lamination. In particular, there are no simple closed curves which have intersection 0 with a minimal foliation.

2.1. **From splitting sequences to mapping classes.** This section sets out the framework connecting mapping class group elements and train track splitting sequences. We refer the reader to [PH] for a detailed treatment of the basic theory of train tracks, and [MM1] for some other concepts we use.

If τ is a train track and F is a foliation, we write $F \prec \tau$ if F *is carried by* τ (compare [PH, Section 1.6], noting that in [PH] the notion of measured geodesic laminations is used in place of foliations. We). We denote by $P(\tau) \subset \mathcal{MF} \setminus \{0\}$ the set of measured foliations which are carried by τ. When it does not cause confusion, we will often identify $P(\tau)$ with the subset of the sphere $\mathcal{PMF}$ of projective measured foliations it defines. The set $P(\tau)$ naturally has the structure of a closed polyhedron, whose faces correspond to the polyhedra $P(\eta)$ of subtracks η of τ.

A train track is called *recurrent*, if for every branch there is a train path which traverses it. It is called *birecurrent* if in addition there is a multicurve hitting the train track efficiently (i.e. without generating bigons) which intersects every branch (compare [PH, Section 1.3] for details on these definitions). From now on, we will usually assume without mention that all train tracks we use are birecurrent. We say that a train track is *large* if every complementary component is simply connected, and *maximal*, if every complementary component is a triangle (which implies largeness).

For maximal, birecurrent train tracks τ, the interior of $P(\tau)$ defines an open set in $\mathcal{PMF}$ [PH, Lemma 3.1.2]. For other train tracks this need not be the case. By the *interior* $\operatorname{int} P(\tau)$ *of* $P(\tau)$ we will always mean the subset of $P(\tau)$ formed by all those measures which assign a positive weight to each branch. We stress again that, in general, this is different from the topological interior of $P(\tau)$ as a subset of $\mathcal{PMF}$ or $\mathcal{MF} \setminus \{0\}$.

Given a train track τ, a branch b is *large*, if every trainpath through either of its endpoints runs through b. Recall that we can perform a *left, right or central split* at a large branch to obtain a new train track τ'. Compare [PH, §2.1] for details on this construction. We recall that a left or right split does not affect the number and type of complementary components of the train track, while a central split can join two complementary components into one.

Let τ be a fixed maximal, birecurrent train track. As noted above, the polyhedron $P(\tau)$ defines an open set in $\mathcal{PMF}$. We let $\mathcal{T}(\tau)$ be the set of all large birecurrent train tracks which can be obtained from τ by any number of splits (left, right, or central). The set $\mathcal{T}(\tau)$ can be stratified in the following way. Put

$\mathcal{T}_0(\tau) = \{\tau\}$, and inductively define $\mathcal{T}_{n+1}(\tau)$ to be the set of large train tracks obtained from each $\sigma \in \mathcal{T}_n(\tau)$ by splitting each large branch once (in one of the up to three possible ways). Note that a central split need not yield a large train track, so not all three possibilities are always allowed.

A large branch b of a large birecurrent train track σ defines a hyperplane H in $P(\sigma)$ cutting $P(\sigma)$ into subpolyhedra P_l, P_r, which are exactly the polyhedra of the left and right splits of σ. The polyhedron of the central split of σ at b is the hyperplane H [PH, Proposition 2.2.2]. Hence, the interiors of the polyhedra $P(\sigma), \sigma \in \mathcal{T}_n(\tau)$ define a decomposition of $P(\tau)$ into disjoint subpolyhedra.

Lemma 2.1. *$U_n(\tau, F)$ is an open neighborhood of F in $\mathcal{PMF}$ for every n.*

Proof. We prove the lemma by induction. For $n = 0$ this is simply openness of $P(\tau)$ ([PH, Lemma 3.1.2]). Suppose now that $U_n(\tau, F)$ is an open neighborhood of F. From the description of the effect of splits on polyhedra given above we conclude that $U_{n+1}(\tau, F)$ is obtained from $U_n(\tau, F)$ by cutting at hyperplanes (corresponding to central splits) and retaining those polyhedra which contain F. If none of these hyperplanes contain F, it is clear that $U_{n+1}(\tau, F)$ is still an open neighbourhood of F. However, suppose that one of them does contain F. This corresponds to the situation in which a track $\eta \in \mathcal{T}_n$ has a large branch so that all three of the left,right and central splits of η along that branch carry F – and therefore all three of these splits will contribute to $U_{n+1}(\tau, F)$, guaranteeing that the latter is still an open neighbourhood of F. $\square$

Now, let $F \in \operatorname{int} P(\tau)$ be given, and let

$$\mathcal{T}(\tau, F) = \{\sigma \in \mathcal{T}(\tau), F \prec \sigma\}$$

be the subset of all those train tracks in $\mathcal{T}(\tau)$ which carry F. We let $\mathcal{T}_n(\tau, F)$ be the set of all those $\sigma \in \mathcal{T}_n(\tau)$ which carry F. For the next lemma, we use the notion of *diagonal extension.* If τ is a train track, then we say that η is a diagonal extension of τ if η is obtained by adding branches inside simply connected complementary components. See [MM1, Section 4.1] for details.

Lemma 2.2. *The sets $\mathcal{T}_n(\tau, F)$ only contain diagonal extensions of the (large) train track $\eta_n \in \mathcal{T}_n(\tau, F)$ with the fewest complementary components.*

Proof. Consider the sequence η_k of train tracks obtained by splitting τ in the direction of F and always choosing a central split when possible. These have the property that they always carry F, and additionally, the weight defined by F is positive on every branch of η_k for all k (*F fills η_k*). Note that for any $\sigma \in \mathcal{T}_n(\tau)$ there exists (at least one) η_k (depending on σ) so that η_k is a subtrack of σ (this follows inductively, since if a foliation is carried by, and fills, a subtrack η of σ and σ splits to σ', then either η or a split of η is a subtrack of σ'). Since F is minimal, and therefore there is no simple closed curve that doesn't intersect F, it can only be carried by large train tracks. Therefore, σ is a diagonal extension of η_k. The lemma now follows, since any set $\mathcal{T}_n(\tau, F)$ contains at most one η_k, and the number of complementary components in a split decreases only during central splits. $\square$

We put

$$U_n(\tau, F) = \bigcup_{\sigma \in \mathcal{T}_n(\tau, F)} \operatorname{int} P(\sigma)$$

A *splitting sequence τ_i in the direction of F* is a sequence τ_i of train tracks with $\tau_0 = \tau$ and so that each τ_i carries F, and τ_{i+1} is obtained from τ_i by splitting exactly one large of τ_i branch once. A *full splitting sequence* instead requires splitting each large branch of τ_i once when passing from τ_i to τ_{i+1}. Hence, if τ_i is a full splitting sequence in the direction of F starting in τ, then $\tau_i \in \mathcal{T}_i(\tau)$ for all i.

If F is a foliation in the minimal stratum (i.e. each singularity is 3–pronged, and there are no saddle connections[2]), then each split in a splitting sequence τ_i is a left or a right split, and furthermore the type is uniquely determined by F. If F has saddle connections or k–prong singularities for $k > 3$, then it is possible that for some n, F is carried by the left, right and central split of τ_n. This is furthermore the last time F is carried in the interior of a maximal train track τ_n along the splitting sequence.

Splitting sequences in the direction of minimal foliations have good contracting properties. In the following theorem, and below, we denote by $\Delta(F)$ the (closed) simplex of projective measured foliations which are topologically equivalent to F.

Theorem 2.3 (compare e.g. [Mos, Theorem 5.1.1])**.** *Suppose that F is a minimal foliation and that τ_i is any splitting sequence in the direction of F. Then*

$$\bigcap_{i=1}^{\infty} P(\tau_i) = \Delta(F).$$

As an immediate corollary, we have

Corollary 2.4. *Let $F \in \operatorname{int} P(\tau)$ be minimal. Then*

$$\bigcap_{n \geq 0} U_n(\tau, F) = \Delta(F).$$

We now describe how to connect splitting sequences to sequences in the mapping class group. The first step is the following lemma.

Lemma 2.5. *There is a finite number of sets*

$$\mathcal{T}^{(1)}, \ldots, \mathcal{T}^{(M)} \tag{1}$$

so that for each maximal train track τ, each minimal F, and each n there is a number $k_{\tau,F,n}$ and a mapping class $f_{\tau,F,n}$ with

$$\mathcal{T}_n(\tau, F) = f_{\tau,F,n}\left(\mathcal{T}^{(k_{\tau,F,n})}\right).$$

The number $k_{\tau,F,n}$ is unique. The mapping class $f_{\tau,F,n}$ is unique up to a finite indeterminacy.

We call the set of $\mathcal{T}^{(i)}$ *standard neighborhood models* and we call the number $k_{\tau,F,n}$ the *type* of $\mathcal{T}_n(\tau, F)$.

Proof of Lemma 2.5. By Lemma 2.2, $\mathcal{T}_n(\tau, F)$ consists of train tracks which are diagonal extensions of some large train track η_n. Since the mapping class group $\mathrm{Mod}(S_g)$ acts on the set of (isotopy classes of) train tracks on S_g with finitely many orbits, there are finitely many choices for such a train track η_n up to the

[2] This means that the corresponding lamination has only triangles as its complementary components.

mapping class group action. Since the number of complementary components of η_n can be bounded from the Euler characteristic of S alone, there are a finite number of diagonal extensions of η_n. This implies that the mapping class group also acts on the sets $\mathcal{T}_n(\tau, F)$ (over all τ, F, n) with finitely many orbits. We can therefore choose the sets $\mathcal{T}^{(i)}$ to be orbit representatives of this action. This shows both the desired existence of $k_{\tau,F,n}$ and $f_{\tau,F,n}$, as well as the uniqueness of $k_{\tau,F,n}$. The (coarse) uniqueness of the $f_{\tau,F,n}$ follows since the set of mapping classes which fix a given train track is finite (compare e.g. [Ham2, Lemma 4.2]), and so the element $f_{n,F}$ is also determined up to a finite choice. □

Let (τ_i) be a full splitting sequence starting in a maximal train track τ towards some minimal foliation F. Then each $\tau_i \in \mathcal{T}(\tau, F)$, and in fact $\tau_i \in \mathcal{T}_i(\tau, F)$. We then get an *associated* Mod-*sequence* (f_i, k_i) by applying Lemma 2.5 to $\mathcal{T}_i(\tau, F)$ for each i. In particular, we then have

$$\mathcal{T}_n(\tau, F) = f_n(\mathcal{T}^{(k_i)}).$$

As before, the numbers k_i are uniquely determined by the splitting sequence, and the mapping classes f_n are determined up to a finite choice. We call the number k_n *the type of the index* n.

Let $\mathcal{U}^{(k)}$ be the neighborhoods associated to our standard models $\mathcal{T}^{(k)}$, i.e.

$$\mathcal{U}^{(k)} = \bigcup_{\sigma \in \mathcal{T}^{(k)}} \mathrm{int} P(\sigma). \tag{2}$$

We call the $\mathcal{U}^{(k)}$ the *standard neighbourhoods*[3]. By the defining property of the associated sequence (f_n, k_n) we can then relate the standard neighbourhoods to the neighbourhoods of F given by the splitting sequence in the following way:

$$U_n(\tau, F) = f_n\left(\mathcal{U}^{(k_n)}\right). \tag{3}$$

The next lemma collects two crucial properties of the associated sequence.

Lemma 2.6. *There is a finite set $M \subset \mathrm{Mod}(S_g)$ with $M = M^{-1}$ and so that the following holds. Suppose that f_n, f_{n+1} are two consecutive terms of an associated* Mod-*sequence. Then we have*

$$f_n^{-1} f_{n+1} \in M \tag{4}$$

Furthermore,

$$f_n^{-1} f_{n+1}\left(\mathcal{U}^{(k_{n+1})}\right) \subset \mathcal{U}^{(k_n)}. \tag{5}$$

Proof. Let T be the (finite) set of all those train tracks which can be obtained from one of the train tracks in $\cup_i \mathcal{T}^{(i)}$ by full splits, and let M_0 be the set of all those mapping classes which map train tracks $\sigma \in T$ to train tracks in any $\cup_j \mathcal{T}^{(j)}$. Note that since T is finite, M_0 is finite by Lemma 2.5. We put $M = M_0 \cup M_0^{-1}$.

To see that it has property (4), observe that if f_n, f_{n+1} are consecutive terms of an associated Mod-sequence, there are train tracks $\eta_n \in \mathcal{T}^{(i_n)}, \eta_{n+1} \in \mathcal{T}^{(i_{n+1})}$, so that $f_{n+1}\eta_{n+1}$ is a full split of $f_n\eta_n$. This implies that $f_n^{-1} f_{n+1}\eta_{n+1}$ is a full split of η_n.

[3] Note that the model neighbourhoods $\mathcal{U}^{(k)}$ need not be contained in the polyhedra $P(\tau_i)$ along the splitting sequence.

In other words, $f_n^{-1} f_{n+1}$ maps a train track in $\mathcal{T}^{(i_{n+1})}$ to one in T, and is therefore an element of M by definition.

Equation (5) follows immediately from the following:

$$f_{n+1}\left(\mathcal{U}^{(k_{n+1})}\right) = U_{n+1}(\tau, F) \subset U_n(\tau, F) = f_n\left(\mathcal{U}^{(k_n)}\right).$$

□

The *type-k subsequence* is the maximal subsequence $f_s^{(k)} = f_{r_s}$ so that $k_{r_s} = k$. We say that type k is *essential for the splitting sequence* (τ_i), if the subsequence $f_s^{(k)}$ is an infinite sequence. At least one type is essential, but we suspect that the type of the initial train track need not repeat infinitely often.

By Lemma 2.6 we have that $f_{i+1} f_i^{-1} \in M$ for all i; but we warn the reader that the elements $f_{s+1}^{(k)}\left(f_s^{(k)}\right)^{-1}$ are not constrained to a finite set in the mapping class group.

2.2. **Minimal Foliations and the Curve Graph.** In this section we will prove that large terms in the associated Mod-sequence for a uniquely ergodic foliation send certain subsets of $\mathcal{PMF}$ into small neighborhoods of the foliation, and will use this to prove contracting properties for associated Mod-sequences. Intuitively, we will show that all curves (and non-minimal foliations) are attracted to the foliations F guiding the splitting sequence, and we will show that the speed of attraction can be controlled for certain geometrically constrained sets of curves.

We begin by rephrasing the contraction exhibited by train track polyhedra under splitting sequences (Theorem 2.3) in terms of associated Mod-sequences.

Corollary 2.7. *Let τ_i be a splitting sequence towards a minimal foliation F and let (f_i, k_i) be an associated* Mod-*sequence. For any essential type k we have that*

$$\bigcap_s f_s^{(k)}(\mathcal{U}^{(k)}) = \Delta(F).$$

Proof. By Corollary 2.4, we have that

$$\bigcap_{n\geq 0} U_n(\tau, F) = \Delta(F),$$

and therefore, by definition of essential type,

$$\bigcap_{n\geq 0, k_n = k} U_n(\tau, F) = \Delta(F).$$

Now, using Equation (3) we see that $U_n(\tau, F) = f_n(\mathcal{U}^{(k)})$ if the index n is of type k, and therefore

$$\bigcap_{n\geq 0, k_n = k} U_n(\tau, F) = \bigcap_s f_s^{(k)}(\mathcal{U}^{(k)}),$$

which shows the corollary. □

In other words, the mapping classes $f_s^{(k)}$ eventually contract $\mathcal{U}^{(k)}$ to a small neighborhood of $\Delta(F)$. The rest of this section is concerned with studying the contraction properties of the mapping classes f_i outside the open sets $\mathcal{U}^{(k)}$.

To this end, we use the geometry of the curve graph. Recall that the *curve graph* $\mathcal{C}(S)$ of a surface is the graph whose vertex set is the set of isotopy classes of essential simple closed curves on S, with edges between classes that admit representatives with intersection 0. We denote by $d_{\mathcal{C}(S)}$ be the resulting metric on $\mathcal{C}(S)$. The core feature of the geometry of the curve graph we need is the following.

Theorem 2.8 (Masur-Minsky [MM1]). *If S is a non-exceptional surface (i.e. $\mathcal{C}(S)$ is connected), then the curve graph is hyperbolic in the sense of Gromov.*

We will need two methods to produce quasigeodesics in the curve graph. The first one is the method employed to show hyperbolicity in [MM1].

Theorem 2.9. *Let S be a surface of finite type. Then there are numbers K, K', depending on S with the following property: suppose that $\rho : \mathbb{R} \to \mathcal{T}(S)$ is a Teichmüller geodesic, and suppose that for each $t \in \mathbb{R}$ the curve α_t has smallest possible extremal length[4] on $\rho(t)$. Then the assigment*

$$t \to \alpha_t$$

is an unparametrised K–quasigeodesic in the curve graph. In particular, for any $t < s$, the set $\{\alpha_r, t \leq r \leq s\}$ has Hausdorff distance at most K' from a curve graph geodesic joining α_t to α_s.

Proof. Theorem 2.3 of [MM1] states that a coarsely transitive path family with the contraction property in a geodesic metric space consists of uniform unparametrised quasigeodesics (for the definitions, compare Section 2.4 of [MM1]). Theorem 2.6 of [MM1] then shows that the family of paths in the curve graph obtained by taking shortest extremal length curves has the contraction property (that these paths are coarsely transitive is easy to see). □

The second, related construction of quasigeodesics uses train tracks. It is proven in [MM2, Theorem 1.3], see also [Ham1, Corollary 2.6]:

Proposition 2.10. *Let S be a surface of finite type. Then there are numbers K, K', depending on S with the following property: suppose that $(\tau_i)_i$ is a splitting sequence and suppose that for each $i \in \mathbb{N}$ the curve α_t is a vertex cycle[5] on τ_i. Then the assigment*

$$i \to \alpha_i$$

is an unparametrised K–quasigeodesic in the curve graph. In particular, for any $t < s$, the set $\{\alpha_r, t \leq r \leq s\}$ has Hausdorff distance at most K' from a curve graph geodesic joining α_t to α_s.

---

[4]See e.g. [Ahl] for a definition. The precise definition of this does not matter too much to understand the theorem; it would remain true also for e.g. the shortest hyperbolic geodesic on $\rho(t)$.

[5]See e.g. [PH] or [MM1, Section 4.1] for a definition of vertex cycle. Again, the precise definition of this does not matter too much; the theorem would remain true for e.g. the shortest trainpath on τ_i.

For a Gromov hyperbolic space, one can define a boundary at infinity, see e.g. [BH, III.H.3][6]. If α_0 is some basepoint, recall the *Gromov product*

$$(x \cdot y)_{\alpha_0} = \frac{1}{2}(d(\alpha_0, x) + d(\alpha_0, y) - d(x, y)).$$

Also note that the Gromov product extends from the space to the boundary at infinity [BH, III.H.3.15]. The Gromov product has the property that

$$|(x \cdot y)_{\alpha_0} - (x' \cdot y')_{\alpha_0}| \leq d(x, x') \tag{6}$$

for any point y and (finite) points x, x'.

In the case of the curve graph, the Gromov boundary can be identified explicitly with a different space. We define the set $\mathcal{EL}(S)$ to be the set of minimal foliations with the measure-forgetting topology. That is, we consider the subset $\mathcal{M} \subset \mathcal{PMF}$ of all minimal foliations, and let $\mathcal{EL}(S)$ be the quotient topological space $\mathcal{M}/\sim$ under the equivalence relation which lets $F \sim F'$ if F, F' are topologically equivalent.

Theorem 2.11 ([Kla, Theorems 1.2, 1.3 and 1.4])**.** *i) The Gromov boundary of $\mathcal{C}(S)$ is homeomorphic to the space $\mathcal{EL}(S)$.*

ii) A sequence α_i of curves (interpreted as points in the curve graph) converges to the point at infinity defined by a minimal foliation F if and only if every accumulation point of $\{\alpha_i, i \in \mathbb{N}\}$ in $\mathcal{PMF}$ is contained in $\Delta(F)$.

iii) Suppose that ρ is a Teichmüller geodesic ray whose vertical foliation is a minimal foliation F, and that for every t, the curve α_t is a curve of smallest extremal length on $\rho(t)$. Then the curves α_t (interpreted as points in the curve graph) converge to F (interpreted as a point in the Gromov boundary)

As a consequence of Theorem 2.11 we have the following characterization of neighborhoods in $\mathcal{PMF}$ using the curve graph.

Lemma 2.12. *Suppose that F is a minimal foliation, and U is an open neighborhood of $\Delta(F)$ in $\mathcal{PMF}$. Let γ be an arbitrary simple closed curve. Then there is a number K with the following property: suppose that β is a simple closed curve so that (as a point in the curve graph) we have*

$$(F \cdot \beta)_\gamma > K.$$

Then β (seen as a projective measured foliation) is contained in U.

Proof. Suppose that the claim were false. Then we would find a sequence (β_i) with $(\beta_i \cdot F)_\gamma > i$ but $\beta_i \notin U$. By the Gromov product condition, (β_i) would then be an admissible sequence converging to the boundary point F. So, by Theorem 2.11 ii), the sequence β_i converges in the measure forgetting topology to F. Since U is an open neighborhood of $\Delta(F)$ this is impossible as $\beta_i \notin U$. □

We also need the following partial converse.

[6]Since the curve graph is locally infinite, some care has to be taken here. The correct definition uses sequences with diverging Gromov products, rather than equivalence classes of quasi-geodesic rays.

Lemma 2.13. *There is a number k_0 with the following property. Suppose that F is a minimal foliation, α is a simple closed curve, and*

$$(F \cdot \alpha)_\gamma > K$$

If μ_i is a sequence of minimal foliations converging to α in $\mathcal{PMF}$, then

$$(F \cdot \mu_i)_\gamma > K - k_0$$

for all large i.

Proof. Denote by $\Phi : \mathcal{T}(S) \to \mathcal{C}(S)$ the map which assigns to a marked hyperbolic surface in Teichmüller space a curve of smallest extremal length[7]. Pick a basepoint X_0 in Teichmüller space for which γ is a curve of smallest extremal length, and consider the Teichmüller geodesic rays ρ_i starting from X_0 in the direction of μ_i. Since the μ_i converge to α in $\mathcal{PMF}$, the rays ρ_i converge uniformly on compact subsets to the Teichmüller geodesic ray ρ_∞ starting in X_0 with vertical foliation α.

Theorem 2.9 implies that there is a constant K so that the images $\Phi \circ \rho_i$ can be reparametrised to be K–quasigeodesics q_i beginning in γ. By Theorem 2.11 iii), the quasigeodesic q_i connects γ to the point μ_i in the Gromov boundary of the curve graph.

There is a constant T_0 so that $\Phi \circ \rho_\infty(t)$ is equal to α for all $t \geq T_0$. As the ρ_i converge to ρ_∞ uniformly on compact sets in Teichmüller space, one concludes that $\Phi \circ \rho_i(T_0) = \alpha$ for all large i. Hence, the q_i pass through α for all large i. This implies that there is a constant k_0, just depending on K and the hyperbolicity constant of the curve graph, so that $(F \cdot \mu_i)_\gamma > (F \cdot \alpha)_\gamma - k_0$, which implies the lemma. □

The next lemma and corollary are well known and standard and included for completeness.

Lemma 2.14. *Let F be a minimal foliation, and K a number. Then suppose that $x, y \in \mathcal{C}(S)$ with*

$$(F \cdot x)_\gamma, (F \cdot y)_\gamma \geq K.$$

Let z be a point on a geodesic between x, y. Then

$$(F \cdot z)_\gamma \geq K - 4\delta,$$

where δ is the hyperbolicity constant of the curve graph.

Proof. First we observe that if x, y, z are three points in $\mathcal{C}(S)$ and z lies on a geodesic between x and y, we have

$$\begin{aligned} 2(x \cdot z)_\gamma = d(\gamma, x) + d(\gamma, z) - d(x, z) &\geq d(\gamma, x) + d(\gamma, y) - d(y, z) - d(x, z) \\ &= d(\gamma, x) + d(\gamma, y) - d(x, y) = 2(x \cdot y)_\gamma. \end{aligned}$$

By δ–hyperbolicity, we have that for all triples a, b, c of points in $\mathcal{C}(S) \cup \partial_\infty \mathcal{C}(S)$

$$(a \cdot c)_\gamma \geq \min\{(a \cdot b)_\gamma, (b \cdot c)_\gamma\} - 2\delta,$$

compare e.g. [BH, III.H.3.17.(4)]. First, apply this to x, F, y to conclude that

$$(x \cdot y)_\gamma \geq K - 2\delta.$$

[7]This curve may not be well-defined, but any two choices have uniformly few intersections due to the collar lemma. Hence, any two choices have uniformly small distance in the curve graph.

Now, apply this same estimate again, to conclude

$$\begin{aligned}
(F \cdot z)_\gamma &\geq \min\{(F \cdot x)_\gamma, (x \cdot z)_\gamma\} - 2\delta \\
&\geq \min\{(F \cdot x)_\gamma, (x \cdot y)_\gamma\} - 2\delta \\
&\geq \min\{K, K - 2\delta\} - 2\delta \\
&\geq K - 4\delta
\end{aligned}$$

which is what we wanted to prove. □

Corollary 2.15. *Let $K, D > 0$ be numbers, F be a minimal foliation. Suppose that $\tilde{x}, \tilde{y}$ are any two points in the curve complex or its boundary, satisfying*

$$(F \cdot \tilde{x})_\gamma, (F \cdot \tilde{y})_\gamma \geq K.$$

Suppose that $z \in \mathcal{C}(S)$ lies on a (possibly infinite) D–quasi-geodesic q with endpoints $\tilde{x}$ and $\tilde{y}$. Then

$$(F \cdot z)_\gamma \geq K - X,$$

where X is a number depending only on the hyperbolicity constant of the curve graph and the quasi-geodesic constant D.

Proof. Choose points $x_i = q(r_i), y_i = q(s_i)$ in the curve complex on the quasi-geodesic q which converge to $\tilde{x}, \tilde{y}$ respectively. If an endpoint of q is finite, we assume that the corresponding sequence is eventually constant. Using [BH, III.H.3.17.(5)] we then conclude from the Gromov product estimate in the prerequisites that

$$\min\{(F \cdot x_i)_\gamma, (F \cdot y_i)_\gamma\} \geq K - 2\delta,$$

for large i. We furthermore assume that i is large enough so that z is contained in the subsegment q_i of q with endpoints x_i, y_i. By δ–hyperbolicity, there is a number B depending only on D, so that the Hausdorff distance between q_i and the geodesic connecting x_i to y_i is at most B. Let z' be a point on that geodesic of distance at most B to z. By Lemma 2.14, we then have

$$(F \cdot z')_\gamma \geq K - 4\delta,$$

and thus

$$(F \cdot z)_\gamma \geq K - 4\delta - B.$$

Hence $X = 4\delta + B$ satisfies the requirement. □

Lemma 2.16. *Let F be a minimal foliation, τ a train track and (τ_i) a splitting sequence in the direction of F and let (f_i, k_i) be an associated* Mod*-sequence. Suppose that (γ_i) is a sequence of simple closed curves so that γ_i is contained in $f_i^{(k_i)}(\mathcal{U}^{(k_i)})$ for every i. Then, for any base point α_0, we have*

$$(\gamma_i \cdot F)_{\alpha_0} \to \infty.$$

Proof. By Corollary 2.4 and the assumption, any accumulation point of the curves γ_i (interpreted as projective measured foliation) is contained in $\Delta(F) \subset \mathcal{PMF}$. By Theorem 2.11 ii), the γ_i therefore converge (interpreted as points in the curve graph) to F in the Gromov boundary. By definition, this implies that the Gromov product condition claimed in the corollary. □

We can use this to show the following contraction behavior for finite-diameter subsets in the curve graph.

Proposition 2.17. *Let F be a minimal foliation, τ a train track and (τ_i) a splitting sequence in the direction of F and let (f_i, k_i) be an associated* Mod*-sequence.*

Consider any neighborhood $\mathcal{V}$ of $\Delta(F)$ in $\mathcal{PMF}$, and let a simple closed curve β_0 and a number $d > 0$ be given.

Then there is a number $N = N(\tau, F, \mathcal{V}, \beta_0, d) > 0$ so that the following holds: If β is any simple closed curve with $d_{\mathcal{C}(S)}(\beta_0, \beta) \leq d$, then

$$f_n(\beta) \in \mathcal{V} \qquad \forall n > N.$$

Proof. As a first reduction, note that by Corollary 2.4 we may assume that $\mathcal{V}$ is of the form $f_s^{(k_s)}(\mathcal{U}^{(k(s))})$ for a large enough s. Fix a basepoint α_0 in the curve graph.

Apply Lemma 2.12 in order to obtain a number $D > 0$ with the property that if γ is any curve so that the Gromov product satisfies

$$(\gamma \cdot F)_{\alpha_0} > D,$$

then $\gamma \in \mathcal{V}$ as an element of $\mathcal{PMF}$.

Now, for each k choose a curve δ_k contained in $\mathcal{U}^{(k)}$ and put $\gamma_n = f_n(\delta_{k(n)})$.

Observe that

$$d_{\mathcal{C}(S)}(f_n(\beta_0), \gamma_n) \leq \max_k d_{\mathcal{C}(S)}(\beta_0, \delta_k) = C_0$$

and, if $d_{\mathcal{C}(S)}(\beta, \beta_0) \leq d$ we therefore have

$$d_{\mathcal{C}(S)}(f_n(\beta), \gamma_n) \leq C_0 + d.$$

Thus, using Equation (6), we see

$$(f_n(\beta) \cdot F)_{\alpha_0} \geq (\gamma_n \cdot F)_{\alpha_0} - d_{\mathcal{C}(S)}(f_n(\beta), \gamma_n) \geq (\gamma_n \cdot F)_{\alpha_0} - (C_0 + d).$$

Applying Lemma 2.16 to the curves γ_n we see that there is a number N so that

$$(\gamma_n \cdot F)_{\alpha_0} > D + C_0 + d \qquad \forall n > N.$$

Together with the previous inequality this implies that

$$(f_n(\beta) \cdot F)_{\alpha_0} > D \qquad \forall n > N,$$

which finishes the proof. □

The next lemma, which requires a definition, will allows us to obtain that large terms in the Mod-sequence to a uniquely ergodic foliation contract certain infinite diameter subsets of the curve graph (thought of as foliations) to a small neghborhood of the uniquely ergodic foliation.

Definition 2.18. *Let D be a number, and ψ a pseudo-Anosov map. A $(D-)$*quasi-axis *is a bi-infinite D–quasi-geodesic $q : \mathbb{R} \to \mathcal{C}(S)$ so that its image $\psi^j q$ has (Hausdorff) distance at most D from the image of q for any power $j \in \mathbb{Z}$.*

Lemma 2.19. *There are constants $D, B > 0$, just depending on the surface, so that every pseudo-Anosov map ψ of S has a D–quasi-axis. Furthermore, any two such quasi-axes have Hausdorff distance at most B.*

Proof. Let $\rho : \mathbb{R} \to \mathcal{T}(S)$ be the Teichmüller geodesic invariant under ψ, i.e. there is some $T > 0$ so that for all t we have $\psi\rho(t) = \rho(t+T)$. For each $t \in [0, T)$, choose a curve α_t of smallest extremal length on $\rho(t)$. For $t \in [i, i+T)$ put $\alpha_t = \psi^i(\alpha_{t-i})$. Then for all t, the curve α_t has smallest extremal length on $\rho(t)$. By Theorem 2.9, the assignment $t \to \alpha_t$ is an (unparametrised) quasigeodesic with quasigeodesic constant just depending on the topological type of the surface. By construction, $t \to \alpha_t$ is invariant under the action of ψ. This shows that quasi-axes exist.

The uniqueness statement follows since any quasiaxis for ψ converges in the Gromov boundary of the curve graph to the stable and unstable foliation of ψ by Theorem 2.11 iii) and two D-quasigeodesics with the same endpoints in a Gromov hyperbolic space have bounded Hausdorff distance. $\square$

In the future, we will choose a D for which Lemma 2.19 holds once and for all, and simply refer to quasi-axes of pseudo-Anosov maps.

Also recall the definition of a *Dehn twist* T_α about a simple closed curve α (compare e.g. [FM1, Section 3.1]). If α is a multicurve, together with a choice of left/right for each component, then we denote by T_α the product of the left/right Dehn twists about the curves in α.

Proposition 2.20. *Let F be a minimal foliation, τ a train track and (τ_i) a splitting sequence in the direction of F and let (f_i, k_i) be an associated* Mod*-sequence.*

Consider any neighborhood $\mathcal{V}$ of $\Delta(F)$ in $\mathcal{PMF}$. Let ψ be a pseudo-Anosov, and let α be a multicurve which is within distance d of its quasi-axis in the curve graph. Let $r > 0$ be any number. Suppose β_0 is a curve.

Then there is a number $N = N(\tau, F, \mathcal{V}, \psi, \alpha, d, r, \beta_0) > 0$ with the following property. Suppose that $n > N$ is given. Then there is a number t_0 (which depends on n), so that for all $t > t_0$ the conjugate $\hat{\psi} = (T_\alpha)^t \circ \psi \circ (T_\alpha)^{-t}$ satisfies the following:

If β is any simple closed curve with $d_{\mathcal{C}(S)}(\beta_0, \beta) \leq d$, then

$$f_n(\hat{\psi}^j \beta) \in \mathcal{V}, \qquad \forall j \in \mathbb{Z}$$

Proof. We follow a similar strategy as in the previous proposition. Apply Lemma 2.12 to find a number D so that if

$$(\gamma \cdot F)_{\alpha_0} > U,$$

then $\gamma \in \mathcal{V}$ as an element of $\mathcal{PMF}$.

Introduce the notation

$$\hat{\psi}_t = (T_\alpha)^t \circ \psi \circ (T_\alpha)^{-t}.$$

We therefore need to show, that there is a number N so that for all $n > N$ there is a t_0 so that

$$(f_n(\hat{\psi}_t^j \beta) \cdot F)_{\alpha_0} > U,$$

for any curve β with $d_{\mathcal{C}(S)}(\beta_0, \beta) \leq d$, and any $t > t_0$, any $j \in \mathbb{Z}$.

The first stage of the proof consists of a (lengthy) reduction of this statement to a similar statement (Equation (7) below) about quasi-axes of the ψ_t. To begin showing this reduction, note that

$$(f_n(\hat{\psi}_t^j \beta) \cdot F)_{\alpha_0} \geq (f_n(\hat{\psi}_t^j \beta_0) \cdot F)_{\alpha_0} - d(\beta, \beta_0) \geq (f_n(\hat{\psi}_t^j \beta_0) \cdot F)_{\alpha_0} - d$$

and therefore it suffices to show

$$(f_n(\hat{\psi}_t^j \beta_0) \cdot F)_{\alpha_0} > U + d.$$

Arguing as above, we have that

$$(f_n(\hat{\psi}_t^j \beta_0) \cdot F)_{\alpha_0} > (f_n(\hat{\psi}_t^j \alpha) \cdot F)_{\alpha_0} - d(\alpha, \beta_0).$$

Hence, it suffices to show that

$$(f_n(\hat{\psi}_t^j \alpha) \cdot F)_{\alpha_0} > U + d + d(\alpha, \beta_0) =: U_1,$$

for any $t > t_0$, any $j \in \mathbb{Z}$.

Now, let ρ be a (D–)quasi-axis for ψ. Since the mapping class group acts as isometries on the curve graph, we have that $f_n T_\alpha^t \rho$ is a (D–)quasi-axis for $f_n \hat{\psi}_t f_n^{-1}$. Furthermore,

$$d(f_n \alpha, f_n T_\alpha^t \rho) = d(\alpha, T_\alpha^t \rho) = d(\alpha, \rho) = A,$$

for all t, since T_α acts as an isometry fixing α. Hence, $f_n \alpha$ is (for all choices of n and t) within A of the D–quasi-axis $f_n T_\alpha^t \rho$ of $f_n \hat{\psi}_t f_n^{-1}$. Let η be a point on $f_n T_\alpha^t \rho$ with $d(f_n \alpha, \eta) \leq A$. The D–quasi-axis property then implies that for any j we have that

$$d((f_n \hat{\psi}_t f_n^{-1})^j \eta, f_n T_\alpha^t \rho) \leq D$$

and therefore

$$d((f_n \hat{\psi}_t f_n^{-1})^j f_n \alpha, f_n T_\alpha^t \rho) \leq A + D$$

As such, we have that

$$d(f_n(\hat{\psi}_t^j \alpha), f_n T_\alpha^t \rho) = d(f_n(\hat{\psi}_t^j f_n^{-1} f_n \alpha), f_n T_\alpha^t \rho) = d((f_n(\hat{\psi}_t f_n^{-1})^j (f_n \alpha), f_n T_\alpha^t \rho) \leq A + D.$$

Therefore, to prove the proposition, it suffices to show that there is a number N, so that for all $n > N$ there is a number t_0, so that for all $t > t_0$:

$$\forall x \in f_n T_\alpha^t \rho : (x \cdot F)_{\alpha_0} > U_1 + A + D =: U_2. \tag{7}$$

Now, use Lemma 2.16 as in the previous proof, to find a number N so that

$$(f_n(\alpha) \cdot F)_{\alpha_0} > 2U_2 + X + k_0 \qquad \forall n > N, \tag{8}$$

where X is the number from Corollary 2.15 and k_0 is the number from Lemma 2.13, applied to the quasi-geodesic constant D. At this point, fix a number $n > N$.

Observe that if μ_+, μ_- are the stable and unstable foliations of ψ, then $T_\alpha^t \mu_+, T_\alpha^t \mu_-$ are the stable and unstable foliations of $\hat{\psi}$. Note that as $t \to \infty$, both of these foliations converge to α in $\mathcal{PMF}$. Consider $f_n \hat{\psi}_t (f_n)^{-1}$, and observe that its stable and unstable foliations therefore converge to $f_n(\alpha)$ in $\mathcal{PMF}$ as the number t increases. By Lemma 2.13, this implies that we can choose t_0 large enough, so that for any $t > t_0$ we have

$$(f_n(T_\alpha^t \mu_+) \cdot F)_{\alpha_0} > U_2 + X$$
$$(f_n(T_\alpha^t \mu_-) \cdot F)_{\alpha_0} > U_2 + X$$

Let now z be any point on a D–quasi-geodesic with endpoints $f_n(T^t \mu_+), f_n(T^t \mu_-)$. Then Corollary 2.15 implies that

$$(z \cdot F)_{\alpha_0} > U_2$$

Since the quasi-axis $f_n T_\alpha^t \rho$ is such a D–quasi-geodesic, the proposition follows. $\square$

In the proof of local path-connectivity, we require uniform control over the constants N appearing in the previous two results (Propositions 2.17 and 2.20) Before stating the corresponding lemma, suppose that $(\tau_i)_i$ is a full splitting sequence in the direction of some minimal foliation F.

Then consider, in Proposition 2.17 or 2.20, a neighbourhood $\mathcal{V} = U_k(\tau, F)$, and observe that it is also a neighbourhood of $\Delta(E)$ for all minimal $E \in U_i(\tau, F), i \geq k$. Additionally, E determines a full splitting sequence starting in τ, whose first i terms are identical with the one defined by F.

Hence, it makes sense to apply Proposition 2.17 or 2.20 for this neighbourhood $\mathcal{V}$, and E in place of F with its full splitting sequence starting in τ. The following lemma shows a boundedness of the resulting numbers N that these propositions produce.

Lemma 2.21. *Suppose that $(\tau_i)_i$ is a full splitting sequence in the direction of some minimal foliation F with $\tau_1 = \tau$. Put $\mathcal{V} = U_k(\tau, F)$ for some k.*

Suppose we are given either

(1) A curve β_0 and a number $d > 0$, or
(2) A pseudo-Ansosov ψ, a curve α, a number $r > 0$ and a curve β_0.

Then there are numbers $M, N > 0$ with the property that the number

(1) $N(\tau, E, \mathcal{V}, \beta_0, d)$ from Proposition 2.17, or
(2) $N(\tau, E, \mathcal{V}, \psi, \alpha, d, r, \beta_0)$ from Proposition 2.20

can be chosen to be smaller than N for all minimal $E \in U_M(\tau, F)$.

Proof. We will describe the case of Proposition 2.20 in detail, the corresponding argument for Proposition 2.17 is similar and simpler.

Recall from the proof of Proposition 2.20 that what one needs to show is the estimate in (7). This in turn is implied simply from (8), which is purely a statement about Gromov product growth of images of α under the associated mapping class group sequence f_n. Hence, to show this lemma, it suffices to show that the number N in (8) can be bounded for associated sequences f'_n as in the statement of this lemma.

If $E \in U_M(\tau, F)$, then by definition the first M terms of the associated Mod-sequence for E and F agree. Hence, to show this lemma, we have to show that the existence of a number N making (8) true can already be guaranteed by knowing a large initial segment of the associated Mod-sequence. The remainder of this proof is concerned with showing that.

Similar to the proof of Proposition 2.17, choose for each k a curve α_k which is carried by each $\sigma \in \mathcal{T}^{(k)}$ as a vertex cycle. By Proposition 2.10, the path $n \mapsto f_n \alpha_{k(n)}$ is then uniformly Hausdorff close to a uniform quasi-geodesic in the curve graph which converges to F.

In particular, this implies that for any K_0 there is an N with the property that

$$(f_n \alpha_{k(n)} \cdot F)_\gamma > K_0 \quad \text{for all } n > N.$$

If now $F' \in U_N(\tau, F)$ and (f'_i) is an associated Mod-sequence for F', then we may assume $f'_i = f_i$ for all $i \leq N$ by definition. Thus, for some uniform constant c

(depending on the quasi-geodesic constant k_1 Proposition 2.10 and the hyperbolicity constant of the curve graph) we have that

$$(f'_n\alpha_{k'(n)} \cdot F)_\gamma > K_0 - c \quad \text{for all } n > N.$$

Since the distance between the curve α and the (finitely many) α_k is bounded, there is a further constant d so that

$$(f'_n\alpha \cdot F)_\gamma > K_0 - c - d \quad \text{for all } n > N.$$

Choosing $K_0 - c - d > 2U_2 + X + k_0$ then yields that the corresponding N works in (8) for the sequences of all $F' \in U_N(\tau, F)$, proving the lemma. □

3. Paths by pushing points

In this section we will construct many special paths of cobounded foliations for punctured surfaces, which will serve as building blocks for all subsequent constructions. The paths we will eventually use to connect uniquely ergodic foliations will be concatenations of paths of this form, except possibly at a countable set of points which will be stable foliations of pseudo-Anosovs (or the endpoints).

The construction described in this section is crucially inspired by the work of Leininger and Schleimer in [LS1], where they build paths of minimal foliations. Our main contribution is that we modify their construction to produce paths of uniquely ergodic (and in fact cobounded) foliations, and obtain some extra control over how these paths follow a "combinatorial skeleton" given by a finite set of curves.

3.1. **Preliminaries on Covers, and on Adding Points.**

Our notation follows [LS1] and we refer the reader to that article for a very good and readable source for background information on the methods used here.

A *smooth surface* will denote a connected, compact, oriented 2–manifold without boundary. All maps between smooth surfaces will be assumed to be smooth unless specified. By a slight abuse of notation, a *(holomorphic) Abelian differential on S* is a smooth 1–form ω which is holomorphic with respect to some complex structure on S (compatible with orientation and smooth structure). We denote by d_ω the (singular) flat metric on the surface defined by integrating ω.

We let $\widetilde{\Omega}(S)$ be the set of all such Abelian differentials. Note that $\widetilde{\Omega}(S)$ is a path-connected set (in fact, a vector bundle over a contractible base; compare [LS1, Section 2.6]).

The quotient

$$\Omega(S) = \widetilde{\Omega}(S)/\mathrm{Diff}_0(S)$$

is the Hodge bundle of Abelian differentials over Teichmüller space of S. We need a variant for surfaces with marked points (which is, crucially, the point of this whole discussion). Namely, if $\mathbf{z} \subset S$ is a finite, ordered set of distinct points, we let $\mathrm{Diff}_0(S, \mathbf{z})$ denote the group of diffeomorphisms of S, fixing each point in $\mathbf{z}$, which are homotopic to the identity through such maps. We let

$$\Omega(S, \mathbf{z}) = \widetilde{\Omega}(S)/\mathrm{Diff}_0(S, \mathbf{z})$$

As in [LS1], the central idea is that any Abelian differential $\omega \in \widetilde{\Omega}(S)$ defines projections $\hat{\omega} \in \Omega(S, \mathbf{z})$ and $\bar{\omega} \in \Omega(S)$ (in the notation of [LS1]).

There is an action of $\mathrm{SL}_2(\mathbb{R})$ on $\widetilde{\Omega}(S)$ defined in the usual way (e.g. by postcomposing canonical flat charts) which descends to the usual $\mathrm{SL}_2(\mathbb{R})$–action on $\Omega(S)$. We denote by g_t the action of diagonal matrices, i.e. Teichmüller geodesic flow.

3.2. **Torus Covers and Badly Approximable Points.** In this section, we begin to construct Abelian differentials with desirable horizontal foliations.

To begin, we say that a Abelian differential $\omega \in \widetilde{\Omega}(S)$ is *(eventually) ϵ–thick* if there exists N so that for all $t > N$ we have that every essential simple closed curve on S has length $\geq \epsilon$ with respect to the singular flat metric $g_t\omega$. We say that ω is *strongly (eventually) ϵ–thick with respect to* $\mathbf{z}$ if the same is true for any arc with endpoints in $\mathbf{z}$. Note that (strong) eventual thickness is invariant under the $\mathrm{Diff}_0(S, z)$–action, and therefore the notion also makes sense for differentials in $\Omega(S, z)$.

The purpose of this section is to give a robust criterion that we will use to construct many paths of thick Abelian differentials.

We make the following (slightly idiosyncratic) definitions, which will be one of the core mechanisms in our construction.

Definition 3.1. *i) Let (X, d) be a metric space and $T : X \to X$ be a dynamical system. We say a pair of points $(x, y) \in X$ is* B-badly approximable *if there exists N so that $k \cdot d(T^k x, y) \geq B$ for all $k \geq N$ and moreover $T^k x \neq y$ for all $k \neq 0$. We may also say that the point y B-*badly approximates *x.*

*ii) We say a rotation R_α of the circle is B-*badly approximable *if the pair (x, x) is B-badly approximable for some (equivalently every) $x \in \mathbb{R}$ for the dynamical system*

$$R_\alpha : \mathbb{R}/\mathbb{Z} \to \mathbb{R}/\mathbb{Z}, \qquad z \mapsto z + \alpha \bmod \mathbb{Z}$$

iii) Similarly, if $F^t : X \to X$ is a measurable flow of a metric space, we say a pair of points (x, y) is B-badly approximable *if there exists N so that $t \cdot d(F^t x, y) \geq B$ for all $t \geq N$ and moreover, $F^t x \neq y$ for all $t \neq 0$. We say a straight line flow on a torus is* B-badly approximable *if the pair (x, x) is B-badly approximable for some x.*

The following lemma shows why we are interested in badly approximable points.

Lemma 3.2. *If q and q' are distinct B-badly approximable points on a torus then any trajectory γ from q to q' has $|g_t\gamma| \geq \sqrt{B}$ for all large enough t.*

Proof. Let t_0 satisfy $d(F^L q, q') > \frac{B}{L}$ for all $L \geq t_0$. Because q and q' are not in the same orbit, by the definition of B-badly approximable, $\lim\limits_{t\to\infty} |g_t\gamma| = \infty$ for every γ a trajectory from q to q'. Thus, we may restrict our attention to the cofinite set of such γ with vertical component at least t_0. Let γ be such a geodesic from q to q'. Because the torus is flat, if the vertical component of γ is a and the horizontal component is b we have that $d(F^a q, q') = b$. Since we assume that (q, q') are B-badly approximable, b is at least $\frac{B}{a}$ if a is large enough. Since the product of the horizontal and vertical components of curves are preserved by g_t, we have $|g_t\gamma|$ is at least $\sqrt{2ab} \geq \sqrt{B}$ for all t. (We are also using the elementary fact that the shortest vector in the positive cone in $\mathbb{R}^2$ with fixed product of horizontal and vertical components has angle $\frac{\pi}{4}$.) $\square$

Definition 3.3. *Let S be a closed surface of genus $g \geq 2$. An Abelian differential $\omega \in \widetilde{\Omega}(S)$ is called* (ϵ, B)–torus good with respect to marked points $q_1, \dots, q_k$ *if there is a regular branched cover, branched over one point,*

$$p : S \to T$$

of S to a torus T and an Abelian differential ω_T on T so that

1. *ω_T is eventually ϵ–thick.*
2. *ω is the pullback of ω_T.*
3. *The images $p(q_i)$ of q_i in T are pairwise B–badly approximable with respect to the flat structure defined by ω_T.*

The associated data *to the (ϵ, B)–torus good ω comprise the cover p and the base differential ω_T.*

The notion of being torus good is invariant under the action of $\mathrm{Diff}_0(S, \{q_1, \dots, q_k\})$ by pulling back differentials, and therefore is also defined for differentials in $\Omega(S, \{q_1, \dots, q_k\})$.

The following proposition shows why we are interested in torus good differentials.

Proposition 3.4. *For any (ϵ, B) and S there is a number $\delta > 0$ with the following property. If ω is (ϵ, B)–torus good with respect to marked points $q_1, \dots, q_k$, then ω is eventually strongly δ–thick with respect to $\mathbf{z} = (q_1, \dots, q_k)$.*

In particular, the horizontal foliation of ω is cobounded as a foliation on $(S, \mathbf{z})$.

Proof. To prove that ω is eventually strongly δ-thick with respect to $\mathbf{z}$, by definition we have to show that there is a t_0 so that:

- if γ is a simple closed curve on ω then $|g_t\gamma| \geq \delta$ for all $t > t_0$ and
- if γ is a trajectory from q_i to q_j with $j \neq i$ we have that $|g_t\gamma| > \delta$ for all $t > t_0$.

The first condition follows for any $\delta \leq \epsilon$ because we are assuming that ω_T is eventually ϵ-thick and any simple closed curve on ω projects to a closed curve of the same length on ω_T because we are branched over a single point. Similarly we have that $\pi(\gamma)$ is a trajectory from $\pi(q_i)$ to $\pi(q_j)$ and $\pi(g_t\gamma) = g_t\pi(\gamma)$ and so any such trajectory has length at least B by the Lemma 3.2. This implies the two conditions above, and therefore eventual strong δ–thickness of ω.

To see the second claim, note that as $t \to \infty$, the differentials $g_t\omega$ all lie in a compact set of the moduli space of flat surfaces by the first part. This in turn implies that Teichmüller flow in the direction of the horizontal foliation of ω also only defines Riemann surfaces which lie in a compact set of the moduli space of $S - \mathbf{z}$. This shows the proposition. □

Next, we will show that these torus good differentials are in fact dense in the set of all differentials. The proof of this uses Schmidt games, a technique from Diophantine approximation, which we briefly define and discuss in the next section.

3.3. **Schmidt game digression.** Suppose we are given a set $E \subset \mathbb{R}^n$. Suppose two players Bob and Alice take turns choosing a sequence of closed Euclidean balls

$$B_0 \supset A_1 \supset B_1 \supset A_2 \supset B_2 \dots$$

(Bob choosing the B_i and Alice the A_i) whose diameters satisfy, for fixed $0 < \alpha, \beta < 1$, and all $i > 0$

$$|B_i| = \beta|A_i| \quad \text{and} \quad |A_{i+1}| = \alpha|B_i|. \tag{9}$$

The only requirement on B_0 is that it has positive diameter. Following Schmidt [Sch] we make the following definition.

Definition 3.5. *We say E is an* (α, β)-winning set *if Alice has a strategy so that no matter what Bob does, $\cap_{i=1}^{\infty} B_i \in E$. It is* α-winning *if it is (α, β)-winning for all $0 < \beta < 1$. E is a* winning set for Schmidt game *if it is α-winning for some $\alpha > 0$.*

A set is called *α-winning* if it is (α, β) winning for all $0 < \beta < 1$.

Because Bob's first move is unconstrained we have:

Lemma 3.6. *If S is an (α, β) winning set for any α, β then S is dense in X.*

Lemma 3.7. *(*[Sch, Theorem 2]*])If $S_1, ..., S_k$ are α-winning sets then $\cap_{i=1}^{k} S_i$ is α-winning.*

In fact the previous lemma is true for countable intersections as well.

Theorem 3.8. [Tse, first line of Section 2] *Let $\xi \in [0, 1)$, R denote rotation by ξ and $x \in [0, 1)$. The set of y so that x, y is $(\frac{\alpha\beta}{4})^3$-badly approximable (for R) is a (α, β) winning set.*

From the previous two results we obtain:

Corollary 3.9. *Given any rotation R and a finite number of points $p_1, ..., p_r$ in $[0, 1)$ we have that the set of q so that p_i, q are $\frac{1}{4 \cdot 4 \cdot 4}^{3r}$-badly approximable for all i with respect to R is (α, β)-winning for some α, β.*

By iterating the previous result and Lemma 3.6 we get:

Corollary 3.10. *Given any rotation on $[0, 1)$, the set of $p_1, ..., p_k$ that are pairwise $(\frac{1}{4 \cdot 4 \cdot 4})^{3(k-1)}$-badly approximable is dense in $[0, 1)^k$.*

3.4. **Density of torus good differentials.**

Proposition 3.11. *Let S be a closed surface of genus $g \geq 2$ and $q_1, \ldots, q_k$ be a set of marked points. Suppose that we fix a regular branched cover, branched over one point,*

$$p : S \to T$$

of S to a torus T and an Abelian differential ω_T on T so that ω is the pullback of ω_T.

Then for every ω_T, every neighborhood U of ω_T in $\widetilde{\Omega}(T)$, and every $\delta > 0$ there exists $\omega'_T \in U$ and points q'_i with $d_\omega(q_i, q'_i) < \delta$, for all i, so that the pullback of ω'_T is (ϵ, B)–torus good with respect to marked points (q'_i). In this, ϵ can be chosen independent of ω and B can be chosen to just depend on k.

Proof. We will work throughout with the canonical flat charts defined by ω, ω_T realizing p as a holomorphic map. We will then show that we can move the q_i by a small amount (in these charts!) and modify ω_T by a small rotation to obtain an (ϵ, B)–torus good differential. This is enough to show the proposition.

Given a straight line flow on a flat torus, there are many (geodesic) transversals so that the first return map of the flow to the transversal is a rotation. Moreover there exists C so that every aperiodic straight line flow on a flat torus of area 1 has infinitely many transversals γ, so that the first return map to γ is a rotation and the return times to γ are between $\frac{1}{C|\gamma|}$ and $\frac{C}{|\gamma|}$. To see this, note that there is a compact set K in the moduli space of flat tori, so that if the orbit $g_t\omega_T$ of a torus ω_T under Teichmüller flow does not diverge to infinity (without recurring), then there exist arbitrarily large t so that $g_t\omega_T \in K$. In the case of an aperiodic straight line flow the first case does not happen. In the second case, a side of the fundamental domain of the torus $g_t\omega_T \in K$ will work as a transversal.

Sublemma: Let p, q be points on a torus T and F^t a minimal straight line flow on T. Suppose that γ is a transversal for F^t, and let T_0 be the minimal first return time of F^t to the transversal γ. Assume further that the first return of F^t defines a rotation R_ξ on γ. Suppose that $s_1, s_2 > 0$ are minimal so that $F^{s_1}p, F^{s_2}q \in \gamma$, and that the points $F^{s_1}p, F^{s_2}q$ are B-badly approximable for R_α. Then p and q are B'-badly approximable for F^t for any $B' < B \cdot T_0$.

Proof of Sublemma. We prove the statement by contradiction. Assume that there exists $\epsilon > 0$ and arbitrarily large L so that

$$d(F^L p, q) < \frac{BT_0 - \epsilon}{L}.$$

Assume that the straight line flow is vertical. We may assume that $F^L p$ is on the same horizontal as q. Let T_1 be the maximal return time of the flow to γ. Then there is some $0 \leq \ell \leq 3T_1$, so that $F^\ell q \in \gamma \setminus \partial\gamma$ (since at most two returns can hit a boundary point of γ). Furthermore, after fixing ℓ, we have that for all large enough L:

$$d(F^L p, q) < d(F^\ell q, \partial\gamma)$$

Let $h \subset \gamma$ be the shortest horizontal segment connecting $F^L p$ to q. Then $F^\ell(h) \subset \gamma$, and it is a horizontal segment of length $d(F^L p, q)$ joining $F^{\ell+L}p$ to $F^\ell q$. Since $F^{\ell+L}p$ and $F^{s_1}p$ are in the same R_α–orbit, there is a power k so that $F^{\ell+L}p = R_\alpha^k F^{s_1}p$. Since s_1 is the first time that the flow line through p hits γ, we know that $k \leq 3 + \frac{L}{T_0}$. In other words, for this k we have:

$$d(R_\alpha^k F^{s_1}p, F^\ell q) = d(F^L p, q)$$

Since rotations are isometries, and $F^\ell q$ is in the R_α orbit of $F^{s_2}q$, there exists some $j \geq 0$ so that

$$d(R_\alpha^{k-j} F^{s_1}p, F^{s_2}q) = d(F^L p, q).$$

If L is large enough (depending on ϵ), we then have a contradiction of our claim that $F^{s_1}p$, $F^{s_2}q$ are B-badly approximable. $\square$

Next, observe that there exists $B > 0$ so that the rotation R_ξ for any ξ whose continued fraction expansion terminates in all 1's is B-badly approximable. Note that such ξ are dense in the reals.

Now, suppose we are given the torus ω_T. Pick a transversal γ so that the first return map for the vertical straight line flow on ω_T defines a rotation on γ, and furthermore the return time is between $\frac{1}{C|\gamma|}$ and $\frac{C}{|\gamma|}$.

By changing the preferred direction on the torus[8] we may assume that this rotation (when rescaling the transversal to have length 1) is in fact B–badly approximable by the density observation above.

These flows are now $\frac{B}{C}$-badly approximable by the Sublemma.

It remains to modify the points. Given $q_1, ..., q_k$ we choose $p_1, ..., p_k$ that are the first times the vertical flows from the q_i intersect our transversal. By Corollary 3.10 we may choose $p'_1, ..., p'_k$ in a δ neighborhood of these points and on the transversal that are $(\frac{1}{4 \cdot 4 \cdot 4})^{3(k-1)}$-badly approximable for the rotation (thought of as being on $[0,1)$). Applying the vertical flow (which is minimal) in the backwards direction, we can obtain $q'_1, .., q'_k$, in a δ neighborhood for $q_1, ..., q_k$, which are pairwise c-badly approximable for the flow for any $c < (\frac{1}{4 \cdot 4 \cdot 4})^{3(k-1)} \frac{1}{C}$ (by our choice of transversal). □

Finally, we need the following density statement for (ϵ, B)–torus good ω:

Proposition 3.12. *The set of (ϵ, B)–torus good ω with respect to q_i is dense in $\widetilde{\Omega}(S, \{q_i\})$.*

Proof. First note that the set of all ω which are lifts of Abelian differentials on tori branched over one point are dense in the space $\widetilde{\Omega}(S, \{q_i\})$. Namely, this notion is invariant under the action of $\mathrm{Diff}(S, \{q_i\})$, and the desired density is true for strata of Abelian differentials in the Hodge bundle over Teichmüller space.

The desired density now follows from Proposition 3.11, since being torus good is invariant under pullback by differentials: if ω is torus good, and ϕ is a diffeomorphism, then $\phi^*\omega$ is also torus good. □

3.5. **Point-pushing and torus good differentials.** Next, we describe constructions which allows us to modify a given (ϵ, B)–torus good ω in a simple way. In its description, we think of simple closed curves as actual maps from $S^1 = \mathbb{R}/\mathbb{Z}$ to S, and not their isotopy classes.

Definition 3.13. *We say that a simple closed curve α on S is* clean *for $\omega \in \widetilde{\Omega}(S)$ if*

(1) α is disjoint from all zeroes of ω.
(2) α is transverse to the horizontal and vertical foliation of ω.

Observe that if α is clean, it intersects every horizontal or vertical segment (in the metric given by ω) in finitely many points, since the angle to the horizontal or vertical direction is bounded away from zero on the compact curve α.

Lemma 3.14. *For every $\omega \in \widetilde{\Omega}(S)$ there is a clean α. Given any clean α there is an open neighbourhood $U_{\omega,\alpha}$ of $\omega \in \widetilde{\Omega}(S)$, so that α is clean for every $\eta \in U_{\omega,\alpha}$.*

Proof. This follows from Lemmas 4.2 and 4.7 of [LS1]. □

Definition 3.15. *Suppose that α is a differentiable, simple closed curve on S which is clean for some $\omega \in \widetilde{\Omega}(S)$. We say that α is parametrised with* constant horizontal

[8]technically, postcomposing the flat charts with a rotation matrix in $\mathrm{SL}_2(\mathbb{R})$ close to the identity.

speed *if, in the flat charts defined by* ω*, the horizontal deriviative* $\alpha(t)$ *is constant in* t*.*

Observe that any clean α admits a parametrisation with constant horizontal speed, since by the definition of clean α is nowhere vertical in the flat charts.

Proposition 3.16. *Suppose that* ω *is* (ϵ, B)*–torus good with respect to marked points* $q_1, \ldots, q_k$*, and that* α *is a simple closed curve on* S *with the following properties*

(1) α *is clean for* ω*, and parametrised with constant horizontal speed.*
(2) There are t_i *so that* $q_i = \alpha(t_i)$*.*

Then for all $B' < B$ *and any* $s \in \mathbb{R}$*, we have that* ω *is* (ϵ, B')*–torus good with respect to the marked points* $\alpha(t_1 + s), \ldots, \alpha(t_k + s)$*.*

Proof. Put $q_i' = \alpha(t_i + s)$. Since the torus is a homogeneous space we may assume without loss of generality $q_1 = q_1'$ and that therefore, by the choice of our parametrization, q_j' is a translate along a vertical leaf from q_j for all $j > 1$; Let γ' be a curve connecting q_i' to q_j' and γ be the curve connecting $q_i = q_i'$ to q_j by traversing γ' and then the vertical segment of length ℓ. Because $|g_u\gamma'| \geq |g_u\gamma| - e^{-\frac{t}{2}}\ell$ we have the the proposition. □

Next, we want to re-interpret the families of (ϵ, B)–torus good differentials constructed in Proposition 3.16 as paths in $\widetilde{\Omega}(S, \mathbf{z})$. It will be useful to describe this construction slightly more generally.

To begin, recall from e.g.[LS1, Section 4.2] that associated to a simple closed curve α there is an isotopy $D_{\alpha,t} : S \to S$ which "pushes along the curve α", i.e. $D_{\alpha,t}(\alpha(s)) = \alpha(t + s)$. Observe that such a diffeomorphism $D_{\alpha,t}$ preserves the curve α setwise. Furthermore, note that any diffeomorphism $F : S \to S$ defines by pullback a map $\widetilde{\Omega}(S) \to \widetilde{\Omega}(S)$, which induces a map

$$\widetilde{\Omega}(S, \mathbf{z}) \to \widetilde{\Omega}(S, F^{-1}(\mathbf{z}))$$

that preserves geometric properties like being (eventually) strongly ϵ–thick, or having a vertical foliation with all leaves closed.

Since $D_{\alpha,t}$ is a smoothly varying family of diffeomorphisms, for any Abelian differential ω, the assignment

$$t \mapsto D_{\alpha,t}^{-1}\omega$$

defines a continuous path $C(\alpha, \omega)$ of Abelian differentials in $\widetilde{\Omega}(S)$. Furthermore, this path depends continuously on the initial differential ω.

As α is a closed curve, the endpoint $D_{\alpha,1}^{-1}$ is actually a diffeomorphism fixing $(q_1, \ldots, q_k)$. Hence, the endpoint of $C(\alpha, \omega)$ is obtained from the initial point by pulling back the differential by that diffeomorphism. This path in $\widetilde{\Omega}(S)$ depends on the choice of the isotopy $D_{\alpha,t}$. Note that the mapping class of $S - \{q_1, \ldots, q_k\}$ defined by $D_{\alpha,1}^{-1}$ depends only on the homotopy class of α relative to the set $\{q_1, \ldots, q_k\}$ and not the actual curve. We call this mapping class a *multi-point-push*, and denote it by P_α. Observe that if α is embedded, then P_α is a product of Dehn twists about curves to either side of α. In particular, results in Section 2 proved for (multi-)Dehn twists also apply for these P_α.

We summarize some more basic properties of these paths in the following proposition.

Proposition 3.17. *Suppose that ω is an Abelian differential on S, and α is a clean simple closed curve which is parametrised with constant horizontal speed. Suppose that $q_i = \alpha(t_i)$ are points on the curve. Then there is a continuous path $c : [0,1] \to \Omega(S, \{q_1, \dots, q_k\})$, whose endpoint $c(1)$ is the image of $c(0)$ under the multi-point-push along α. Furthermore, we have*

(1) If ω is (ϵ, B)–torus good with respect to $q_1, \dots, q_k$, then any point on c is eventually strongly δ–thick with respect to $q_1, ..., q_k$.
(2) If ω has vertical foliation a weighted multicurve, then the same is true for every point on c.
(3) The path c depends (for a fixed α[9]) continuously on the initial differential ω.

Proof. (1) If we suppose that $\omega, q_1, \dots, q_k$ and α satisfy the requirements of Proposition 3.16 then, for any t, the differential $D_{\alpha,t}^{-1}\omega$ is (ϵ, B)–torus good with respect to $(q_1, \dots, q_k)$[10], by Proposition 3.16.
(2) If the vertical foliation of ω is a multicurve C, and ϕ is any diffeomorphism, then the vertical foliation of $\phi^*\omega$ is $\phi(C)$. In particular, the vertical foliation of $\phi^*\omega$ is also a multicurve. This implies the desired statement.
(3) This follows because the diffeomorphisms $D_{\alpha,t}^{-1}$ act continuously on $\widetilde{\Omega}(S)$, vary smoothly in t, and the map $\widetilde{\Omega}(S) \to \Omega(S, \{q_1, \dots, q_k\})$ is continuous. □

Lemma 3.18. *Suppose that q_i, α are as in Proposition 3.17, but that ω is an Abelian differential whose vertical foliation is a multicurve δ. Consider the path in $\Omega(S, \{q_1, \dots, q_k\})$ from Proposition 3.17. Then, only a finite number of weighted multicurves appear along this path as vertical foliations.*

Proof. The fact that every vertical foliation along the path is a (weighted) multicurve follows from Proposition 3.17. Pick regular leaves $\gamma_1, \dots, \gamma_n$, so that the vertical foliation of ω consists exactly of cylinders around the γ_i (seen as a foliation on $S - \{q_1, \dots, q_k\}$). Consider a time s so that

$$D_{\alpha,s}^{-1}(\cup\gamma_i) \cap \{q_1, \dots, q_k\} \neq \emptyset.$$

Since $D_{\alpha,s}$ preserves α, and all q_i are contained in α, for such an s there has to be a point $\alpha(t_0) \in \gamma_i$ (for a suitable i) so that $\alpha(t_0 + s) = q_j$ (for a suitable j). Hence, each such time s corresponds to one of the finitely many intersection points of α with $\cup\gamma_i$ and a choice of q_j – in particular there are only finitely many such times, say $s_j, j = 1, \dots, J$. Observe that for $t \in (s_j, s_{j+1})$ the multicurves

$$t \to D_{\alpha,t}^{-1}(\cup\gamma_i), \quad t \in (s_j, s_{j+1})$$

are then all disjoint from the points q_i by definition of the times s_j, and thus freely homotopic multicurves on the surface $S - \{q_1, \dots, q_k\}$

[9]Strictly, the curve is only fixed up to reparametrisation; for any ω one has to choose a constant horizontal speed parametrisation.

[10]The covers certifying torus goodness vary in t, by pullback under the $D_{\alpha,t}^{-1}$

The multicurve defined by the vertical foliation of $D_{\alpha,t}^{-1}\omega$ is exactly $D_{\alpha,t}^{-1}(\cup\gamma_i)$, seen as a foliation on $S-\{q_1,\dots,q_k\}$, and it is therefore constant (up to homotopy) for all $t\in(s_j,s_{j+1})$.

This shows that the multicurve defined by the vertical foliation of $D_{\alpha,t}^{-1}\omega$ can only change at $t=s_j$ for some j, and therefore only takes finitely many values. $\square$

Definition 3.19. *A* twisting pair *for an Abelian differential ω and a number of marked points $q_1,\dots,q_k$ is a pair of curves α,β so that*

(1) α,β fill S.
(2) α,β are clean.
(3) There are numbers t_i,s_i so that $q_i=\alpha(t_i)=\beta(s_i)$.

Arguing exactly as in the proof of [LS1, Lemma 4.2], we see the following.

Lemma 3.20. *Let S be a closed surface of genus at least 2, and with some number of marked points $q_1,\dots,q_k$. For every $\omega\in\widetilde{\Omega}(S)$ there is a twisting pair (α,β). Given any twisting pair (α,β) there is an open neighbourhood $U_{\omega,\alpha,\beta}$ of $\omega\in\widetilde{\Omega}(S)$, so that (α,β) is a twisting pair for every $\eta\in U_{\omega,\alpha,\beta}$.*

As an immediate consequence of Proposition 3.17 and the fact that the product of multi-point-pushes around filling curves are pseudo-Anosov, we have the following result, analogous to [LS1, Lemma 4.5].

Corollary 3.21. *Let $(q_1,\dots,q_k)$ be points on S. Suppose that ω is an Abelian differential and α,β is a twisting pair for ω. Then, for any j, define the diffeomorphism $\psi^{(j)}=P_\alpha^j(P_\alpha P_\beta^{-1})P_\alpha^{-j}$.*

Let F be the vertical foliation of ω. Then, there is a point push path

$$P(F,\psi^{(j)}F)$$

joining F to $\psi^{(j)}F$. If ω is (ϵ,B)–torus good with respect to $(q_1,\dots,q_k)$, then any point on $P(F,\psi^{(j)}F)$ is a cobounded foliation. If F is a multicurve, then there is a finite set of multicurves F_i so that every point on $P(F,\psi^{(j)}F)$ consists of a weighted multicurve homotopic to one of the F_i (with varying weights).

Proof. The idea is to apply Proposition 3.17 $2j+2$ times to join ω to $\Psi^{(j)}\omega$ (where $\Psi^{(j)}$ is a diffeomorphism defining the multi-point-push $\psi^{(j)}$), and then obtain P as the associated path of vertical foliations. To make this precise, denote by P_α,P_β point pushing diffeomorphisms around α,β, and denote by $C(\eta,P_*\eta)$ the path of Abelian differentials guaranteed by applying Proposition 3.17. We now form the concatenated path

$$\begin{aligned}C:=C(\omega,P_\alpha\omega)*C(P_\alpha\omega,P_\alpha^2\omega)*\dots\\ C(P_\alpha^j\omega,P_\alpha^{j+1}\omega)*C(P_\alpha^{j+1}\omega,P_\alpha^{j+1}P_\beta^{-1}\omega)*\\ C(P_\alpha^{j+1}P_\beta^{-1}\omega,P_\alpha^{j+1}P_\beta^{-1}P_\alpha^{-1}\omega)*\dots\\ C(P_\alpha^{j+1}P_\beta^{-1}P_\alpha^{-j+1}\omega,P_\alpha^{j+1}P_\beta^{-1}P_\alpha^{-j}\omega).\end{aligned}$$

Taking the vertical foliations, we then obtain a path

$$
\begin{aligned}
P := P(F, P_\alpha F) * P(P_\alpha F, P_\alpha^2 F) * \dots \\
P(P_\alpha^j F, P_\alpha^{j+1} F) * P(P_\alpha^{j+1} F, P_\alpha^{j+1} P_\beta^{-1} F) * \\
P(P_\alpha^{j+1} P_\beta^{-1} F, P_\alpha^{j+1} P_\beta^{-1} P_\alpha^{-1} F) * \dots \\
P(P_\alpha^{j+1} P_\beta^{-1} P_\alpha^{-j+1} F, P_\alpha^{j+1} P_\beta^{-1} P_\alpha^{-j} F)
\end{aligned}
\tag{10}
$$

of foliations joining the vertical foliation F of ω to $\psi^{(j)}(F)$. Proposition 3.17, if ω was (ϵ, B)–torus good, the same is true for any point on the path C, hence P consists of cobounded foliations. If F was a multicurve, the claim follows from Lemma 3.18. □

Corollary 3.22. *Suppose that $q_i, \omega, \alpha, \beta$ and $\psi^{(j)}$ are as in Corollary 3.21, and suppose that ω is (ϵ, B)–torus good. Then the concatenation*

$$P(F, \psi^{(j)} F) * \psi^{(j)} P(F, \psi^{(j)} F) * \left(\psi^{(j)}\right)^2 P(F, \psi^{(j)} F) * \dots$$

extends to a continuous path of cobounded foliations connecting F to the stable foliation of $\psi^{(j)}$.

Proof. First observe that $P(F, \psi^{(j)} F)$ is disjoint from the unstable foliation of $\psi^{(j)}$ for all j. Namely, the unstable foliation of $\psi^{(j)}$ has an angle-π singularity, since it is a point-pushing map (compare [LS1, Lemma 2.2]), whereas F (and any point push of it) as a lift of a foliation on a torus has no such singularities. Now, the corollary is an immediate consequence of the fact that pseudo-Anosov maps act on $\mathcal{PMF}$ with north-south dynamics. □

Proposition 3.23. *Let $(q_1, \dots, q_k)$ be points on S. Suppose that ω is an Abelian differential whose vertical foliation is a multicurve, and let α, β be a twisting pair for ω. Then, for any j, define the mapping class $\psi^{(j)} = P_\alpha^j (P_\alpha P_\beta^{-1}) P_\alpha^{-j}$.*

(1) There is a constant $C = C(\omega, \alpha, \beta) > 0$, so that the union of the sets of multicurves appearing in paths $P(F, \psi^{(j)} F)$ from Corollary 3.21 (over all j) has diameter at most C in the curve graph.

(2) If ω_n is a sequence of Abelian differentials converging to ω, with vertical foliations F_n, then the paths $P(F_n, \psi^{(j)} F_n)$ converge to $P(F, \psi^{(j)} F)$.

Proof. (1) We inductively consider the terms used in the proof of Equation (10) in Corollary 3.21. Let δ be one of the curves in the multicurve F. By Lemma 3.18 only finitely many curves appear in $P(F, P_\alpha F)$; call that set of curves G_0. In the next terms

$$P(P_\alpha^i F, P_\alpha^{i+1} F) = P_\alpha^i P(F, P_\alpha F)$$

the curves that appear are the images of G_0 under powers of a Dehn multitwist, and as these act on the curve graph by isometries with fixed points (*elliptically*), all curves that appear are contained in a C_0–neighbourhood of δ. The curves appearing in the next two terms:

$$P(P_\alpha^j F, P_\alpha^{j+1} F) * P(P_\alpha^{j+1} F, P_\alpha^{j+1} P_\beta^{-1} F) = P_\alpha^j (P(F, P_\alpha F) * P(P_\alpha F, P_\alpha P_\beta^{-1} F))$$

are images under P_α^j of the (finitely many) curves appearing in $P(F, P_\alpha F) * P(P_\alpha F, P_\alpha P_\beta^{-1} F)$, and are therefore also contained in some C_1–neighbourhood of δ. Finally, in the terms of the third type

$$P(P_\alpha^{j+1} P_\beta^{-1} P_\alpha^{-i} F, P_\alpha^{j+1} P_\beta^{-1} P_\alpha^{-i-1} F) = P_\alpha^{j+1} P_\beta^{-1} P_\alpha^{-i} P(F, P_\alpha^{-1} F)$$

we argue similarly. The path $P(F, P_\alpha^{-1} F)$ involves finitely many curves, which remain in a C_3–neighbourhood of δ by application of any power P_α^{-i}. Let $C_4 = 2C_3 + d(\delta, P_\alpha^{-1} P_\beta \delta)$. Then the image of the C_3–neighbourhood around δ is mapped by the pseudo-Anosov $P_\alpha^{-1} P_\beta$ into the C_4–neighbourhood around δ.

Finally, letting $C_5 = 2C_4 + d(\delta, \alpha)$, the C_4-neighbourhood around δ is sent to a C_5–neighbourhood around δ by the application of any further power of P_α^j. Hence, C_5 has the desired property.

(2) This is a consequence of the fact that diffeomorphisms act continuously on the space of Abelian differentials; compare Proposition 3.17 (3).

□

4. Paths in the punctured case, and Islands of point-pushes

We now come to the main technical connectivity result for (ϵ, B)–torus good foliations. Let S be a surface, and fix a finite set of curves $\mathbf{z} \neq \emptyset$.

Suppose (ϵ, B) are given so that there is a (ϵ, B)–torus good ω on S with respect to $\mathbf{z}$. We begin by defining $\mathcal{TG}$ to be the set of all vertical foliations on S of all ω which are (ϵ, B)–torus good. Since being (ϵ, B)–torus good is invariant under diffeomorphisms preserving $\mathbf{z}$, $\mathcal{TG}$ is a mapping-class-group invariant set, and thus dense in $\mathcal{F}(S, \mathbf{z})$.

We begin with the following easy lemma, which is completely analogous to the argument used to prove Theorem 1.1 for $\mathbf{z} \neq \emptyset$ in [LS1, Section 4.4].

Lemma 4.1. *Suppose that $F, F' \in \mathcal{TG}$ are arbitrary. Then, there are*

(1) A finite number of simple closed curves $\alpha_i, \beta_i, i = 1, \ldots, N-1$,
(2) numbers $\epsilon, B > 0$,
(3) Abelian differentials $\omega_1, \ldots, \omega_N$,
(4) Abelian differentials $\hat\omega_1, \ldots, \hat\omega_N$,
(5) and multicurves $\delta_1, \ldots, \delta_N$,

so that the following hold:

i) The vertical foliation of ω_1 is F, and the vertical foliation of ω_N is F'.
ii) The curves (α_i, β_i) are a twisting pair for both $\omega_i, \hat\omega_i$ and $\omega_{i+1}, \hat\omega_{i+1}$.
iii) All ω_i are (B, ϵ)–torus good, given by pullbacks of ω_T^i along a cover $p_i : S \to T$,
iv) all $\hat\omega_i$ are pullbacks of Abelian differentials $\hat\omega_T^i$ whose vertical foliation is a single cylinder.
v) the multicurves δ_i are the core curves of the vertical cylinders of the $\hat\omega_i$.

Proof. Let $F, F' \in \mathcal{TG}$ be given. By definition of $\mathcal{TG}$, they are vertical foliations of (ϵ, B)–torus good Abelian differentials ω_1, ω_N. Note that these satisfy i) and iii) by choice. Choose a path $\gamma(t)$ from ω_1 to ω_N in $\widetilde{\Omega}(S)$. For every $\gamma(t)$ there is a twisting pair $\alpha(t), \beta(t)$ for $\gamma(t)$ by Lemma 3.20. In fact, by the same lemma, there is a small open neighbourhood U_t so that the curves $\alpha(t), \beta(t)$ are twisting pairs for all differentials in U_t. By compactness of the path γ, a finite number $U_1, \ldots, U_N$ of such neighborhoods suffice to cover the path γ. We let $\omega_i \subset U_i \cap U_{i-1}$ be (ϵ, B)–torus good (which is possible since (ϵ, B)–torus good differentials are dense), and $p_i : S \to T$ the defining covering. This implies the existence of the desired objects (1) through (3) with properties i) through iii).

It remains to show the existence of Abelian differentials and curves as in (4),(5) with properties iv) and v). Let ω_T^i be the differential so that ω_i is the pullback of ω_T^i by p_i. As the U_i are open, and cylinder directions are dense, there is a differential $\hat{\omega}_T^i$ on T which has vertical direction a simple closed curve $\delta_i' \subset T$, and so that the pullback $p_i^*\hat{\omega}_T^i = \hat{\omega}_i$ is also contained in $U_i \cap U_{i-1}$. By the definition of that set, (α_i, β_i) and $(\alpha_{i-1}, \beta_{i-1})$ are then twisting pairs for $\hat{\omega}_i$, so they satisfy iii). Furthermore, the vertical foliation of $\hat{\omega}_i$ is the multicurve $p_i^{-1}(\delta_i') = \delta_i$. Hence, $\hat{\omega}_i, \delta_i$ satisfy properties iv) and v). □

Using the output of Lemma 4.1, we can construct paths between torus good foliations in the following way.

Definition 4.2. *Let $\omega_i, (\alpha_i, \beta_i), p_i, \omega_T^i$ be as in Lemma 4.1. For each i choose a number K_i and a corresponding mapping class*

$$\psi_i^{(K_i)} = P_{\alpha_i}^{K_i}(P_{\alpha_i}P_{\beta_i}^{-1})P_{\alpha_i}^{-K_i}$$

which we call the peak pseudo-Anosovs*, and for each $i = 2, \ldots, N-1$ choose a cobounded foliation F_i which is a lift of a foliation under p_i, which we call the* base foliations*, so that the (α_i, β_i) are a twisting pair for that lift. Put $F = F_1, F' = F_N$.*

The associated push-and-peak-path *is then the path γ obtained as a concatenation*

$$\gamma = \gamma_1^+ * \overline{\gamma_2^-} * \gamma_2^+ \cdots * * \gamma_{N-1}^+ * \overline{\gamma_N^-},$$

where $\bar{\cdot}$ denotes the path with opposite orientation, in the following way:

(1) γ_i^+ is the path starting in F_i, and ending in the stable foliation of $\psi_i^{(K_i)}$ which is obtained as the concatenation

$$P(F_i, \psi_i^{(K_i)}F_i) * \psi_i^{(K_i)}P(F_i, \psi_i^{(K_i)}F_i) * \left(\psi_i^{(K_i)}\right)^2 P(F_i, \psi_i^{(K_i)}F_i) * \ldots$$

of images of the point-push path $P(F_i, \psi_i^{(K_i)}F_i)$ (compare Corollary 3.21 and 3.22) under $\psi_i^{(K_i)}$.

(2) γ_i^- is the path starting in F_i, and ending in the stable foliation of $\psi_{i-1}^{(K_{i-1})}$ which is similarly obtained as the concatenation

$$P(F_i, \psi_{i-1}^{(K_{i-1})}F_i) * \psi_{i-1}^{(K_{i-1})}P(F_i, \psi_{i-1}^{(K_{i-1})}F_i) * \left(\psi_{i-1}^{(K_{i-1})}\right)^2 P(F_i, \psi_{i-1}^{(K_{i-1})}F_i) * \cdots$$

Observe that peak-and-push paths are defined using Abelian differentials, but really depend only on the choice of suitable $F_1, ..., F_N$; $\alpha_1, \beta_1, ..., \alpha_N, \beta_N$ and $K_1, ..., K_N$.

Corollary 4.3. *Any two points in $\mathcal{TG} \subset \mathcal{PMF}$ can be joined by a path of cobounded foliations.*

Proof. Let $F, F' \in \mathcal{TG}$ be given. Apply Lemma 4.1 to obtain $\omega_i, (\alpha_i, \beta_i), p_i, \omega_i^T$. Construct the push-and-peak path as in Definition 4.2. First observe that by using

Corollary 3.22, we see that this is indeed a continous path of cobounded foliations. It joins F to F' by construction. □

We now use the machinery developed in Section 2.2 in order to contract suitable point-push-paths into small neighbourhoods of uniquely ergodic foliations. We briefly recall the setup from that section. Namely, suppose that (τ_n) is a full splitting sequence in the direction of a uniquely ergodic foliation F, and let (f_m, k_m) be an associated Mod-sequence.

Recall from Section 2 that there are (nested) neighbourhoods $U_n(F, \tau)$, so that

$$\bigcap_n U_n(\tau, F) = \{F\}$$

and finitely many "model neighbourhoods" $\mathcal{U}^{(k)}$, so that

$$f_n\left(\mathcal{U}^{(k(n))}\right) = U_n(\tau, F).$$

The following theorem is concerned with finding paths P which connect two points in a model neighbourhood $\mathcal{U}^{(k)}$, and which are also moved by the f_n into smaller and smaller neighbourhoods of F (even though the path P may leave $\mathcal{U}^{(k)}$!).

Theorem 4.4. *Suppose that (τ_n) is a full splitting sequence in the direction of a uniquely ergodic foliation F, and let (f_m, k_m) be an associated* Mod*-sequence.*

Fix an essential type k, and let $F, F' \in \mathcal{TG} \cap \mathcal{U}^{(k)}$ be two foliations defined by torus good Abelian differentials ω, ω'. Furthermore let δ, δ' be lifts of simple closed curves on the base tori. Assume that

- $(*)$**:** *$\mathcal{U}^{(k)}$ contains every foliation which is a lift of the torus covers defined by ω, ω'.*

Then for any n there is a number m_0 with the following property. For any $m > m_0$ with $k_m = k$ there is a peak-and-push path γ connecting F to F', so that $f_m\gamma$ is completely contained in $U_n(\tau, F)$.

Without property $(*)$ the conclusion remains true for F, F' which are sufficiently close (depending on m) to the curves δ, δ'.

Proof. We begin by noting that due to property $(*)$, the initial segment γ_1^+ and terminal segment γ_N^- are automatically contained in $\mathcal{U}^{(k)}$, independent of all other choices. Hence, for any $m > n$ with $k_m = k$, the images of these segments under f_m are contained in $U_n(\tau, F)$, by Equation (3) of the associated sequence. If $(*)$ does not hold, we will argue for the initial/terminal segment exactly as below.

We will now explain how to construct the path segments γ_i^+ of the push-and-peak-path; the segments γ_i^- will be constructed analogously. Whenever a constant K_i is chosen, it needs to be chosen to be large enough for the construction of both γ_i^+ and γ_{i+1}^-.

Let $\omega_i, (\alpha_i, \beta_i), \delta_i$ be the objects guaranteed by Lemma 4.1 applied to F, F'. Consider the point-pushing pseudo-Anosov map

$$\psi_i^{(K_i)} = P_{\alpha_i}^{K_i} P_{\alpha_i} P_{\beta_i}^{-1} P_{\alpha_i}^{-K_i}.$$

By Proposition 3.23 there are numbers C_i so that for any choice of the numbers K_i in the construction of the peak-and-push paths, every point on the paths $P(\delta_i, \psi_i^{(K_i)}\delta_i)$

corresponds to a multicurve which is contained in the C_i–neighbourhood of δ_i in the curve graph. Let G_i be the set of all multicurves appearing on such paths. The (finite) union

$$G = \bigcup_{i=1}^{N} G_i$$

also has finite diameter. We can therefore choose a number d large enough so that for all i, every curve in G has distance at most d from α_i.

By increasing d, we may also assume that (for any choice of powers K_i in the push-and-peak-paths), the quasi-axes of $\psi_i^{(K_i)}$ pass within distance d of α_i as well. Indeed, if ρ is a quasi-axis for $\psi_i^{(0)}$ then $P_{\alpha_i}^{K_i}\rho$ is a quasi axis for $\psi_i^{(K_i)}$, and the claim follows since P_{α_i} fixes α_i.

Apply Proposition 2.20 with this d to $P_{\alpha_i}P_{\beta_i}^{-1}$ as the pseudo-Anosov, and $\mathcal{V} = U_n(\tau, F)$ as the neighbourhood and any curve in G as the curve β_0 for every i to get a constant $N = N_i$. Let m_0 be the maximum of these constants.

Let now $m > m_0$ be given. Then Proposition 2.20 yields[11] that if we choose the powers K_i in the definition of $\psi_i^{(K_i)}$ large enough, the images of the point-pushing paths $\left(\psi_i^{(K_i)}\right)^j P(\delta_i, \psi_i^{(K_i)}\delta_i), \left(\psi_{i-1}^{(K_{i-1})}\right)^j P(\delta_i, \psi_{i-1}^{(K_{i-1})}\delta_i)$ under f_m are contained in $U_n(\tau, F)$ for all j.

As we let $K_i \to \infty$, the stable foliation of $\psi_i^{(K_i)}$ converges to α_i. Hence, we can choose numbers K_i large enough, so that the stable foliation of $\psi_i^{(K_i)}$ is sent into $U_n(\tau, F)$ by f_m (in addition to the previous constraints).

Since the pseudo-Anosov $\psi_i^{(K_i)}$ acts on $\mathcal{PMF}$ with north-south-dynamics, and the (compact) path $P(\delta_i, \psi_i^{(K_i)}\delta_i)$ does not intersect the unstable foliation of $\psi_i^{(K_i)}$ (as the path consists only of multicurves), there is a number $\epsilon > 0$ and $J > 0$ so that the ϵ–neighbourhood of $P(\delta_i, \psi_i^{(K_i)}\delta_i)$ is mapped into $U_n(\tau, F)$ by $f_m\left(\psi_i^{(K_i)}\right)^j$ for all $j > J$.

By continuity of the maps $f_m\psi_i^{(K_i)}, ..., f_m\left(\psi_i^{(K_i)}\right)^J$, we may therefore choose F_i close enough to δ_i so that in fact the path $P(F_i, \psi_i^{(K_i)}F_i)$ is contained in the ϵ–neighbourhood of $P(\delta_i, \psi_i^{(K_i)}\delta_i)$, and therefore

$$f_m\left(\psi_i^{(K_i)}\right)^j P(F_i, \psi_i^{(K_i)}F_i)$$

is contained in $U_n(\tau, F)$ for all $j > J$.

By the continuity of the maps $\psi_i^{(K_i)}, \left(\psi_i^{(K_i)}\right)^2, ..., \left(\psi_i^{(K_i)}\right)^J$ we can choose the foliation F_i even closer to δ_i, to ensure that for all $j \geq 0$, since the paths $f_m\left(\psi_i^{(K_i)}\right)^j P(\delta_i, \psi_i^{(K_i)}\delta_i)$ are all contained in $U_n(\tau, F)$. Repeating the same argument for all i, and analogously for the paths for γ_i^- finishes the argument. □

[11] noting that since multi-point pushing maps are multitwists, we can apply that Proposition in this situation

The following corollary will allow us to use the paths given by the previous Theorem to build paths of foliations on surfaces without punctures.

Corollary 4.5. *Suppose that S is a surface with marked points, which is a branched cover over a torus. Suppose further that $\hat{S}$ is a closed surface and $p : \hat{S} \to S$ is a properly branched cover, with branching set $\mathbf{z}$ equal to the marked points of S. Suppose that τ_n is a splitting sequence of train tracks on $\hat{S}$ in the direction of a uniquely ergodic foliation $\hat{E}$. Let $f_1, \dots$ be an associated* Mod*-sequence.*

Fix an essential type k, and let $\hat{F} = p^{-1}(F), \hat{F}' = p^{-1}(F') \in \mathcal{U}^{(k)}$ be lifts under p of torus good foliations F, F' on $(S, \mathbf{z})$, defined by Abelian differentials ω, ω', and let δ, δ' be lifts of simple closed curves on the base tori. Assume that

> $(*)$**:** *$\mathcal{U}^{(k)}$ contains every lift under p of a foliation on S which is a lift of the torus covers defined by ω, ω'.*

Then for any n there is a number m_0 with the following property. For any $m > m_0$ with $k_m = k$ there is a peak-and-push path γ connecting F to F', which lifts under p to a path $\hat{\gamma}$ of cobounded foliations, and so that $f_m\hat{\gamma}$ is completely contained in $U_n(\tau, F)$.

Without property $(*)$ the conclusion remains true for $\hat{F}, \hat{F}'$ which are sufficiently close (depending on τ) to lifts $\hat{\delta}, \hat{\delta}'$ of the curves δ, δ'.

Proof. This follows exactly like the previous proof, using that the lifting map $\mathcal{PMF}(S) \to \mathcal{PMF}(\hat{S})$ is continuous. □

As an application of Theorem 4.4, we can now prove the main theorem in the case of punctured surfaces.

Theorem 4.6. *Suppose that Σ is a surface of genus $g \geq 2$ and with $p \geq 3$ punctures. Then the set of uniquely ergodic foliations on Σ is path-connected.*

To prove the theorem, the main step is to show that one can connect an arbitrary uniquely ergodic foliation F to a torus good foliation. In order to do this, we use the connection to splitting sequences described in Section 2.

To this end, let τ be a maximal train track carrying F, and τ_s a full splitting sequence in direction of F. We let (f_n, k_n) be an associated Mod-sequence. First, we need the following statement, purely about the model neighbourhoods.

Lemma 4.7. *Given any k there is a torus cover $p_k : \Sigma \to T$, so that the lift of every foliation from T via p_k is contained in $\mathcal{U}^{(k)}$.*

Proof. Let $p : \Sigma \to T$ be any branched torus cover, and let $L \subset \mathcal{PMF}$ be the set of all lifts of foliations on T via p. Precomposing the cover p by a mapping class φ^{-1} replaces L by $\varphi(L)$.

Choose a pseudo-Anosov φ whose attracting foliation is contained in the (open) set $\mathcal{U}^{(k)}$, and whose repelling foliation is not contained in L. As pseudo-Anosovs act on $\mathcal{PMF}$ with north-south dynamics, there is a power N so that $\varphi^N(L) \subset \mathcal{U}^{(k)}$, which shows the existence of the desired cover. □

From now on, we fix for each k covers p_k as in Lemma 4.7. Furthermore we choose, once and for all, torus good foliations $F^{(k)} \in \mathcal{U}^{(k)}$ which are defined by these covers p_k.

Recall that the associated Mod-sequence has the property that

$$U_s(\tau, F) = f_s(\mathcal{U}^{(k_s)}).$$

In particular, the (torus good) foliations $f_s\left(F^{(k_s)}\right)$ converge to F. Our strategy will be to find paths γ_s of cobounded foliations which connect $f_s\left(F^{(k_s)}\right)$ to $f_{s+1}\left(F^{(k_{s+1})}\right)$, so that the concatenated paths

$$c_n = \gamma_1 * \gamma_2 * \cdots \gamma_n$$

converge, as $n \to \infty$, to a path connecting the torus good foliation $f_1\left(F^{(k_1)}\right)$ to F.

Recall from Lemma 2.6 that there is a finite set M of mapping classes, so that for all n we have

$$f_n^{-1} f_{n+1} \in M.$$

The following corollary of Theorem 4.4 is what makes our construction of paths work:

Corollary 4.8. *Given any n, there is a number m with the following property: if $s > m$, then there is a path γ_s with the following properties:*

(1) γ_s joins $f_s\left(F^{(k_s)}\right)$ to $f_{s+1}\left(F^{(k_{s+1})}\right)$,
(2) γ_s consists only of cobounded foliations, and
(3) $\gamma_s \subset U_n(\tau, F)$.

Proof. Since we only make a claim about large s, we may assume without loss of generality that every type k_s for $s > m$ is essential.

We will then find γ_s as

$$\gamma_s = f_s \iota_s.$$

In order to satisfy (1), the path ι_s needs to join $F^{(k_s)}$ to $f_s^{-1} f_{s+1}\left(F^{(k_{s+1})}\right)$. Note that by the second claim of Lemma 2.6 (Equation (5)) we have

$$f_s^{-1} f_{s+1}\left(\mathcal{U}^{(k_{s+1})}\right) \subset \mathcal{U}^{(k_s)}$$

and therefore we have that

$$F^{(k_s)}, f_s^{-1} f_{s+1}\left(F^{(k_{s+1})}\right) \in \mathcal{U}^{(k_s)}.$$

In fact, as the foliations $F^{(k_s)}$ are defined by the covers from Lemma 4.7, the foliations $F = F^{(k_s)}, F' = f_s^{-1} f_{s+1}\left(F^{(k_{s+1})}\right)$ are defined by Abelian differentials ω, ω' which satisfy condition $(*)$ in Theorem 4.4 by the comment right after the proof of Lemma 4.7.

Hence, for any essential type k and s with $k_s = k$, we can apply Theorem 4.4 to $U_n(\tau, F)$ and pairs of foliations $(F^{(k)}, f_s^{-1} f_{s+1}\left(F^{(k_{s+1})}\right))$, to obtain thresholds $m_0(k, F^{(k)}, f_s^{-1} f_{s+1}\left(F^{(k_{s+1})}\right))$. Note that since there are finitely many $F^{(i)}$ and for all s, $f_s^{-1} f_{s+1} \in M$ for the finite set M from Lemma 2.6, there is a number

$$m = \max m_0(k, F^{(k)}, f_s^{-1} f_{s+1}\left(F^{(k_{s+1})}\right)).$$

We claim that this has the desired property. Namely, suppose that $s > m$. Then, let $k = k_s$ be the type of the index s. By our choice of m the foliations $F^{(k_s)}, f_s^{-1} f_{s+1}\left(F^{(k_{s+1})}\right)$ and the number s then satisfy the prerequisites of Theorem 4.4, and we can choose ι_s to be the path guaranteed by that theorem. Since peak-and-push-paths consist only of cobounded foliations, and this property is invariant under the mapping class group, $f_s \iota_s$ then satisfies (1) and (2). Property (3) is directly guaranteed by Theorem 4.4. $\square$

Proof of Theorem 4.6. In order to show the theorem, in light of Corollary 4.3 it suffices to show that any uniquely ergodic foliation F can be joined to a torus good foliation. We will do this by using the construction outlined above.

Namely, apply Corollary 4.8 for every n to get a sequence m_n of threshold indices. We may assume without loss of generality that m_n is increasing in n. For $s \leq m_1$, choose γ_s to be any path of cobounded foliations connecting $f_s\left(F^{(k_s)}\right)$ to $f_{s+1}\left(F^{(k_{s+1})}\right)$ (which is possible by Corollary 4.3). For $m_{n+1} \geq s > m_n$, let γ_s be the result of applying Corollary 4.8. We then have that $\gamma_s \subset U_n(\tau, F)$ for $m_{n+1} \geq s > m_n$.

Consider now the paths
$$c_r = \gamma_1 * \cdots * \gamma_r,$$
and note that they join the torus good foliation $f_1\left(F^{(k_1)}\right)$ to $f_{r+1}\left(F^{(k_{r+1})}\right)$. For any $s < r$, let
$$i_{s,r} = \gamma_{s+1} * \cdots * \gamma_r,$$
so that
$$c_r = c_s * i_{s,r}.$$
By our construction of the γ_s, we have that for any n there is some m_n, so that for all $r > s > m_n$:
$$i_{s,r} \subset U_n(\tau, F)$$
As by Corollary 2.4 we have that
$$\bigcap_n U_n(\tau, F) = \{F\},$$
this shows that since $c_r \subset U_n$ for all $r > m_n$, the infinite concatenation
$$c_\infty = c_1 * c_2 * \cdots * c_n * \cdots$$
extends to a continuous path with endpoints $f_1\left(F^{(k_1)}\right), F$, finishing the proof. $\square$

5. Paths in the closed case, and Islands of branched covers

Theorem 5.1. *Suppose that Σ is a closed surface of genus $g \geq 5$. Then the set of uniquely ergodic foliations on Σ is path-connected.*

To prove this theorem, we want to run the strategy of the proof of Theorem 4.6, with the addition of using branched covers to lift paths from punctured to closed surfaces.

The first ingredient is the following theorem, which follows from the methods developed in [LS1].

Proposition 5.2. *Suppose that $g \geq 5$. Then there is an involution σ of the closed surface Σ_g with the following properties.*

i) Σ_g/σ is a surface of genus at least 2 with several marked points.

ii) For any conjugate $\hat{\sigma}$ of σ in the mapping class group there is a sequence σ_i so that

$$\sigma = \sigma_1, \ldots, \sigma_n = \hat{\sigma},$$

and for any i the group $G_i = \langle \sigma_i, \sigma_{i+1} \rangle$ is a finite group so that Σ_g / G_i is a torus with four marked points. In that case we also say that $\sigma, \hat{\sigma}$ are a good pair*.*

In the proof we need the notion of *Humphries generators* for the mapping class group. We refer the reader to [FM1, Chapter 4] for a detailed discussion, and only recall the definition for convenience. Namely, a Humphries generating set for the mapping class group of a genus g surface consists of Dehn twists about curves[12] $\alpha_i, i = 1, \ldots, 2g+1$ so that

- $\alpha_1, \ldots, \alpha_{2g}$ form a chain, i.e. α_i, α_j intersect in one point if $|i-j| = 1$, and are disjoint otherwise.
- α_{2g+1} is disjoint from all α_i except α_4, which it intersects in a single point.

The crucial result [FM1, Theorem 4.14] is that Dehn twists about any such set of curves generate the mapping class group.

Proof of Proposition 5.2. When g is even this is [LS1, Theorem 5.3]. The case of odd genus is a fairly straightforward modification which is below.

The strategy is as follows. We show that for $f_1, \ldots, f_n$ a suitably chosen generating set for $\mathrm{Mod}(\Sigma_g)$ and σ a suitably chosen involution we have that $\sigma, f_i \sigma f_i^{-1}$ are a good pair. Since whenever σ, σ' are good pair, $g\sigma g^{-1}$ and $g\sigma' g^{-1}$ are as well, we have that by induction of the word length in $f_1, \ldots, f_n$, σ can be joined to $f\sigma f^{-1}$ for any mapping class f.

To construct σ and σ', we use the following setup (compare Figure 1). We realise

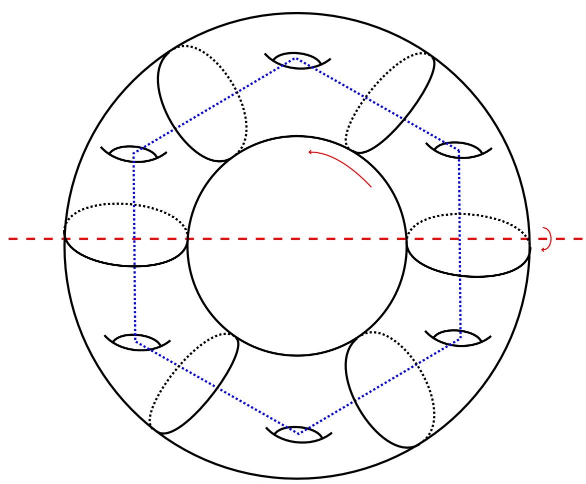

FIGURE 1. The setup for Proposition 5.2: realising the dihedral group action

[12]In the terminology of [FM1, Theorem 4.14], also referring to [FM1, Figure 4.5], the curves $\alpha_1, \ldots, \alpha_{2g}$ are the curves $m_1, a_1, c_1, a_2, c_2, \ldots, c_{g-1}, a_g$, and the curve α_{2g+1} is the curve m_2.

the surface S of genus $2k+1$ as a union

$$S = \bigcup_{i=0}^{2k-1} H_i$$

where each H_i is a torus with two boundary components, and the two boundaries of H_i are glued to H_{i+1}, H_{i-1} in a ring (compare Figure 1). Denote by $\delta_0, \ldots, \delta_{2k-1}$ the boundary curves of the H_i, so that $\partial H_i = \delta_i \cup \delta_{i+1}$. The dihedral group of order $4k$ embeds into the mapping class group of S, generated by an order $2k$ element r and an order 2 element σ. We have that $r(H_i) = H_{i+1}, r(\delta_i) = \delta_{i+1}$ (where indices are taken mod $2k$), and σ can be described in the following way: the curves δ_0, δ_k cut S into two subsurfaces S_+, S_-, each of which has genus $(g-1)/2$ and has two boundary components. The involution σ will exchange S^+ and S^- and fix both boundary components of S^+ setwise.

Intuitively, we imagine S as a symmetric, thickened $2k$-gon in three-space, with a torus in each corner. The element r then rotates the $2k$-gon by π/k around its center, while σ rotates by π about an axis through δ_0, δ_k (compare Figure 1).

We then define $\sigma' = r\sigma r^{-1}$. We claim that $\Sigma_g/\langle \sigma, \sigma' \rangle$ is a torus with four marked points. Indeed, $\langle \sigma, \sigma' \rangle$ contains r^2 (recall that σ, r generate a dihedral group), and thus

$$H_0 \cup H_1 \to \Sigma_g/\langle \sigma, \sigma' \rangle$$

is already surjective. Since σ' exchanges H_0 and H_1, even

$$H_0 \to \Sigma_g/\langle \sigma, \sigma' \rangle$$

is already surjective. In fact, $\Sigma_g/\langle \sigma, \sigma' \rangle$ is obtained from H_0 by identifying two halves of δ_0 with each other (via the action of σ) and identifying two halves of δ_1 with each other (via the action of σ'). This shows that $\Sigma_g/\langle \sigma, \sigma' \rangle$ is indeed a torus with four marked points (coming from the fixed points of σ, σ' in H_0).

Next, we claim that there are simple closed curves α_i with the following properties:

a) Dehn twists about the α_i form a (Humphries) generating set for the mapping class group of Σ_g.
b) Each α_i is either contained in one of the $S^\pm$, or is invariant under σ.
c) If $\alpha_i \subset S^\pm$, then it is nonseparating in that subsurface.
d) There is one α_{j_0} which is contained in S^- and which is invariant under σ'.

That such a set of curves exists is an exercise using Figure 2.

Now, from property c) we get the following:

$$\forall \alpha_i \text{ not invariant under } \sigma \quad \exists \phi_i \in \mathrm{Mcg}(\Sigma_g) : [\phi_i, \sigma] = 1, \phi_i(\alpha_{j_0}) = \alpha_i. \tag{11}$$

Namely, suppose first that $\alpha_i \subset S^-$. Then, since both α_i, α_{j_0} are nonseparating in S^-, there is a mapping class f of S^- fixing ∂S^- which sends α_{j_0} to α_i. Extend f to a mapping class ϕ_i of S by setting it to be $\sigma f \sigma$ on S^+. This has the desired property. In the case where $\alpha_i \subset S^+$, we start with f which sends $\sigma\alpha_i$ to α_{j_0} as above, and let ϕ_i be σf on S^- and $f\sigma$ on S^+.

We claim that for any of the Humphries generators $T = T_{\alpha_i}$ we can connect σ to $T\sigma T^{-1}$ with a path as in ii) of the statement of the Proposition.

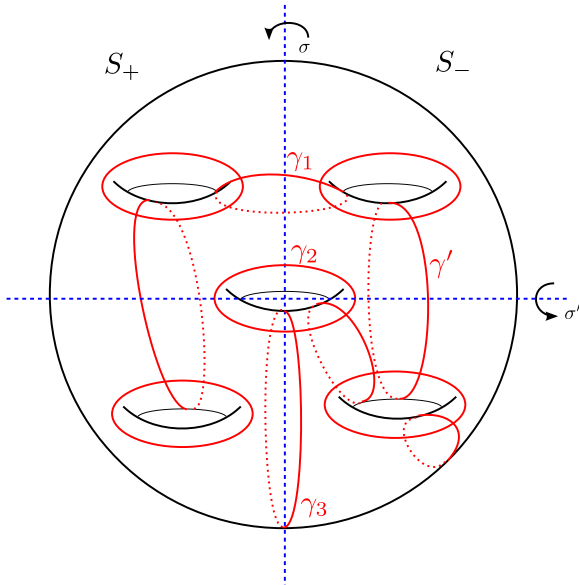

FIGURE 2. The setup for Proposition 5.2: $\gamma_1, \gamma_2, \gamma_3$ are the curves which are not in $S^\pm$ and they are invariant under σ. The curve γ' is invariant under σ'.

For twists about curves α_i which are invariant under σ there is nothing to show, as such twists commute with σ, and therefore the trivial path connects σ and $T_{\alpha_i}\sigma T_{\alpha_i}^{-1} = \sigma$. If α_i is not invariant, let ϕ_i be the mapping class guaranteed by (11). We claim that

$$\sigma_1 = \sigma,$$

$$\sigma_2 = \phi_i\sigma'\phi_i^{-1},$$

$$\sigma_3 = T_{\alpha_i}\sigma T_{\alpha_i}^{-1}$$

is a path as desired. To begin with, note that

$$G_1 = \langle \sigma, \phi_i\sigma'\phi_i^{-1}\rangle = \langle \phi_i\sigma\phi_i^{-1}, \phi_i\sigma'\phi_i^{-1}\rangle = \phi_i\langle\sigma,\sigma'\rangle\phi_i^{-1},$$

since ϕ_i commutes with σ. As by assumption σ, σ' is a good pair, G_1 is a group as desired.

Next, observe that

$$\phi_i T_{\alpha_{j_0}}\phi_i^{-1} = T_{\phi_i\alpha_{j_0}} = T_{\alpha_i}$$

and therefore

$$[T_{\alpha_i}, \phi_i\sigma'\phi_i^{-1}] = [\phi_i T_{\alpha_{j_0}}\phi_i^{-1}, \phi_i\sigma'\phi_i^{-1}] = \phi_i[T_{\alpha_{j_0}}, \sigma']\phi_i^{-1} = 1,$$

since σ' preserves α_{j_0} and therefore commutes with the Dehn twist about α_{j_0}. As G_1 is generated by a good pair, so is

$$T_{\alpha_{j_0}}G_1T_{\alpha_{j_0}}^{-1} = \langle T_{\alpha_{j_0}}\sigma T_{\alpha_{j_0}}^{-1}, T_{\alpha_{j_0}}\phi_i\sigma'\phi_i^{-1}T_{\alpha_{j_0}}^{-1}\rangle = \langle T_{\alpha_{j_0}}\sigma T_{\alpha_{j_0}}^{-1}, \phi_i\sigma'\phi_i^{-1}\rangle = \langle\sigma_3,\sigma_2\rangle.$$

Hence, $\sigma_1, \sigma_2, \sigma_3$ is indeed a path as desired. $\square$

For the remainder of this section, we fix σ to be as in the conclusion of Proposition 5.2. Say that a foliation F is *lifted torus good*, if F is the lift of a torus good foliation on $\Sigma_g/\hat{\sigma}$ for $\hat{\sigma}$ a conjugate of σ in $\mathrm{Mod}(S)$ (possibly by the identity).

The following will replace Lemma 4.1.

Lemma 5.3. *Suppose that F, F' are lifted torus good. Then there are*

(1) Involutions $\sigma_1, \ldots, \sigma_N$, which are conjugate to σ,
(2) Abelian differentials $\omega_i, i = 1, \ldots, N$ on Σ_g,

so that the following hold:

i) For any i, the group $\langle \sigma_i, \sigma_{i+1} \rangle$ is finite and $T_i = \Sigma_g/\langle \sigma_i, \sigma_{i+1} \rangle$ is a torus with four marked points.
ii) The differential ω_i is a lift of a torus good differential on the torus T_i (with marked points).

Proof. Suppose that F is a lift of a foliation on Σ_g/σ and F' is a lift of a foliation on Σ_g/σ'. Apply Proposition 5.2 to σ, σ' to find the involutions σ_i with property i). The differentials ω_1, ω_N are chosen to be the ones defining F, F'; the other ω_i can be chosen as arbitrary lifts of torus good differentials on T_i. □

Finally, the following will replace Theorem 4.4.

Theorem 5.4. *Suppose that (τ_n) is a full splitting sequence in the direction of a uniquely ergodic foliation F, and let f_m be an associated* Mod*-sequence.*

Fix an essential type k and let $E, E' \in \mathcal{U}^{(k)}$ be two lifted torus good foliations lifted from covers $\Sigma_g/\sigma, \Sigma_g/\sigma'$. Assume that

$(*)$: *$\mathcal{U}^{(k)}$ contains every foliation which is a lift of the cover defined by $\Sigma_g/\sigma, \Sigma_g/\sigma'$.*

Then for any n there is a number m_0 with the following property. Suppose that $m > m_0$ and that $k_m = k$. Then there is an path γ connecting F to F', so that $f_m\gamma$ is completely contained in $U_n(\tau, F)$, and consists only of cobounded foliations.

Proof. Suppose that E, E' are given as in the theorem. First, apply Lemma 5.3 to obtain a sequence of involutions $\sigma_1, ..., \sigma_N$. We now have two sequences of covers

$$p_i : \Sigma_g \to \Sigma_g/\sigma_i$$

and

$$t_i : \Sigma_g \to \Sigma_g/\langle \sigma_i, \sigma_{i+1} \rangle$$

which are compatible in the sense that t_i factors through both p_i and p_{i+1}:

$$\begin{tikzcd}
& \Sigma_g \arrow[dl, "p_i"'] \arrow[dr, "p_{i+1}"] \arrow[dd, "t_i"] & \\
\Sigma_g/\sigma_i \arrow[dr, dashed] & & \Sigma_g/\sigma_{i+1} \arrow[dl, dashed] \\
& \Sigma_g/\langle \sigma_i, \sigma_{i+1} \rangle &
\end{tikzcd}$$

Now for each i, let δ_i be a lift of a simple closed curve on the four times punctured torus $\Sigma_g/\langle\sigma_i,\sigma_{i+1}\rangle$ by the map t_i, and let μ_i be a lift of a simple closed curve from Σ_g/σ_i by the map p_i. We will next construct lifted torus good foliations B_i, I_j^+, I_j^-, and the desired path as a concatenation

$$\gamma = \gamma_1^+ * \overline{\gamma_2^-} * \gamma_2^0 * \gamma_2^+ * \cdots * \overline{\gamma_{N-1}^-} * \gamma_{N-1}^0 * \gamma_{N-1}^+ * \overline{\gamma_N^-}$$

where $\bar{\cdot}$ denotes the path with opposite orientation, and

(1) γ_j^+ is a path starting in I_j^+, and ending in B_{j+1},
(2) γ_j^0 is a path starting in I_j^-, and ending in I_j^+,
(3) γ_j^- is a path starting in I_j^-, and ending in B_{j-1},

All γ_j^* will be produced by using Corollary 4.5.

Namely, put $I_0^+ = E, I_N^- = E'$, and choose B_i, I_j^+, I_j^- lifted torus good foliations (for the covers p_i, t_j, t_j respectively) close enough to δ_i, μ_j so that Corollary 4.5. (to Theorem 4.4), in the version without $(*)$ except at the endpoints applies in all three cases mentioned above. This can e.g. be achieved by starting with any lifted torus good foliations, and Dehn twisting them about the δ_i, μ_j.

Now, Corollary 4.5 yields, for a pair i, j, paths $\gamma_j^+, \gamma_j^0, \gamma_j^-$ which are contracted into $U_n(\tau, F)$ by f_m for $m > m_0(i, j, *)$, $* \in \{-, +, 0\}$. By choosing m_0 to be the largest of those finitely many $m_0(i, j, *)$, the desired property holds. □

With this in place, we can finish the proof of Theorem 5.1 exactly as in the case of Theorem 4.6.

In fact, the proof shows something a little bit stronger, which will be useful to show local path-connectivity.

Corollary 5.5. *Suppose τ is a train track carrying a uniquely ergodic foliation F, and suppose that τ_n is a splitting sequence in the direction of F. Then for any n there is a $m = m(\tau, n, F)$ with the following property. If E is any uniquely ergodic foliation contained in $U_m(\tau, F)$, then there is a path of uniquely ergodic laminations connecting F to E completely contained in $U_n(\tau, F)$.*

Proof. In the case of a punctured surface, i.e. Theorem 4.6, all bounds on m come from applying Proposition 2.17 or 2.20 within the proof of Theorem 4.4. By Lemma 2.21 we can choose these bounds to be independent of the actual foliation guiding the splitting sequence, as long as the foliation is contained in $U_k(\tau, F)$ for k large enough. The bounds in Theorem 5.1 come from applying Theorem 4.4 and its Corollary 4.5, and so the same is true there. □

6. Local Path Connectivity

In this section, we improve the Theorem from the last section to the following.

Theorem 6.1. *If $g \geq 5$ or $g \geq 2, p \geq 3$, the set of uniquely ergodic foliations on $S_{g,p}$ is locally path-connected.*

Given a uniquely ergodic foliation F and a full splitting sequence $(\tau_s)_s$ towards F. For any n, we let $m(\tau, n, F)$ the number guaranteed by Corollary 5.5. Define

$$\mathcal{G}_n(\tau, F) = U_{m(\tau,n,F)}(\tau, F).$$

Corollary 5.5 guarantees that for any $F' \in \mathcal{G}_n(\tau, F)$ there exists a path $P_{F,F'}$ of cobounded foliations joining F to F', which is contained in $U_n(\tau, F)$.

Let $\hat{\mathcal{G}}_n(\tau, F)$ be the intersection of $\mathcal{G}_n(\tau, F)$ with the set of uniquely ergodic foliations. For any point $p \in P_{F,F'}$, we can define a neighbourhood

$$\mathcal{G}_n(p, \tau)$$

as above, i.e. with the property that p can be joined to any $p' \in \mathcal{G}_n(p, \tau)$ by a path of cobounded foliations which is contained in $U_n(\tau, F)$.

Also observe that

$$U_i(\tau, p) \subset U_i(\tau, F) \tag{12}$$

for all $i \leq n$.

Define

$$N^{(1)}(F, n) := \bigcup_{F' \in \hat{\mathcal{G}}_n(\tau, F)} \bigcup_{p \in P_{F,F'}} \mathcal{G}_n(p, \tau).$$

Inductively, put

$$N^{(r+1)}(F, n) = \bigcup_{p \in N^{(r)}(F,n)} N^{(1)}(p, n).$$

Also observe that we have $N^{(r)}(F, n) \subset U_n(\tau, F)$ by Equation (12), whenever $F' \in \mathcal{G}_n(\tau, F)$.

Proposition 6.2. *Any point in $N^{(r)}(F, n)$ is connected to F by a path of uniquely ergodic foliations, which is contained in in $N^{(r+1)}(F, n)$.*

Proof. We prove this by induction.

Base case: If $p \in N^{(1)}(F, n)$ then we can connect it to F by a path in $N^{(2)}(F, n)$.

Proof. If $p \in P_{F,F'}$ this is obvious. Otherwise $p \in \mathcal{G}_n(\hat{p}, \tau)$ for some $\hat{p} \in P_{F,F'}$ where $F' \in \mathcal{G}_n(\tau, F)$. By definition we have that there exists a path of cobounded foliations contained in $\mathcal{G}_n(\hat{p}, \tau)$ connecting p to $\hat{p}$. Concatenating this with the segment of $P_{F,F'}$ connecting $\hat{p}$ to F connects p to F. The first segment of the path is in $N^{(1)}(\hat{p}, n)$ and so the whole thing is in $N^{(2)}(F, n)$. □

Inductive step: Assume $p \in N^{(r)}(F, n)$ and that any point in $N^{(r-1)}(F, n)$ is connected to F by a path of cobounded foliations in $N^{(r)}(F, n)$. We will now show that p is path connected by cobounded foliations in $N^{(r+1)}(F, n)$ to F.

Proof. Because $p \in N^{(r)}(F, n)$ we know (by definition of $N^{(r+1)}$) $p \in N^{(1)}(\hat{p}, n)$ for some $\hat{p} \in N^{(r-1)}(F, n)$. By the base case of induction applied to $\hat{p}$ it is connected to $\hat{p}$ by a path in $N^{(2)}(\hat{p}, n) = \cup_{p' \in N^{(1)}(\hat{p},n)} N^{(1)}(p', n)$. This is contained in $\cup_{p' \in N^{(r)}(F,n)} N^{(1)}(p', n) = N^{(r+1)}(F, n)$. To finish linking p to F we use our inductive assumption to link $\hat{p}$ to F by a path in $N^{(r-1+1)}(F, n)$. □

□

Corollary 6.3. *For any uniquely ergodic foliation F, and any n, the set*

$$\left(\bigcup_{r \geq 1} N^{(r)}(F, n) \right) \cap \mathcal{UE}$$

is an open neighbourhood of F in $\mathcal{UE}$, which is path-connected and contained in $U_n(\tau, F)$.

Proof. The set is open as a union of open subsets. It is contained in $U_n(\tau, F)$, since all $N^{(r)}(F, n)$ have this property. It is path-connected by Proposition 6.2. □

By Corollary 2.4, the $U_n(\tau, F)$ are a basis for neighbourhoods of F in $\mathcal{UE}$, and thus this finishes the proof of Theorem 6.1.

References

[Ahl] Lars V. Ahlfors. *Conformal invariants.* AMS Chelsea Publishing, Providence, RI, 2010. Topics in geometric function theory, Reprint of the 1973 original, With a foreword by Peter Duren, F. W. Gehring and Brad Osgood.

[BH] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences].* Springer-Verlag, Berlin, 1999.

[CH] Jon Chaika and Sebastian Hensel. The set of uniquely ergodic interval exchange transformations is path-connected. *Ergodic Theory and Dynamical Systems* (2017), 1–46.

[FM1] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series.* Princeton University Press, Princeton, NJ, 2012.

[FM2] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.* **6**(2002), 91–152.

[Gab] David Gabai. Almost filling laminations and the connectivity of ending lamination space. *Geom. Topol.* **13**(2009), 1017–1041.

[Ham1] Ursula Hamenstädt. Train tracks and the Gromov boundary of the complex of curves. In *Spaces of Kleinian groups*, volume 329 of *London Math. Soc. Lecture Note Ser.*, pages 187–207. Cambridge Univ. Press, Cambridge, 2006.

[Ham2] Ursula Hamenstädt. Geometry of the mapping class groups. I. Boundary amenability. *Invent. Math.* **175**(2009), 545–609.

[Kla] Erica Klarreich. The boundary at infinity of the curve complex and the relative Teichmüller space. Preprint, 1999.

[LS1] Christopher J. Leininger and Saul Schleimer. Connectivity of the space of ending laminations. *Duke Math. J.* **150**(2009), 533–575.

[LS2] Christopher J. Leininger and Saul Schleimer. Hyperbolic spaces in Teichmüller spaces. *J. Eur. Math. Soc. (JEMS)* **16**(2014), 2669–2692.

[Lev] Gilbert Levitt. Foliations and laminations on hyperbolic surfaces. *Topology* **22**(1983), 119–135.

[Mas] Howard Masur. Interval exchange transformations and measured foliations. *Ann. of Math. (2)* **115**(1982), 169–200.

[MS] Howard Masur and John Smillie. Hausdorff dimension of sets of nonergodic measured foliations. *Ann. of Math. (2)* **134**(1991), 455–543.

[MM1] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.* **138**(1999), 103–149.

[MM2] Howard A. Masur and Yair N. Minsky. Quasiconvexity in the curve complex. In *In the tradition of Ahlfors and Bers, III*, volume 355 of *Contemp. Math.*, pages 309–320. Amer. Math. Soc., Providence, RI, 2004.

[Mos] Lee Mosher. Train track expansions of measured foliations. Preprint, 2003.

[PH] R. C. Penner and J. L. Harer. *Combinatorics of train tracks*, volume 125 of *Annals of Mathematics Studies.* Princeton University Press, Princeton, NJ, 1992.

[Sch] Wolfgang M. Schmidt. On badly approximable numbers and certain games. *Trans. Amer. Math. Soc.* **123**(1966), 178–199.

[Tse] Jimmy Tseng. Badly approximable affine forms and Schmidt games. *J. Number Theory* **129**(2009), 3020–3025.

[Vee] William A. Veech. The Teichmüller geodesic flow. *Ann. of Math. (2)* **124**(1986), 441–530.