

quatrième série - tome 47 fascicule 2 mars-avril 2014

*ANNALES
SCIENTIFIQUES
de
L'ÉCOLE
NORMALE
SUPÉRIEURE*

Yair MINSKY & Barak WEISS

*Cohomology classes represented by measured foliations, and Mahler's
question for interval exchanges*

SOCIÉTÉ MATHÉMATIQUE DE FRANCE

COHOMOLOGY CLASSES REPRESENTED BY MEASURED FOLIATIONS, AND MAHLER'S QUESTION FOR INTERVAL EXCHANGES

BY YAIR MINSKY AND BARAK WEISS

ABSTRACT. – A translation structure on (S, Σ) gives rise to two transverse measured foliations \mathcal{F}, \mathcal{G} on S with singularities in Σ , and by integration, to a pair of relative cohomology classes $[\mathcal{F}], [\mathcal{G}] \in H^1(S, \Sigma; \mathbb{R})$. Given a measured foliation \mathcal{F} , we characterize the set of cohomology classes \mathbf{b} for which there is a measured foliation \mathcal{G} as above with $\mathbf{b} = [\mathcal{G}]$. This extends previous results of Thurston [19] and Sullivan [18].

We apply this to two problems: unique ergodicity of interval exchanges and flows on the moduli space of translation surfaces. For a fixed permutation $\sigma \in \mathcal{S}_d$, the space \mathbb{R}_+^d parametrizes the interval exchanges on d intervals with permutation σ . We describe lines ℓ in \mathbb{R}_+^d such that almost every point in ℓ is uniquely ergodic. We also show that for $\sigma(i) = d+1-i$, for almost every $s > 0$, the interval exchange transformation corresponding to σ and (s, s^2, \dots, s^d) is uniquely ergodic. As another application we show that when $k = |\Sigma| \geq 2$, the operation of “moving the singularities horizontally” is globally well-defined. We prove that there is a well-defined action of the group $B \times \mathbb{R}^{k-1}$ on the set of translation surfaces of type (S, Σ) without horizontal saddle connections. Here $B \subset \mathrm{SL}(2, \mathbb{R})$ is the subgroup of upper triangular matrices.

RÉSUMÉ. – Une structure de translation sur une surface marquée (S, Σ) donne lieu à deux feuilletages mesurés \mathcal{F}, \mathcal{G} sur S à singularités dans Σ et, par intégration, à un couple de classes de cohomologie relative $[\mathcal{F}], [\mathcal{G}] \in H^1(S, \Sigma; \mathbb{R})$. Étant donné un feuilletage mesuré \mathcal{F} , nous caractérisons l'ensemble des classes de cohomologie \mathbf{b} pour lesquelles il existe un feuilletage mesuré \mathcal{G} comme ci-dessus tel que $\mathbf{b} = [\mathcal{G}]$. Cela généralise des résultats antérieurs de Thurston [19] et Sullivan [18].

Nous appliquons ce résultat à deux problèmes : l'unique ergodicité des échanges d'intervalles et les flots sur l'espace des modules des surfaces de translation. Étant donnée une permutation $\sigma \in \mathcal{S}_d$, l'ensemble \mathbb{R}_+^d paramètre les échanges d'intervalles sur d intervalles de permutation associée σ . Nous décrivons les droites ℓ de \mathbb{R}_+^d dont presque tout point est uniquement ergodique. Nous démontrons aussi que si σ est donnée par $\sigma(i) = d+1-i$, pour presque tout $s > 0$, l'échange d'intervalles correspondant à σ et à (s, s^2, \dots, s^d) est uniquement ergodique. Une autre application est que lorsque $k = |\Sigma| \geq 2$, l'opération consistant à « déplacer horizontalement les singularités » est bien définie. En notant B le sous-groupe des matrices triangulaires supérieures de $\mathrm{SL}(2, \mathbb{R})$, nous prouvons qu'il y a une action bien définie du groupe $B \times \mathbb{R}^{k-1}$ sur l'ensemble des surfaces de translation de type (S, Σ) sans connexion horizontale.

1. Introduction

1.1. Motivating questions and nonsensical pictures

To introduce the problems discussed in this paper, consider some pictures. Suppose that $\mathbf{a} = (a_1, \dots, a_d)$ is a vector with positive entries, $I = [0, \sum a_i]$ is an interval, σ is a permutation on d symbols, and $\mathcal{T} = \mathcal{T}_\sigma(\mathbf{a}) : I \rightarrow I$ is the interval exchange obtained by cutting up I into segments of lengths a_i and permuting them according to σ . A fruitful technique for studying the dynamical properties of \mathcal{T} is to consider it as the return map to a transverse segment along the vertical foliation in a translation surface, i.e., a union of polygons with edges glued pairwise by translations. See Figure 1.1 for an example with one polygon; note that the interval exchange determines the horizontal coordinates of vertices, but there are many possible choices of the vertical coordinates.




FIGURE 1.1. Masur's construction ([9], according to Masur this was suggested by the referee of [9]): an interval exchange embedded in a one-polygon translation surface.

Given a translation surface q with a transversal, one may deform it by applying the horocycle flow, i.e., deforming the polygon with the linear map

$$(1) \quad h_s = \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix}.$$

The return map to a transversal in $h_s q$ depends on s , so we get a one-parameter family \mathcal{T}_s of interval exchange transformations (Figure 1.2). For sufficiently small s , one has $\mathcal{T}_s = \mathcal{T}_\sigma(\mathbf{a}(s))$, where $\mathbf{a}(s) = \mathbf{a} + s\mathbf{b}$ is a line segment, whose derivative $\mathbf{b} = (b_1, \dots, b_d)$ is determined by the *heights* of the vertices of the polygon. We will consider an inverse problem: given a line segment $\mathbf{a}(s) = \mathbf{a} + s\mathbf{b}$, does there exist a translation surface q such that for all sufficiently small s , $\mathcal{T}_\sigma(\mathbf{a}(s))$ is the return map along vertical leaves to a transverse segment in $h_s q$? Attempting to interpret this question with pictures, we see that some choices of \mathbf{b} lead to a translation surface while others lead to nonsensical pictures—see Figure 1.3. The solution to this problem is given by Theorem 5.3.

Now consider a translation surface q with two singularities. We may consider the operation of moving one singularity horizontally with respect to the other. That is, at time s , the line segments joining one singularity to the other are made longer by s , while line segments




FIGURE 1.2. The horocycle action gives a linear one parameter family of interval exchanges




FIGURE 1.3. The choice $\mathbf{b} = (2, 1, -1, -2)$ (left) gives a translation surface but what about $\mathbf{b} = (0, 3, -1, -2)$?

joining a singularity to itself are unchanged. For small values of s , one obtains a new translation surface q_s by examining the picture. But for large values of s , some of the segments in the figure cross each other and it is not clear whether the operation defined above gives rise to a well-defined surface. Our Theorem 11.2 shows that the operation of moving the zeroes is well-defined for all values of s , provided one rules out the obvious obstruction that two singularities connected by a horizontal segment collide.




FIGURE 1.4. The singularity \circ is moved to the right with respect to \bullet , and the picture becomes nonsensical.

1.2. Main geometrical result

Let S be a compact oriented surface of genus $g \geq 2$ and $\Sigma \subset S$ a finite subset. A *translation surface structure* on (S, Σ) is an atlas of charts into the plane, whose domains cover $S \setminus \Sigma$,

and such that the transition maps are translations. Such structures arise naturally in complex analysis and in the study of interval exchange transformations and polygonal billiards and have been the subject of intensive research, see the recent surveys [12, 25].

Several geometric structures on the plane can be pulled back to $S \setminus \Sigma$ via the atlas, among them the foliations of the plane by horizontal and vertical lines. We call the resulting oriented foliations of $S \setminus \Sigma$ the *horizontal and vertical foliation* respectively. Each can be completed to a singular foliation on S , with a pronged *singularity* at each point of Σ . Label the points of Σ by ξ_1, \dots, ξ_k and fix natural numbers r_1, \dots, r_k . We say that the translation surface is of *type* $\mathbf{r} = (r_1, \dots, r_k)$ if the horizontal and vertical foliations have a $2(r_j + 1)$ -pronged singularity at each ξ_j .

By pulling back dy (resp. dx) from the plane, the horizontal (vertical) foliation arising from a translation surface structure is equipped with a *transverse measure*, i.e., a family of measures on each arc transverse to the foliation which is invariant under holonomy along leaves. We will call an oriented singular foliation on S , with singularities in Σ , which is equipped with a transverse measure a *measured foliation on* (S, Σ) . We caution the reader that we deviate from the convention adopted in several papers on this subject, by considering the number and orders of singularities as part of the structure of a measured foliation; we call these the *type* of the foliation. In other words, we do not consider two measured foliations which differ by a Whitehead move to be the same.

Integrating the transverse measures gives rise to two well-defined cohomology classes in the relative cohomology group $H^1(S, \Sigma; \mathbb{R})$. That is we obtain a map

$$\text{hol} : \{\text{translation surfaces on } (S, \Sigma)\} \rightarrow (H^1(S, \Sigma; \mathbb{R}))^2.$$

This map is a local homeomorphism and serves to give coordinate charts to the set of translation surfaces (see §2.1 for more details), but it is not globally injective. For example, precomposing with a homeomorphism which acts trivially on homology may change a marked translation surface structure but does not change its image in under hol ; see [14] for more examples. On the other hand it is not hard to see that the pair of horizontal and vertical measured foliations associated to a translation surface uniquely determine it, and hence the question arises of reconstructing the translation surface from just the cohomological data recorded by hol . Our main theorems give results in this direction.

To state them we define the set of (*relative*) *cycles carried by* \mathcal{F} , denoted $H_+^{\mathcal{F}}$, to be the image in $H_1(S, \Sigma; \mathbb{R})$ of all (possibly atomic) transverse measures on \mathcal{F} (see §2.5).

THEOREM 1.1. – *Suppose \mathcal{F} is a measured foliation on (S, Σ) , and $\mathbf{b} \in H^1(S, \Sigma; \mathbb{R})$. Then the following are equivalent:*

1. *There is a measured foliation \mathcal{G} on (S, Σ) , everywhere transverse to \mathcal{F} and of the same type, such that \mathcal{G} represents \mathbf{b} .*
2. *Possibly after replacing \mathbf{b} with $-\mathbf{b}$, for any $\delta \in H_+^{\mathcal{F}}$, $\mathbf{b} \cdot \delta > 0$ ($\mathbf{b} \cdot \delta$ denotes the natural pairing of homology and cohomology).*

After proving Theorem 1.1 we learned from F. Bonahon that it has a long and interesting history. Similar result were proved by Thurston [19] in the context of train tracks and measured laminations, and by Sullivan [18] in a very general context involving foliations. Bonahon neglected to mention his own contribution [1]. Our result is a ‘relative version’ in that we

control the type of the foliation, and need to be careful with the singularities. This explains why our definition of $H_+^{\mathcal{F}}$ includes the relative cycles carried by critical leaves of \mathcal{F} . The proof we present here is close to the one given by Thurston.

The arguments proving Theorem 1.1 imply the following stronger statement (see §2 for detailed definitions):

THEOREM 1.2. – *Given an oriented topological singular foliation \mathcal{F} on (S, Σ) , let $\tilde{\mathcal{H}}(\mathcal{F})$ denote the set of marked translation surfaces whose vertical foliation is topologically equivalent to \mathcal{F} . Let $\mathbb{A}(\mathcal{F}) \subset H^1(S, \Sigma; \mathbb{R})$ denote the set of cohomology classes corresponding to non-atomic transverse measures on \mathcal{F} . Let $\mathbb{B}(\mathcal{F}) \subset H^1(S, \Sigma; \mathbb{R})$ denote the set of cohomology classes that pair positively with all elements of $H_+^{\mathcal{F}} \subset H_1(S, \Sigma)$. Then*

$$\text{hol} : \tilde{\mathcal{H}}(\mathcal{F}) \rightarrow \mathbb{A}(\mathcal{F}) \times \mathbb{B}(\mathcal{F})$$

is a homeomorphism.

1.3. Applications

We present two applications of Theorem 1.1. The first concerns the generic properties of interval exchange transformations. Let σ be a permutation on d symbols and let \mathbb{R}_+^d be the vectors $\mathbf{a} = (a_1, \dots, a_d)$ for which $a_i > 0, i = 1, \dots, d$. The pair \mathbf{a}, σ determines an interval exchange $\mathcal{T}_\sigma(\mathbf{a})$ by subdividing the interval $I_{\mathbf{a}} = [0, \sum a_i)$ into d subintervals of lengths a_i , which are permuted according to σ . In 1982 Masur [9] and Veech [21] confirmed a conjecture of Keane, proving (assuming that σ is irreducible, i.e., $\sigma(\{1, \dots, k\}) \neq \{1, \dots, k\}$ for $k < d$) that almost every \mathbf{a} , with respect to Lebesgue measure on \mathbb{R}_+^d , is *uniquely ergodic*, i.e., the only invariant measure for $\mathcal{T}_\sigma(\mathbf{a})$ is Lebesgue measure. On the other hand Masur and Smillie [11] showed that the set of non-uniquely ergodic interval exchanges is large in the sense of Hausdorff dimension. A basic problem is to understand the finer structure of the set of non-uniquely ergodic interval exchanges. Specifically, motivated by analogous developments in diophantine approximation, we will ask: For which curves $\ell \subset \mathbb{R}_+^d$ is the non-uniquely ergodic set of zero measure, with respect to the natural measure on the curve? Which properties of a measure μ on \mathbb{R}_+^d guarantee that μ -a.e. \mathbf{a} is uniquely ergodic?

In this paper we obtain several results in this direction, involving three ingredients: a curve in \mathbb{R}_+^d ; a measure supported on the curve; and a dynamical property of interval exchanges. The goal will be to understand the dynamical properties of points in the support of the measure. We state these results, and some open questions in this direction, in §6. To illustrate them we state a special case, which may be thought of as an interval exchanges analogue of a famous result of Sprindzhuk (Mahler's conjecture, see e.g., [6, §4]):

THEOREM 1.3. – *For $d \geq 2$, let*

$$(2) \quad \mathbf{a}(x) = (x, x^2, \dots, x^d)$$

and let $\sigma(i) = d + 1 - i$. Then $\mathbf{a}(x)$ is uniquely ergodic for Lebesgue a.e. $x > 0$.

Identifying two translation structures which differ by a precomposition with an orientation preserving homeomorphism of S which fixes each point of Σ we obtain the *stratum* $\mathcal{H}(\mathbf{r})$ of translation surfaces of type \mathbf{r} . There is an action of $G = \text{SL}(2, \mathbb{R})$ on $\mathcal{H}(\mathbf{r})$, and its restriction to the subgroup $\{h_s\}$ as in (1) is called the *horocycle flow*. To prove our results on

unique ergodicity we employ the strategy, introduced in [5], of lifting interval exchanges to translation surfaces, and studying the dynamics of the G -action on $\mathcal{H}(\mathbf{r})$. Specifically we use quantitative nondivergence estimates [15] for the horocycle flow. Theorem 1.1 is used to characterize the lines in \mathbb{R}_+^d which may be lifted to horocycle paths.

The second application concerns an operation of ‘moving singularities with respect to each other’ which has been discussed in the literature under various names (cf. [25, §9.6]) and which we now define. Let $\tilde{\mathcal{H}}(\mathbf{r})$ be the *stratum of marked translation surfaces of type \mathbf{r}* , i.e., two translation surface structures are equivalent if they differ by precomposition by a homeomorphism of S which fixes Σ and is isotopic to the identity relative to Σ . Integrating transverse measures as above induces a well-defined map $\tilde{\mathcal{H}}(\mathbf{r}) \rightarrow H^1(S, \Sigma; \mathbb{R}^2)$ which can be used to endow $\tilde{\mathcal{H}}(\mathbf{r})$ (resp. $\mathcal{H}(\mathbf{r})$) with the structure of an affine manifold (resp. orbifold), such that the natural map $\tilde{\mathcal{H}}(\mathbf{r}) \rightarrow \mathcal{H}(\mathbf{r})$ is an orbifold cover. We describe foliations on $\tilde{\mathcal{H}}(\mathbf{r})$ which descend to well-defined foliations on $\mathcal{H}(\mathbf{r})$. The two summands in the splitting

$$(3) \quad H^1(S, \Sigma; \mathbb{R}^2) \cong H^1(S, \Sigma; \mathbb{R}) \oplus H^1(S, \Sigma; \mathbb{R})$$

induce two foliations on $\tilde{\mathcal{H}}(\mathbf{r})$, which we call the *real foliation* and *imaginary foliation* respectively. Also, considering the exact sequence in cohomology

$$(4) \quad H^0(S; \mathbb{R}^2) \rightarrow H^0(\Sigma; \mathbb{R}^2) \rightarrow H^1(S, \Sigma; \mathbb{R}^2) \xrightarrow{\text{Res}} H^1(S; \mathbb{R}^2) \rightarrow \{0\},$$

we obtain a natural subspace $\ker \text{Res} \subset H^1(S, \Sigma; \mathbb{R}^2)$, consisting of the cohomology classes which vanish on the subspace of ‘absolute periods’ $H_1(S) \subset H_1(S, \Sigma)$. The foliation induced on $\tilde{\mathcal{H}}(\mathbf{r})$ is called the *REL foliation* or *kernel foliation*. Finally, intersecting the leaves of the real foliation with those of the REL foliation yields the *real REL foliation*. It has leaves of dimension $k - 1$ (where $k = |\Sigma|$). Two nearby translation surfaces q and q' are in the same plaque if the integrals of the flat structures along all closed curves are the same on q and q' , and if the integrals of curves joining *distinct* singularities only differ in their horizontal component. Intuitively, q' is obtained from q by fixing one singularity as a reference point and moving the other singularities horizontally. Understanding this foliation is important for the study of the dynamics of the horocycle flow. It was studied in some restricted settings in [4, 2], where it was called *Horiz*.

The leaves of the kernel foliation, and hence the real REL foliation, are equipped with a natural translation structure, modeled on the vector space $\ker \text{Res} \cong H^0(\Sigma; \mathbb{R})/H^0(S, \mathbb{R})$. One sees easily that the leaf of q is incomplete if, when moving along the leaf, a saddle connection on q is made to have length zero, i.e., if ‘singularities collide’. Using Theorems 1.1 and 1.2 we show in Theorem 11.2 that this is the only obstruction to completeness of leaves. This implies that on a large set, the leaves of real REL are the orbits of an action. More precisely, let $Q \subset \mathcal{H}$ be the set of translation surfaces with no horizontal saddle connections. This is a set of full measure which is invariant under the group B of upper triangular matrices in G . We show that it coincides with the set of complete real REL leaves. Let F denote the group $B \ltimes \mathbb{R}^{k-1}$, where B acts on \mathbb{R}^{k-1} via

$$\begin{pmatrix} a & b \\ 0 & 1/a \end{pmatrix} \cdot \vec{v} = a\vec{v}.$$

We prove:

THEOREM 1.4. – *The group F acts on \mathcal{Q} continuously and affinely, preserving the natural measure, and leaving invariant the subset of translation surfaces of area one. The action of B is the same as that obtained by restricting the G -action, and the \mathbb{R}^{k-1} -action is transitive on each real REL leaf in \mathcal{Q} .*

Note that while the F -action is continuous, \mathcal{Q} is not complete: it is the complement in \mathcal{H} of a dense countable union of proper affine submanifolds with boundary. Also note that the leaves of the real foliation or the kernel foliation are not orbits of a group action on \mathcal{H} —but see [4] for a related discussion of pseudo-group-actions.

1.4. Organization of the paper

We present the proof of Theorem 1.1 in §3 and of Theorem 1.2, in §4. We interpret these theorems in the language of interval exchanges in §5. This interpretation furnishes a link between line segments in the space of interval exchanges, and horocycle paths in a corresponding stratum of translation surfaces: it turns out that the line segments which may be lifted to horocycle paths form a cone in the tangent space to interval exchange space, and this cone can be explicitly described in terms of a bilinear form studied by Veech. We begin §6 with a brief discussion of Mahler's question in diophantine approximation, and the question it motivates for interval exchanges. We then state in detail our results for generic properties of interval exchanges. The proofs of these results occupy §7–10. Nondivergence results for horocycles make it possible to analyze precisely the properties of interval exchanges along a line segment, in the cone of directions described in §12. To obtain information about curves we approximate them by line segments, and this requires the quantitative nondivergence results obtained in [15]. In §11 we prove our results concerning real REL. These sections may be read independently of §6–10. We conclude with a discussion which connects real REL with some of the objects encountered in §7–10.

1.5. Acknowledgements

We thank John Smillie for many valuable discussions. We thank Francis Bonahon for pointing out the connection between our Theorem 1.1 and previous work of Sullivan and Thurston. The authors were supported by BSF grant 2004149, ISF grant 584/04 and NSF grant DMS-0504019.

2. Preliminaries

In this section we recall some standard facts and set our notation. For more information we refer the reader to [12, 25] and the references therein.

2.1. Strata of translation surfaces

Let $S, \Sigma = (\xi_1, \dots, \xi_k)$, and $\mathbf{r} = (r_1, \dots, r_k)$ be as in the introduction. A *translation structure* (resp., a *marked translation structure*) of type \mathbf{r} on (S, Σ) is an equivalence class of $(U_\alpha, \varphi_\alpha)$, where:

- $(U_\alpha, \varphi_\alpha)$ is an atlas of charts for $S \setminus \Sigma$;
- the transition functions $\varphi_\beta \circ \varphi_\alpha^{-1}$ are of the form

$$(5) \quad \mathbb{R}^2 \ni \vec{x} \mapsto \vec{x} + c_{\alpha, \beta};$$

- around each $\xi_j \in \Sigma$ the charts glue together to form a cone point with cone angle $2\pi(r_j + 1)$.

By definition $(U_\alpha, \varphi_\alpha), (U'_\beta, \varphi'_\beta)$ are equivalent if there is an orientation preserving homeomorphism $h : S \rightarrow S$ (for a marked structure, isotopic to the identity via an isotopy fixing Σ), fixing all points of Σ , such that $(U_\alpha, \varphi_\alpha)$ is compatible with $(h(U'_\beta), \varphi'_\beta \circ h^{-1})$. Thus the equivalence class \mathbf{q} of a marked translation surface is a subset of that of the corresponding translation surface q , and we will say that \mathbf{q} is obtained from q by *specifying a marking*, q from \mathbf{q} by *forgetting the marking*, and write $q = \pi(\mathbf{q})$. Note that our convention is that singularities are labelled.

Pulling back dx and dy from the coordinate charts we obtain two well-defined closed 1-forms, which we can integrate along any path α on S . If α is a cycle or has endpoints in Σ (a relative cycle), then the result, which we denote by

$$\text{hol}(\alpha, \mathbf{q}) = \begin{pmatrix} x(\alpha, \mathbf{q}) \\ y(\alpha, \mathbf{q}) \end{pmatrix} \in \mathbb{R}^2,$$

depends only on the homology class of α in $H_1(S, \Sigma)$. We let $\text{hol}(\mathbf{q}) = \text{hol}(\cdot, \mathbf{q})$ be the corresponding element of $H^1(S, \Sigma; \mathbb{R}^2)$, with coordinates $x(\mathbf{q}), y(\mathbf{q})$ in $H^1(S, \Sigma; \mathbb{R})$.

A *saddle connection* for q is a straight segment which connects singularities and does not contain singularities in its interior.

The set of all (marked) translation surfaces on (S, Σ) of type \mathbf{r} is called the *stratum* of (marked) translation surface of type \mathbf{r} and is denoted by $\mathcal{H}(\mathbf{r})$ (resp. $\tilde{\mathcal{H}}(\mathbf{r})$). We have suppressed the dependence on Σ from the notation since for a given type \mathbf{r} there is an isomorphism between the corresponding set of translation surfaces on (S, Σ) and on (S, Σ') for any other finite subset $\Sigma' = (\xi'_1, \dots, \xi'_k)$.

The map $\text{hol} : \tilde{\mathcal{H}} \rightarrow H^1(S, \Sigma; \mathbb{R}^2)$ just defined gives local charts for $\tilde{\mathcal{H}}$, endowing it (resp. \mathcal{H}) with the structure of an affine manifold (resp. orbifold). To see how this works, fix a triangulation τ of S with vertices in Σ . Then $\text{hol}(\mathbf{q})$ associates a vector in the plane to each oriented edge in τ , and hence associates an oriented Euclidean triangle to each oriented triangle of τ . If all the orientations are consistent, then a translation structure with the same holonomy as \mathbf{q} can be realized explicitly by gluing the Euclidean triangles to each other. Let $\tilde{\mathcal{H}}_\tau$ be the set of all translation structures obtained in this way (we say that τ is *realized geometrically* in such a structure). Then the restriction $\text{hol} : \tilde{\mathcal{H}}_\tau \rightarrow H^1(S, \Sigma; \mathbb{R}^2)$ is injective and maps onto an open subset. Conversely every \mathbf{q} admits some geometric triangulation (e.g., a Delaunay triangulation as in [11]) and hence \mathcal{H} is covered by the \mathcal{H}_τ , and so these provide an atlas for a linear manifold structure on $\tilde{\mathcal{H}}$. We should remark that a topology on $\tilde{\mathcal{H}}$ can

be defined independently of this, by considering nearly isometric comparison maps between different translation structures, and that this topology is the same as that induced by the charts of hol.

Let $\text{Mod}(S, \Sigma)$ denote the mapping class group, i.e., the orientation preserving homeomorphisms of S fixing Σ pointwise, up to an isotopy fixing Σ . The map hol is $\text{Mod}(S, \Sigma)$ -equivariant. This means that for any $\varphi \in \text{Mod}(S, \Sigma)$, $\text{hol}(\mathbf{q} \circ \varphi) = \varphi_* \text{hol}(\mathbf{q})$, which is nothing more than the linearity of the holonomy map with respect to its first argument.

One can show that the $\text{Mod}(S, \Sigma)$ -action on $\tilde{\mathcal{H}}$ is properly discontinuous. Thus $\mathcal{H} = \tilde{\mathcal{H}}/\text{Mod}(S, \Sigma)$ is a linear orbifold and $\pi : \tilde{\mathcal{H}} \rightarrow \mathcal{H}$ is an orbifold covering map. We have

$$(6) \quad \dim \tilde{\mathcal{H}} = \dim \mathcal{H} = \dim H^1(S, \Sigma; \mathbb{R}^2) = 2(2g + k - 1).$$

The Poincaré-Hopf index theorem implies that

$$(7) \quad \sum r_j = 2g - 2.$$

There is an action of $G = \text{SL}(2, \mathbb{R})$ on \mathcal{H} and on $\tilde{\mathcal{H}}$ by post-composition on each chart in an atlas. The projection $\pi : \tilde{\mathcal{H}} \rightarrow \mathcal{H}$ is G -equivariant. The G -action is linear in the homology coordinates, namely, given a marked translation surface structure \mathbf{q} and $\gamma \in H_1(S, \Sigma)$, and given $g \in G$, we have

$$(8) \quad \text{hol}(\gamma, g\mathbf{q}) = g \cdot \text{hol}(\gamma, \mathbf{q}),$$

where on the right hand side, g acts on \mathbb{R}^2 by matrix multiplication.

We will write

$$g_t = \begin{pmatrix} e^{t/2} & 0 \\ 0 & e^{-t/2} \end{pmatrix}, \quad r_\theta = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}.$$

2.2. Interval exchange transformations

Suppose σ is a permutation on d symbols. For each

$$\mathbf{a} \in \mathbb{R}_+^d = \{(a_1, \dots, a_d) \in \mathbb{R}^d : \forall i, a_i > 0\}$$

we have an interval exchange transformation $\mathcal{T}_\sigma(\mathbf{a})$ defined by dividing the interval

$$I = I_{\mathbf{a}} = \left[0, \sum a_i\right)$$

into subintervals of lengths a_i and permuting them according to σ . It is customary to take these intervals as closed on the left and open on the right, so that the resulting map has $d - 1$ discontinuities and is left-continuous.

More precisely, set $x_0 = x'_0 = 0$ and for $i = 1, \dots, d$,

$$(9) \quad x_i = x_i(\mathbf{a}) = \sum_{j=1}^i a_j, \quad x'_i = x'_i(\mathbf{a}) = \sum_{j=1}^i a_{\sigma^{-1}(j)} = \sum_{\sigma(k) \leq i} a_k;$$

then for every $x \in I_i = [x_{i-1}, x_i)$ we have

$$(10) \quad \mathcal{T}(x) = \mathcal{T}_\sigma(\mathbf{a})(x) = x - x_{i-1} + x'_{\sigma(i)-1} = x - x_i + x'_{\sigma(i)}.$$

In particular, if (following Veech [20]) we let Q be the alternating bilinear form given by

$$(11) \quad Q(\mathbf{e}_i, \mathbf{e}_j) = \begin{cases} 1 & i > j, \sigma(i) < \sigma(j) \\ -1 & i < j, \sigma(i) > \sigma(j) \\ 0 & \text{otherwise} \end{cases}$$

where $\mathbf{e}_1, \dots, \mathbf{e}_d$ is the standard basis of \mathbb{R}^d , then

$$\mathcal{T}(x) - x = Q(\mathbf{a}, \mathbf{e}_i).$$

An interval exchange $\mathcal{T} : I \rightarrow I$ is said to be *minimal* if there are no proper closed \mathcal{T} -invariant subsets of I . We say that \mathcal{T} is *uniquely ergodic* if the only invariant measure for \mathcal{T} , up to scaling, is Lebesgue measure. We will say that $\mathbf{a} \in \mathbb{R}_+^d$ is minimal or uniquely ergodic if $\mathcal{T}_\sigma(\mathbf{a})$ is.

Below we will assume that σ is *irreducible*, i.e., there is no $k < d$ such that σ leaves the subset $\{1, \dots, k\}$ invariant, and *admissible* (in Veech's sense), see §2.3. For the questions about interval exchanges which we will study, these hypotheses on σ entail no loss of generality.

It will be helpful to consider a more general class of maps which we call *generalized interval exchanges*. Suppose J is a finite union of intervals. A generalized interval exchange $\mathcal{T} : J \rightarrow J$ is an orientation preserving piecewise isometry of J , i.e., it is a map obtained by subdividing J into finitely many subintervals and re-arranging them to obtain J . These maps are not often considered because studying their dynamics easily reduces to studying interval exchanges. However they will arise naturally in our setup.

2.3. Measured foliations, transversals, and interval exchange induced by a translation surface

Given a surface S and a finite $\Sigma \subset S$, a *singular foliation (on S with singularities in Σ)* is a foliation \mathcal{F} on $S \setminus \Sigma$ such that for any $z \in \Sigma$ there is $k = k_z \geq 3$ such that \mathcal{F} extends to form a k -pronged singularity at z . A singular foliation \mathcal{F} is orientable if there is a continuous choice of a direction on each leaf. If \mathcal{F} is orientable then k_z is even for all z . Leaves which meet the singularities are called *critical*. A *transverse measure* on a singular foliation \mathcal{F} is a family of measures defined on arcs transverse to the foliation and invariant under restriction to subsets and isotopy through transverse arcs. A *measured foliation* is a singular foliation equipped with a transverse measure, which we further require has no atoms and has full support (no open arc has measure zero). We will only consider orientable singular foliations which can be equipped with a transverse measure. This implies (see e.g., [12, Thm. 6.6]) that the surface S is decomposed into finitely many domains on each of which the foliation is either minimal (any ray is dense) or periodic (any leaf is periodic). A periodic component is also known as a *cylinder*. These components are separated by saddle connections.

Given a flat surface structure \mathbf{q} on S , pulling back via charts the vertical and horizontal foliations on \mathbb{R}^2 gives oriented singular foliations on S called the vertical and horizontal foliations, respectively. Transverse measures are defined by integrating the pullbacks of dx and dy , i.e., they correspond to the holonomies $x(\mathbf{q})$ and $y(\mathbf{q})$. Conversely, given two oriented everywhere transverse measured foliations on S with singularities in Σ , one obtains an atlas of charts as in §2.1 by integrating the measured foliation. I.e., for each $z \in S \setminus \Sigma$, a local coordinate system is obtained by taking a simply connected neighborhood $U \subset S \setminus \Sigma$ of z and defining the two coordinates of $\varphi(w) \in \mathbb{R}^2$ to be the integral of the measured foliations along

some path connecting z to w (where the orientation of the foliations is used to determine the sign of the integral). One can verify that this procedure produces an atlas with the required properties.

We will often risk confusion by using the symbol \mathcal{F} to denote both a measured foliation and the corresponding singular foliation supporting it. A singular foliation is called *minimal* if any noncritical leaf is dense, and *uniquely ergodic* if there is a unique (up to scaling) transverse measure on S which is supported on noncritical leaves. Where confusion is unavoidable we say that q is minimal or uniquely ergodic if its vertical foliation is.

At a singular point $p \in \Sigma$ with k prongs, a small neighborhood of p divides into k foliated disks, glued along singular leaves of \mathcal{F} , which we call *foliated half-disks*. Since the minimal and periodic components are separated by saddle connections, a sufficiently small foliated half-disk is either contained in a single periodic component or in a minimal component.

Now let \mathcal{F} be a singular foliation on a surface S with singularities in Σ . We will consider three kinds of transversals to \mathcal{F} .

- We define a *transverse system* to be an injective map $\gamma : J \rightarrow S$ where J is a finite union of intervals J_i , the restriction of γ to each interval J_i is a smooth embedding, the image of the interior of γ intersects every non-critical leaf of \mathcal{F} transversally, and does not intersect Σ .
- We define a *judicious curve* to be a transverse system $\gamma : J \rightarrow S$ with J connected, such that γ begins and ends at singularities, and the interior of γ intersects all leaves including critical ones.
- We say a transverse system γ is *special* if all its components are of the following types (see Figure 2.3):
 - For every foliated half-disk D of a singularity $p \in \Sigma$ which is contained in a minimal component, there is a component of γ whose interior intersects D and terminating at p . This component of γ meets Σ at only one endpoint.
 - For every periodic component (cylinder) P of \mathcal{F} , γ contains one arc crossing P and joining two singularities on opposite sides of P .

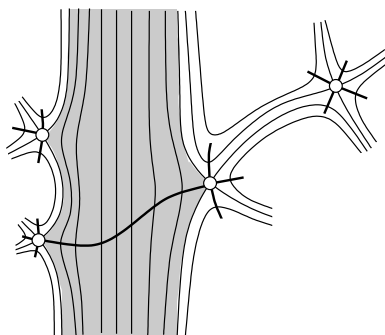


FIGURE 2.1. A *special* transverse system cuts across periodic components and into minimal components.

Note that since every non-critical leaf in a minimal component is dense, the non-cylinder edges of a special transverse system can be made as short as we like, without destroying the property that they intersect every non-critical leaf.

In each of these cases, we can parametrize points of γ using the transverse measure, and consider the first return map to γ when moving up along vertical leaves. When γ is a judicious curve, this is an interval exchange transformation which we denote by $\mathcal{T}(\mathcal{F}, \gamma)$, or by $\mathcal{T}(q, \gamma)$ when \mathcal{F} is the vertical foliation of q . Then there is a unique choice of σ and \mathbf{a} with $\mathcal{T}(\mathcal{F}, \gamma) = \mathcal{T}_\sigma(\mathbf{a})$, and with σ an irreducible admissible permutation. The corresponding number of intervals is

$$(12) \quad d = 2g + |\Sigma| - 1;$$

note that $d = \dim H^1(S, \Sigma; \mathbb{R}) = \frac{1}{2} \dim \mathcal{H}$ if $q \in \mathcal{H}$. The return map to a transverse system is a generalized interval exchange. We denote it also by $\mathcal{T}(\mathcal{F}, \gamma)$. In each of the above cases, any non-critical leaf returns to γ infinitely many times.

If γ is a transversal to \mathcal{F} then $\mathcal{T}(\mathcal{F}, \gamma)$ completely determines the transverse measure on \mathcal{F} . In particular the vertical foliation of q is uniquely ergodic (minimal) if and only if $\mathcal{T}(q, \gamma)$ is for some (any) transverse system γ .

There is an inverse construction which associates with an irreducible permutation σ a surface S of genus g , and a k -tuple \mathbf{r} satisfying (7), such that the following holds. For any $\mathbf{a} \in \mathbb{R}_+^d$ there is a translation surface structure $q \in \mathcal{H}(\mathbf{r})$, and a transversal γ on S such that $\mathcal{T}_\sigma(\mathbf{a}) = \mathcal{T}(q, \gamma)$. Variants of this construction can be found in [24, 9, 21]. Veech's admissibility condition amounts to requiring that there is no transverse arc on S for which $\mathcal{T}(q, \gamma)$ has fewer discontinuities. Fixing σ , we say that a flat structure q on S is a *lift* of \mathbf{a} if there is a judicious curve γ on S such that $\mathcal{T}(q, \gamma) = \mathcal{T}_\sigma(\mathbf{a})$. It is known that for any σ , there is a stratum \mathcal{H} such that all lifts of all \mathbf{a} lie in \mathcal{H} . We call it the *stratum corresponding to σ* .

2.4. Decomposition associated with a transverse system

Suppose \mathcal{F} is an oriented singular foliation on (S, Σ) and $\gamma : J \rightarrow S$ is a transverse system to \mathcal{F} . There is an associated cellular decomposition $\mathcal{B} = \mathcal{B}(\gamma)$ of (S, Σ) defined as follows. Let $\mathcal{T} = \mathcal{T}(\mathcal{F}, \gamma)$ be the generalized interval exchange corresponding to γ .

The 2-cells in \mathcal{B} correspond to the intervals of continuity of \mathcal{T} . For each such interval I , the corresponding cell consists of the union of interiors of leaf intervals beginning at I and ending at $\mathcal{T}(I)$. Hence it fibers over I and hence has the structure of an open topological rectangle. The boundary of a 2-cell lies in γ and in certain segments of leaves, and the union of these form the 1-skeleton. The 0-skeleton consists of points of Σ , endpoints of γ , and points of discontinuity of \mathcal{T} and \mathcal{T}^{-1} . Edges of the 1-skeleton lying on γ will be called *transverse edges* and edges lying on \mathcal{F} will be called *leaf edges*. Leaf edges inherit an orientation from \mathcal{F} and transverse edges inherit the transverse orientation induced by \mathcal{F} .

Note that opposite boundaries of a 2-cell could come from the same points in S : a particular example occurs for a special transverse system, where if there is a transverse edge crossing a cylinder, that cylinder is obtained as a single 2-cell with its bottom and top edges identified. Such a 2-cell is called a *cylinder cell*.

It is helpful to consider a *spine* for \mathcal{B} , which we denote $\chi = \chi(\gamma)$, and is composed of the 1-skeleton of \mathcal{B} together with one leaf ℓ_R for every rectangle R , traversing R from bottom to top. The spine is closely related to Thurston's train tracks; indeed, if we delete from χ the singular points Σ and the leaf edges that meet them, and collapse each element of the transversal γ to a point, we obtain a train track that 'carries \mathcal{F} ' in Thurston's sense (see [16] for details about train tracks). But note that keeping the deleted edges allows us to keep track of information relative to Σ , and in particular, to keep track of the saddle connections in \mathcal{F} .

2.5. Transverse cocycles, homology and cohomology

We now describe cycles supported on a foliation \mathcal{F} and their dual cocycles.

We will see that a transverse measure μ on \mathcal{F} defines an element $[c_\mu] \in H_1(S, \Sigma)$, expressed concretely as a cycle c_μ in the spine $\chi(\gamma)$ of a transverse system. Poincaré duality identifies $H_1(S, \Sigma)$ with $H^1(S \setminus \Sigma)$, and the dual $[d_\mu]$ of $[c_\mu]$ is represented by the cochain corresponding to integrating the measure μ .

If μ has no atoms, then in fact we obtain $[c'_\mu] \in H_1(S \setminus \Sigma)$, and its dual $[d'_\mu]$ lies in $H^1(S, \Sigma)$. The natural maps $H_1(S \setminus \Sigma) \rightarrow H_1(S, \Sigma)$ and $H^1(S, \Sigma) \rightarrow H^1(S \setminus \Sigma)$ take $[c'_\mu]$ to $[c_\mu]$ and $[d'_\mu]$ to $[d_\mu]$ respectively.

We will now describe these constructions in more detail.

Let γ be a transverse system and $\chi(\gamma)$ the spine of its associated complex $\mathcal{B}(\gamma)$ as above. Given μ we define a 1-chain on χ as follows. For each rectangle R whose bottom side is an interval κ in γ , set $\mu(R) = \mu(\text{int } \kappa)$ (using the interior is important here because of possible atoms in the boundary). For each leaf edge f of \mathcal{B} , set $\mu(\{f\})$ to be the transverse measure of μ across f (which is 0 unless the leaf f is an atom of μ). The 1-chain $x = \sum_R \mu(R)\ell_R + \sum_f \mu(\{f\})f$ may not be a cycle, but we note that invariance of μ implies that, on each component of γ , the sum of measures taken with sign (ingoing vs. outgoing) is 0, so that ∂x restricted to each component is null-homologous. Hence by 'coning off' ∂x in each component of γ we can obtain a cycle of the form:

$$(13) \quad c_\mu = \sum_{R \text{ rectangle}} \mu(R)\ell_R + \sum_{f \text{ leaf edge of } \mathcal{B}} \mu(\{f\})f + z,$$

where z is a 1-chain supported in γ such that $\partial z = -\partial x$. Invariance and additivity of μ imply that the homology class in $H_1(S, \Sigma)$ is independent of the choice of γ .

The cochain d_μ is constructed as follows: in any product neighborhood U for \mathcal{F} in $S \setminus \Sigma$, integration of μ gives a map $U \rightarrow \mathbb{R}$, constant along leaves (but discontinuous at atomic leaves). On any oriented path in U with endpoints off the atoms, the value of the cochain is obtained by mapping endpoints to \mathbb{R} and subtracting. Via subdivision this extends to a cochain defined on 1-chains whose boundary misses atomic leaves. This cochain is a cocycle via additivity and invariance of the measures, and suffices to give a cohomology class (or one may extend it to all 1-chains by a suitable chain-homotopy perturbing vertices on atomic leaves slightly).

A cycle corresponding to a transverse measure will be called a (*relative*) *cycle carried by \mathcal{F}* . The set of all (relative) cycles carried by \mathcal{F} is a convex cone in $H_1(S, \Sigma; \mathbb{R})$ which we denote by $H_+^\mathcal{F}$. Since we allow atomic measures, we can think of (positively oriented) saddle connections or closed leaves in \mathcal{F} as elements of $H_+^\mathcal{F}$.

In the case with no atoms, we note that the expression for c_μ has no terms of the form $\mu(\{f\})f$, and hence we get a cycle in $S \setminus \Sigma$. The definition of the cochain extends in that case to neighborhoods of singular points, and evaluates consistently on relative 1-chains, giving a class in $H^1(S, \Sigma)$. We let $\mathbb{A}(\mathcal{F}) \subset H^1(S, \Sigma)$ denote the classes obtained in this way.

LEMMA 2.1. – *Every element of $H_+^{\mathcal{F}}$ can be expressed as a convex combination of saddle connections and elements of $H_+^{\mathcal{F}} \cap H_1(S)$.*

Proof. – Note that $H_1(S)$ naturally embeds in $H_1(S, \Sigma)$ via the exact sequence $H_1(\Sigma) \rightarrow H_1(S) \rightarrow H_1(S, \Sigma)$.

If μ has no atoms, as observed above c_μ gives us a cycle in $S \setminus \Sigma$. The inclusion into S yields an element of $H_+^{\mathcal{F}} \cap H_1(S)$.

If μ does have atoms, they are supported on saddle connections and closed leaves. The latter are already in $H_1(S)$. We may subtract these from μ , and what remains is still a non-negative transverse measure without atoms, whose homology class therefore lies in $H_1(S)$. \square

The cycles we construct here are closely related to Schwartzman's asymptotic cycles [17], but we will not need to use this point of view explicitly.

2.6. Intersection pairing

Via Poincaré duality, the canonical pairing on $H^1(S, \Sigma) \times H_1(S, \Sigma)$ becomes the intersection pairing on $H_1(S \setminus \Sigma) \times H_1(S, \Sigma)$. In the former case we denote this pairing by $(d, c) \mapsto d(c)$, and in the latter, by $(c, c') \mapsto c \cdot c'$. Suppose \mathcal{F} and \mathcal{G} are two mutually transverse oriented singular foliations, with transverse measures μ and ν respectively. If we allow μ but not ν to have atoms, then $[c_\mu] \in H_1(S, \Sigma)$ and $[c_\nu] \in H_1(S \setminus \Sigma)$ so we have the intersection pairing

$$(14) \quad c_\nu \cdot c_\mu = d_\nu(c_\mu) = \int_S \nu \times \mu.$$

In other words we integrate the transverse measure of \mathcal{G} along the leaves of \mathcal{F} , and then integrate against the transverse measure of \mathcal{F} . (The sign of $\nu \times \mu$ should be chosen so that it is positive when the orientation of \mathcal{G} agrees with the transverse orientation of \mathcal{F}). We can see this explicitly by choosing a transversal γ for \mathcal{F} lying in the leaves of \mathcal{G} . Then the cochain representing d_ν is 0 along γ , and using the form (13) we have

$$(15) \quad d_\nu(c_\mu) = \sum_R \mu(R) \int_{\ell_R} \nu + \sum_f \mu(\{f\}) \int_f \nu.$$

2.6.1. *Judicious case.* – Now suppose γ is a judicious transversal. In this case the pairing of $H^1(S, \Sigma)$ and $H_1(S, \Sigma)$ has a concrete form which we will use in § 5.

There is a cell decomposition \mathcal{D} of S that is dual to \mathcal{B} , defined as follows. Because γ intersects all leaves, and terminates at Σ on both ends, each rectangle R of \mathcal{B} has exactly one point of Σ on each of its leaf edges. Connect these two points by a transverse arc in R and let \mathcal{D}^1 be the union of these arcs. \mathcal{D}^1 cuts S into a disk \mathcal{D}^2 , bisected by γ . Indeed, upward flow from γ encounters \mathcal{D}^1 in a sequence of edges which is the *upper boundary* of the disk, and downward flow encounters the *lower boundary* which goes through the edges of \mathcal{D}^1 in a permuted order, in fact exactly the permutation σ of the interval exchange $\mathcal{T}(\mathcal{F}, \gamma)$.

A class in $H^1(S, \Sigma)$ is determined by its values on the (oriented) edges of \mathcal{D}^1 , and in fact this gives a basis, which we can label by the intervals of continuity of $\mathcal{T}(\mathcal{F}, \gamma)$ (the condition that the sum is 0 around the boundary of the disk is satisfied automatically). The Poincaré dual basis for $H_1(S \setminus \Sigma)$ is given by the loops $\hat{\ell}_R$ obtained by joining the endpoints of ℓ_R along γ .

The pairing restricted to non-negative homology is computed by the form Q of (11). Ordering the rectangles R_1, \dots, R_d according to their bottom arcs along γ and writing $\hat{\ell}_i = \hat{\ell}_{R_i}$, we note that $\hat{\ell}_i$ and $\hat{\ell}_j$ have nonzero intersection number precisely when the order of i and j is reversed by σ (i.e., $(i - j)(\sigma(i) - \sigma(j)) < 0$), and in particular (accounting for sign),

$$(16) \quad \hat{\ell}_i \cdot \hat{\ell}_j = Q(\mathbf{e}_i, \mathbf{e}_j).$$

In other words, Q is the intersection pairing on $H_1(S \setminus \Sigma) \times H_1(S \setminus \Sigma)$ with the given basis (note that this form is degenerate, as the map $H_1(S \setminus \Sigma) \rightarrow H_1(S, \Sigma)$ has a kernel).

3. The lifting problem

In this section we prove Theorem 1.1.

Proof of Theorem 1.1. – We first explain the easy direction (1) \implies (2). Since \mathcal{F} is everywhere transverse to \mathcal{G} , and the surface is connected, reversing the orientation of \mathcal{G} if necessary we can assume that positively oriented paths in leaves of \mathcal{F} always cross \mathcal{G} from left to right. Therefore, using (14) and (15), we find that $\mathbf{b}(\delta) > 0$ for any $\delta \in \mathcal{H}_+^{\mathcal{F}}$.

Before proving the converse we indicate the idea of proof. We will consider a sequence of finer and finer cell decompositions associated to a shrinking sequence of special transverse systems. As the transversals shrink, the associated train tracks split, each one being carried by the previous one. We examine the weight that a representative of \mathbf{b} places on the vertical leaves in the cells of these decompositions (roughly speaking the branches of the associated train tracks). If any of these remain non-positive for all time, then a limiting argument produces an invariant measure on \mathcal{F} which has non-positive pairing with \mathbf{b} , a contradiction. Hence eventually all cells have positive ‘heights’ with respect to \mathbf{b} , and can be geometrically realized as rectangles. The proof is made complicated by the need to keep track of the singularities, and in particular by the appearance of cylinder cells in the decomposition.

Fix a special transverse system γ for \mathcal{F} (see § 2.3), and let $\mathcal{B} = \mathcal{B}(\gamma)$ be the corresponding cell decomposition as in § 2.4. Given a path α on S which is contained in a leaf of \mathcal{F} and begins and ends in transverse edges of \mathcal{B} , we will say that α is *parallel to saddle connections* if there is a continuous family of arcs α_s contained in leaves, where $\alpha_0 = \alpha$, α_1 is a union of saddle connections in \mathcal{F} , and the endpoints of each α_s are in γ .

We claim that for any N there is a special transverse system $\gamma_N \subset \gamma$ such that the following holds: for any leaf edge e of $\mathcal{B}_N = \mathcal{B}(\gamma_N)$, either e is parallel to saddle connections or, when moving along e from bottom to top, we return to $\gamma \setminus \gamma_N$ at least N times before returning to γ_N . Indeed, if the claim is false then for some N and any special $\gamma' \subset \gamma$ there is a leaf edge e' in \mathcal{B}_N which is not parallel to saddle connections, and starts and ends at points x, y in γ' , making at most N crossings with $\gamma \setminus \gamma'$. Now take γ'_j to be shorter and shorter special transverse subsystems of γ , denote the corresponding edge by e_j and the points by x_j, y_j . Passing to a subsequence, we find that x_j, y_j converge to points in Σ and e_j converges to a concatenation of at most N saddle connections joining these points. In particular, for large enough j , e_j is parallel to saddle connections, a contradiction proving the claim.

Since γ is special, each periodic component of \mathcal{F} consists of one 2-cell of \mathcal{B} , called a *cylinder cell*, with its top and bottom boundaries identified along an edge traversing the component, which we call a *cylinder transverse edge*. The same holds for \mathcal{B}_N , and our construction ensures that \mathcal{B} and \mathcal{B}_N contain the same cylinder cells and the same cylinder transverse edges.

Let β be a singular (relative) 1-cocycle in (S, Σ) representing \mathbf{b} . We claim that we may choose β so that it vanishes on non-cylinder transverse edges. Indeed, each such edge meets Σ only at one endpoint, so they can be deformation-retracted to Σ , and pulling back a cocycle via this retraction gives β . Note that β assigns a well-defined value to the periodic leaf edges, namely the value of \mathbf{b} on the corresponding loops.

In general β may assign non-positive heights to rectangles, and thus it does not assign any reasonable geometry to \mathcal{B} . We now claim that there is a positive N such that for any leaf edge $e \in \mathcal{B}_N$,

$$\beta(e) > 0.$$

Since a saddle connection in \mathcal{F} represents an element of $H_+^{\mathcal{F}}$, our assumption implies that $\beta(e) > 0$ for any leaf edge of \mathcal{B}_N which is a saddle connection. Moreover by construction, if e is parallel to saddle connections, then $\beta(e) = \sum \beta(e_i)$ for saddle connections e_i , so again $\beta(e) > 0$. Now suppose by contradiction that for any N we can find a leaf edge e_N in \mathcal{B}_N which is not parallel to saddle connections and such that $C_N = |e_N \cap \gamma| \geq N$ and

$$\limsup_{N \rightarrow \infty} \frac{\beta(e_N)}{C_N} \leq 0.$$

Then defining measures ν_N by

$$\nu_N(I) = \frac{|e_N \cap I|}{C_N}, \text{ where } I \subset \gamma \text{ is an interval,}$$

we can pass to a subsequence and obtain a measure μ on γ as a weak-* limit. This measure is invariant under the return map to γ (since the return map changes ν_N by at most $1/C_N$ on any subinterval), and thus defines a transverse measure on \mathcal{F} representing a class $[c_\mu] \in H_+^{\mathcal{F}}$. Moreover by construction it has no atoms and gives measure zero to the cylinder cells. We will evaluate $\beta(c_\mu)$.

For each rectangle R in \mathcal{B} , let $\theta_R \subset \gamma$ be the transverse arc on the bottom of R and ℓ_R a leaf segment going through R from bottom to top, as in §2.4. Since e_N is not parallel to

saddle connections, its intersection with each R is a union of arcs parallel to ℓ_R . Since β gives all such arcs the same values $\beta(\ell_R)$, we have

$$\beta(e_N) = \beta \left(\sum_R |e_N \cap \theta_R| \ell_R \right) = \sum_R |e_N \cap \theta_R| \beta(\ell_R).$$

By (13), since μ has no atoms we can write

$$c_\mu = \sum_R \mu(\theta_R) \ell_R + z$$

where $z \subset \gamma$. Since β vanishes along γ we have

$$\mathbf{b}(c_\mu) = \sum_R \mu(\theta_R) \beta(\ell_R) = \sum_R \lim_N \frac{|e_N \cap \theta_R|}{C_N} \beta(\ell_R) = \lim_N \frac{\beta(e_N)}{C_N} \leq 0.$$

This contradicts the hypothesis, proving the claim.

The claim implies that the topological rectangles in \mathcal{B}_N can be given a compatible Euclidean structure, using the transverse measure of \mathcal{F} and β to measure respectively the horizontal and vertical components of all relevant edges. Note that all non-cylinder cells become metric rectangles, and the cylinder cells become metric parallelograms. Thus we have constructed a translation surface structure on (S, Σ) whose horizontal foliation \mathcal{G} represents β , as required. \square

4. The homeomorphism theorem

We now prove Theorem 1.2, which states that

$$\text{hol} : \tilde{\mathcal{H}}(\mathcal{F}) \rightarrow \mathbb{A}(\mathcal{F}) \times \mathbb{B}(\mathcal{F})$$

is a homeomorphism, where $\tilde{\mathcal{H}}(\mathcal{F})$ is the set of marked translation surface structures with vertical foliation topologically equivalent to \mathcal{F} , $\mathbb{A}(\mathcal{F}) \subset H^1(S, \Sigma)$ is the set of cohomology classes obtained from transverse measures on \mathcal{F} without atoms (see §2.5), and $\mathbb{B}(\mathcal{F}) \subset H^1(S, \Sigma)$ is the set of \mathbf{b} such that $\mathbf{b}(\alpha) > 0$ for all $\alpha \in H_+^\mathcal{F}$.

Proof. – The fact that hol maps $\tilde{\mathcal{H}}(\mathcal{F})$ to $\mathbb{A}(\mathcal{F}) \times \mathbb{B}(\mathcal{F})$ is an immediate consequence of the definitions, and of the easy direction of Theorem 1.1. That it is continuous is also clear from definitions. That $\text{hol}|_{\tilde{\mathcal{H}}(\mathcal{F})}$ maps onto $\mathbb{A}(\mathcal{F}) \times \mathbb{B}(\mathcal{F})$ is the hard direction of Theorem 1.1. Injectivity is a consequence of the following:

LEMMA 4.1. – *Let $\tilde{\mathcal{H}}$ be a stratum of marked translation surfaces of type (S, Σ) . Fix a singular measured foliation on (S, Σ) , and let $\mathbf{b} \in H^1(S, \Sigma; \mathbb{R})$. Then there is at most one $\mathbf{q} \in \tilde{\mathcal{H}}$ with vertical foliation \mathcal{F} and horizontal foliation representing \mathbf{b} .*

Proof of Lemma 4.1. – Suppose that \mathbf{q}_1 and \mathbf{q}_2 are two marked translation surfaces, such that the vertical measured foliation of both is \mathcal{F} , and the horizontal measured foliations $\mathcal{G}_1, \mathcal{G}_2$ both represent \mathbf{b} . We need to show that $\mathbf{q}_1 = \mathbf{q}_2$. Let γ be a special transverse system to \mathcal{F} (as in §2.3 and the proof of Theorem 1.1). Recall that the non-cylinder edges of γ can be made as small as we like. Since \mathcal{F} and \mathcal{G}_1 are transverse, we may take each non-cylinder segment of γ to be contained in leaves of \mathcal{G}_1 .

In a sufficiently small neighborhood U of any $p \in \Sigma$, we may perform an isotopy of \mathbf{q}_2 , preserving the leaves of \mathcal{F} , so as to make \mathcal{G}_2 coincide with \mathcal{G}_1 . This follows from the fact that in \mathbb{R}^2 , the leaves of any foliation transverse to the vertical foliation can be expressed as graphs over the horizontal direction. Having done this, we may choose the non-cylinder segments of γ to be contained in such neighborhoods, and hence simultaneously in leaves of \mathcal{G}_1 and \mathcal{G}_2 .

Now consider the cell decomposition $\mathcal{B} = \mathcal{B}(\gamma)$, and let $\beta_i = [\mathcal{G}_i]$ be the 1-cocycle on \mathcal{B} obtained by integrating \mathcal{G}_i . For a transverse non-cylinder edge e we have $\beta_1(e) = \beta_2(e) = 0$ since e is a leaf for both foliations. If e is contained in a leaf of \mathcal{F} , we may join its endpoints to Σ by paths d and f along γ . Then $\delta = d + e + f$ represents an element of $H_1(S, \Sigma)$ so that $\beta_1(\delta) = \beta_2(\delta) = \mathcal{B}(\delta)$. Since $\beta_i(d) = \beta_i(f) = 0$ we have $\beta_i(e) = \mathcal{B}(\delta)$. For a cylinder edge e , its endpoints are already on Σ so $\beta_1(e) = \beta_2(e)$. We have shown that $\beta_1 = \beta_2$ on all edges of \mathcal{B} .

Recall from the proof of Theorem 1.1 that \mathbf{q}_i may be obtained explicitly by giving each cell the structure of a Euclidean rectangle or parallelogram (the latter for cylinder cells) as determined by \mathcal{F} and β_i on the edges. Therefore $\mathbf{q}_1 = \mathbf{q}_2$. \square

Finally we need to show that the inverse of hol is continuous. This is an elaboration of the well-known fact that hol is a local homeomorphism, which we can see as follows. Let $\mathbf{q} \in \tilde{\mathcal{H}}$, and consider a geometric triangulation τ of \mathbf{q} with vertices in Σ (e.g., a Delaunay triangulation [11]). The shape of each triangle is uniquely and continuously determined by the hol image of each of its edges. Hence if we choose a neighborhood \mathcal{U} of \mathbf{q} small enough so that none of the triangles becomes degenerate, we have a homeomorphism $\text{hol} : \mathcal{U} \rightarrow \mathcal{V}$ where $\mathcal{V} = \text{hol}(\mathcal{U}) \subset H^1(S, \Sigma; \mathbb{R}^2)$.

If for $\mathbf{q}' \in \mathcal{U}$, the first coordinate $x(\mathbf{q}')$ of $\text{hol}(\mathbf{q}')$ lies in $\mathbb{A}(\mathcal{F})$, then we claim that $\mathbf{q}' \in \tilde{\mathcal{H}}(\mathcal{F})$. This is because the vertical foliation is determined by the weights that $x(\mathbf{q})$ assigns to edges of the triangulation. By Lemma 4.1, \mathbf{q}' is the unique preimage of $\text{hol}(\mathbf{q}')$ in $\tilde{\mathcal{H}}(\mathcal{F})$. Hence $(\text{hol}|_{\mathcal{U}})^{-1}$ and $(\text{hol}|_{\tilde{\mathcal{H}}(\mathcal{F})})^{-1}$ coincide on their overlap, so continuity of one implies continuity of the other. \square

5. Positive pairs

We now reformulate Theorem 1.1 in the language of interval exchanges, and derive several useful consequences. For this we need some more definitions. Let the notation be as in §2.2, so that σ is an irreducible and admissible permutation on d elements. The tangent space $T\mathbb{R}_+^d$ has a natural product structure $T\mathbb{R}_+^d = \mathbb{R}_+^d \times \mathbb{R}^d$ and a corresponding affine structure. Given $\mathbf{a} \in \mathbb{R}_+^d$, $\mathbf{b} \in \mathbb{R}^d$, we can think of (\mathbf{a}, \mathbf{b}) as an element of $T\mathbb{R}_+^d$. We will be using the same symbols \mathbf{a}, \mathbf{b} which were previously used to denote cohomology classes; the reason for this will become clear momentarily. Let $\mathcal{I} = \mathcal{I}_\sigma(\mathbf{a})$ be the interval exchange associated with σ and \mathbf{a} .

For $\mathbf{b} \in \mathbb{R}^d$, in analogy with (9), define $y_i(\mathbf{b}), y'_i(\mathbf{b})$ via

$$(17) \quad y_i = y_i(\mathbf{b}) = \sum_{j=1}^i b_j, \quad y'_i = y'_i(\mathbf{b}) = \sum_{j=1}^i b_{\sigma^{-1}(j)} = \sum_{\sigma(k) \leq i} b_k.$$

In the case of Masur's construction (Figure 1.1), the y_i are the heights of the points in the upper boundary of the polygon, and the y'_j in the lower.

Consider the following step functions $f, g, L : I \rightarrow \mathbb{R}$, depending on \mathbf{a} and \mathbf{b} :

$$(18) \quad \begin{aligned} f(x) &= y_i && \text{for } x \in I_i = [x_{i-1}, x_i) \\ g(x) &= y'_i && \text{for } x \in I'_i = [x'_{i-1}, x'_i) \\ L(x) &= f(x) - g(\mathcal{T}(x)). \end{aligned}$$

Note that for Q as in (11) and $x \in I_i$ we have

$$(19) \quad L(x) = Q(\mathbf{e}_i, \mathbf{b}).$$

If there are $i, j \in \{0, \dots, d - 1\}$ (not necessarily distinct) and $m > 0$ such that $\mathcal{T}^m(x_i) = x_j$ we will say that (i, j, m) is a *connection* for \mathcal{T} . We denote the set of invariant non-atomic probability measures for \mathcal{T} by $\mathcal{M}_{\mathbf{a}}$, and the set of connections by $\mathcal{L}_{\mathbf{a}}$.

DEFINITION 5.1. – We say that $(\mathbf{a}, \mathbf{b}) \in \mathbb{R}_+^d \times \mathbb{R}^d$ is a positive pair if

$$(20) \quad \int L d\mu > 0 \text{ for any } \mu \in \mathcal{M}_{\mathbf{a}}$$

and

$$(21) \quad \sum_{n=0}^{m-1} L(\mathcal{T}^n x_i) > y_i - y_j \text{ for any } (i, j, m) \in \mathcal{L}_{\mathbf{a}}.$$

As explained in §2.3, following [24] one can construct a surface S with a finite subset Σ , a foliation \mathcal{F} on (S, Σ) , and a judicious curve γ on S such that $\mathcal{T}_{\sigma}(\mathbf{a}) = \mathcal{T}(\mathcal{F}, \gamma)$ (we identify γ with I via the transverse measure).

Moreover, as in §2.6 there is a complex \mathcal{D} with a single 2-cell \mathcal{D}^2 containing γ as a properly embedded arc, and whose boundary is divided by $\partial\gamma$ into two arcs each of subjects to the 1-skeleton of \mathcal{D} . The upper arc is divided by Σ into d oriented segments K_1, \dots, K_d images of the segments I_i under flow along \mathcal{F} . The vector \mathbf{b} can be interpreted as a class in $H^1(S, \Sigma)$ by assigning b_i to the segment K_i . The Poincaré dual of \mathbf{b} in $H_1(S \setminus \Sigma)$ is written $\beta = \sum b_i \hat{\ell}_i$, where $\hat{\ell}_i$ are as in §2.6.

Using these we show:

PROPOSITION 5.2. – (\mathbf{a}, \mathbf{b}) is a positive pair if and only if $\mathbf{b}(\alpha) > 0$ for any $\alpha \in H_+^{\mathcal{F}}$.

Proof. – We show that implication \implies , the converse being similar. It suffices to consider the cases where $\alpha \in H_+^{\mathcal{F}}$ corresponds to $\mu \in \mathcal{M}_{\mathcal{F}}$ or to a positively oriented saddle connection in \mathcal{F} , because a general element in $H_+^{\mathcal{F}}$ is a convex combination of these.

In the first case, define $a'_k = \mu(I_k)$ and $\mathbf{a}' = \sum a'_k \mathbf{e}_k$. The corresponding homology class in $H_1(S \setminus \Sigma)$ is $\alpha = \sum a'_k \hat{\ell}_k$. Hence, as we saw in (16),

$$\mathbf{b}(\alpha) = \beta \cdot \alpha = Q(\mathbf{a}', \mathbf{b}) = \sum a'_k Q(\mathbf{e}_k, \mathbf{b}) \stackrel{(19)}{=} \sum \mu(I_k) L|_{I_k} = \int L d\mu > 0.$$

Now consider the second case. Given a connection (i, j, m) for \mathcal{T} , the corresponding saddle connection α meets the disk \mathcal{D}^2 in a union of leaf segments η_1, \dots, η_m where each η_n is the leaf segment in \mathcal{D}^2 intersecting the interval $\gamma = I$ in the point $\mathcal{T}^n(x_i)$. This point lies in some interval I_r and some interval I'_s . If we let $\hat{\eta}_n$ be a line segment connecting the left

endpoint of I'_s to the left endpoint of I_r then the chain $\sum \hat{\eta}_n$ is homologous to $\sum \eta_n$ (note that the first endpoint of η_1 and the last endpoint of η_m do not change). Now we apply our cocycle \mathbf{b} to each $\hat{\eta}_n$ to obtain

$$\mathbf{b}(\hat{\eta}_n) = y_r - y'_s = f(\mathcal{F}^n(x_i)) - g(\mathcal{F}^n(x_i)).$$

Summing, we get

$$\begin{aligned} \mathbf{b}(\alpha) &= \mathbf{b}\left(\sum \eta_n\right) = \mathbf{b}\left(\sum \hat{\eta}_n\right) \\ &= \sum f(\mathcal{F}^n(x_i)) - g(\mathcal{F}^n(x_i)) \\ &= -f(x_i) + f(\mathcal{F}^m(x_i)) + \sum_{n=0}^{m-1} f(\mathcal{F}^n(x_i)) - g(\mathcal{F}^{n+1}(x_i)) \\ &= -y_i + y_j + \sum_{n=0}^{m-1} L(\mathcal{F}^n(x_i)) > 0. \end{aligned} \quad \square$$

Now we can state the interval exchange version of Theorem 1.1:

THEOREM 5.3. – *Let $\tilde{\mathcal{H}}$ be the stratum of marked translation surfaces corresponding to σ . Then for any positive pair $(\mathbf{a}_0, \mathbf{b}_0)$ there is a neighborhood \mathcal{U} of $(\mathbf{a}_0, \mathbf{b}_0)$ in $T\mathbb{R}_+^d$, and a map $\mathbf{q} : \mathcal{U} \rightarrow \tilde{\mathcal{H}}$ such that the following hold:*

- (i) \mathbf{q} is an affine map and a local homeomorphism.
- (ii) For any $(\mathbf{a}, \mathbf{b}) \in \mathcal{U}$, $\mathbf{q}(\mathbf{a}, \mathbf{b})$ is a lift of \mathbf{a} .
- (iii) Any (\mathbf{a}, \mathbf{b}) in \mathcal{U} is positive.
- (iv) Suppose $(\mathbf{a}, \mathbf{b}) \in \mathcal{U}$ and $\varepsilon_0 > 0$ is small enough so that $(\mathbf{a} + s\mathbf{b}, \mathbf{b}) \in \mathcal{U}$ for $|s| \leq \varepsilon_0$. Then $h_s \mathbf{q}(\mathbf{a}, \mathbf{b}) = \mathbf{q}(\mathbf{a} + s\mathbf{b}, \mathbf{b})$ for $|s| \leq \varepsilon_0$.

Proof. – Above and in §2.6 we identified \mathbb{R}^d with $H^1(S, \Sigma; \mathbb{R})$, obtaining an injective affine map

$$T\mathbb{R}_+^d \cong \mathbb{R}_+^d \times \mathbb{R}^d \rightarrow H^1(S, \Sigma; \mathbb{R})^2 \cong H^1(S, \Sigma; \mathbb{R}^2)$$

and we henceforth identify $T\mathbb{R}_+^d$ with its image.

Given a positive pair $(\mathbf{a}_0, \mathbf{b}_0)$, as discussed at the end of §2.3 we obtain a surface (S, Σ) together with a measured foliation $\mathcal{F} = \mathcal{F}(\mathbf{a}_0)$ and a judicious transversal γ_0 , so that the return map $\mathcal{F}(\mathcal{F}, \gamma)$ is equal to $\mathcal{F}_\sigma(\mathbf{a}_0)$, where I parametrizes γ_0 via the transverse measure.

The positivity condition, together with Proposition 5.2 and Theorem 1.1 give us a translation surface structure $\mathbf{q} = \mathbf{q}(\mathbf{a}_0, \mathbf{b}_0)$ whose vertical foliation is \mathcal{F} and whose image under hol is $(\mathbf{a}_0, \mathbf{b}_0)$.

Let τ be a geometric triangulation on \mathbf{q} . Assume for the moment that τ has no vertical edges, and in particular that each edge for τ is transverse to \mathcal{F} . Now consider \mathbf{a} very close to \mathbf{a}_0 . The map $\mathcal{F}_\sigma(\mathbf{a})$ is close to $\mathcal{F}_\sigma(\mathbf{a}_0)$ and hence induces a foliation $\mathcal{F}(\mathbf{a})$ whose leaves are nearly parallel to those of $\mathcal{F}(\mathbf{a}_0)$. More explicitly, $\mathcal{F}(\mathbf{a})$ is obtained by modifying $\mathcal{F}(\mathbf{a}_0)$ slightly in a small neighborhood of γ so that it remains transverse to both γ and τ , and so that the return map becomes $\mathcal{F}_\sigma(\mathbf{a})$. This can be done if \mathbf{a} is in a sufficiently small neighborhood of \mathbf{a}_0 .

Now if (\mathbf{a}, \mathbf{b}) is sufficiently close to $(\mathbf{a}_0, \mathbf{b}_0)$, the pair (\mathbf{a}, \mathbf{b}) assigns to edges of τ vectors which retain the orientation induced by $(\mathbf{a}_0, \mathbf{b}_0)$. Hence we obtain a new geometric triangulation, for which $\mathcal{F}(\mathbf{a})$ is transverse to the edges, has transverse measure agreeing with \mathbf{a} on the edges, and hence is still realized by the vertical foliation. Moreover γ_0 is still transverse to the new foliation and the return map is the correct one. That is what we wanted to show.

Returning to the case where τ is allowed to have vertical edges: note that at most one edge in a triangle can be vertical. Hence, if we remove the vertical edges we are left with a decomposition whose cells are Euclidean triangles and quadrilaterals, and whose edges are transverse to \mathcal{F} . The above argument applies equally well to this decomposition.

We have shown that in a neighborhood of $(\mathbf{a}_0, \mathbf{b}_0)$, the map \mathbf{q} maps to $\tilde{\mathcal{H}}_\tau$, and is a local inverse for hol. Hence it is affine and a local homeomorphism, establishing (i). Part (ii) is by definition part of our construction. Part (iii) follows from the implication (1) \implies (2) in Theorem 1.1. In verifying (iv) we use (8). \square

The following useful observation follows immediately:

COROLLARY 5.4. – *The set of positive pairs is open.*

We also record the following corollary of our construction, coming immediately from the description we gave of the variation of the triangles in the geometric triangulation τ :

COROLLARY 5.5. – *Suppose $\sigma, \tilde{\mathcal{H}}$, a positive pair $(\mathbf{a}_0, \mathbf{b}_0)$, and $\mathbf{q} : \mathcal{U} \rightarrow \tilde{\mathcal{H}}$ are as in Theorem 5.3. The structures $\mathbf{q}(\mathbf{a}, \mathbf{b})$ can be chosen in their isotopy class so that the following holds: A single curve γ is a judicious transversal for the vertical foliation of $\mathbf{q}(\mathbf{a}, \mathbf{b})$ for all $(\mathbf{a}, \mathbf{b}) \in \mathcal{U}$; the flat structures vary continuously with (\mathbf{a}, \mathbf{b}) , meaning that the charts in the atlas, modulo translation, vary continuously; the return map to γ satisfies $\mathcal{F}(\mathbf{q}, \gamma) = \mathcal{F}_\sigma(\mathbf{a})$. In particular, for any $\mathbf{q}' = \mathbf{q}(\mathbf{a}, \mathbf{b})$, there is $\varepsilon > 0$ such that for $|s| < \varepsilon$, $\mathbf{a}(s) = \mathbf{a} + s\mathbf{b}$, we have*

$$\mathcal{F}_\sigma(\mathbf{a}(s)) = \mathcal{F}(h_s \mathbf{q}', \gamma).$$

6. Mahler's question for interval exchanges and its generalizations

A vector $\mathbf{x} \in \mathbb{R}^d$ is *very well approximable* if for some $\varepsilon > 0$ there are infinitely many $\mathbf{p} \in \mathbb{Z}^d, q \in \mathbb{N}$ satisfying $\|q\mathbf{x} - \mathbf{p}\| < q^{-(1/d+\varepsilon)}$. It is a classical fact that almost every (with respect to Lebesgue measure) $\mathbf{x} \in \mathbb{R}^d$ is not very well approximable, but that the set of very well approximable vectors is large in the sense of Hausdorff dimension. Mahler conjectured in the 1930's that for almost every (with respect to Lebesgue measure on the real line) $x \in \mathbb{R}$, the vector $\mathbf{a}(x) = (x, x^2, \dots, x^d)$ as in (2), is not very well approximable. This famous conjecture was settled by Sprindzhuk in the 1960's and spawned many additional questions of a similar nature. A general formulation of the problem is to describe measures μ on \mathbb{R}^d for which almost every \mathbf{x} is not very well approximable. See [6] for a survey.

In this section we apply Theorems 1.1 and 5.3 to analogous problems concerning interval exchange transformations. Fix a permutation σ on d symbols which is irreducible and admissible. In answer to a conjecture of Keane, it was proved by Masur [9] and Veech [21] that almost every \mathbf{a} (with respect to Lebesgue measure on \mathbb{R}_+^d) is uniquely ergodic. On the other hand Masur and Smillie [11] showed that the set of non-uniquely ergodic interval exchanges

is large in the sense of Hausdorff dimension. In this paper we consider the problem of describing measures μ on \mathbb{R}_+^d such that μ -a.e. \mathbf{a} is uniquely ergodic. In a celebrated paper [5], it was shown that for certain σ and certain line segments $\ell \subset \mathbb{R}_+^d$ (arising from a problem in billiards on polygons), for μ -almost every $\mathbf{a} \in \ell$, $\mathcal{T}_\sigma(\mathbf{a})$ is uniquely ergodic, where μ denotes Lebesgue measure on ℓ . This was later abstracted in [22], where the same result was shown to hold for a general class of measures in place of μ . Our strategy is strongly influenced by these papers.

Before stating our results we introduce more terminology. Let $B(x, r)$ denote the interval $(x - r, x + r)$ in \mathbb{R} . We say that a finite regular Borel measure μ on \mathbb{R} is *decaying and Federer* if there are positive C, α, D such that for every $x \in \text{supp } \mu$ and every $0 < \varepsilon, r < 1$,

$$(22) \quad \mu(B(x, \varepsilon r)) \leq C\varepsilon^\alpha \mu(B(x, r)) \quad \text{and} \quad \mu(B(x, 3r)) \leq D\mu(B(x, r)).$$

It is not hard to show that Lebesgue measure, and the coin-tossing measure on Cantor's middle thirds set, are both decaying and Federer. More constructions of such measures are given in [22, 8]. Let \dim denote Hausdorff dimension, and for $x \in \text{supp } \mu$ let

$$\underline{d}_\mu(x) = \liminf_{r \rightarrow 0} \frac{\log \mu(B(x, r))}{\log r}.$$

Now let

$$(23) \quad \varepsilon_n(\mathbf{a}) = \min \left\{ \left| \mathcal{T}^k(x_i) - \mathcal{T}^\ell(x_j) \right| : 0 \leq k, \ell \leq n, 1 \leq i, j \leq d-1, (i, k) \neq (j, \ell) \right\},$$

where $\mathcal{T} = \mathcal{T}_\sigma(\mathbf{a})$, and the x_i are defined via (9). We say that \mathbf{a} is *of recurrence type* if $\limsup n\varepsilon_n(\mathbf{a}) > 0$ and *of bounded type* if $\liminf n\varepsilon_n(\mathbf{a}) > 0$. It is known by work of Masur, Boshernitzan, Veech and Cheung that if \mathbf{a} is of recurrence type then it is uniquely ergodic, but that the converse does not hold—see §7 below for more details.

We have:

THEOREM 6.1 (Lines). – *Suppose (\mathbf{a}, \mathbf{b}) is a positive pair. Then there is $\varepsilon_0 > 0$ such that the following hold for $\mathbf{a}(s) = \mathbf{a} + s\mathbf{b}$ and for every decaying and Federer measure μ with $\text{supp } \mu \subset (-\varepsilon_0, \varepsilon_0)$:*

- (a) *For μ -almost every s , $\mathbf{a}(s)$ is of recurrence type.*
- (b) $\dim \{s \in \text{supp } \mu : \mathbf{a}(s) \text{ is of bounded type}\} \geq \inf_{x \in \text{supp } \mu} \underline{d}_\mu(x)$.
- (c) $\dim \{s \in (-\varepsilon_0, \varepsilon_0) : \mathbf{a}(s) \text{ is not of recurrence type}\} \leq 1/2$.

THEOREM 6.2 (Curves). – *Let I be an interval, let μ be a decaying and Federer measure on I , and let $\beta : I \rightarrow \mathbb{R}_+^d$ be a C^2 curve, such that for μ -a.e. $s \in I$, $(\beta(s), \beta'(s))$ is positive. Then for μ -a.e. $s \in I$, $\beta(s)$ is of recurrence type.*

Following some preliminary work, we will prove Theorem 6.1 in §8 and Theorem 6.2 in §10. In §9 we will prove a strengthening of Theorem 6.1(a).

7. Saddle connections, compactness criteria

A link between the G -action and unique ergodicity questions was made in the following fundamental result.

LEMMA 7.1 (Masur [10]). – *If $q \in \mathcal{H}$ is not uniquely ergodic then the trajectory $\{g_t q : t \geq 0\}$ is divergent, i.e., for any compact $K \subset \mathcal{H}$ there is t_0 such that for all $t \geq t_0$, $g_t q \notin K$.*

Masur's result is in fact stronger as it provides divergence in the moduli space of quadratic differentials. The converse statement is not true, see [3]. It is known (see [23, Prop. 3.12] and [22, §2]) that \mathbf{a} is of recurrence (resp. bounded) type if and only if the forward geodesic trajectory of any of its lifts returns infinitely often to (resp. stays in) some compact subset of \mathcal{H} . It follows using Lemma 7.1 that if \mathcal{T} is of recurrence type then it is uniquely ergodic. In this section we will prove a quantitative version of these results, linking the behavior of G -orbits to the size of the quantity $n\varepsilon_n(\mathbf{a})$.

We denote the set of all saddle connections for a marked translation surface \mathbf{q} by $\mathcal{L}_{\mathbf{q}}$. There is a natural identification of $\mathcal{L}_{\mathbf{q}}$ with $\mathcal{L}_{g\mathbf{q}}$ for any $g \in G$. We define

$$\phi(q) = \min \{ \ell(\alpha, \mathbf{q}) : \alpha \in \mathcal{L}_{\mathbf{q}} \},$$

where $\mathbf{q} \in \pi^{-1}(q)$ and $\ell(\alpha, \mathbf{q}) = \max\{|x(\alpha, \mathbf{q})|, |y(\alpha, \mathbf{q})|\}$. Let \mathcal{H}_1 be the area-one sublocus in \mathcal{H} , i.e., the set of $q \in \mathcal{H}$ for which the total area of the surface is one. A standard compactness criterion for each stratum asserts that a set $X \subset \mathcal{H}_1$ is compact if and only if

$$\inf_{q \in X} \phi(q) > 0.$$

Thus, for each $\varepsilon > 0$,

$$K_\varepsilon = \{q \in \mathcal{H}_1 : \phi(q) \geq \varepsilon\}$$

is compact, and $\{K_\varepsilon\}_{\varepsilon > 0}$ forms an exhaustion of \mathcal{H}_1 . We have:

PROPOSITION 7.2. – *Suppose γ is judicious for q . Then there are positive κ, c_2, n_0 such that for $\mathcal{T} = \mathcal{T}(q, \gamma)$ we have*

– *If $n \geq n_0$, $\zeta \geq n\varepsilon_n(\mathcal{T})$, and $e^{t/2} = n\sqrt{2c_2/\zeta}$, then*

$$(24) \quad \phi(g_t q) \leq \kappa\sqrt{\zeta}.$$

– *If $n = \lfloor \kappa e^{t/2} \rfloor$, then*

$$(25) \quad n\varepsilon_n(\mathcal{T}) \leq \kappa\phi(g_t q).$$

Moreover, κ, c_2, n_0 may be taken to be uniform for q ranging over a compact subset of \mathcal{H} and γ ranging over smooth curves of uniformly bounded length, with return times to the curve bounded above and below.

Proof. – We first claim that

$$(26) \quad \varepsilon_n(\mathbf{a}) = \min \{ |x_i - \mathcal{T}^r x_j| : 1 \leq i, j \leq d-1, |r| \leq n, (j, r) \neq (i, 0) \}.$$

Indeed, if the minimum in (23) is equal to $|\mathcal{T}^k x_i - \mathcal{T}^\ell x_j|$ with $\ell \geq k \geq 1$, then the interval between $\mathcal{T}^k x_i$ and $\mathcal{T}^\ell x_j$ does not contain discontinuities for \mathcal{T}^{-k} (if it did the minimum could be made smaller). This implies that \mathcal{T}^{-k} acts as an isometry on this interval

so that $|x_i - \mathcal{F}^{\ell-k} x_j| = \varepsilon_n(\mathcal{F})$. Similarly, if the minimum in (26) is obtained for i, j and $r = -k \in [-n, 0]$ then the interval between x_i and $\mathcal{F}^{-k} x_j$ has no discontinuities of \mathcal{F}^k , so that the same value is also obtained for $|\mathcal{F}^k x_i - x_j|$. Hence the minimum in (26) equals the minimum in (23).

Let $\mathbf{q} \in \pi^{-1}(q)$ and let $n_0 \geq 1$. Suppose the return times to γ along the vertical foliation are bounded below and above by c_1 and c_2 , respectively. Making c_2 larger we can also assume that the total variation in the vertical direction along γ is no more than c_2 . Write $n\varepsilon_n(\mathcal{F}) = n|x_i - \mathcal{F}^r x_j| \leq \zeta$ and let t be as in (24). Let σ_i and σ_j be the singularities of \mathbf{q} lying vertically above x_i and x_j . Let α be the path moving vertically from σ_j along the forward trajectory of x_j until $\mathcal{F}^r x_j$, then along γ to x_i , and vertically up to σ_i . Then $|x(\mathbf{q}, \alpha)| = \varepsilon_n(\mathcal{F}) \leq \zeta/n$ and $|y(\mathbf{q}, \alpha)| \leq c_2 r + c_2 \leq 2nc_2$ for $n \geq n_0$. Therefore, since $e^{t/2} = n\sqrt{2c_2/\zeta}$, we have

$$\begin{aligned} |x(g_t \mathbf{q}, \alpha)| &= e^{t/2} |x(\mathbf{q}, \alpha)| \leq \sqrt{2c_2 \zeta}, \\ |y(g_t \mathbf{q}, \alpha)| &= e^{-t/2} |y(\mathbf{q}, \alpha)| \leq \sqrt{2c_2 \zeta}, \end{aligned}$$

so $\ell(\alpha, g_t \mathbf{q}) \leq \kappa\sqrt{\zeta}$, where $\kappa = \sqrt{2c_2}$. A shortest representative for α with respect to $g_t \mathbf{q}$ is a concatenation $\bar{\alpha}$ of saddle connections. Since α travels monotonically along both horizontal and vertical foliations of \mathbf{q} , we claim that $\bar{\alpha}$ does the same. This will imply that the coordinates of $\bar{\alpha}$'s saddle connections have consistent signs, so that the same bound holds for each of them, and (24) follows.

To prove the claim, note that after restricting to subpaths we may assume that α and $\bar{\alpha}$ have disjoint interiors and bound a disk D . Suppose that two successive saddle connections of $\bar{\alpha}$ move with opposite orientation with respect to, say, the vertical direction. Then their common endpoint p is contained in the interior of a horizontal leaf segment ℓ , such that $\ell - p$ is in the interior of the disk D . Extending ℓ in both directions until it meets the boundary of D , we see that it must meet the boundary at two points of α since $\bar{\alpha}$ is a geodesic. But this implies that α is not monotonic in the vertical direction, establishing the claim.

Now we establish (25). Let α be a saddle connection minimizing $\ell(\cdot, g_t \mathbf{q})$, and write $x_t = x(g_t \mathbf{q}, \alpha)$ and $y_t = y(g_t \mathbf{q}, \alpha)$. Without loss of generality (reversing the orientation of α if necessary) we may assume that $x_t \geq 0$. Minimality means

$$\phi = \phi(g_t \mathbf{q}) = \max(x_t, |y_t|).$$

In \mathbf{q} , the coordinates of α satisfy

$$x_0 = e^{-t/2} x_t \leq e^{-t/2} \phi$$

and

$$|y_0| = e^{t/2} |y_t| \leq e^{t/2} \phi.$$

Let \mathcal{U} be the strip $[0, x_0] \times \mathbb{R}$ in \mathbb{R}^2 , and let $v \subset \mathcal{U}$ be the line segment connecting $v_- = (0, 0)$ to $v_+ = (x_0, y_0)$. A neighborhood of v in \mathcal{U} embeds in S by a local isometry that preserves horizontal and vertical directions. We can extend this to an isometric immersion $\psi : \mathcal{U}' \rightarrow S$, where \mathcal{U}' has the following form: There is a discrete set $\widehat{\Sigma} \subset \mathcal{U} \setminus \text{int}(v)$, and for each $\sigma = (x, y) \in \widehat{\Sigma}$ a vertical ray R_σ of the form $\{x\} \times (y, \infty)$ ("upward pointing") or $\{x\} \times (-\infty, y)$ ("downward pointing"), so that the rays are pairwise disjoint, disjoint from v , and $\mathcal{U}' = \mathcal{U} \setminus \bigcup_{\sigma \in \widehat{\Sigma}} R_\sigma$ (see Figure 7.1). The map ψ takes $\widehat{\Sigma}$ to Σ , and it is defined by




FIGURE 7.1. The vertical strip \mathcal{U} minus rays R_σ immerses isometrically in S .

extending the embedding at each $p \in v$ maximally along the vertical line through p (in both directions) until the image encounters a singularity in Σ . (We include v_- and v_+ in $\hat{\Sigma}$, and for these two points delete both an upward and a downward pointing ray.)

Let $\hat{\gamma}$ be the preimage $\psi^{-1}(\gamma)$. This is a union of arcs properly embedded in \mathcal{U}' , and transverse to the vertical foliation in \mathbb{R}^2 . By definition of c_1 and c_2 , each vertical line in \mathcal{U}' is cut by $\hat{\gamma}$ into segments of length at least c_1 and at most c_2 . Moreover the total vertical extent of each component of $\hat{\gamma}$ is at most c_2 .

Consider γ_1 the component of $\hat{\gamma}$ that meets the downward ray based at v_+ at the highest point \hat{r} . The other endpoint \hat{p} of γ_1 lies on some other ray R_σ .

The width of γ_1 is at most x_0 , so the image points $r = \psi(\hat{r})$ and $p = \psi(\hat{p})$ satisfy $|p - r| \leq x_0$, with respect to the induced transverse measure on γ . We now check that p and r are images of discontinuity points of \mathcal{T} , by controlled powers of \mathcal{T} .

By choice of γ_1 , the upward leaf emanating from r encounters the singularity $\psi(v_+)$ before it returns to γ , and hence r itself is a discontinuity point x_i .

For p , let us write $\sigma = (x, y)$ and $p = (x, y')$. Suppose first that $y_0 \geq 0$. There are now two cases. If R_σ lies above v (and hence is upward pointing): we have $y' \geq y \geq 0$, and moreover (since the vertical variation of γ_1 is bounded) $y' \leq y_0 + c_2$. The segment of R_σ between σ and p is cut by $\hat{\gamma}$ (incident from the right) into at most $(y' - y)/c_1$ pieces, and this implies that there is some $k \geq 0$ bounded by

$$k \leq \frac{y_0 + c_2}{c_1} \leq \frac{e^{t/2}\phi + c_2}{c_1}$$

such that $p = T^k x_j$ for some discontinuity x_j .

If R_σ lies below v and is downward pointing: we have $y' \leq y \leq y_0$ and $y' \geq y_0 - 2c_2$, so that by the same logic there is $k \geq 0$ with

$$k \leq \frac{2c_2}{c_1}$$

such that $T^k p = x_j$ for some discontinuity x_j .

Hence in either case we have

$$(27) \quad |T^m x_j - x_i| \leq x_0 \leq e^{-t/2} \phi$$

where $-2c_2/c_1 \leq m \leq (y_0 + c_2)/c_1$.

If $y_0 < 0$ there is a similar analysis, yielding the bound (27) where now

$$(y_0 - 2c_2)/c_1 \leq m \leq c_2/c_1.$$

Noting that $\phi < 1$ by area considerations, if we take $n = \lfloor \kappa e^{t/2} \rfloor$, where $\kappa = 1 + (1 + 2c_2)/c_1$, then we guarantee $|m| \leq n$, and hence get

$$n\varepsilon_n \leq \kappa e^{t/2} e^{-t/2} \phi \leq \kappa \phi. \quad \square$$

8. Mahler's question for lines

In this section we will derive Theorem 6.1 from Theorem 5.3 and earlier results of [5, 10, 7, 8]. We will need the following observation, which gives an explicit sense in which a horocycle and a circle look the same from far away.

LEMMA 8.1. – *For any $|\theta| < \pi/2$, there is a bounded subset $\Omega \subset G$ such that for any $t \geq 0$ there is $w \in \Omega$ such that $g_t h_{-\tan \theta} = w g_t r_\theta$.*

Proof. – The assertion is equivalent to saying that the set $\Omega = \{g_t h_{-\tan \theta} r_{-\theta} g_{-t} : t \geq 0\}$ is bounded in G . Letting

$$x = h_{-\tan \theta} r_{-\theta} = \begin{pmatrix} \cos \theta & 0 \\ -\sin \theta & 1/\cos \theta \end{pmatrix} \in G,$$

we note that $g_t x g_{-t}$ converges in G as $t \rightarrow \infty$, so the claim follows. \square

Proof of Theorem 6.1. – Let (\mathbf{a}, \mathbf{b}) be positive, let \mathcal{U} be a neighborhood of (\mathbf{a}, \mathbf{b}) in $\mathbb{R}_+^d \times \mathbb{R}^d$, let $\mathbf{q} : \mathcal{U} \rightarrow \mathcal{H}$ as in Theorem 5.3, let $q = \pi \circ \mathbf{q}$ where $\pi : \tilde{\mathcal{H}} \rightarrow \mathcal{H}$ is the natural projection, and let $\varepsilon_0 > 0$ so that $\mathbf{a}(s) = \mathbf{a} + s\mathbf{b} \in \mathcal{U}$ for all $s \in (-\varepsilon_0, \varepsilon_0)$. Making ε_0 smaller if necessary, let γ be a judicious curve for q such that $\mathcal{I}_\sigma(\mathbf{a}(s)) = \mathcal{I}(h_s q, \gamma)$ for all $s \in (-\varepsilon_0, \varepsilon_0)$. By Theorem 5.3 and Lemma 7.2, $\mathbf{a}(s)$ is of recurrence (resp. bounded) type if and only if there is a compact subset $K \subset \mathcal{H}$ such that $\{t > 0 : g_t h_s q \in K\}$ is unbounded (resp., is equal to $(0, \infty)$). The main result of [5] is that for any q , for Lebesgue-a.e. $\theta \in (-\pi, \pi)$ there is a compact $K \subset \mathcal{H}$ such that $\{t > 0 : g_t r_\theta q \in K\}$ is unbounded. Thus (a) (with μ equal to Lebesgue measure) follows via Lemma 8.1. For a general measure μ , the statement will follow from Corollary 9.3 below. Similarly (b) follows from [7] for μ equal to Lebesgue measure, and from [8] for a general decaying Federer measure, and (c) follows from [10]. \square

9. Quantitative nondivergence for horocycles

In this section we will recall a quantitative nondivergence result for the horocycle flow, which is a variant of results in [15], and will be crucial for us. At the end of the section we will use it to obtain a strengthening of Theorem 6.1(a).

Given positive constants C, α, D , we say that a regular finite Borel measure μ on \mathbb{R} is (C, α) -decaying and D -Federer if (22) holds for all $x \in \text{supp } \mu$ and all $0 < \varepsilon, r < 1$. For an interval $J = B(x, r)$ and $c > 0$ we write $cJ = B(x, cr)$. Let \mathcal{H}_1 and K_ε be as in §7, and let $\tilde{\mathcal{H}}_1 = \pi^{-1}(\mathcal{H}_1)$. We recall that a *half-translation surface* is a generalization of a translation surface, and is defined similarly to a translation surface, but with the transition maps as in (5) of the form $\vec{x} \mapsto \pm \vec{x} + c_{\alpha, \beta}$. In the language of complex analysis, translation surfaces correspond to abelian differentials while half-translation surfaces correspond to quadratic differentials. Although we are not concerned with half-translation surfaces in this paper, we will state the theorem below in its more general form in light of potential applications.

THEOREM 9.1. – *Given a stratum \mathcal{H} of translation or half-translation surfaces, there are positive constants C_1, λ, ρ_0 , such that for any (C, α) -decaying and D -Federer measure μ on an interval $B \subset \mathbb{R}$, the following holds. Suppose $J \subset \mathbb{R}$ is an interval with $3J \subset B$, $0 < \rho \leq \rho_0$, $\mathbf{q} \in \tilde{\mathcal{H}}_1$, and suppose*

$$(28) \quad \forall \delta \in \mathcal{L}_{\mathbf{q}}, \sup_{s \in J} \ell(\delta, h_s \mathbf{q}) \geq \rho.$$

Then for any $0 < \varepsilon < \rho$:

$$(29) \quad \mu(\{s \in J : h_s \pi(\mathbf{q}) \notin K_\varepsilon\}) \leq C' \left(\frac{\varepsilon}{\rho}\right)^{\lambda \alpha} \mu(J),$$

where $C' = C_1 2^{\alpha} CD$.

The proof is similar to that of [15, Thm. 6.10], but with the assumption that μ is Federer substituting for condition (36) of that paper. The theorem was stated without proof in [8, Prop. 8.3]. To keep the paper self-contained we give a detailed proof.

Proof. – For an interval $J \subset \mathbb{R}$, let $|J|$ denote its length. For a function $f : \mathbb{R} \rightarrow \mathbb{R}_+$ and $\varepsilon > 0$, let

$$J_{f, \varepsilon} = \{x \in J : f(x) < \varepsilon\} \text{ and } \|f\|_J = \sup_{x \in J} f(x).$$

For $\delta \in \mathcal{L}_{\mathbf{q}}$ let ℓ_δ be the function $\ell_\delta(s) = \ell(\delta, h_s \mathbf{q})$. Suppose, for $\mathbf{q} \in \tilde{\mathcal{H}}$, $\delta \in \mathcal{L}_{\mathbf{q}}$ and an interval J , that $\|\ell_\delta\|_J \geq \rho$. Since the function ℓ_δ is the norm of a linear function of s , an elementary computation (see [15, Lemma 4.4]) shows that $J_\varepsilon = J_{\ell_\delta, \varepsilon}$ is a subinterval of J and

$$(30) \quad |J_\varepsilon| \leq \frac{2\varepsilon}{\rho} |J|.$$

Moreover by another elementary calculation (see [15, Lemma 6.4]), if $b > 0$ and $s \in \mathbb{R}$, for saddle connections δ, δ' we have

$$(31) \quad \ell_{\mathbf{q}, \delta}(s), \ell_{\mathbf{q}, \delta'}(s) < b/3 \implies J_{\ell_\delta, b/3} \subset J_{\ell_{\delta'}, b}$$

(possibly after exchanging δ with δ').

Suppose that μ is (C, α) -decaying and D -Federer, and $\text{supp } \mu \cap \tilde{J}_\varepsilon \neq \emptyset$. Let $x \in \text{supp } \mu \cap J_\varepsilon$. Note that $J_\varepsilon \subset B(x, |J_\varepsilon|)$ and $B(x, |J|) \subset 3J$. One has

$$\begin{aligned} \mu(J_\varepsilon) &\leq \mu(B(x, |J_\varepsilon|)) \\ &\stackrel{\text{decay, (30)}}{\leq} C \left(\frac{2\varepsilon}{\rho}\right)^\alpha \mu(B(x, |J|)) \\ &\leq 2^\alpha C \left(\frac{\varepsilon}{\rho}\right)^\alpha \mu(3J) \leq C'' \left(\frac{\varepsilon}{\rho}\right)^\alpha \mu(J), \end{aligned}$$

where $C'' = 2^\alpha CD$. This shows that if J is an interval, $\mathbf{q} \in \tilde{\mathcal{H}}$, and $\delta \in \mathcal{L}_\mathbf{q}$ is such that $\|\ell_\delta\|_J \geq \rho$, then for any $0 < \varepsilon < \rho$,

$$(32) \quad \frac{\mu(J_{\ell_\delta, \varepsilon})}{\mu(J)} \leq C'' \left(\frac{\varepsilon}{\rho}\right)^\alpha.$$

From this estimate we obtain that if $M < \infty$ and $\mathcal{A} \subset \mathcal{L}_\mathbf{q}$ such that

$$(33) \quad \forall t \in J, \#\{\delta \in \mathcal{A} : \ell_\delta(t) < \rho\} \leq M,$$

then

$$(34) \quad \mu\left(\bigcup_{\delta \in \mathcal{A}} J_{\ell_\delta, \varepsilon}\right) \leq MC'' \left(\frac{\varepsilon}{\rho}\right)^\alpha \mu(J).$$

Indeed, using the estimate (32) for each $J_{\ell_\delta, \rho}$, we have

$$\begin{aligned} \mu\left(\bigcup_{\delta \in \mathcal{A}} J_{\ell_\delta, \varepsilon}\right) &\leq \sum_{\delta \in \mathcal{A}} \mu(J_{\ell_\delta, \varepsilon}) \\ &\stackrel{(28), (32)}{\leq} C'' \left(\frac{\varepsilon}{\rho}\right)^\alpha \sum_{\delta \in \mathcal{A}} \mu(J_{\ell_\delta, \rho}) \\ &\stackrel{(33)}{\leq} MC'' \left(\frac{\varepsilon}{\rho}\right)^\alpha \mu(J), \end{aligned}$$

as claimed.

Let M denote the maximal number of disjoint saddle connections for a surface in \mathcal{H} ([15, Prop. 4.7]). For $r = 1, \dots, M$, $\mathbf{q} \in \tilde{\mathcal{H}}_1$ and a collection E of r disjoint saddle connections for \mathbf{q} , let

$$\ell_{\mathbf{q}, E}(t) = \max_{\delta \in E} \ell_{\mathbf{q}, \delta}(t), \quad \text{and} \quad \hat{\alpha}_r(t) = \hat{\alpha}_{\mathbf{q}, r}(t) = \inf \ell_{\mathbf{q}, E}(t),$$

where the infimum ranges over all possible E . In particular $\hat{\alpha}_1(t)$ is the same as $\phi(h_t \mathbf{q})$. For each E , let $S(E)$ denote the closure of the simply connected components of $S \setminus E$. The following two facts were proved by [5] (see also [15, § 6]):

- (i) There exists $\rho_0 > 0$ such that for every \mathbf{q} , every r and every E for which $S(E) = S$ we have $\ell_{\mathbf{q}, E}(0) \geq \rho_0$.
- (ii) Suppose E consists of $r < M$ disjoint saddle connections for $\mathbf{q} \in \tilde{\mathcal{H}}_1$, and for some θ we have

$$\ell_{\mathbf{q}, E}(0) < \frac{\theta}{3\sqrt{2}} \quad \text{and} \quad \hat{\alpha}_{r+1}(0) \geq \theta.$$

Then for any two saddle connections δ, δ' with nontrivial intersection, such that δ is part of the boundary of $S(E)$, we have

$$\ell_{\mathbf{q},\delta'}(0) \geq \frac{\sqrt{2}}{3}\theta.$$

We will prove the theorem with ρ_0 as in (i), and

$$\lambda = \frac{1}{M-1}.$$

Suppose $\mathbf{q}, \rho < \rho_0$ and J satisfy (28). By enlarging C_1 if necessary, it will suffice to prove the theorem for all

$$0 < \varepsilon < \left(\frac{\sqrt{2}}{9}\right)^{M-1} \rho.$$

Let

$$V_\varepsilon = \{t \in J : \hat{\alpha}_1(t) < \varepsilon\}.$$

For $k = 1, \dots, M-1$ define

$$L_k = \varepsilon \left(\frac{\rho}{\varepsilon}\right)^{(k-1)/(m-1)}, \text{ so that } \frac{L_k}{L_{k+1}} = \left(\frac{\varepsilon}{\rho}\right)^{1/(m-1)} < \frac{\sqrt{2}}{9}.$$

For every $t \in V_\varepsilon$ let $r(t) = \max\{k : \hat{\alpha}_k(t) < L_k\}$. Then $r(t) \leq M-1$ by the choice of ρ_0 , and

$$\hat{\alpha}_{r(t)}(t) < L_{r(t)}, \quad \hat{\alpha}_{r(t)+1}(t) > L_{r(t)+1}.$$

Since V_ε is the disjoint union of the sets

$$V_k = \{t : r(t) = k\},$$

there is $r \leq M-1$ such that

$$(35) \quad \mu(V_r) \geq \frac{\mu(V_\varepsilon)}{M-1},$$

and we set $L = L_r, U = L_{r+1}$. Also, for any $\delta \in \mathcal{L}_{\mathbf{q}}$ we set

$$H(\delta) = \left\{ t \in J : \ell_{\mathbf{q},\delta}(t) < L, \left[\delta' \cap \delta \neq \emptyset \implies \ell_{\mathbf{q},\delta'}(t) \geq \frac{\sqrt{2}U}{3} \right] \right\}.$$

We claim that (33) holds for

$$\mathcal{A} = \left\{ \delta \in \mathcal{L}_{\mathbf{q}} : V_r \cap H(\delta) \neq \emptyset, \ell_\delta(t) \leq \frac{\sqrt{2}U}{9} \right\},$$

and also that

$$(36) \quad V_r \subset \bigcup_{\delta} H(\delta).$$

To see (33) for \mathcal{A} , note that if $\delta, \delta' \in \mathcal{A}$ then applying (31) with $b = \frac{\sqrt{2}U}{3}$, possibly after replacing δ and δ' , we have that for some $t \in H(\delta)$, $\ell_{\delta'}(t) < b$. The definition of $H(\delta)$ now implies that δ and δ' are disjoint. So the bound follows from the definition of M . To justify (36) we take, for each $t \in V_r$, a saddle connection $\delta \in E$ where E achieves the minimum in the definition of $\hat{\alpha}_r(t)$, and use (ii), with $\theta = U$.

With this in hand, we have:

$$\begin{aligned} \mu(V_\varepsilon) &\stackrel{(35)}{\leq} (M-1)\mu(V_r) \\ &\stackrel{(36)}{\leq} (M-1)\mu\left(\bigcup_\delta H(\delta)\right) \\ &\stackrel{(34)}{\leq} (M-1)MC''\left(\frac{L}{U}\right)^\alpha \mu(J). \end{aligned}$$

This implies (29). \square

COROLLARY 9.2. – *For any stratum \mathcal{H} of translation surfaces and any C, α, D there is a compact $K \subset \mathcal{H}_1$ such that for any $q \in \mathcal{H}_1$, any unbounded $\mathcal{T} \subset \mathbb{R}_+$ and any (C, α) -decaying and D -Federer measure μ on an interval $J \subset \mathbb{R}$, for μ -a.e. $s \in J$ there is a sequence $t_n \rightarrow \infty$, $t_n \in \mathcal{T}$ such that $g_{t_n} h_s q \in K$.*

Proof. – Given C, α, D , let λ, ρ_0, C' be as in Theorem 9.1. Let ε be small enough so that

$$C' \left(\frac{\varepsilon}{\rho_0} \right)^{\lambda\alpha} < 1,$$

and let $K = K_\varepsilon$. Suppose to the contrary that for some (C, α) -decaying and D -Federer measure μ on some interval J_0 we have

$$\mu(A) > 0, \text{ where } A = \{s \in J_0 : \exists t_0 \forall t \in \mathcal{T} \cap (t_0, \infty), g_t h_s q \notin K\}.$$

Then there is $A_0 \subset A$ and $t_0 > 0$ such that $\mu(A_0) > 0$ and

$$(37) \quad s \in A_0, t \in \mathcal{T} \cap (t_0, \infty) \implies g_t h_s q \notin K.$$

By a general density theorem, see e.g., [13, Cor. 2.14], there is an interval J with $3J \subset J_0$ such that

$$(38) \quad \frac{\mu(A_0 \cap J)}{\mu(J)} > C' \left(\frac{\varepsilon}{\rho_0} \right)^{\lambda\alpha}.$$

We claim that by taking $t > t_0$ sufficiently large we can assume that for all $\delta \in \mathcal{L}_q$ there is $s \in J$ such that $\ell(\delta, g_t h_s q) \geq \rho_0$. This will guarantee that (28) holds for the horocycle $s \mapsto g_t h_s q = h_{e^t s} g_t q$, and conclude the proof since (37) and (38) contradict (29) (with $g_t q$ in place of q).

It remains to prove the claim. Let $\zeta = \phi(q)$ so that for any $\delta \in \mathcal{L}_q$,

$$\ell(\delta, q) = \max\{|x(\delta, q)|, |y(\delta, q)|\} \geq \zeta.$$

If $|x(\delta, q)| \geq \zeta$, then $|x(\delta, g_t q)| = e^{t/2}|x(\delta, q)| \geq \zeta e^{t/2}$, and if $|y(\delta, q)| \geq \zeta$ then the function

$$s \mapsto x(\delta, g_t h_s q) = e^{t/2}(x(\delta, q) + sy(\delta, q))$$

has slope $|e^{t/2}y(\delta, q)| \geq e^{t/2}\zeta$, hence $\sup_{s \in J} |x(\delta, g_t h_s q)| \geq \zeta e^{t/2} \frac{|J|}{2}$. Thus the claim holds when $\zeta e^{t/2} \geq \max\left\{\rho_0, \rho_0 \frac{|J|}{2}\right\}$. \square

This yields a strengthening of Theorem 6.1(a).

COROLLARY 9.3. – Suppose (\mathbf{a}, \mathbf{b}) is positive, and write $\mathbf{a}(s) = \mathbf{a} + s\mathbf{b}$. There is $\varepsilon_0 > 0$ such that given C, α, D there is $\zeta > 0$ such that if μ is (C, α) -decaying and D -Federer, and $\text{supp } \mu \subset (-\varepsilon_0, \varepsilon_0)$, then for μ -almost every s , $\limsup n\varepsilon_n(\mathbf{a}(s)) \geq \zeta$.

Proof. – Repeat the proof of Theorem 6.1, using Corollary 9.2 and Proposition 7.2 instead of [5]. □

10. Mahler's question for curves

In this section we prove Theorems 1.3 and 6.2 by deriving them from a stronger statement.

THEOREM 10.1. – Let $J \subset \mathbb{R}$ be a compact interval, let $\beta : J \rightarrow \mathbb{R}_+^d$ be a C^2 curve, let μ be a decaying Federer measure on J , and suppose that for every $s_1, s_2 \in J$, $(\beta(s_1), \beta'(s_2))$ is a positive pair. Then there is $\zeta > 0$ such that for μ -a.e. $s \in J$, $\limsup_{n \rightarrow \infty} n\varepsilon_n(\beta(s)) \geq \zeta$.

Derivation of Theorem 6.2 from Theorem 10.1. – If Theorem 6.2 is false then there is $A \subset I$ with $\mu(A) > 0$ such that for all $s \in A$, $\beta(s)$ is not of recurrence type but $(\beta(s), \beta'(s))$ is positive. Let $s_0 \in A \cap \text{supp } \mu$ so that $\mu(A \cap J) > 0$ for any open interval J containing s_0 . Since the set of positive pairs is open (Corollary 5.4), there is an open J containing s_0 such that $(\beta(s_1), \beta'(s_2))$ is positive for every $s_1, s_2 \in J$, so Theorem 10.1 implies that $\beta(s)$ is of recurrence type for almost every $s \in J$, a contradiction. □

Proof of Theorem 1.3. – Let $\mathbf{a}(x)$ be as in (2) and let $\| \cdot \|_1$ be the 1-norm on \mathbb{R}^d . Since unique ergodicity is unaffected by dilations, it is enough to verify the conditions of Theorem 6.2 for the permutation $\sigma(i) = d + 1 - i$, a decaying Federer measure μ , and for

$$\beta(s) = \frac{\mathbf{a}(s)}{\|\mathbf{a}(s)\|_1} = \frac{1}{s + \dots + s^d} (s, s^2, \dots, s^d).$$

For any connection (i, j, m) the set $\{\mathbf{a} \in \Delta : (i, j, m) \in \mathcal{L}_{\mathbf{a}}\}$ is a proper affine subspace of \mathbb{R}_+^d transversal to $\{(x_1, \dots, x_d) : \sum x_i = 1\}$, and since $\beta(s)$ is analytic and not contained in any such affine subspace, the set $\{s \in I : \beta(s) \text{ has connections}\}$ is countable, so $\beta(s)$ is without connections for μ -a.e. s .

Letting $R = R(s) = s + \dots + s^d$, we have

$$\beta'(s) = \frac{1}{R^2} (\gamma_1(s), \dots, \gamma_d(s)), \quad \text{where } \gamma_i(s) = \sum_{\ell=i}^{i+d-1} (2i - \ell - 1) s^\ell.$$

Then for $j = 1, \dots, d-1$ and $\ell = 1, \dots, j+d-1$, setting $a = \max\{1, \ell+1-d\}$, $b = \min\{j, \ell\}$ we find:

$$\begin{aligned} R^2 y_j &= \sum_{i=1}^j \gamma_i(s) = \sum_{\ell=1}^{j+d-1} \left(\sum_{i=a}^b (2i - \ell - 1) \right) s^\ell \\ &= \sum_{\ell=1}^{j+d-1} [b(b+1) - a(a-1) - (b-a+1)(\ell+1)] s^\ell. \end{aligned}$$

Considering separately the 3 cases $1 \leq \ell \leq j$, $j < \ell \leq d$, $d < \ell$ one sees that in every case $y_j < 0$. Our choice of σ insures that for $j = 1, \dots, d-1$, $y'_j = -y_{d-j} > 0$. This implies via (18) that $L < 0$ on I , thus for all s for which $\beta(s)$ is without connections, $(\beta(s), -\beta'(s))$

is positive. Define $\hat{\beta}(s) = \beta(-s)$, so that $(\hat{\beta}(s), \hat{\beta}'(s)) = (\beta(-s), -\beta'(-s))$ is positive for a.e. $s < 0$. Thus Theorem 6.2 applies to $\hat{\beta}$, proving the claim. \square

Proof of Theorem 10.1. – Let \mathcal{H} be the stratum corresponding to σ . For a (C, α) -decaying and D -Federer measure μ , let C', ρ_0, λ be the constants as in Theorem 9.1. Choose $\varepsilon > 0$ small enough so that

$$(39) \quad B = C' \left(\frac{\varepsilon}{\rho_0} \right)^{\lambda\alpha} < 1,$$

and let $K = K_\varepsilon$.

By making J smaller if necessary, we can assume that for all $s_1, s_2 \in J$, there is a translation surface $q(s_1, s_2) = \pi \circ \mathbf{q}(\beta(s_1), \beta'(s_2))$ corresponding to the positive pair $(\beta(s_1), \beta'(s_2))$ via Theorem 5.3. That is $Q = \{q(s_1, s_2) : s_i \in J\}$ is a bounded subset of \mathcal{H} and $q(s_1, s_2)$ depends continuously on s_1, s_2 . By appealing to Corollary 5.5, we can also assume that there is a fixed curve γ so that $\mathcal{T}_\sigma(\beta(s_1)) = \mathcal{T}(q(s_1, s_2), \gamma)$, for $s_1, s_2 \in J$. Define $q(s) = q(s, s)$. By rescaling we may assume with no loss of generality that $q(s) \in \mathcal{H}_1$ for all s , and by making K larger let us assume that $Q \subset K$. By continuity, the return times to γ along vertical leaves for $q(s_1, s_2)$ are uniformly bounded from above and below, and the length of γ with respect to the flat structure given by $q(s_1, s_2)$ is uniformly bounded.

We claim that there is C_1 , depending only on Q , such that for any interval $J_0 \subset \mathbb{R}$ with $0 \in J_0$, any $t > 0$, any $q \in Q$, any $\mathbf{q} \in \pi^{-1}(q)$ and any $\delta \in \mathcal{L}_{\mathbf{q}}$, we have

$$\sup_{s \in J_0} \ell(\delta, g_t h_s \mathbf{q}) \geq C_1 |J_0| e^{t/2}.$$

Here $|J_0|$ is the length of J_0 .

Indeed, let $\theta = \inf\{\phi(q) : q \in Q\}$, which is a positive number since Q is bounded. Let $C_1 = \min\left\{\frac{\theta}{|J_0|}, \frac{\theta}{2}\right\}$, and let $q \in Q$, $\mathbf{q} \in \pi^{-1}(q)$, $\delta \in \mathcal{L}_{\mathbf{q}}$. Then $\max\{|x(\delta, \mathbf{q})|, |y(\delta, \mathbf{q})|\} \geq \theta$. Suppose first that $|x(\delta, \mathbf{q})| \geq \theta$, then

$$\sup_{s \in J_0} \ell(\delta, g_t h_s \mathbf{q}) \geq |x(\delta, g_t \mathbf{q})| = e^{t/2} |x(\delta, \mathbf{q})| \geq \theta e^{t/2} \geq C_1 |J_0| e^{t/2}.$$

Now if $|y(\delta, \mathbf{q})| \geq \theta$ then the function

$$s \mapsto x(\delta, g_t h_s \mathbf{q}) = e^{t/2} (x(\delta, \mathbf{q}) + sy(\delta, \mathbf{q}))$$

has slope $|e^{t/2} y(\delta, \mathbf{q})| \geq e^{t/2} \theta$, hence

$$\sup_{s \in J_0} \ell(\delta, g_t h_s \mathbf{q}) \geq \sup_{s \in J_0} |x(\delta, g_t h_s \mathbf{q})| \geq e^{t/2} \theta |J_0| / 2 \geq C_1 |J_0| e^{t/2}.$$

This proves the claim.

For each $s_0, s \in J$ let $\mathbf{a}^{(s_0)}(s) = \beta(s_0) + \beta'(s_0)(s - s_0)$ be the linear approximation to β at s_0 . Using the fact that β is a C^2 -map, there is \tilde{C} such that

$$(40) \quad \max_{s \in J_0} \|\beta(s) - \mathbf{a}^{(s_0)}(s)\| < \tilde{C} |J_0|^2$$

whenever $J_0 \subset J$ is a subinterval centered at s_0 .

Let κ and c_2 be as in Proposition 7.2, let ε be as chosen in (39), let

$$C_2 = \frac{\rho_0}{2C_1}, \quad \zeta_1 < \left(\frac{\varepsilon}{\kappa}\right)^2 \quad \text{and} \quad \zeta = \zeta_1 \cdot \frac{c_2}{d^2 \tilde{C}}.$$

If the theorem is false then $\mu(A) > 0$, where

$$A = \{s \in J : \limsup n\varepsilon_n(\beta(s)) < \zeta\}.$$

Moreover there is N and $A_0 \subset A$ such that $\mu(A_0) > 0$ and

$$(41) \quad n \geq N, s \in A_0 \implies n\varepsilon_n(\beta(s)) < \zeta.$$

Using [13, Cor. 2.14] let $s_0 \in A_0$ be a density point, so that for any sufficiently small interval J_0 centered at s_0 we have

$$(42) \quad \mu(A_0 \cap J_0) > B\mu(J_0),$$

where B is as in (39).

For $t > 0$ we will write

$$c(t) = C_2 e^{-t/2} \quad \text{and} \quad J_t = B(s_0, c(t)).$$

Let $\mathbf{a}(s) = \mathbf{a}^{(s_0)}(s)$ and let $\tilde{q}(s) = h_{s-s_0}q(s_0)$, which is the surface $\pi \circ \mathbf{q}(\mathbf{a}(s), \beta'(s_0))$. The trajectory $s \mapsto g_t \tilde{q}(s) = g_t h_{s-s_0}q(s_0) = h_{e^t s} g_t h_{-s_0}q(s_0)$ is a horocycle path, which, by the claim and the choice of C_2 , satisfies (28) with $\rho = \rho_0$ and $J = J_t$ for all $t > 0$. Therefore

$$(43) \quad \mu \{s \in J_t : g_t \tilde{q}(s) \notin K\} \leq B\mu(J_t).$$

Now for a large $t > 0$ to be specified below, let

$$(44) \quad n_1 = \sqrt{\frac{\zeta_1}{2c_2}} e^{t/2} \quad \text{and} \quad n_2 = \frac{2c_2}{d^2 \tilde{C}} n_1.$$

By making \tilde{C} larger we can assume that $n_2 < n_1$. For large enough t we will have $n_2 > n_0$ (as in Proposition 7.2) and $n_2 > N$ (as in (41)).

We now claim

$$(45) \quad s \in J_t, n_2 \varepsilon_{n_2}(\beta(s)) < \zeta \implies n_1 \varepsilon_{n_1}(\mathbf{a}(s)) < \zeta_1.$$

Assuming this, note that by (24), (44) and the choice of ζ_1 , if $n_1 \varepsilon_{n_1}(\mathbf{a}(s)) < \zeta_1$ then $g_t \tilde{q}(s) \notin K$. Combining (41) and (45) we see that $A_0 \cap J_t \subset \{s \in J_t : g_t \tilde{q}(s) \notin K\}$ for all large enough t . Combining this with (42) we find a contradiction to (43).

It remains to prove (45). Let $r \leq n_2$, let $\mathcal{T} = \mathcal{T}_\sigma(\mathbf{a}(s))$, $\mathcal{P} = \mathcal{T}_\sigma(\beta(s))$, let x_i, x_j be discontinuities of \mathcal{T} and let x'_i, x'_j be the corresponding discontinuities of \mathcal{P} . By choice of \tilde{C} we have

$$\|\beta(s) - \mathbf{a}(s)\| < \tilde{C}e^{-t},$$

where $\|\cdot\|$ is the max norm on \mathbb{R}^d . Suppose first that x_i and x'_i have the same itinerary under \mathcal{T} and \mathcal{P} until the r th iteration; i.e., $\mathcal{T}^k x_i$ is in the ℓ th interval of continuity of \mathcal{T} if and only if $\mathcal{P}^k x'_i$ is in the ℓ th interval of continuity of \mathcal{P} for $k \leq r$. Then one sees from (9) and (10) that

$$|\mathcal{T}^r x_i - \mathcal{P}^r x'_i| \leq \sum_0^r d^2 \|\beta(s) - \mathbf{a}(s)\|.$$

Therefore, using (44),

$$\begin{aligned} |x_i - \mathcal{F}^r x_j - (x'_i - \mathcal{F}^r x'_j)| &\leq |x_i - x'_i| + |\mathcal{F}^r x_j - \mathcal{F}^r x'_j| \\ &\leq 2 \sum_1^r d^2 \|\beta(s) - \mathbf{a}(s)\| \\ &\leq 2d^2 n_2 \tilde{C} e^{-t} = \frac{\zeta_1}{2n_1}. \end{aligned}$$

If $n_1 \varepsilon_{n_1}(\mathbf{a}(s)) \geq \zeta_1$ then by the triangle inequality we find that $|x'_i - \mathcal{F}^r x'_j| < \zeta_1/2n_1 > 0$. In particular this shows that for all i , x_i and x'_i do have the same itinerary under \mathcal{F} and \mathcal{F} , and moreover

$$\begin{aligned} \varepsilon_{n_2}(\beta(s)) &\geq \varepsilon_{n_2}(\mathbf{a}(s)) - \frac{\zeta_1}{2n_1} \\ &\geq \varepsilon_{n_1}(\mathbf{a}(s)) - \frac{\zeta_1}{2n_1} \\ &> \frac{\zeta_1}{2n_1} = \zeta_1 \frac{n_2}{2n_1} \cdot \frac{1}{n_2} = \frac{\zeta}{n_2}, \end{aligned}$$

as required. □

11. Real REL

This section contains our results concerning the real REL foliation, whose definition we now recall. Let \mathbf{q} be a marked translation surface of type \mathbf{r} , where k (the number of singularities) is at least 2. Recall that \mathbf{q} determines a cohomology class $\text{hol}(\mathbf{q})$ in $H^1(S, \Sigma; \mathbb{R}^2)$, where for a relative cycle $\gamma \in H_1(S, \Sigma)$, the value \mathbf{q} takes on γ is $\text{hol}(\gamma, \mathbf{q})$. Also recall that there are open sets $\tilde{\mathcal{H}}_\tau$ in $\tilde{\mathcal{H}}(\mathbf{r})$, corresponding to a given triangulation τ of (S, Σ) , such that the map hol restricted to $\tilde{\mathcal{H}}_\tau$ endows $\tilde{\mathcal{H}}$ with a linear manifold structure. Now recall the map Res as in (4), let V_1 be the first summand in the splitting (3), and let $W = V_1 \cap \ker \text{Res}$, so that $\dim W = k - 1$, where $k = |\Sigma|$. The REL foliation is modeled on $\ker \text{Res}$, the real foliation is modeled on V_1 , and the real REL foliation is modeled on W . That is, a ball $\mathcal{U} \subset \tilde{\mathcal{H}}_\tau$ provides a product neighborhood for these foliations, where $\mathbf{q}, \mathbf{q}_1 \in \mathcal{U}$ belong to the same plaque for REL, real, or real REL, if $\text{hol}(\mathbf{q}) - \text{hol}(\mathbf{q}_1)$ belongs respectively to $\ker \text{Res}, V_1$, or W .

Recall that \mathcal{H} is an orbifold and the orbifold cover $\pi : \tilde{\mathcal{H}} \rightarrow \mathcal{H}$ is defined by taking a quotient by the $\text{Mod}(S, \Sigma)$ -action. Since hol is equivariant with respect to the action of the group $\text{Mod}(S, \Sigma)$ on $\tilde{\mathcal{H}}$ and $H^1(S, \Sigma; \mathbb{R}^2)$, and (by naturality of the splitting (3) and the sequence (4)) the subspaces V_1, W , and $\ker \text{Res}$ are $\text{Mod}(S, \Sigma)$ -invariant, the foliations defined by these subspaces on $\tilde{\mathcal{H}}$ descend naturally to \mathcal{H} . More precisely leaves in $\tilde{\mathcal{H}}$ descend to ‘orbifold leaves’ on \mathcal{H} , i.e., leaves in $\tilde{\mathcal{H}}$ map to immersed sub-orbifolds in \mathcal{H} .

PROPOSITION 11.1. – *The REL and real REL leaves on both $\tilde{\mathcal{H}}$ and \mathcal{H} have a well-defined translation structure.*

Proof. – A translation structure on a leaf \mathcal{L} amounts to saying that there is a fixed vector space V and an atlas of charts on \mathcal{L} taking values in V , so that transition maps are translations. We take $V = \ker \text{Res}$ for the REL leaves, and $V = W$ for the real REL leaves. For each $\mathbf{q}_0 \in \tilde{\mathcal{H}}$, the atlas is obtained by taking the chart $\mathbf{q} \mapsto \text{hol}(\mathbf{q}) - \text{hol}(\mathbf{q}_0)$. By the definition of the corresponding foliations, these are homeomorphisms in a sufficiently small neighborhood of \mathbf{q}_0 , and the fact that transition maps are translations is immediate. In order to check that this descends to \mathcal{H} , recall that $\mathcal{H} = \tilde{\mathcal{H}} / \text{Mod}(S, \Sigma)$, and let $\varphi \in \text{Mod}(S, \Sigma)$. We need to show that if $\hat{\mathbf{q}} = \mathbf{q} \circ \varphi$, $\hat{\mathbf{q}}_0 = \mathbf{q}_0 \circ \varphi$, then $\text{hol}(\mathbf{q}) - \text{hol}(\mathbf{q}_0) = \text{hol}(\hat{\mathbf{q}}) - \text{hol}(\hat{\mathbf{q}}_0)$. Since hol is $\text{Mod}(S, \Sigma)$ -equivariant, this amounts to checking that φ acts trivially on $\ker \text{Res}$. Invoking (4), this follows from our convention that any $\varphi \in \text{Mod}(S, \Sigma)$ fixes each point of Σ , so acts trivially on $H^0(\Sigma)$. \square

It is an interesting question to understand the geometry of individual leaves. For the REL foliation this is a challenging problem, but for the real REL foliation, our main theorems give a complete answer. Given a marked translation surface \mathbf{q} we say that a saddle connection δ for \mathbf{q} is *horizontal* if $y(\mathbf{q}, \delta) = 0$. Note that for a generic \mathbf{q} there are no horizontal saddle connections, and in any stratum there is a uniform upper bound on their number.

Let $\mathcal{O}_{\mathbf{q}} \subset W$ denote those elements whose x -coordinate is positive on every horizontal saddle connection of \mathbf{q} . This is an open convex cone in W , and equal to all of W if \mathbf{q} has no horizontal saddle connections. Note that $\text{hol}(\mathbf{q}) \in \mathcal{O}_{\mathbf{q}}$ by the easy direction in Theorem 1.1. Let $\mathcal{V}_{\mathbf{q}} = \mathcal{O}_{\mathbf{q}} - \text{hol}(\mathbf{q})$.

THEOREM 11.2. – *Let $\mathbf{q} \in \tilde{\mathcal{H}}$. Then there is a continuous map $\psi : \mathcal{V}_{\mathbf{q}} \rightarrow \tilde{\mathcal{H}}$ such that for any $\mathbf{c} \in \mathcal{V}_{\mathbf{q}}$,*

$$(46) \quad \text{hol}(\psi(\mathbf{c})) = \text{hol}(\mathbf{q}) + \mathbf{c},$$

and the horizontal foliation of $\psi(\mathbf{c})$ is the same as that of \mathbf{q} . Moreover the image of ψ is contained in the REL leaf of \mathbf{q} .

Proof. – Let \mathcal{F} and \mathcal{G} be the vertical and horizontal foliations of \mathbf{q} respectively. We will apply Theorem 1.2, reversing the roles of the horizontal and vertical foliations of \mathbf{q} ; that is we use \mathcal{G} in place of \mathcal{F} . Writing $\text{hol}(\mathbf{q}) = (x(\mathbf{q}), y(\mathbf{q}))$ and $\mathbf{c} = (x_{\mathbf{c}}, 0) \in \mathcal{V}_{\mathbf{q}}$, we have

$$\text{hol}(\mathbf{q}) + \mathbf{c} = (x(\mathbf{q}) + x_{\mathbf{c}}, y(\mathbf{q})).$$

We must check that this cohomology class is contained in $\mathbb{B}(\mathcal{G}) \times \mathbb{A}(\mathcal{G})$ (notation as in Theorem 1.2), which will enable us to conclude that

$$\psi(\mathbf{c}) \equiv \text{hol}^{-1}(x(\mathbf{q}) + x_{\mathbf{c}}, y(\mathbf{q}))$$

is continuous and satisfies (46).

Clearly $y(\mathbf{q})$, the cohomology class represented by \mathcal{G} , is in $\mathbb{A}(\mathcal{G})$. To check that $x(\mathbf{q}) + x_{\mathbf{c}} \in \mathbb{B}(\mathcal{G})$, we need to show that for any element $\delta \in H_+^{\mathcal{G}}$, $x(\mathbf{q}, \delta) + x_{\mathbf{c}}(\delta) > 0$. To see this, we treat separately the cases when δ is a horizontal saddle connection, and when δ is represented by an element of $H_1(S, \Sigma)$ which lies in $H_1(S)$. This suffices because by Lemma 2.1 every element of $\mathcal{H}_+^{\mathcal{G}}$ is a convex combination of these two cases.

If δ is represented by a saddle connection, then by definition of $\mathcal{V}_{\mathbf{q}}$ we have $x(\mathbf{q}, \delta) + x_{\mathbf{c}}(\delta) > 0$.

If δ is in $H_1(S)$, since $\mathbf{c} \in W \subset \ker \text{Res}$ we have $x_{\mathbf{c}}(\delta) = 0$, and the easy direction in Theorem 1.1 implies

$$x(\mathbf{q}, \delta) + x_{\mathbf{c}}(\delta) = x(\mathbf{q}, \delta) > 0.$$

To check $\psi(\mathbf{c})$ is contained in the real REL leaf of \mathbf{q} , it suffices to note that according to (46), in any local chart provided by hol , $\psi(\mathbf{c}_t)$ moves along plaques of the foliation. \square

Since the leaves are modeled on W , in the case when $\mathcal{V}_{\mathbf{q}} = W$ we obtain:

COROLLARY 11.3. – *If \mathbf{q} has no horizontal saddle connections, then there is a homeomorphism $\psi : W \rightarrow \tilde{\mathcal{H}}$ onto the leaf of \mathbf{q} .*

Moreover the above maps are compatible with the transverse structure for the real REL foliation. Namely, let $\psi_{\mathbf{q}}$ be the map in Theorem 11.2. We have:

COROLLARY 11.4. – *For any $\mathbf{q}_0 \in \tilde{\mathcal{H}}$ there is a neighborhood \mathcal{V} of 0 in W , and a submanifold $\mathcal{U} \subset \tilde{\mathcal{H}}$ everywhere transverse to real-REL leaves and containing \mathbf{q}_0 such that:*

1. *For any $\mathbf{q} \in \mathcal{U}$, $\psi_{\mathbf{q}}$ is defined on \mathcal{V} .*
2. *The map $\mathcal{U} \times \mathcal{V} \rightarrow \tilde{\mathcal{H}}$ defined by $(\mathbf{q}, \mathbf{c}) \mapsto \psi_{\mathbf{q}}(\mathbf{c})$ is a homeomorphism onto a neighborhood \mathcal{N} of \mathbf{q}_0 .*
3. *Each plaque of the real REL foliation in \mathcal{N} is of the form $\psi_{\mathbf{q}}(\mathcal{V})$.*

Proof. – Let \mathcal{N}_0 be a bounded neighborhood of \mathbf{q}_0 on which hol is a local homeomorphism and let $\mathcal{U}_0 \subset \mathcal{N}_0$ be any submanifold everywhere transverse to the real REL leaves and of complementary dimension. For example we can take \mathcal{U}_0 to be the pre-image under hol of a small ball around $\text{hol}(\mathbf{q}_0)$ in an affine subspace of $H^1(S, \Sigma; \mathbb{R}^2)$ which is complementary to W .

Since \mathcal{N}_0 is bounded there is a uniform lower bound $r > 0$ on the lengths of saddle connections for $\mathbf{q} \in \mathcal{N}_0$. Thus a ball B of radius $r/2$ around 0 in W is contained in $\mathcal{V}_{\mathbf{q}}$ for each $\mathbf{q} \in \mathcal{N}_0$.

Thus, for any $\mathbf{q} \in \mathcal{U}_0$ we can apply Theorem 11.2 to define $\psi_{\mathbf{q}}$ on B , and after choosing sufficiently small $\mathcal{U} \subset \mathcal{U}_0$ and $\mathcal{V} \subset B$, we may assume that $\psi_{\mathbf{q}}(\mathcal{V}) \subset \mathcal{N}_0$ for all $\mathbf{q} \in \mathcal{U}$.

From (46) and the choice of \mathcal{U} it follows that $\{\psi_{\mathbf{q}}(\mathbf{c}) : \mathbf{q} \in \mathcal{U}, \mathbf{c} \in \mathcal{V}\}$ is a neighborhood of $\text{hol}(\mathbf{q}_0)$ in $H^1(S, \Sigma; \mathbb{R}^2)$. Assertions (2), (3) now follow from the fact that $\text{hol}|_{\mathcal{N}_0}$ is a homeomorphism. \square

Proof of Theorem 1.4. – Let $\mathcal{H}, \tilde{\mathcal{H}}$ be as above, let $\pi : \tilde{\mathcal{H}} \rightarrow \mathcal{H}$ be the projection, let \tilde{Q} denote the subset of translation surfaces in $\tilde{\mathcal{H}}$ without horizontal saddle connections, and let $Q = \pi(\tilde{Q})$. Note that Q and \tilde{Q} are B -invariant, where B is the subgroup of upper triangular matrices in G , acting on $\mathcal{H}, \tilde{\mathcal{H}}$ in the usual way. We will extend the B -action to an action of $F = B \times W$.

The action of W is defined as follows. For each $\mathbf{q} \in \tilde{Q}$, $\mathcal{V}_{\mathbf{q}} = W$ and Theorem 11.2 allows us to define $\mathbf{c}\mathbf{q} = \psi_{\mathbf{q}}(\mathbf{c})$. The group action law

$$(47) \quad (\mathbf{c}_1 + \mathbf{c}_2)\mathbf{q} = \mathbf{c}_1(\mathbf{c}_2\mathbf{q})$$

follows from (46), associativity of addition in $H^1(S, \Sigma; \mathbb{R}^2)$, and the uniqueness in Lemma 4.1. Thus we have defined an action of W on \tilde{Q} , and by Proposition 11.1 this descends to a well-defined action on Q . To see that the action map $W \times \tilde{Q} \rightarrow \tilde{Q}$ is continuous, we take $w_n \rightarrow w$ in W , $\mathbf{q}^{(n)} \rightarrow \mathbf{q}$ in \tilde{Q} , and need to show that

$$(48) \quad w_n \mathbf{q}^{(n)} \rightarrow w \mathbf{q}.$$

Corollary 11.4 implies that for any t there are neighborhoods \mathcal{U} of $(tw)\mathbf{q}$ and \mathcal{V} of 0 in W such that the map

$$\mathcal{U} \times \mathcal{V} \rightarrow \tilde{\mathcal{H}}, \quad (\mathbf{q}, \mathbf{c}) \mapsto \mathbf{c}\mathbf{q}$$

is continuous. By compactness we can find $\mathcal{U}_i, i = 1, \dots, k$, a partition $0 = t_0 < \dots < t_k = 1$, and a fixed open $\mathcal{V} \subset W$ such that

- $\{(tw)\mathbf{q} : t \in [t_{i-1}, t_i]\} \subset \mathcal{U}_i$.
- $(t_i - t_{i-1})w \in \mathcal{V}$.

It now follows by induction on i that $t_i w^{(n)} \mathbf{q}_n \rightarrow (t_i w)\mathbf{q}$ for each i , and putting $i = k$ we get (48).

The action is affine and measure preserving since in the local charts given by $H^1(S, \Sigma; \mathbb{R}^2)$, it is defined by vector addition. Since the area of a surface can be computed in terms of its absolute periods alone, this action preserves the subset of area-one surfaces. A simple calculation using (8) shows that for $\mathbf{q} \in \tilde{Q}$, $\mathbf{c} \in W$ and

$$g = \begin{pmatrix} a & b \\ 0 & a^{-1} \end{pmatrix} \in B,$$

we have $g\mathbf{c}\mathbf{q} = (ac)g\mathbf{q}$. That is, the actions of B and W respect the commutation relation defining F , so that we have defined an F -action on \tilde{Q} . To check continuity of the action, let $f_n \rightarrow f$ in F and $\mathbf{q}_n \rightarrow \mathbf{q} \in \tilde{Q}$. Since F is a semi-direct product we can write $f_n = w_n b_n$, where $w_n \rightarrow w$ in W and $b_n \rightarrow b$ in B . Since the B -action is continuous, $b_n \mathbf{q}_n \rightarrow b\mathbf{q}$, and since (as verified above) the W -action on \tilde{Q} is continuous,

$$f_n \mathbf{q}_n = w_n (b_n \mathbf{q}_n) \rightarrow w (b\mathbf{q}) = f\mathbf{q}. \quad \square$$

12. Cones of good directions

Suppose σ is an irreducible and admissible permutation, and recall that $(\mathbf{a}, \mathbf{b}) \in \mathbb{R}_+^d \times \mathbb{R}^d$ is a positive pair if

$$(49) \quad \int L d\mu > 0 \text{ for any } \mu \in \mathcal{M}_{\mathbf{a}}$$

and

$$(50) \quad \sum_{n=0}^{m-1} L(\mathcal{I}^n x_i) > y_i - y_j \text{ for any } (i, j, m) \in \mathcal{L}_{\mathbf{a}}.$$

Let

$$(51) \quad \mathcal{C}_{\mathbf{a}} = \{\mathbf{b} \in \mathbb{R}^d : (\mathbf{a}, \mathbf{b}) \text{ is positive}\} \subset T_{\mathbf{a}}\mathbb{R}_+^d.$$

As we showed in Corollary 5.5, this is the set of good directions at the tangent space of \mathbf{a} , i.e., the directions of lines which may be lifted to horocycles. In this section we will relate the

cones $\mathcal{C}_{\mathbf{a}}$ with the bilinear form Q as in (11), and show there are ‘universally good’ directions for σ , i.e., specify certain \mathbf{b} such that $\mathbf{b} \in \mathcal{C}_{\mathbf{a}}$ for all \mathbf{a} which are without connections. We will also find ‘universally bad’ directions, i.e., directions which do not belong to $\mathcal{C}_{\mathbf{a}}$ for any \mathbf{a} ; these will be seen to be related to real REL.

Set $\mathcal{C}_{\mathbf{a}}^+ = \{\mathbf{b} \in \mathbb{R}^d : (49) \text{ holds}\}$, so that $\mathcal{C}_{\mathbf{a}} \subset \mathcal{C}_{\mathbf{a}}^+$ for all \mathbf{a} , and $\mathcal{C}_{\mathbf{a}}^+ = \mathcal{C}_{\mathbf{a}}$ when \mathbf{a} has no connections. Now let

$$(52) \quad \mathcal{C} = \{\mathbf{b} \in \mathbb{R}^d : \forall i, Q(\mathbf{e}_i, \mathbf{b}) > 0\}.$$

We have:

PROPOSITION 12.1. – \mathcal{C} is a nonempty open convex cone, and $\mathcal{C} = \bigcap_{\mathbf{a}} \mathcal{C}_{\mathbf{a}}^+$.

Proof. – It is clear that \mathcal{C} is open and convex. It follows from (19) that

$$\mathcal{C} = \{\mathbf{b} \in \mathbb{R}^d : \forall x \in I, \forall \mathbf{a} \in \mathbb{R}_+^d, L_{\mathbf{a}, \mathbf{b}}(x) > 0\}.$$

The irreducibility of σ implies that $\mathbf{b}_0 = (b_1, \dots, b_d)$ defined by $b_i = \sigma(i) - i$ (as in [9]) satisfies $y_i(\mathbf{b}_0) > 0 > y'_i(\mathbf{b}_0)$ for all i , so by (18), $L_{\mathbf{a}, \mathbf{b}_0}$ is everywhere positive irrespective of \mathbf{a} . This shows that \mathbf{b}_0 belongs to \mathcal{C} , and moreover that \mathcal{C} is contained in $\mathcal{C}_{\mathbf{a}}^+$ for any \mathbf{a} .

For the inclusion $\bigcap_{\mathbf{a}} \mathcal{C}_{\mathbf{a}}^+ \subset \mathcal{C}$, suppose $\mathbf{b} \notin \mathcal{C}_{\mathbf{a}}^+$, so that for some $\mu \in \mathcal{M}_{\mathbf{a}}$ we have $\int L d\mu \leq 0$. Writing $\mathbf{a}' = (a'_1, \dots, a'_d)$, where $a'_j = \mu(I_j)$, we have

$$Q(\mathbf{a}', \mathbf{b}) = \sum a'_i Q(\mathbf{e}_i, \mathbf{b}) = \int L d\mu \leq 0,$$

so that $\mathbf{b} \notin \mathcal{C}$. □

Note that in the course of the proof of Theorem 1.3, we actually showed that $\mathbf{a}'(s) \in -\mathcal{C}$ for all $s > 0$. Indeed, given a curve $\{\alpha(s)\} \subset \mathbb{R}_+^d$, the easiest way to show that $\alpha(s)$ is uniquely ergodic for a.e. s , is to show that $\alpha'(s) \in \pm \mathcal{C}$ for a.e. s .

Let \mathbf{R} denote the null-space of Q , that is

$$\mathbf{R} = \{\mathbf{b} \in \mathbb{R}^d : Q(\cdot, \mathbf{b}) \equiv 0\}.$$

PROPOSITION 12.2. – We have $\mathbf{R} = \mathbb{R}^d \setminus \bigcup_{\mathbf{a}} \pm \mathcal{C}_{\mathbf{a}}$.

Proof. – By (19), $\mathbf{R} = \{\mathbf{b} \in \mathbb{R}^d : L_{\mathbf{a}, \mathbf{b}}(x) \equiv 0\}$, so that containment \subset is clear. Now suppose $\mathbf{b} \notin \mathbf{R}$, that is $Q(\mathbf{e}_i, \mathbf{b}) \neq 0$ for some i , and by continuity there is an open subset \mathcal{U} of \mathbb{R}_+^d such that $Q(\mathbf{a}, \mathbf{b}) \neq 0$ for $\mathbf{a} \in \mathcal{U}$. Now taking $\mathbf{a} \in \mathcal{U}$ which is uniquely ergodic and without connections we have $\mathbf{b} \in \pm \mathcal{C}_{\mathbf{a}}$. □

Consider the map $(\mathbf{a}, \mathbf{b}) \mapsto \mathbf{q}(\mathbf{a}, \mathbf{b})$ defined in Theorem 5.3. It is easy to see that the image of an open subset of $\mathbb{R}_+^d \times \{\mathbf{b}\}$ is a plaque for the real foliation. Additionally, recalling that $Q(\mathbf{a}, \mathbf{b})$ records the intersection pairing on $H_1(S, \Sigma) \times H_1(S \setminus \Sigma)$, and that the intersection pairing gives the duality $H_1(S \setminus \Sigma) \cong H^1(S, \Sigma)$, one finds that the image of $\mathbf{R} \times \{\mathbf{b}\}$ is a plaque for the real REL foliation. That is, Proposition 12.2 says that the tangent directions in $T\mathbb{R}_+^d$ which can never be realized as horocycle directions, are precisely the directions in the real REL leaf.

BIBLIOGRAPHY

- [1] F. BONAHOON, Shearing hyperbolic surfaces, bending pleated surfaces and Thurston's symplectic form, *Ann. Fac. Sci. Toulouse Math.* **5** (1996), 233–297.
- [2] K. CALTA, K. WORTMAN, On unipotent flows in $\mathcal{H}(1, 1)$, *Ergodic Theory Dynam. Systems* **30** (2010), 379–398.
- [3] Y. CHEUNG, H. MASUR, A divergent Teichmüller geodesic with uniquely ergodic vertical foliation, *Israel J. Math.* **152** (2006), 1–15.
- [4] A. ESKIN, J. MARKLOF, D. WITTE MORRIS, Unipotent flows on the space of branched covers of Veech surfaces, *Ergodic Theory Dynam. Systems* **26** (2006), 129–162.
- [5] S. KERCKHOFF, H. MASUR, J. SMILLIE, Ergodicity of billiard flows and quadratic differentials, *Ann. of Math.* **124** (1986), 293–311.
- [6] D. KLEINBOCK, Some applications of homogeneous dynamics to number theory, in *Smooth ergodic theory and its applications (Seattle, WA, 1999)*, Proc. Sympos. Pure Math. **69**, Amer. Math. Soc., 2001, 639–660.
- [7] D. KLEINBOCK, B. WEISS, Bounded geodesics in moduli space, *Int. Math. Res. Not.* **2004** (2004), 1551–1560.
- [8] D. KLEINBOCK, B. WEISS, Badly approximable vectors on fractals, *Israel J. Math.* **149** (2005), 137–170.
- [9] H. MASUR, Interval exchange transformations and measured foliations, *Ann. of Math.* **115** (1982), 169–200.
- [10] H. MASUR, Hausdorff dimension of the set of nonergodic foliations of a quadratic differential, *Duke Math. J.* **66** (1992), 387–442.
- [11] H. MASUR, J. SMILLIE, Hausdorff dimension of sets of nonergodic measured foliations, *Ann. of Math.* **134** (1991), 455–543.
- [12] H. MASUR, S. TABACHNIKOV, Rational billiards and flat structures, in *Handbook of dynamical systems, Vol. 1A*, North-Holland, 2002, 1015–1089.
- [13] P. MATTILA, *Geometry of sets and measures in Euclidean spaces*, Cambridge Studies in Advanced Math. **44**, Cambridge Univ. Press, 1995.
- [14] C. T. MCMULLEN, Foliations of Hilbert modular surfaces, *Amer. J. Math.* **129** (2007), 183–215.
- [15] Y. MINSKY, B. WEISS, Nondivergence of horocyclic flows on moduli space, *J. reine angew. Math.* **552** (2002), 131–177.
- [16] R. C. PENNER, J. L. HARER, *Combinatorics of train tracks*, Annals of Math. Studies **125**, Princeton Univ. Press, 1992.
- [17] S. SCHWARTZMAN, Asymptotic cycles, **3** (2008), no. 12, 5463, http://www.scholarpedia.org/article/Asymptotic_cycles.
- [18] D. SULLIVAN, Cycles for the dynamical study of foliated manifolds and complex manifolds, *Invent. Math.* **36** (1976), 225–255.
- [19] W. P. THURSTON, Minimal stretch maps between hyperbolic surfaces, preprint arXiv:math.GT/9801039.
- [20] W. A. VEECH, Interval exchange transformations, *J. Analyse Math.* **33** (1978), 222–272.

- [21] W. A. VEECH, Gauss measures for transformations on the space of interval exchange maps, *Ann. of Math.* **115** (1982), 201–242.
- [22] W. A. VEECH, Measures supported on the set of uniquely ergodic directions of an arbitrary holomorphic 1-form, *Ergodic Theory Dynam. Systems* **19** (1999), 1093–1109.
- [23] Y. B. VOROBETS, Plane structures and billiards in rational polygons: the Veech alternative, *Russian Math. Surveys* **51** (1996), 779–817.
- [24] A. N. ZEMLJAKOV, A. B. KATOK, Topological transitivity of billiards in polygons, *Mat. Zametki* **18** (1975), 291–300.
- [25] A. ZORICH, Flat surfaces, in *Frontiers in number theory, physics, and geometry. I* (P. Cartier, B. Julia, P. Moussa, P. Vanhove, eds.), Springer, 2006, 437–583.

(Manuscrit reçu le 28 juillet 2011 ;
accepté, après révision, le 12 avril 2013.)

Yair MINSKY
Yale University
Mathematics Department
P.O. Box 208283
New Haven, CT 06520-8283, U.S.A.
E-mail: yair.minsky@yale.edu

Barak WEISS
Ben Gurion University
Department of Mathematics
Be'er Sheva 84105, Israel
E-mail: barakw@math.bgu.ac.il