FUNDAMENTA
MATHEMATICAE
180 (2003)

Identifying points of a pseudo-Anosov homeomorphism

by

Gavin Band (Coventry and Warszawa)

Abstract. We investigate the question, due to S. Smale, of whether a hyperbolic automorphism T of the n-dimensional torus can have a compact invariant subset homeomorphic to a compact manifold of positive dimension, other than a finite union of subtori. In the simplest case such a manifold would be a closed surface. A result of Fathi says that T can sometimes have an invariant subset which is a finite-to-one image of a closed surface under a continuous map which is locally injective except possibly at a finite number of points, these being the singularities of the invariant foliations of a suitable pseudo-Anosov homeomorphism. For a class of pseudo-Anosov homeomorphisms whose invariant foliations are of a particularly simple type, we show that this map is never locally injective at the singularities. The proof involves finding pairs of points having lifts in the universal abelian cover whose orbits are similar, and in fact we find whole pairs of horseshoes worth of such points.

1. Introduction. One of the basic questions one can ask about a dynamical system is: what are its invariant subsets? In [Hir70] this question was asked of hyperbolic dynamical systems, and in particular of the hyperbolic toral automorphisms. Several people worked in this area, but one of the questions that still remains is

Question 1.1. *Can a hyperbolic toral automorphism have a compact invariant topological submanifold of positive dimension, different from a finite union of subtori?*

The simplest such submanifolds would be closed orientable surfaces of genus $g > 1$. Let us say straightaway that any such surface would have to be embedded in a rather complicated way: like most invariant subsets of Anosov diffeomorphisms, it could not contain a nontrivial Lipschitz arc [Man78].

A good way to begin studying Question 1.1 was provided in [Fat88]. Suppose f is a pseudo-Anosov homeomorphism of the closed orientable surface $M = M_g$ of genus $g > 1$ whose stable and unstable foliations are orientable. Raising f to some power if necessary we may assume that f

2000 *Mathematics Subject Classification*: 37B05, 37D10, 37D20, 37E30.

preserves the orientations of its invariant foliations. In this situation, the expansion factor $\lambda > 1$ of f and its inverse λ^{-1} are each known to appear as simple eigenvalues of the map f_* induced by f in first homology. They are, in particular, algebraic integers, and there are two distinct possibilities: either λ and λ^{-1} are conjugate over $\mathbb{Q}$, or not. In the first case we let P be the minimum polynomial of λ (over $\mathbb{Z}$), which is also the minimum polynomial of λ^{-1}; in the second, we let P be the product of the minimum polynomial of λ and that of λ^{-1}. In either case P is a factor of the characteristic polynomial $\chi(f_*)$ of f_* having both λ and λ^{-1} as roots, and whose degree n satisfies $2 \leq n \leq 2g$.

Suppose that P has no roots of modulus 1. Fathi applied a theorem of J. Franks ([Fra70]) to show that there exists a nontrivial continuous map $\varphi : M \to \mathbb{T}^n$ which makes the diagram

$$\begin{array}{ccc} M & \xrightarrow{f} & M \\ {\scriptstyle\varphi}\downarrow & & \downarrow{\scriptstyle\varphi} \\ \mathbb{T}^n & \xrightarrow{T} & \mathbb{T}^n \end{array}$$

commute. Here T is a hyperbolic automorphism of $\mathbb{T}^n$ whose characteristic polynomial is P, and φ is surjective on first homology groups. We ask

QUESTION 1.2. *Is φ injective?*

If φ is injective, then it is a homeomorphism, and $\operatorname{im}\varphi$ is a compact invariant subset for T answering Question 1.1 affirmatively. However, to the best of our knowledge no examples are known for which φ is injective.

In [Fat88] it is proved that φ is locally injective away from the singularities of the invariant foliations of the pseudo-Anosov map f (and is consequently finite-to-one). The purpose of this paper is to show that φ need not be locally injective at the singularities of the foliations.

Results. The paper is structured as follows. In Section 2 we define an equivalence relation $\sim$ between points of M and show that equivalent points must be identified by Fathi's map φ. By an *f-horseshoe* we will mean a map $\epsilon : \Sigma \to M$ satisfying $\epsilon\sigma = f\epsilon$, where σ is the left-shift map of the two-sided shift space $\Sigma = \{0,1\}^{\mathbb{Z}}$. In Section 3 we associate f-horseshoes to certain rectangles in M, the horseshoes fitting into the rectangles in a nice way. By comparing paths which are the diagonals of such rectangles, we show how pairs of f-horseshoes can arise, which are equivalent in the sense that corresponding points of their images are equivalent. Then in Section 4 we identify a class $\mathcal{C}$ of pseudo-Anosov maps whose invariant foliations are of a particularly simple type, having in particular only one singularity, and exploit the special geometric features of maps in this class to prove the following theorem.

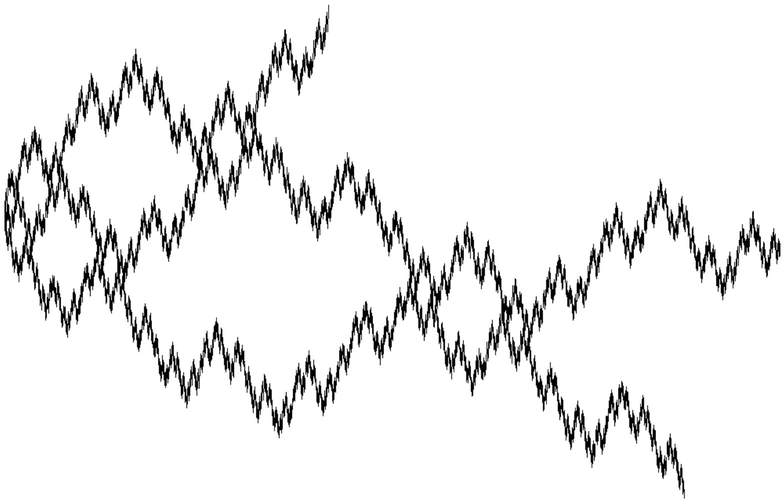

Fig. 1.1. The behaviour of φ for $f \in \mathcal{C}$

THEOREM 1.3. *Let $f \in \mathcal{C}$. For every stable prong s_1 of the singularity p there is a stable prong $s_2 \neq s_1$ of p and f-horseshoes ϵ_1 and ϵ_2 corresponding to rectangles $\Lambda_1, \Lambda_2 \subset M$ such that*

(1) *either the bottom edge or the top edge of Λ_i is a subinterval of s_i having p as one endpoint*, *for $i = 1, 2$;*

(2) *for all $z \in \Sigma$ we have $\epsilon_1(z) \sim \epsilon_2(z)$ but $\epsilon_1(z) \neq \epsilon_2(z)$*, *except that if z is one of the two fixed points of σ then $\epsilon_1(z) = \epsilon_2(z) = p$.*

In particular Fathi's map φ (if it exists for f) is not locally injective at p.

The behaviour produced by Theorem 1.3 is illustrated in Figure 1.1. This figure depicts the image under φ of the three outwardly-oriented stable prongs of the singularity p, corresponding to a certain pseudo-Anosov map f of M_2 which is described in more detail in [Ban03]. The relevant torus has dimension 4 and the two dimensions depicted are those of the stable manifold of the fixed point $\varphi(p)$ (which is the leftmost point in the figure).

Acknowledgements. I would like to express my thanks to my PhD supervisor Anthony Manning, who suggested Question 1.2 to me and taught me how to study it. I would also like to thank Feliks Przytycki, Godofredo Iommi and the referee for their many helpful comments.

2. Equivalent points and identifications. We work with a pseudo-Anosov map f of the smooth closed orientable surface $M = M_g$ of genus $g \geq 1$. Thus f preserves a transverse pair $(\mathcal{F}^{\mathrm{u}}, \mathcal{F}^{\mathrm{s}})$ of singular measured foliations, expanding uniformly along $\mathcal{F}^{\mathrm{u}}$ by a factor of $\lambda > 1$ and contracting uniformly along $\mathcal{F}^{\mathrm{s}}$ by a factor of λ^{-1}. We denote the finite set of singularities of these foliations by Z. We will assume throughout that $\mathcal{F}^{\mathrm{u}}$ and $\mathcal{F}^{\mathrm{s}}$ are orientable. Since Fathi's map φ (if it exists) is not changed by raising f to some power, we may also assume without loss of generality that f preserves the orientation of its invariant foliations, and also when necessary that f has a fixed point (to be used as a base point).

Rather than working directly with transverse measures we will work with "transverse" smooth real 1-forms ω^{s} and ω^{u}, each vanishing on Z. These are defined by the requirement that for any smooth path α in M, the integrals

$$\langle \omega^{\mathrm{s}}, \alpha\rangle = \int_{\alpha} \omega^{\mathrm{s}} \quad \text{and} \quad \langle \omega^{\mathrm{u}}, \alpha\rangle = \int_{\alpha} \omega^{\mathrm{u}}$$

are the measures of α, measured, respectively, transverse to $\mathcal{F}^{\mathrm{s}}$ and to $\mathcal{F}^{\mathrm{u}}$; the sign of these integrals reflects the orientation of α relative to that of $\mathcal{F}^{\mathrm{s}}$ and $\mathcal{F}^{\mathrm{u}}$. For vectors v and w in the tangent space to M at a point x, the expression $(v, w) = \omega^{\mathrm{u}}(v)\omega^{\mathrm{u}}(w) + \omega^{\mathrm{s}}(v)\omega^{\mathrm{s}}(w)$ defines a "singular" Riemannian metric on M which makes $M \setminus Z$ a flat Riemannian manifold, and which induces a metric d on the whole of M.

We denote by $\widetilde{M}_0$ the *universal abelian cover* of M, this being the covering space of M with group of deck transformations

$$\pi_1(M)/[\pi_1(M), \pi_1(M)] \cong H_1(M; \mathbb{Z}).$$

The Riemannian metric on M lifts to one on $\widetilde{M}_0$, and we denote again by d the corresponding metric.

The following definition is analogous to that of global shadowing (attributed in [Han85] to A. Katok), which, however, works in the universal covering space rather than in $\widetilde{M}_0$.

DEFINITION 2.1. For $x, y \in M$, we write $x \sim y$ (for the homeomorphism f) if x and y have lifts $\widetilde{x}$ and $\widetilde{y}$ in $\widetilde{M}_0$ such that $d(\widetilde{f}^k\widetilde{x}, \widetilde{f}^k\widetilde{y})$, $k \in \mathbb{Z}$, is bounded.

Here $\widetilde{f}$ denotes a lift of f to $\widetilde{M}_0$; since any two lifts differ by an isometry covering the identity, the above definition is independent of the choice of $\widetilde{f}$.

The usefulness of the definition for our purposes comes from the following lemma. We suppose ourselves in the situation of the introduction, so we have a hyperbolic automorphism T of $\mathbb{T}^n$ and a map $\varphi : M \to \mathbb{T}^n$ with $\varphi f = T\varphi$.

LEMMA 2.2. *If $x \sim y$ then $\varphi(x) = \varphi(y)$.*

Proof. Choose lifts $\widetilde{x}$ and $\widetilde{y}$ of x and y in $\widetilde{M}_0$ such that $d(\widetilde{f}^k\widetilde{x}, \widetilde{f}^k\widetilde{y})$ is bounded, by $D > 0$, say. Since the fundamental group of $\mathbb{T}^n$ is abelian, φ lifts to a map $\widetilde{\varphi} : \widetilde{M}_0 \to \mathbb{R}^n$ of covering spaces. Choose the lift $\widetilde{T}$ of T to $\mathbb{R}^n$ for which $\widetilde{\varphi}\widetilde{f} = \widetilde{T}\widetilde{\varphi}$. If $\widetilde{p}$ is the lift to $\widetilde{M}_0$ of our chosen base point which is fixed by $\widetilde{f}$, then $\widetilde{T}$ is the lift of T which fixes $\widetilde{\varphi}\widetilde{p}$.

Since φ is uniformly continuous, so is $\widetilde{\varphi}$. Choose a small $\varepsilon > 0$ and choose $\delta > 0$ such that any δ-ball in $\widetilde{M}_0$ is mapped by $\widetilde{\varphi}$ into an ε-ball in $\mathbb{R}^n$. There is an integer $K > 0$ such that for each $k > 0$ we may choose K points $p_0 = \widetilde{f}^k\widetilde{x}, p_1, \ldots, p_K = \widetilde{f}^k\widetilde{y}$ such that $d(p_i, p_{i+1}) < \delta$ for each $0 \le i < K$. Indeed any $K > D/\delta$ will do. These points may be chosen, for example, to lie in some path α from $\widetilde{f}^k\widetilde{x}$ to $\widetilde{f}^k\widetilde{y}$ of length at most D.

Then in $\mathbb{R}^n$ we have

$$\|\widetilde{T}^k\widetilde{\varphi}(\widetilde{x}) - \widetilde{T}^k\widetilde{\varphi}(\widetilde{y})\| = \|\widetilde{\varphi}(\widetilde{f}^k\widetilde{x}) - \widetilde{\varphi}(\widetilde{f}^k\widetilde{y})\| < \varepsilon K$$

for all $k \in \mathbb{Z}$. Since $\widetilde{T}$ is a hyperbolic linear map of $\mathbb{R}^n$ this is impossible unless $\widetilde{\varphi}(\widetilde{x}) = \widetilde{\varphi}(\widetilde{y})$ and hence $\varphi(x) = \varphi(y)$ as required. ∎

3. Equivalent horseshoes. Recall that by an f-horseshoe we mean a map $\epsilon : \Sigma \to M$ which satisfies $\epsilon\sigma = f\epsilon$, where σ is the left-shift map of $\Sigma = \{0,1\}^{\mathbb{Z}}$. In this section we will see how, in principle, pairs of f-horseshoes can be found, corresponding image points of which are equivalent in the sense of the previous section. The basic idea is to find pairs of rectangles of a certain kind in M, which are similar from the point of view of homology.

A *rectangle* Λ is a closed subset of M whose interior is diffeomorphic to some rectangle $(0, w) \times (0, h)$ in the x-y plane, by a diffeomorphism ψ which satisfies $\psi^* dx = \omega^{\mathrm{u}}$ and $\psi^* dy = \omega^{\mathrm{s}}$. The numbers w and h are the *width* and *height* of Λ. A rectangle can contain no point of Z in its interior, since ω^{s} and ω^{u} vanish on Z. A rectangle Λ has a (closed) top edge $T(\Lambda)$, a bottom edge $B(\Lambda)$, and left and right edges $L(\Lambda)$ and $R(\Lambda)$, defined in the obvious way. Similarly we will frequently refer to the bottom-left point, the top-right point, etc., of Λ.

Call a rectangle Λ *horseshoe-like* if either its bottom-left and top-right points, or its top-left and bottom-right points, are fixed points, and moreover $f(T(\Lambda)) \subset T(\Lambda)$ and $f(B(\Lambda)) \subset B(\Lambda)$. Any rectangle which has periodic points at opposite corners is horseshoe-like for some power of f. To a horseshoe-like rectangle Λ we can associate an f-horseshoe whose image is a Cantor set of points inside Λ, as follows. Let p_0 and p_1 be the fixed points lying at the corners of Λ. Denote by Λ^0 and Λ^1 the connected components of $f^{-1}\Lambda \cap \Lambda$ containing p_0 and p_1 respectively. These are each subrectangles of Λ of height $\lambda^{-1} \cdot \mathrm{height}(\Lambda)$ and the same width as Λ. Then f sends them to their images, also subrectangles of Λ, by an affine map; and as for Smale's horseshoe it follows that the intersection

$$\bigcap_{k\in\mathbb{Z}} f^{-k}(\Lambda^0 \cup \Lambda^1)$$

is a Cantor set whose points are uniquely coded by their itineraries (except that if $p_0 = p_1$, then this point has two codes). For any $z = (\ldots z_{-1}z_0z_1\ldots)$

$\in \Sigma$ we define $\epsilon(z)$ to be the unique point of this intersection which satisfies $f^k\epsilon(z) \in \Lambda^{z_k}$ for all $k \in \mathbb{Z}$; then ϵ is an f-horseshoe which we will call the f-horseshoe associated to Λ.

The diagonals of a rectangle are the straight-line paths between opposite corners. For a horseshoe-like rectangle Λ we will say "the diagonal of Λ" to mean the diagonal of Λ between the two fixed points. We have

LEMMA 3.1. *Suppose Λ_1 and Λ_2 are two rectangles in M, horseshoe-like for some power f^t of f, with corners at the same two fixed points p_0, p_1 of f^t. Suppose the diagonals γ_1 and γ_2 of Λ_1 and Λ_2 satisfy $[\gamma_1] = [\gamma_2]$ in the relative homology group $H_1(M, \{p_0, p_1\}; \mathbb{Z})$. If ϵ_1 and ϵ_2 are the associated f^t-horseshoes, then $\epsilon_1(z) \sim \epsilon_2(z)$ (for f) for every $z \in \Sigma$.*

Proof. Fix $z \in \Sigma$ and choose a lift $\widetilde{f}$ of f to $\widetilde{M_0}$. To fix ideas, we work in the case where Λ_1 and Λ_2 have fixed points at their bottom-left and top-right corners. Thus p_0 is the bottom-left point of Λ_i for $i = 1, 2$, and p_1 the top-right point. Choose a lift $\widetilde{p}_0$ of p_0 in $\widetilde{M_0}$, and for $i = 1, 2$ let q_i be the corresponding lift of $\epsilon_i(z)$; this means that q_i lies in the lift of Λ_i whose bottom-left point is $\widetilde{p}_0$. We claim that for all $m \in \mathbb{Z}$, the lifts of Λ_1 and Λ_2 containing $\widetilde{f}^{tm}q_1$ and $\widetilde{f}^{tm}q_2$ respectively have the same bottom-left point, some lift $\widetilde{p}_0(m)$ of p_0. Since $[\gamma_1] = [\gamma_2]$ this is the same thing as claiming that they have the same top-right point, some lift $\widetilde{p}_1(m)$ of p_1. Of course for $m = 0$ the claim is true with $\widetilde{p}_0(0) = \widetilde{p}_0$.

To prove the claim, suppose it is satisfied for some $k \geq 0$, and for $i = 1, 2$ let $\widetilde{\Lambda}_i$ be the relevant lift of Λ_i. It has bottom-left point $\widetilde{p}_0(k)$ and top-right point $\widetilde{p}_1(k)$, independent of i. If $z_k = 0$ then $\widetilde{f}^{tk}q_i$ lies in the subrectangle $\widetilde{\Lambda}_i^0$ of $\widetilde{\Lambda}_i$ obtained by lifting Λ_i^0. Since f^t sends Λ_i^0 to its image, also a subrectangle of Λ_i, by an affine map fixing p_0, it follows that $\widetilde{f}^{t(k+1)}q_i$ lies in the lift of Λ_i whose bottom-left point is $\widetilde{p}_0(k+1) := \widetilde{f}^t(\widetilde{p}_0(k))$. On the other hand if $z_k = 1$ then $\widetilde{f}^{tk}q_i$ lies in the subrectangle $\widetilde{\Lambda}_i^1$ of $\widetilde{\Lambda}_i$ obtained by lifting Λ_i^1. Arguing as above we find that $\widetilde{f}^{t(k+1)}q_i$ lies in the lift of Λ_i whose top-right point is $\widetilde{p}_1(k+1) := \widetilde{f}^t\widetilde{p}_1(k)$. In either case we have shown that the claim holds for $m = k+1$. A similar argument works for f^{-t} (for which Λ_1 and Λ_2 are still horseshoe-like), so the claim holds for all $m \in \mathbb{Z}$ by induction.

In particular we have

$$\begin{aligned} d(\widetilde{f}^{tm}q_1, \widetilde{f}^{tm}q_2) &\leq d(\widetilde{f}^{tm}q_1, \widetilde{p}_0(m)) + d(\widetilde{p}_0(m), \widetilde{f}^{tm}q_2) \\ &\leq \operatorname{diam}(\Lambda_1) + \operatorname{diam}(\Lambda_2) \end{aligned}$$

for all $m \in \mathbb{Z}$; i.e. $\epsilon_1(z) \sim \epsilon_2(z)$ for f^t. Since $\widetilde{f}$ expands distances by a factor of at most λ, it immediately follows that we have $\epsilon_1(z) \sim \epsilon_2(z)$ for f, as required. ∎

4. Finding equivalent horseshoes. In this section we will identify a class $\mathcal{C}$ of pseudo-Anosov maps for which the pairs of equivalent horseshoes described in the previous section actually occur, proving Theorem 1.3 of the introduction. The class consists of all those pseudo-Anosovs whose invariant foliations have a particularly simple type, and the proof involves exploiting concrete geometrical features of these foliations.

To express the "type" of the foliations, we will consider combinatorial properties of the interval exchange map induced by the unstable foliation on a segment of a stable leaf. We begin by recalling some material from [Rau79] and [Vee82].

4.1. *Admissible segments and partitions.* As above we work with a pseudo-Anosov map f of the surface M whose invariant foliations are orientable. Let s be an initial segment of an outwardly-oriented stable prong of some singularity p. That is, s is a subset of a stable prong of p homeomorphic to the closed unit interval, whose left-hand endpoint is p. We say s is *admissible* if the other endpoint x of s lies in an unstable prong of a singularity p' such that the (closed) segment $u(x)$ of this prong joining x to p' does not intersect the interior of s. We regard $u(x)$ as an oriented line segment and say it is positively or negatively oriented according to whether it agrees or disagrees with the orientation of $\mathcal{F}^{\mathrm{u}}$.

It is well known that the first return map F of the unstable foliation to any such segment s is an interval exchange map. A discontinuity d of F is a point of intersection of the interior of s with an inwardly-oriented prong of a singularity q such that the open segment of this prong from d to q does not intersect the interior of s. Our assumption that s is admissible is a nondegeneracy assumption: it means that x lies in an unstable prong in such a way that it does not give rise to an "extra" discontinuity of F.

Let us write $p = d_0, d_1, \ldots, d_k = x$ for the discontinuities of F taken together with the endpoints of s, written in the order they occur in s. For each $1 \leq i \leq k$ the open subinterval (d_{i-1}, d_i) of s is maximal with respect to the following property: we may homotope (d_{i-1}, d_i) by a nontrivial holonomy along $\mathcal{F}^{\mathrm{u}}$, in the direction given by orientation of $\mathcal{F}^{\mathrm{u}}$, until we obtain a new subinterval of s. This is of course the subinterval $F(d_{i-1}, d_i)$. The closure of the region swept out by this holonomy is a rectangle which we denote by P_i. These rectangles have pairwise disjoint interiors, and because the unstable foliation of a pseudo-Anosov map is uniquely ergodic (see [MR879], [Vee82]), they cover M. Our choice of the d_i's ensures that the left edge of each P_i contains a singularity, and the right edge of each P_i either contains a singularity or the point x (or both).

We write $\mathcal{P}_s = \{P_i \mid 1 \leq i \leq k\}$, and call $\mathcal{P}_s$ the *partition of M associated to s*.

4.2. *Combinatorics of $\mathcal{P}_s$.* To describe the class $\mathcal{C}$ we need a method of recording how the rectangles in the partition $\mathcal{P}_s$ are "stuck together" along s. We do this by defining a permutation π_s of $\{1, \ldots, k\}$ as follows. The subinterval $F(d_i, d_{i+1})$ of s is the interior of the top edge of the rectangle P_i, and these subintervals are pairwise disjoint and lie in some order in s. For $1 \leq i \leq k$ we define $\pi_s(i) = j$ if $F(d_i, d_{i+1})$ is the jth such interval in this ordering.

From the permutation π_s one can calculate directly the number and orders of singularities of the stable and unstable foliations, and hence the genus g of M. Intuitively, this calculation can be performed by "walking" once around each singularity, recording the rectangles one crosses in doing so. This gives one a collection of cycles of indices of rectangles, one for each singularity, the length of a cycle reflecting the order of the corresponding singularity. See [Vee82] for details.

In fact π_s encodes information about $\mathcal{F}^{\mathrm{u}}$ and $\mathcal{F}^{\mathrm{s}}$ which is more subtle than the number and orders of singularities. Namely, it encodes in which connected component, in the moduli space of flat singular Riemannian metrics on M_g having the given number and orders of singularities, the metric defined by $\mathcal{F}^{\mathrm{s}}$ and $\mathcal{F}^{\mathrm{u}}$ lies. Although we will not need this theory, which is developed in [Vee90] and [KZ03], we will refer to it briefly to help us indicate what the class $\mathcal{C}$, defined below, comprises.

4.3. *The class $\mathcal{C}$.* For any $k \geq 2$ we denote by $\mathcal{C}_k$ the collection of all those pseudo-Anosov maps which have an admissible segment s whose associated permutation π_s is the permutation π_k given by

$$\pi_k(i) = k + 1 - i. \tag{4.1}$$

We define $\mathcal{C}$ to be the union of $\mathcal{C}_k$ over all *even* $k \geq 4$. That $\mathcal{C}$ is nonempty follows from the techniques of [Vee82]; in fact, each C_k contains infinitely many distinct pseudo-Anosovs (up to taking powers). The reason we consider only even k is that, for odd $k > 4$, the invariant foliations of any $f \in \mathcal{C}_k$ have two singularities, whereas for even k they have exactly one; this property will be important below. Indeed in the latter case the single singularity has order $k-1$ (so it has $2k-2$ stable prongs and $2k-2$ unstable prongs) and so f is a homeomorphism of the surface $M_{k/2}$ of genus $k/2$.

Homeomorphisms in $\mathcal{C}$ have special geometric features which we will now exploit to prove Theorem 1.3. However, not every pseudo-Anosov homeomorphism whose invariant foliations are orientable and have only one singularity lies in $\mathcal{C}$: see Remark 4.4 below.

4.4. *The geometry of $\mathcal{P}_s$.* We shall henceforth assume that $f \in \mathcal{C}_k$ for some even $k \geq 4$. Let s, $\mathcal{P}_s$ and $\pi_s = \pi_k$ be as above. We now describe the geometry of the partition $\mathcal{P}_s$ in this special case in more detail. One such

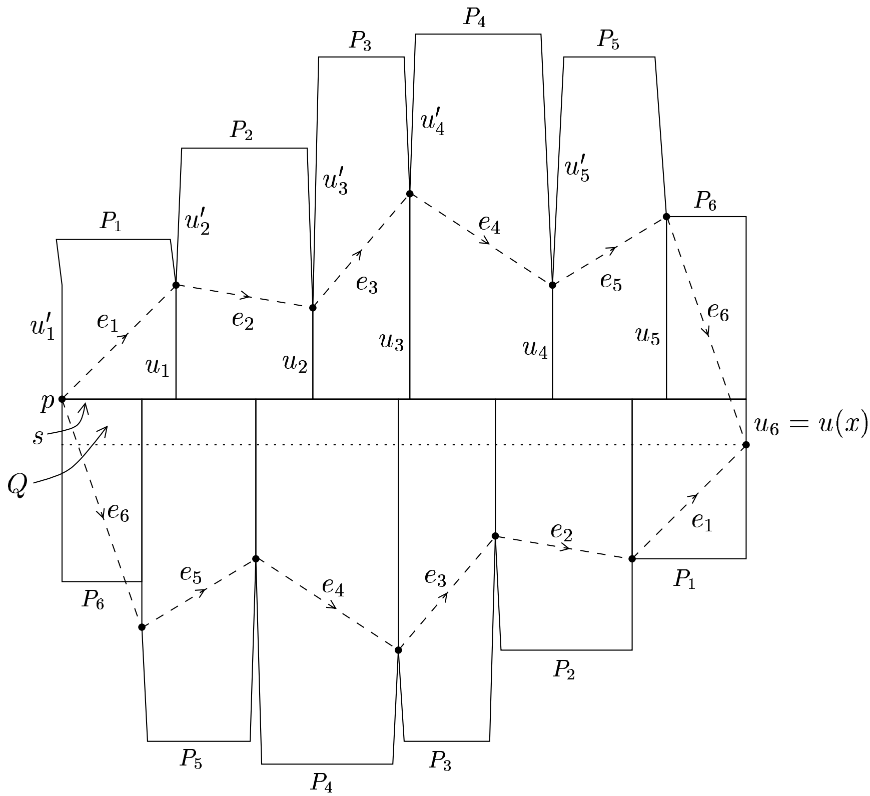

Fig. 4.1. A partition for a pseudo-Anosov map in the class $\mathcal{C}_6$

partition is depicted in Figure 4.1 in the case $k = 6$. The diagram is drawn in the universal cover of the surface M, and two lifts of each rectangle P_i are depicted; to form M, one must identify them, preserving horizontal and (almost) vertical lines. Here the segment $u(x)$ coming from the admissibility of s is negatively oriented. Our aim in this section is to show that $\mathcal{P}_s$ has, in general, the features seen in Figure 4.1.

Denote by p the singularity of the stable and unstable foliations. By definition, for each $0 \le i \le k$ there is a (closed) segment u_i of an unstable prong connecting d_i to p (for $i = 0$ it is just the one-point set $\{p\}$). Indeed, for $0 \le i < k$ this segment is contained in the left edge $L(P_{i+1})$ of P_{i+1}, whereas $u_k = u(x)$. For $1 \le i \le k$ we denote by u'_i the segment of $L(P_i)$ from p to the top-left point of P_i; it is the closure of $L(P_i) \setminus u_{i-1}$. We define also $u'_0 = -u_k$, that is, u_k taken with the opposite orientation.

The assumption $\pi_s = \pi_k$ implies that the u_i's and the u'_i's are rather symmetrically arranged in M. To express this, let us define $h_i = \langle \omega^{\mathrm{s}}, u_i \rangle$ and $h'_i = \langle \omega^{\mathrm{s}}, u'_i \rangle$. The following inequalities now follow from the above definitions and because we have $\pi_s(1) = k$ and $\pi_s(k) = 1$:

$$0 < h_i \le \text{height}(P_{i+1}), \quad 0 < h'_i \le \text{height}(P_i) \text{ for } 1 \le i < k, \tag{4.2}$$

$$h_0 = 0, \quad h'_k = 0, \tag{4.3}$$

$$-\text{height}(P_1) < h_k < \text{height}(P_k), \quad h'_0 = -h_k. \tag{4.4}$$

Some similar formulae (with different notation) can be found on p. 206 of [Vee82]. We have the following lemma.

LEMMA 4.1. *For all* $0 \le i < j \le k$ *we have*

$$h_j - h_i = h'_j - h'_i. \tag{4.5}$$

Moreover, *for* $1 \le i < k$ *we have* $h_i > h_k$ *and* $h'_i > -h_k = h'_0$.

Proof. Note that $u_{i-1} \cup u'_i = L(P_i)$ for all $1 \le i \le k$, and thus $h_{i-1} = \text{height}(P_i) - h'_i$. On the other hand, for all $1 < i \le k$ we have $\pi_s(i-1) = \pi_s(i)+1$, meaning that the top-right point of P_i is the top-left point of P_{i-1}. If $i < k$, or if $i = k$ and $u(x)$ is positively oriented, then $R(P_i)$ contains p, and it follows that $R(P_i) = u_i \cup u'_{i-1}$. Hence $h_i = \text{height}(P_i) - h'_{i-1}$ in this case. In fact, this formula also holds for the case $i = k$ and $u(x)$ negatively oriented because, in this case, u'_{k-1} is the union of $u(x) = u_k$ and $R(P_k)$. The formula also holds for $i = 1$ because of our choice that $u'_0 = -u_k$.

Now let $0 \le i < j \le k$. We have

$$\begin{aligned} h_j - h_i &= \sum_{r=i+1}^{j} (h_r - h_{r-1}) \\ &= \sum_{r=i+1}^{j} [(\text{height}(P_r) - h'_{r-1}) - (\text{height}(P_r) - h'_r)] \\ &= \sum_{r=i+1}^{j} (h'_r - h'_{r-1}) = h'_j - h'_i. \end{aligned}$$

The other inequalities follow by applying this together with (4.2) and (4.3) above. ∎

From the last two inequalities in the lemma we get the following. If $u(x)$ is negatively oriented, so $h_k < 0$, there is a rectangle Q in M of height $-h_k$ whose top edge is s, whose bottom-right point is p, and whose right edge lies inside $R(P_1)$. If instead $h_k > 0$, there is a rectangle Q in M of height h_k whose bottom edge is s, whose top-right point is p, and whose right edge lies inside $R(P_k)$. We will use this rectangle Q below.

4.5. *Paths for rectangles.* For $1 < i < k$ let us define e_i to be the loop at p which is a straight-line path traversing P_i once from its left edge to its right edge. This makes sense because p lies in both the left and the right edge of P_i. If $i = 1$ or $i = k$ we instead define e_i to be the path composition $\alpha \cdot u(x)$, where α denotes the diagonal of P_i from p to x. Thus e_i is again a loop at p, and Lemma 4.1 implies that we may homotope e_i fixing endpoints, until it is a straight-line path in $P_i \cup Q$. These paths are depicted by dashed lines in Figure 4.1.

DEFINITION 4.2. A *diagonal pair* is a pair (i, j), $1 \leq i < j \leq k$, with the property that for all $i \leq r < j$ we have $h_r > \max(h_{i-1}, h_j)$.

We are now ready to prove the following theorem, from which Theorem 1.3 of the introduction follows directly, by applying Lemma 3.1.

THEOREM 4.3. (a) *If* (i, j) *is a diagonal pair other than* $(1, k)$, *then there are distinct rectangles* Λ_1 *and* Λ_2 *in* M, *horseshoe-like for some power* f^t *of* f, *whose diagonals* γ_1 *and* γ_2 *satisfy*

$$[\gamma_1] = [\gamma_2] = \sum_{r=i}^{j} [e_r] \tag{4.6}$$

in $H_1(M, \{p\}; \mathbb{Z})$. *If* ϵ_1 *and* ϵ_2 *are the corresponding* f^t*-horseshoes*, *then for all* $z \in \Sigma$ *we have* $\epsilon_1(z) \neq \epsilon_2(z)$, *except that if* z *is one of the two fixed points of the shift map then* $\epsilon_1(z) = \epsilon_2(z) = p$.

(b) *There are exactly* $k - 2$ *diagonal pairs other than* $(1, k)$. *Moreover each* i *with* $1 < i < k$, *and at least one of* 1 *or* k, *occurs in some such pair.*

Proof. (a) Let (i, j) be a diagonal pair, $1 \leq i < j \leq k$, and suppose first of all that $j \neq k$. We denote by Λ_1 the rectangle of height $|h_j - h_{i-1}|$ and width $\sum_{r=i}^{j} \operatorname{width}(P_r)$, defined as follows: we require $L(\Lambda_1) \subset L(P_i)$ and $R(\Lambda_1) \subset R(P_j)$, and the bottom-left or top-left point of Λ_1 is p depending on whether $h_j > h_{i-1}$ or $h_j < h_{i-1}$. Because (i, j) is diagonal, this rectangle is well defined, and it can be written as the union $\bigcup_{r=i}^{j} Q_r$, where Q_r is a subrectangle of P_r whose left edge lies in u_{r-1} (except possibly for $r = i$) and whose right edge lies in u_r (except possibly for $r = j$).

We denote by Λ_2 the rectangle of the same width and height as Λ_1 which has $L(\Lambda_2) \subset L(P_j)$ and $R(\Lambda_2) \subset R(P_i)$, and whose bottom-left or top-left point is p depending on whether $h'_j > h'_{i-1}$ or $h'_j < h'_{i-1}$. Again Λ_2 is well defined, because since (i, j) is diagonal and by Lemma 4.1, we have $h'_r > \max(h'_{i-1}, h'_j)$ for all $i \leq r < j$. In fact Λ_2 can be written as the union $\bigcup_{r=i}^{j} Q'_r$, where Q'_r is a subrectangle of P_r whose left edge lies in u'_r (except possibly for $r = j$), and whose right edge lies in u'_{r-1} (except possibly for $r = i$).

If $j = k$ and $h_k < 0$ then Λ_2 is defined as above, but Λ_1 no longer fits in the same way into the rectangles P_r. However, since by Lemma 4.1 we have $h_k < h_r$ for all $1 \leq r < k$, it follows that Λ_1 can again be constructed as a union $\bigcup_{r=i}^{j} Q_r$, where Q_r is a subrectangle not of P_r, but of $P_r \cup Q$. A similar argument when $h_k > 0$ is used to construct Λ_2.

If γ_1 and γ_2 are the diagonals of Λ_1 and Λ_2, we have $[\gamma_1] = [\gamma_2] = \sum_{r=i}^{j} [e_i]$ in $H_1(M, \{p\}; \mathbb{Z})$ by construction. Indeed, the homotopy class of γ_1 is that of the path composition $e_i e_{i+1} \dots e_j$ and the homotopy class of γ_2 is that of the path composition $e_j e_{j-1} \dots e_i$ (where we make some choice of bracketing in these compositions to make them well defined). Moreover unless $(i, j) = (1, k)$, Λ_1 and Λ_2 are distinct because their left edges lie in different prongs of the singularity p, as do their right edges (note however that $\Lambda_1 \cap \Lambda_2 \setminus \{p\}$ may be nonempty).

The rectangles Λ_1 and Λ_2 are horseshoe-like for any power f^t of f which fixes all stable prongs, and have the same width and height. Since such an f^t expands unstable leaves and contracts stable leaves everywhere by the same factor λ^t, it follows that (with the notation of Section 3) the actions of f^t on $\Lambda_1^0 \cup \Lambda_1^1$ and on $\Lambda_2^0 \cup \Lambda_2^1$ are conjugate by the isometry $\Phi : \Lambda_1 \to \Lambda_2$ which preserves the stable and unstable foliations and their orientations. In particular $\epsilon_2(z) = \Phi\epsilon_1(z)$ for all $z \in \Sigma$. Since Φ sends stable (resp. unstable) leaves to stable (resp. unstable) leaves isometrically, it is determined by its value at any point of Λ_1 other than the singularity p. In particular if Φ fixes any point of Λ_1 other than p, it must fix every point of Λ_1, which is impossible since Λ_1 and Λ_2 have distinct left edges. This finishes the proof of (a).

(b) Now we will count the number of diagonal pairs. For $1 \le i \le j \le k$, denote by $S(i, j)$ the set of diagonal pairs (i', j') satisfying $i \le i' < j' \le j$. We will prove that if (i, j) is diagonal then $\#S(i, j) = j - i$. The proof is by induction on $j - i$. Assume that the formula holds for all diagonal pairs (a, b) with $b - a < j - i$ (it certainly holds when $b - a = 1$). Pick the $i \le i' < j$ which makes $h_{i'}$ as small as possible (there is only one such i' because the stable foliation has no closed leaves or saddle connections). Then the pairs (i, i') and $(i' + 1, j)$ are each diagonal, with the obvious exceptions when $i' = i$ and when $i' + 1 = j$, and the diagonality of (i, j) ensures that $S(i, j) = S(i, i') \cup S(i' + 1, j) \cup \{(i, j)\}$. Applying our induction hypothesis we get $\#S(i, j) = i' - i + j - (i' + 1) + 1 = j - i$, as required. In particular $\#S(1, k) = k - 1$.

For any $1 \le i \le k$, one checks that i is the first element of some diagonal pair if and only if $h_{i-1} < h_i$, and that i is the second element of some diagonal pair if and only if $h_{i-1} > h_i$. We cannot have $h_{i-1} = h_i$; thus every i occurs in some diagonal pair. If $(1, k)$ is the only pair containing 1 then we must have $h_2 > h_1$, so $(2, k)$ is a diagonal pair containing k. This completes the proof. ∎

REMARK 4.4. An obvious question is whether there are pseudo-Anosovs whose invariant foliations are orientable and have exactly one singularity, but which do not lie in the class $\mathcal{C}$. The answer turns out to be "no" on M_2, but "yes" on surfaces of higher genus. To construct such pseudo-Anosovs, one applies the construction of Section 8 of [Vee82] starting with a permutation π which does not lie in the so-called "extended Rauzy class" of π_k, but which corresponds as above (Section 4.2) to foliations on M_g with one singularity. The existence of such permutations follows from the classification, in [KZ03], of connected components of strata of the moduli space of flat singular Riemannian metrics on M_g. One such permutation is

$$\theta_k = \begin{pmatrix} 1 & 2 & 3 & \dots & k-3 & k-2 & k-1 & k \\ k & k-3 & k-4 & \dots & 2 & k-1 & k-2 & 1 \end{pmatrix}$$

for $k = 2g$ and $g > 2$.

For $g = 2$, on the other hand, no such permutations exist (since there are only $4! = 24$ permutations of $\{1, 2, 3, 4\}$ this can easily be verified directly). Consequently, the construction of [Vee82] and the theorem on p. 327 of [Rau79], which describes the structure of the Rauzy class of π_k, together imply that any pseudo-Anosov homeomorphism of M_2 whose invariant foliations have only one singularity, must lie in $\mathcal{C}$.

References

[Ban03] G. Band, *How to identify points of a pseudo-Anosov homeomorphism*, PhD thesis, Univ. of Warwick, 2003.

[Fat88] A. Fathi, *Some compact invariant sets for hyperbolic linear automorphisms of torii [tori]*, Ergodic Theory Dynam. Systems 8 (1988), 191–204.

[Fra70] J. Franks, *Anosov diffeomorphisms*, in: Global Analysis (Berkeley, CA, 1968), Proc. Sympos. Pure Math. 14, Amer. Math. Soc., Providence, RI, 1970, 61–93.

[Han85] M. Handel, *Global shadowing of pseudo-Anosov homeomorphisms*, Ergodic Theory Dynam. Systems 5 (1985), 373–377.

[Hir70] M. W. Hirsch, *On invariant subsets of hyperbolic sets*, in: Essays on Topology and Related Topics (Mémoires dédiés à Georges de Rham), Springer, New York, 1970, 126–135.

[KZ03] M. Kontsevich and A. Zorich, *Connected components of the moduli spaces of Abelian differentials with prescribed singularities*, Invent. Math. 153 (2003), 631–678.

[Man78] R. Mañé, *Invariant sets of Anosov's diffeomorphisms*, ibid. 46 (1978), 147–152.

[MR879] *Travaux de Thurston sur les surfaces*, Astérisque 66 (1979).

[Rau79] G. Rauzy, *Échanges d'intervalles et transformations induites*, Acta Arith. 34 (1979), 315–328.

[Vee82] W. A. Veech, *Gauss measures for transformations on the space of interval exchange maps*, Ann. of Math. (2) 115 (1982), 201–242.

[Vee90] W. A. Veech, *Moduli spaces of quadratic differentials*, J. Anal. Math. 55 (1990), 117–171.

Department of Mathematics
University of Warwick
Coventry CV4 7AL
United Kingdom
E-mail: gavinband@tiscali.co.uk

Institute of Mathematics
Polish Academy of Sciences
P.O. Box 21
00-956 Warszawa, Poland
E-mail: gav@impan.gov.pl

Received 28 August 2003;
in revised form 27 November 2003