Pergamon

Topology Vol. 36, No. 1, pp. 209–227, 1997

0040-9383/96/$15.00 + 0.00

0040-9383(95)00066-6

FINITE FOLIATIONS AND SIMILARITY INTERVAL EXCHANGE MAPS

D. COOPER†, D. D. LONG‡
and A. W. REID§

(*Received* 10 *February* 1995; *received for publication* 30 *November* 1995)

1. INTRODUCTION

This paper continues the study, initiated in [1], of understanding surfaces immersed transverse to the suspension flow in a hyperbolic surface bundle over the circle. Briefly, our context is the following. Suppose that $\theta : F \to F$ is an orientation-preserving pseudo-Anosov homeomorphism of a closed surface. Then we may form the mapping torus $M = M(\theta)$ which results of Thurston show is hyperbolic. This mapping torus is equipped with an obvious one-dimensional foliation, denoted throughout as $\mathscr{L}$. We shall consider surfaces $g : S \looparrowright M$ which are immersed into M so as to be transverse to $\mathscr{L}$. It is shown by Mangum in [2] that such surfaces are automatically incompressible. It follows that $g_*(\pi_1(S))$ is a surface subgroup of $\pi_1(M)$; this Kleinian group may be geometrically finite in which case we refer to $g : S \looparrowright M$ as a *geometrically finite immersion.* Although this is not used in an essential way here, we note that a combination of deep results due to Marden [3] and Bonahon [4] shows that in this context, the geometrically finite case is precisely the case when the surface group is *quasi-fuchsian*, that is to say its limit set is a topological circle.

Since the map θ is pseudo-Anosov, it preserves a pair of singular foliations, $\mathscr{F}^{\pm}$ of F and these give rise to a pair of codimension one singular foliations of M which we denote by $\Omega\mathscr{F}^{\pm}$. The fact that $g(S)$ is transverse to $\mathscr{L}$ means that these foliations pull back to give a pair of singular foliations $\mathscr{F}_S^{\pm}$ of S. In these terms we have:

THEOREM 1.1. (Cooper *et al.* [1]). *The immersion $g : S \looparrowright M$ is geometrically finite if and only if one (and hence both) of $\mathscr{F}_S^{\pm}$ is a finite foliation.*

We shall use the term *foliation* of a surface F to mean the integral curves of a singular line field with a finite number of singularities, all of prong type and recall that such a foliation is said to be *finite* if it contains a finite number of closed leaves and both ends of every nonclosed leaf spiral onto one of these closed leaves. In the context of general foliations, this could also be called depth one. Understanding whether a given foliation is

---

†Partially supported by the NSF.
‡Partially supported by the A.P. Sloan Foundation and the NSF.
§Supported by The Royal Society.

finite or not seems to be a subtle problem. We shall show:

THEOREM 1.2. *It is a decidable question if a surface immersed transverse to a suspension flow is geometrically finite or not.*

We shall use this result to exhibit what is apparently the first nontrival example of an immersed geometrically infinite surface. It is in fact already known that, in principle, the question resolved by Theorem 1.2 is decidable for any incompressible surface at all [5]; however, that algorithm is impractical, relying as it does on an enumeration of finite coverings of the ambient hyperbolic manifold. Moreover, our proof in this case has the benefit of additional theoretic insight as well as providing a convenient tool for computing the associated dilatation and invariant measure.

The tool we introduce arises as follows. We sketch the main ingredients here, deferring some of the details until later. If one fixes an arc $\alpha \subset S$ transverse to one of the foliations $\mathscr{F}_S$, which for convenience we assume orientable, then we can consider the first return map $\eta : \alpha \to \alpha$ which this foliation induces. It is in fact a possibility at this stage that some leaf never returns to α, in which case we are in the situation of a finite foliation, since we know the dichotomy provided by Theorem 1.1 (see [1] for the proof that this is in fact a dichotomy) is that either $\mathscr{F}_S$ has all leaves dense or it is finite. We can suppose then that η is a well-defined map away from a finite number of points corresponding to where the separatrices of $\mathscr{F}_S$ meet α for the first time. Thus, we obtain a map which is closely related to an interval exchange map, the only difference being that some intervals may be stretched or contracted. We define such a map to be a *(geometric) similarity interval exchange map.* In contrast to the more usual notion of interval exchange map, Lebesgue measure is not typically invariant for η; indeed, if the foliation is finite the only invariant probability measures will involve a finite number of atoms. One can also define similarity exchange maps in a completely general setting. They seem to be of interest both as dynamical systems and for the connections such understanding offers to the theory of affine foliations (see [6]).

In Section 3 we construct a bundle and use the idea of [1] to make an immersed incompressible surface transverse to the flow, then use our methods to prove that this surface is geometrically infinite. This appears to be the first explicitly constructed example of an immersion of a geometrically infinite surface. Moreover, part of the power of our method is that one can calculate dilatations and measures associated with the immersion.

More specifically, we recall from [1] that given any pseudo-Anosov map as above, there is always some finite sheeted covering and a lift of the map $\tilde{\theta} : \tilde{F} \to \tilde{F}$ so that some simple closed curve $C \subset \tilde{F}$ is carried disjoint from its image by $\tilde{\theta}$. We can use this to construct a surface immersed transverse to the flow, denoted $S(C, \tilde{\theta}(C), \tilde{F}) = S$; namely remove annular neighbourhoods of C and $\tilde{\theta}(C)$ and reglue them in a carefully chosen way (see Section 3 for details) The example of Section 3 is of this type. One of the surprising features (somewhat concealed here) is that it seems to be extremely difficult to produce geometrically infinite surfaces in this simple-minded way and it is of some interest to understand this difficulty, both for its own sake and because of the following theorem (notation as in Section 3).

THEOREM 1.3. *Suppose that $\theta : F \to F$ is a pseudo-Anosov map for which we can find a simple closed curve C with* (a) *$C \cup \theta(C)$ is nonseparating, and* (b) *C, $\theta(C)$ and $\theta^2(C)$ are all disjoint. Then either*

(1) *the immersion $g : S = S(C, \theta(C), F) \looparrowright M(\theta)$ cannot be lifted to become an embedding in any finite covering of $M(\theta)$,*

or

(2) *there is a finite sheeted covering $p: F_1 \to F$, a map θ_1 covering θ and a lift C_1 of C so that the immersion $S = S(C_1, \theta_1(C_1), F_1) \subset M(\theta)$ is geometrically infinite.*

In particular, with our current lack of understanding, it does not seem impossible that for sufficiently complicated monodromies, it might be that one can never obtain a geometrically infinite surface by a single operation of cut and cross join. Such a bundle would therefore contain a nonseparable geometrically finite surface group.

2. SIMILARITY INTERVAL EXCHANGES AND THE ALGORITHM

We begin with some preliminary observations. The first of these is the definition of a similarity interval exchange map. Recall that an interval exchange map can be informally described as cutting up the unit interval into a finite number of half open intervals and permuting them. Our notion is a generalisation of this in that we also allow intervals to be expanded or contracted by powers of a fixed real number s. Here is the formal definition (cf. [7] or [8]).

Definition 2.1. Suppose that $(\alpha_1, \dots, \alpha_n)$ is a vector of real numbers with $\alpha_i > 0$ and $\sum_{i=1}^n \alpha_i = 1$. Suppose further that $s > 1$ is some fixed real number with the property that there are integers $k_1, \dots, k_n$ so that $\sum_{i=1}^n s^{k_i}\alpha_i = 1$. Set $\beta_0 = 0$ and $\beta_i = \sum_{j=1}^i \alpha_j$.

Then given a permutation σ on $\{1, \dots, n\}$ we set $\alpha^\sigma = (s^{k_{\sigma^{-1}(1)}}\alpha_{\sigma^{-1}(1)}, \dots, s^{k_{\sigma^{-1}(n)}}\alpha_{\sigma^{-1}(n)})$. Our hypothesis guarantees that this is a probability vector and we form a corresponding vector β_i^σ.

We define the *similarity interval exchange map* $\eta: I \to I$ associated with this data by

$$\eta(x) = s^{k_{i-1}}(x - \beta_{i-1}) + \beta^\sigma_{\sigma(i)-1}, \quad x \in [\beta_{i-1}, \beta_i).$$

The effect of this definition is to map the interval $[\beta_{i-1}, \beta_i)$ to the interval $[\beta^\sigma_{\sigma(i)-1}, \beta^\sigma_{\sigma(i)})$; this has involved a scaling by factor s^{k_i}. Of course, when all the k_i's are zero, this reduces to the more usual notion of interval exchange map.

These arise naturally in our context in the following way. In the notation established in the introduction, let $\alpha \subset S$ be a short arc lying inside a leaf of the foliation $\mathscr{F}_S^-$. Temporarily, we assume that the foliations $\mathscr{F}_S^\pm$ are globally orientable, that is, each is the integral curve of a vector field with zeroes. Further, we assume that α has been chosen so that it contains no singularity. In particular, the vector field has no zeroes on α.

The arc α is transverse to $\mathscr{F}^+$ so that, given $p \in \alpha$, we may flow along the leaf of $\mathscr{F}_S^+$ in the direction given by the vector field until the first return to (the other side of) α. It is possible that there are points for which the leaf never returns to α, but Theorem 1.1 implies that this can only happen in the presence of a closed leaf, so we shall assume that this is not the case and that we have defined a map $\eta: \alpha \to \alpha$. Strictly speaking, the map is not quite well defined; exactly as in the case of an interval exchange map, there are a finite number of points corresponding to where the separatrices first meet α where η is potentially multivalued, but (exactly as in the interval exchange case) we adopt some convention so that the map becomes well defined and only discontinuous at those points.

There is also another map which arises by flowing in the opposite direction along the vector field and similar considerations apply, so that we may as well assume that η is surjective. We claim:

LEMMA 2.2. *The map $\eta:\alpha \to \alpha$ is a similarity interval exchange map.*

Proof. The foliations $(\mathscr{F}^{\pm}, \mu^{\pm})$ induce on F a metric, ρ_μ, which is flat away from the singularities (see [9]). We have supposed that the arc α is very short; in particular, we can assume that it is actually a subset of a leaf of $\mathscr{F}_S^-$ and that the metric induces a length along α which is just the measure μ^+. Then the flow map is piecewise linear, as in the case of measured foliations, the difference being that if this arc flows around the monodromy factor a nonzero number of times algebraically, its μ^+-measure changes by some power of the stretching factor. Q.E.D

Definition 2.3. A similarity interval exchange map which arises in this way will be called *geometric.*

Remark. (i) This can also be thought of in terms of the affinely measured foliations of [1, 10].

(ii) The requirement that the foliations be orientable is no real restriction since there is always a double branched covering of the surface so that the lifted foliations have this property and the presence or otherwise of a closed leaf in $\mathscr{F}_S$ will be unaffected by this.

Theorem 1.1 has the following interpretation in this language. Finite foliations which arise in the context of bundles contain a finite number of closed leaves which correspond to periodic points of the associated similarity interval exchange. In general, the closed leaves may attract on one side and repel on the other, but in the case of globally orientable foliations one sees easily that each such periodic point is either attracting or repelling.

In general, it seems that similarity interval exchanges can have rather complicated dynamics; however we easily have:

LEMMA 2.4. *A geometric similarity interval exchange either has an invariant probability measure of full support (in which case this measure is unique) or the only invariant probability measures are atomic.*

Proof. The given dichotomy corresponds to the geometrically finite and infinite dichotomy provided by Theorem 1.1.

If the immersion is geometrically infinite, then any η-invariant measure gives rise to an invariant measure on one of the measured foliations associated with the monodromy of the fibre S. We briefly recall the construction: Any sufficiently short arc a on S transverse to the foliation may be isotoped (without changing which leaves are crossed) so it lies in the arc α. We use the measure given on α to define the measure of a. Of course, there may be two ways of so isotoping the arc, but these give the same answer by η invariance. The measure of a general arc can then be defined by integration; it is built into the construction that this yields an invariant measure on the foliation on S and these are well known to be uniquely ergodic [9].

On the other hand, if the surface is geometrically finite, then the foliation is finite. By passing to a power if necessary, we can suppose that the periodic points of the similarity exchange are all fixed points.

In particular, any subarc of α can be subdivided until each subarc has the property that either it contains a fixed point, or it lies entirely in the basin of attraction of some fixed point. If β is one of the arcs of the latter type, it cannot have positive measure, since any neighbourhood of the fixed point to which it is attracted contains countably many disjoint images of this arc, all of the same measure. It follows that the only possibility is a measure supported in the fixed set. Q.E.D.

Remark. Examples show that a general similarity interval exchange map may have periodic orbits which are not attractors and this means that more complicated measures may exist. Moreover, it is also known that an interval exchange map may not be uniquely ergodic [7, 8] so that even if there are no closed orbits, the invariant measure may not be projectively unique. To illustrate the dynamics of similarity exchange maps, we offer the following simple example.

Example. Fix real numbers $0 \leqslant a, s \leqslant 1$ so that $a + a \cdot s < 1$. We define a function $f: I \to I$ as follows:

$$f(x) = \begin{cases} s \cdot x + (1 - a - s \cdot a), & 0 \leqslant x < a \\ (x - a)/s + (1 - a), & a \leqslant x < a + a \cdot s \\ x - a - a \cdot s & a + a \cdot s \leqslant x \leqslant 1. \end{cases}$$

It is easily seen that this is a similarity interval exchange map, with three intervals being permuted and with stretching factor s. A systematic study of such maps has not yet been made, but we have illustrated some of the dynamics in Figs 1 and 2 for fixed $a = \frac{1}{3}$ and two values of the parameter s. These show pictures of the measures μ_K defined by

$$\mu_K(\alpha) = (\text{number of times } f^i(x) \in \alpha \text{ for } i \leqslant K)/K$$

for some fixed $x \in [0, 1]$ (Figs 1 and 2 have $K = 30{,}000$). In particular, we have not established that this apparent qualitative difference is in fact genuine.

We now begin the description of the algorithm alluded to in the introduction. This has two ingredients; firstly an algorithm to decide if an affine lamination has a closed leaf and secondly an algorithm to decide if there is a certain self-similarity of the dynamical system defined by the similarity interval exchange. Both of these parts use the fact that the similarity interval exchange is geometric. We do not know how to decide if the general similarity interval exchange contains no periodic point.

Let $\theta: F \to F$ be a pseudo-Anosov map with associated invariant foliations $(\mathscr{F}^{\pm}, \mu^{\pm})$ and for the purposes of describing these we assume that a Markov partition $\mathscr{M} = R_1 \cup \cdots \cup R_n$ has been constructed (see [9] for a description of Markov partitions and their properties). Let $M = M(\theta)$ be the mapping torus and suppose that $g: S \looparrowright M$ is a surface immersed transverse to the flow. (As a notational convenience, we suppress the map g where this will cause no confusion.) It has been shown by Mangum [2] that such a surface is automatically incompressible. By lifting to the universal covering, we see that the Markov partition induces a dissection of the universal covering of the surface S into rectangles. Of course, in general this will not give a flat structure on S since if a bunch of leaves crosses over the fibre F, then its measure is multiplied by a factor of λ or λ^{-1}; however, we do obtain a singular affine structure (see [6]). It will be necessary to do exact calculations and to this end we observe the following lemma.

LEMMA 2.5. *Suppose that α and β are algebraic over $\mathbb{Q}$ whose minimal polynomials are known. Then one can compute the minimum polynomial of $\alpha + \beta$ and $\alpha \cdot \beta$.*

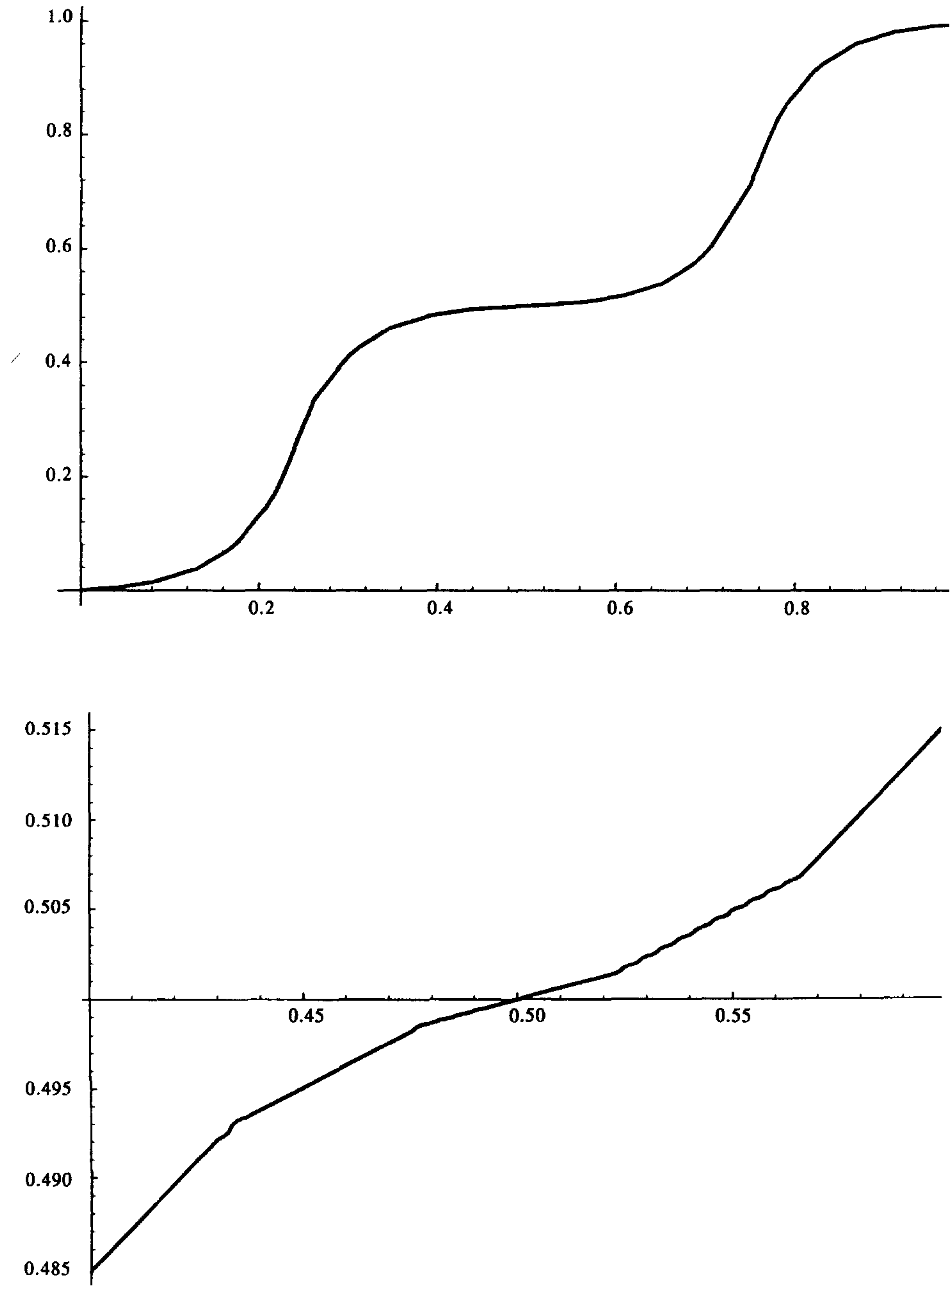

Fig. 1.

Proof. It will become clear from the proof that one can actually compute the irreducible polynomial of $\sum_{i,j} a_{i,j}, \alpha^i \beta^j$ for any rationals $a_{i,j}$. Let $f(X)$ be the irreducible polynomial for α and $g(Y)$ be the irreducible polynomial for β. Form $P(Z, X, Y) = Z - \sum_{i,j} a_{i,j} X^i Y^j$; by taking resultants (see [11]) we can construct a polynomial with integer coefficients $Q(Z)$ which can be factorised into $\mathbb{Z}$-irreducible polynomials, one of which is the irreducible polynomial for $z = \sum_{i,j} a_{i,j} \alpha^i \beta^j$. It is well known that factorisation can be done algorithmically. One then computes the number $\sum_{i,j} a_{i,j} \alpha^i \beta^j$ with sufficient accuracy that one can decide which irreducible polynomial is the relevant factor. Q.E.D.

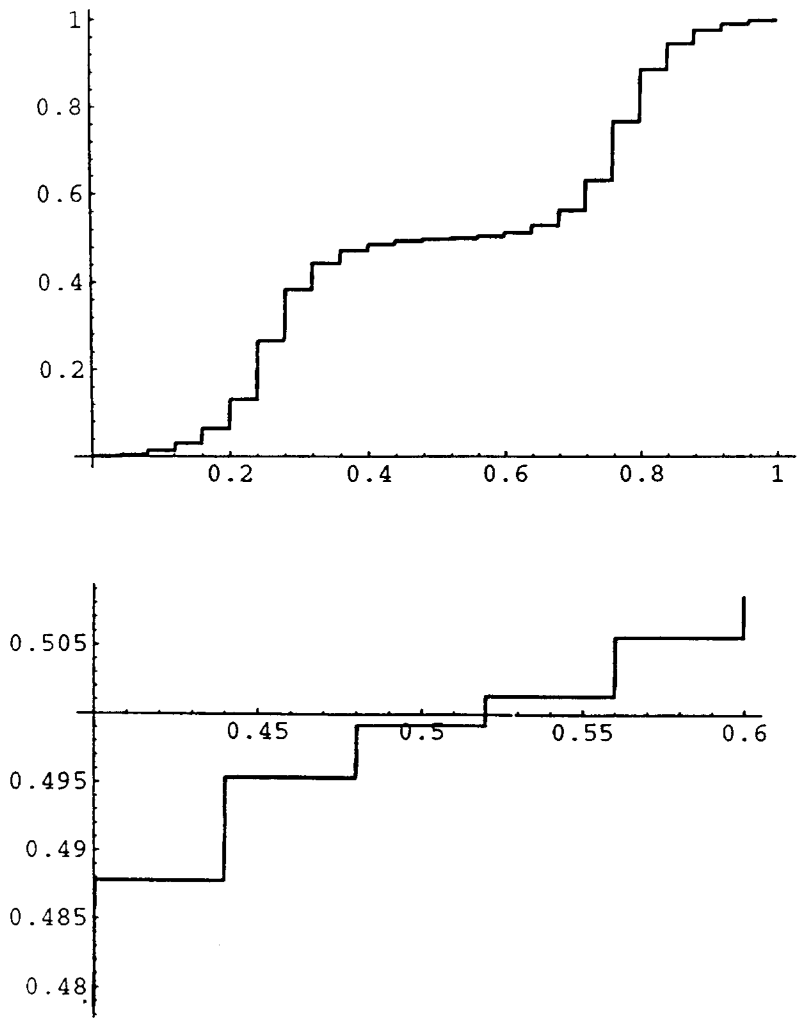

Fig. 2.

We note that the side lengths in the Markov parition $\mathscr{M}$ are the positive eigenvector of an integer matrix where the relevant eigenvalue is an algebraic integer, i.e. the dilatation. Thus, by taking a large number of resultants, we can compute the irreducible polynomials of all side lengths. In this sense, we know all side lengths exactly.

Moreover, the defining property of the Markov partition is that when the map is applied, the image rectangles either run over a rectangle in the set $\mathscr{M}$ completely or not at all. We refer to Fig. 3(a), which shows how rectangles may meet in F. By taking a large power of the map, we see that arcs of the type β and γ of the figure are each the unions of arcs each of which has measure $\lambda^k \cdot r_i$, where r_i runs over measures of rectangles R_i. For example, this situation is illustrated in Fig. 3(b), which exhibits the measure β as the sum of 3 algebraic numbers. In particular, the measure of any such arc is algebraic and exactly computable.

Our objective is to track the paths of leaves in S; we do this by "shadowing the path" down in F, using measure as our parametrisation of height in any rectangle. Thus, we may regard the bundle M as $F \times [0, 1]$ with top and bottom levels identified via the monodromy and a cut-apart version of S embedded transverse to the I fibres. A typical leaf in S shadows a leaf which runs from rectangle to rectangle in F (we refer again to Fig. 3(a)) but

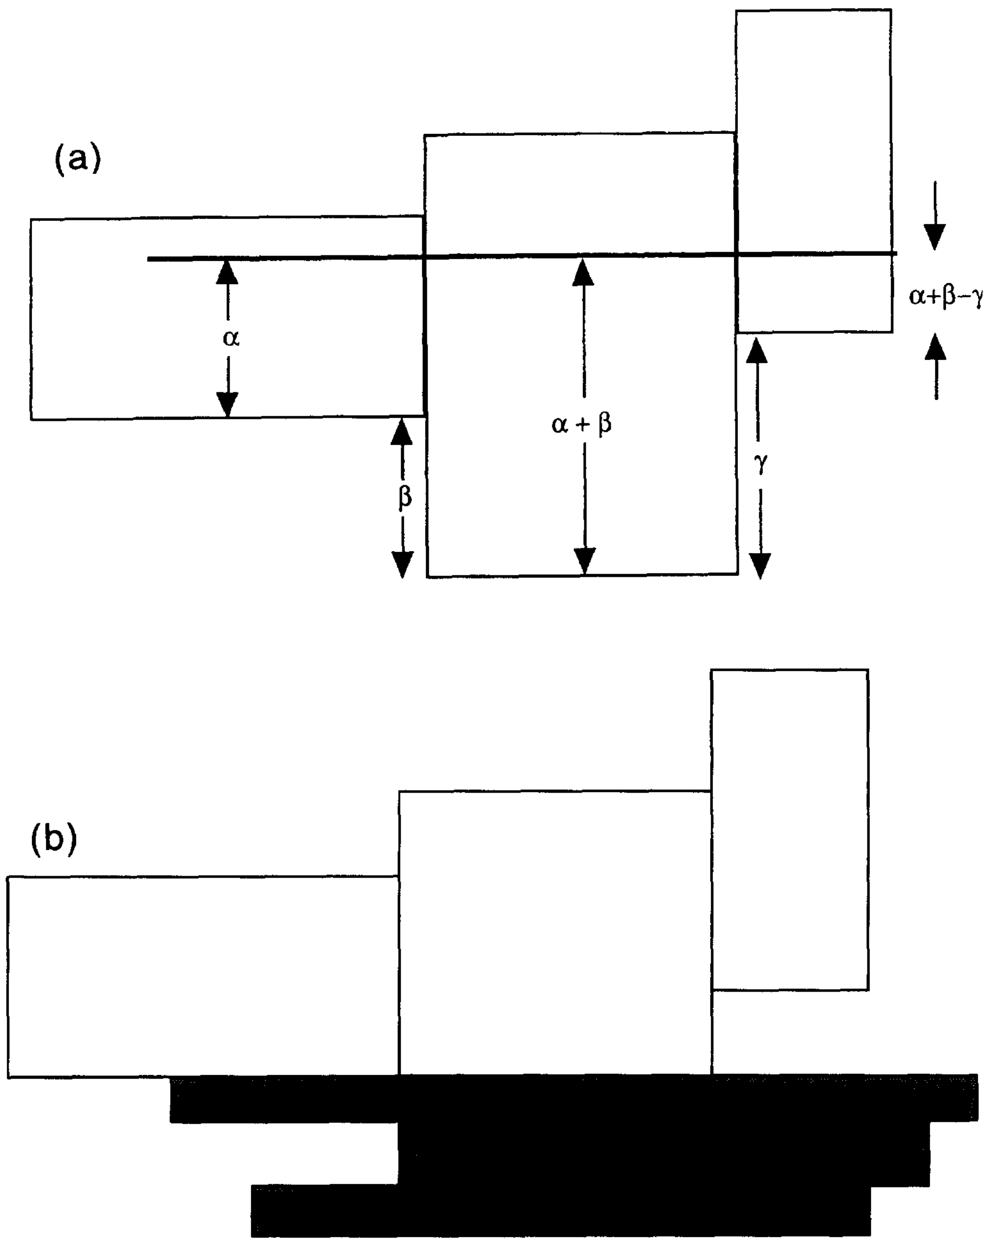

Fig. 3.

occasionally it will be necessary to cross the levels $F \times \{0\}$ or $F \times \{1\}$, and this entails an application of the monodromy; this kind of jump is illustrated in Figs 4(a) and (b) which exhibits a leaf-shadow in F which in S is crossing from a rectangle to its image. As the arc crosses from top to bottom, it may enter the image rectangle in the middle of some R_a; however this shift is accounted for by the image measures k_1 and k_2 of other rectangles. Now we note that, as we observed above, the measures k_1 and k_2 are algebraic, so that if the leaf-shadow in F is at a height which is algebraic (in the notation of Fig. 4(a), this is the statement $\alpha + \sum_i \beta_i$ is algebraic) then its height in the new rectangle (exhibited in Fig. 4(b) as $k_1 + k_2 + \lambda(\alpha + \sum_i \beta_i)$) in F is also algebraic. This observation suffices for the following lemma.

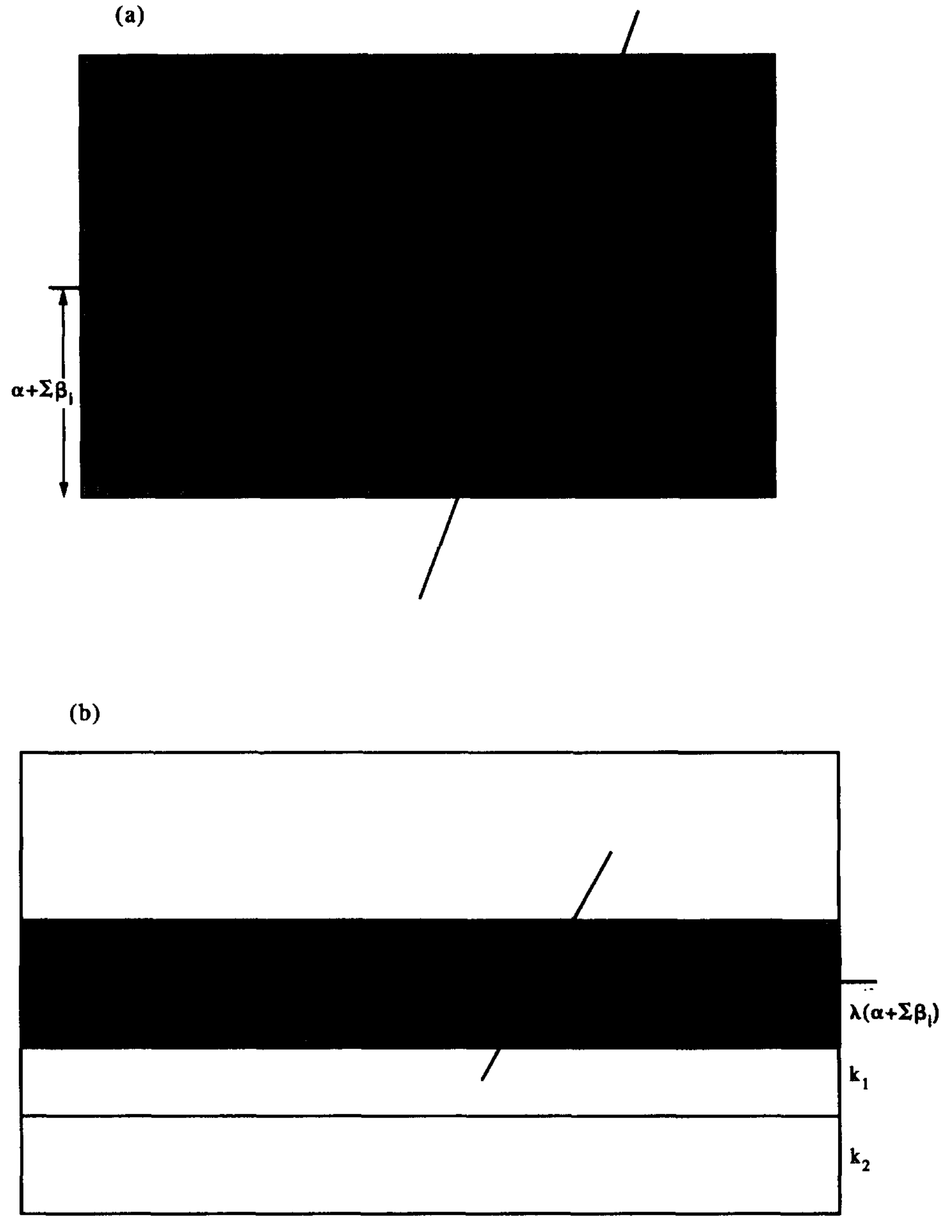

Fig. 4.

PROPOSITION 2.6. *Suppose that C is a closed leaf on S passing through rectangle $R \subset F$ at height (in the sense of measure) α. Then α is an algebraic number.*

Proof. Suppose that we are given a path of a leaf in S which closes up. Then as we remarked above, the path of the leaf in question shadows the path of some leaf in F as this leaf moves from rectangle to rectangle in F, except when the leaf crosses the 0 or 1 levels, when there is an application of the monodromy whose effect on the coordinates is multiplication by λ or its reciprocal. Given a point $p \in R$ which is at some height α (in the sense of measure) as shown, the height transitions from rectangle to rectangle in F are as shown in Figs 3 and 4. Using this description, we can follow the path of the leaf in S, and we see that the condition that it closes up is a linear equation in $\alpha, \beta_1, \ldots, \beta_m$ with integer coefficients of the form $f(\lambda)\alpha + g(\lambda, \beta_1, \ldots, \beta_m) = \alpha$. The coefficient $f(\lambda)$ cannot be 1 as this would imply that the closed leaf ran over the fibre algebraically zero times. This is impossible as it is shown in [1] the closed leaf can never be freely homotoped into the fibre group. Q.E.D.

COROLLARY 2.7. *If the foliation $\mathscr{F}_S$ contains a closed leaf then this is decidable.*

Proof. In purely formal terms, one proceeds as follows: For each rectangle enumerate the algebraic heights. Fix one such height, α. The above considerations show that we can follow the path of the leaf at that height exactly. We then examine each return to the rectangle in question and see whether the associated linear equation has the number α as a solution. It does if and only if there is a closed leaf at this height. If $\mathscr{F}_S$ has a closed leaf, this process will eventually find it.

Remark. It seems worthwhile to point out that, in practice, the process is simpler and faster than this. Theorem 1.1 shows that any leaf in S spirals onto one of the closed leaves in $\mathscr{F}_S$. Accordingly, we choose any leaf and move forward and wait until this leaf comes very close to itself, say some computer calculation repeats a height to a large number of decimal places. This yields a segment of leaf and thus a candidate path for the closed leaf, which in turn yields a linear equation of the type promised by the above proposition. One then verifies that the exact solution to this equation follows the purported path. This technique is particularly powerful in the presence of the branched flat structures of [12]. See Section 3 where we construct the closed leaf in this fashion.

The second part of the algorithm involves establishing if there is no closed leaf. Here again we shall strongly use the geometric nature of our situation. The key fact here is that if the surface is to be geometrically infinite, it is embeddable in some finite sheeted covering where it is transverse to the lift of the flow $\mathscr{L}$ (see [4]). It follows from the results of Fried [13] that the monodromy of S comes from the first return map of $\mathscr{L}$ on S in this finite sheeted covering. Let ρ be a closed flowline in $\mathscr{L}$; this corresponds to a periodic point of the monodromy θ, by passing to a finite sheeted cover (and taking the preimage of S) we may suppose that ρ meets F once, that is to say that $\rho \cap F$ is a fixed point of θ.

Fix some point of intersection $q \in \rho \cap S$ and choose an arc I lying in $\mathscr{F}^-$ containing q in its interior. Application of powers of the monodromy θ to this arc maps it inside itself and the observation of the above paragraph implies that if the surface S is geometrically infinite, then some power of the monodromy θ is the restriction of the monodromy associated with S applied to the arc α. This has the consequence that if the surface S is geometrically infinite, the similarity interval exchange map associated with the first return of $\mathscr{F}_S^+$ on the arc I is isomorphic to the similarity interval exchange associated with the first return on $\theta^k(I)$ for some $k > 0$. Our notion of isomorphism here is the obvious one: Two similarity exchanges α_1 and α_2 are isomorphic if there is a homeomorphism $h: I \to I$ so that $\alpha_1 = h\alpha_2 h^{-1}$.

In fact, the above argument shows more than just isomorphism; the interval $\theta^k(I) \subset I$ and there is an algebraic number so that the homothety centred at q which maps $\theta^k(I) \to I$ conjugates one similarity exchange to the other. We refer to such an isomorphism as an *algebraic homothety*.

Moreover, we claim that once such an algebraic homothety is established, this also proves the surface must in fact be geometrically infinite. The reason is that by Theorem 1.1, the alternative is that the similarity exchange has a finite number of periodic points, corresponding to where the closed leaves meet I. However, the existence of an algebraic homothety means that the presence of any periodic point other than q will force infinitely many periodic points, a contradiction. However, q cannot be a periodic point for the similarity exchange, as [1] shows that (in the terms of the description that we use here) the closed leaf on S contains no point on F which is θ-periodic. Therefore, the final ingredient in our algorithm is:

PROPOSITION 2.8. *In the above notation, it is possible to compute the geometric similarity interval exchange map exactly. Moreover, given two geometric similarity interval exchanges, one can decide if they are conjugate by an algebraic homothety.*

Proof. The proof of Proposition 2.6 actually shows a little more: If a leaf is at algebraic height in some rectangle, the entire forwards and backwards history of the leaf is at algebraic heights in every rectangle. Since singularities of the foliation lie in the frontiers of the rectangles of $\mathscr{M}$ (i.e. at height zero or algebraic) this implies that all separatrices of the foliation meet every rectangle at algebraic heights. We construct the similarity exchange as follows: Fix some small arc I lying entirely inside one rectangle and follow out each separatix until the first hit on I. Note that in our case, every leaf is dense; otherwise there is a closed leaf, so we may assume that there is always such a hit. Exactly as in the case of an interval exchange map, these hits define the exchange. Moreover, all the hits on I are at algebraic heights so that all the intervals have algebraic widths. As above all computations can be made exactly.

The question of whether two given similarity exchanges are conjugate by algebraic homothety can also be decided as this is exactly the question of whether one set of algebraic numbers is a constant multiple of another set. Q.E.D

COROLLARY 2.9. *If $\mathscr{F}_S$ contains no closed leaf, then this is decidable.*

Proof. Theorem 1.1 shows that this happens if and only if the surface is geometrically infinite, so there is a monodromy map and, as explained above, an isomorphism of the similarity interval exchange on I with the similarity interval exchange on $\theta^j(I)$ for some j.

The procedure to decide this is the following. We compute exactly the similarity exchanges on each of the arcs in the descending sequence $I \supset \theta(I) \supset \theta^2(I) \supset \cdots$ and at each stage test for isomorphism. The proposition shows that all of these calculations can be done exactly. Q.E.D.

In the case that one constructs an algebraic homothety of the similarity exchange, this in fact constructs the monodromy, which is precisely the map which carries rectangles to the isomorphic copies of themselves. In particular, we obtain a Markov partition which can be used to compute the dilatation and defines the invariant measure.

3. AN EXAMPLE

In this section we use the ideas developed earlier to exhibit an example of a surface obtained by a single cut and cross join inside a hyperbolic bundle which is geometrically infinite. With our current lack of understanding, such examples appear to be rather hard to construct.

We begin with some generalities from [1]. In our standard notation, assume that C is a simple closed curve on F which is disjoint from θC and that $C \cup \theta C$ does not separate. Let

F_- be the surface obtained by cutting F open along C and θC and then compactifying. Thus, ∂F_- has four components $C_+, C_-, \theta C_+, \theta C_-$ where the signs are chosen so that θ takes the $+$ side of C to the $+$ side of θC. Now define S to be the surface obtained from F_- by identifying C_+ with θC_- via θ and similarly identifying C_- with θC_+. Thus, S is an orientable connected surface. Informally, we refer to this construction as a cut and cross join. It is easily seen that:

LEMMA 3.1. *There is an immersion $g: S \looparrowright M$ which is transverse to $\mathscr{L}$.*

To generate examples, we use the method for constructing pseudo-Anosov maps, described in [12], which is a slight generalisation of an idea due to Thurston (see [9]).

Consider the pair of one submanifolds $\mathscr{C}$ and $\mathscr{D}$ of the surface F depicted in Fig. 5. These fill the surface and so define a flat struture with a single singularity, shown in Fig. 6. If X is any simple closed curve on this surface, we shall denote the Dehn twist about X by T_X. In this notation, we define maps $T(\mathscr{C}) = T_{C_1}^2 T_{C_2} T_{C_3}$ and $T(\mathscr{D}) = T_{D_1} T_{D_2} T_{D_3}^2$. If we regard each of the rectangles in Fig. 6 as a square, we see that both these maps are affine maps of the flat structure, inducing the linear maps

$$[T(\mathscr{C})] = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad [T(\mathscr{D})] = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix}.$$

We construct pseudo-Anosov maps by exhibiting affine maps of the flat structure whose local behaviour is given by a 2×2-matrix of determinant 1 and trace bigger than 2. To this end, we introduce another affine map $\xi: F \to F$. Consider Fig. 6; each corner in this picture carries a label and if we develop the picture of these corners around the singularity, we see the picture given in Fig. 7. In particular, one can check from this picture that the map which rotates clockwise around the singularity carrying corner 1 to corner 18 extends to an affine map $\xi: F \to F$ which induces the liner map

$$[\xi] = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

One also sees that $\xi(C_1) = D_3$ with the orientations shown in Fig. 8. In this notation, we consider the family of maps, $\theta(k) = \xi \circ T(\mathscr{C})^k$ which appear in the affine structure as

$$[\xi] \cdot [T(\mathscr{C})]^k = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 2k \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -1 & -2k \end{pmatrix}.$$

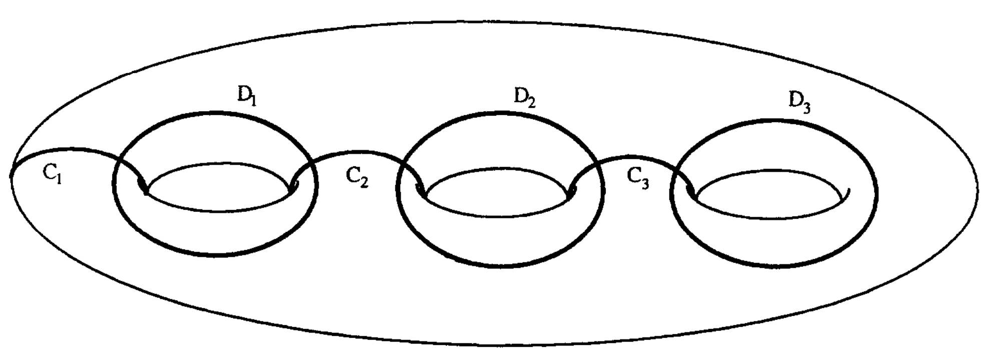

Fig. 5.

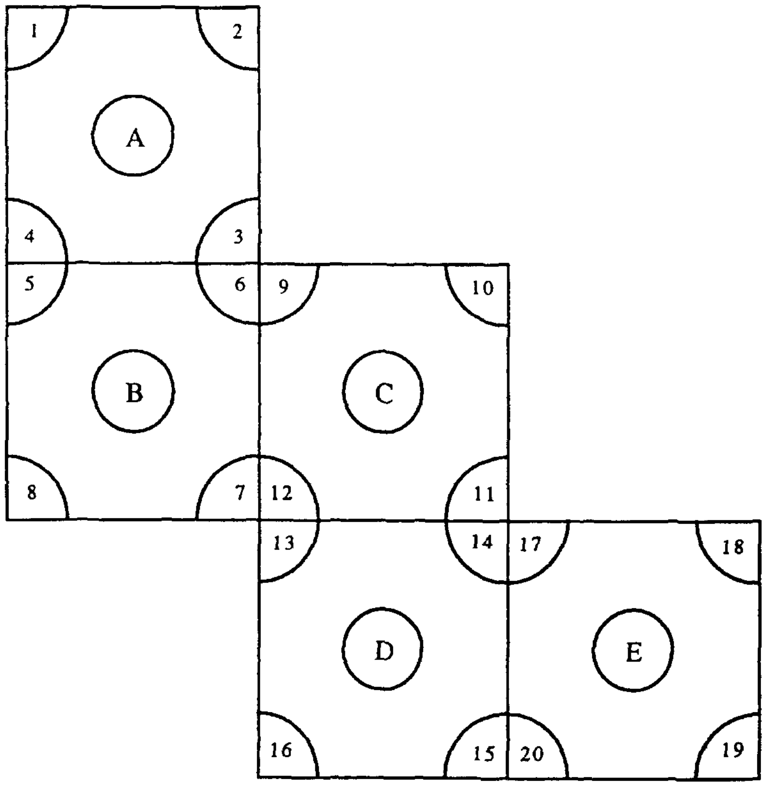

Fig. 6.

These maps all have the property that they carry C_1 to D_3 with the orientation shown in Fig. 8 and if $|k| > 1$ they are pseudo-Anosov. Therefore, each bundle $M(\theta(k))$ contains a surface $S(k)$ obtained from the fibre F by doing a single cut and cross join using the pair $\{C_1, \theta(k)(C_1) = D_3\}$. We claim that the methods developed above show:

THEOREM 3.2. (1) *The surface* $S(-3)$ *is geometrically finite.*
(2) *The surface* $S(-2)$ *is geometrically infinite.*

Computer experimentation suggests the following conjecture may be true. Our results show that it can be verified on a case by case basis, but no theoretic proof has yet been found.

CONJECTURE 1. *With the notation established above, if* $|k| > 1$, *the surface* $S(k)$ *is geometrically finite when* k *is odd.*

It is in fact easy to verify for low values of odd k that this conjecture is true. By contrast, the geometrically infinite case is much more difficult to deal with and only the case $k = -2$ has been thoroughly checked. Nonetheless, we conjecture (with somewhat less confidence):

CONJECTURE 2. *The surface* $S(k)$ *is geometrically infinite when* k *is even.*

The proof of Theorem 3.2 goes as follows. One computes that if $k = -3$, then the invariant foliations for $\theta(-3)$ have slopes $3 \pm \sqrt{8}$. One can work with either of these

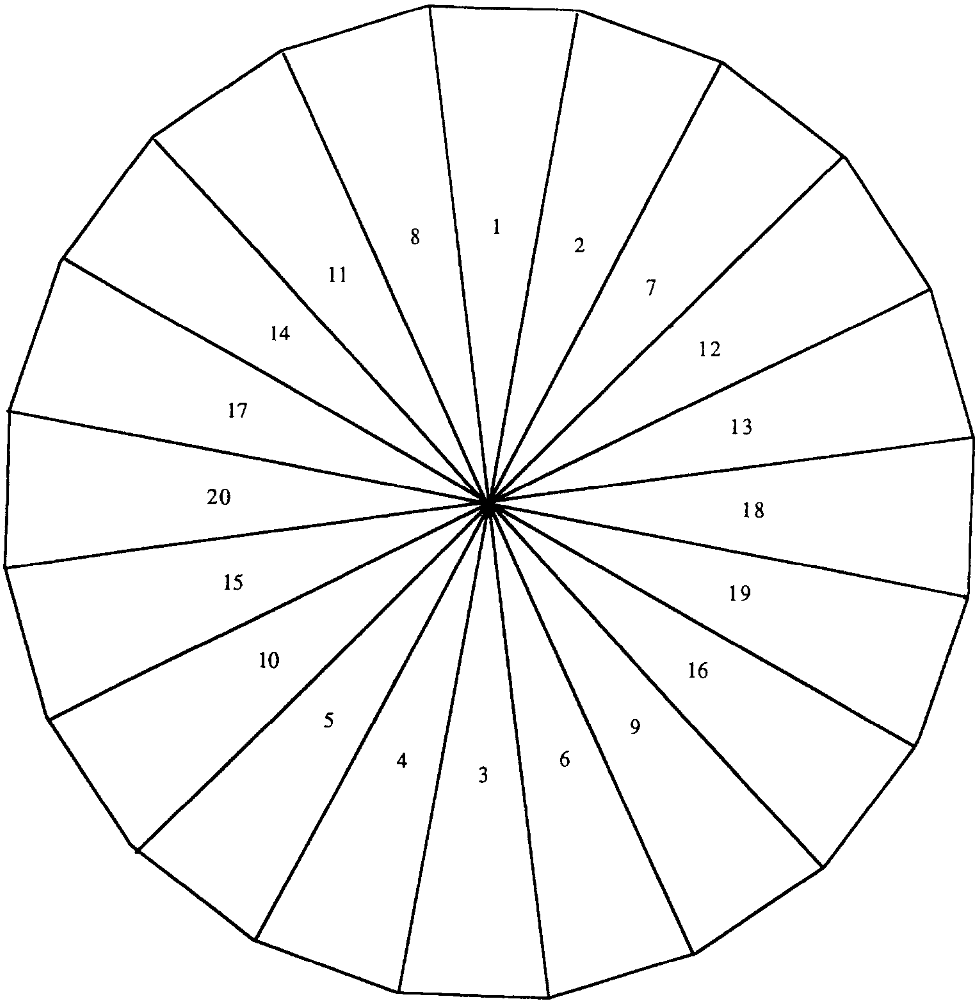

Fig. 7.

slopes; if we choose the positive sign, then as indicated above, one easily finds an approximation to the closed leaf.

The candidate path is shown in Fig. 9. In this context, one quickly shows by linear algebra and trigonometry that the solution is an exact one. This uses the fact that all the squares have side length one. We indicate how this can be done. For brevity, we denote the slope of the foliation by $1/s$ and measure lengths along each horizontal edge. Denote the starting point 1 by length α. The point 2 then has coordinate $\alpha + 2s$, the point 3 has coordinate $\alpha + 4s$. Then some trigonometry reveals that point 4 has horizontal coordinate $1 - (1 - \alpha - 4s)/s = 5 - 1/s + \alpha/s$. Continuing in this way we see that the other end of the leaf has coordinate $-1 + 4s + \alpha \cdot s + 16s^2 + 2s^3$ so that the condition that the leaf is closed becomes $-1 + 4s + \alpha \cdot s + 16s^2 + 2s^3 = \alpha$. After simplifying with $s = 3 - \sqrt{8}$, this is a linear equation for α which one readily checks has a solution.

The case $k = -2$ involves a more complicated calculation; however, matters are simplified by the presence of a point L fixed by the monodromy $\theta(-2) = \theta$. In this case, the foliations have slopes $2 \pm \sqrt{3} = \mathscr{F}^{\pm}$ and Fig. 10 shows an arc $\alpha \subset \mathscr{F}^{-}$ lying between points a and b, which are defined as the first points of intersection of the $\mathscr{F}^{+}$-singular arcs shown with a short $\mathscr{F}^{-}$arc through L.

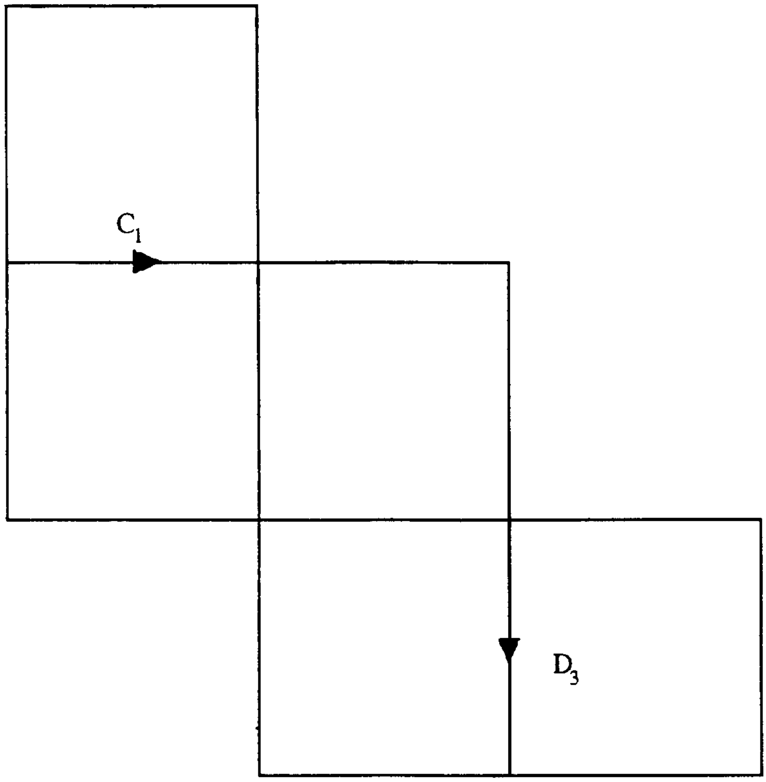

Fig. 8.

One then computes the first return map for the $\mathscr{F}^+$ foliation on the arc α; this is the similarity interval exchange shown in Fig. 11. (As a convenience, the edge α is parametrised by the hits of the arcs u_i and v_i on the horizontal edge just below α; here $s = 2 - \sqrt{3}$ is the reciprocal of the slope of $\mathscr{F}^+$. This is the table on the left.) The monodromy θ acts as a five cycle on the separatrices, so this suggests that one computes the first return maps for the images $\theta^{5i}(\alpha)$. If $i = 1$ there is no algebraic homothety, but if $i = 2$, one finds that the similarity interval exchange coming from $\theta^{10}(\alpha)$ is isomorphic to that on α, with parameters shown in the right-hand table. This shows that the surface $S(-2)$ is geometrically infinite and gives a complete description of the monodromy map. This completes the proof of Theorem 3.2. Q.E.D.

We conclude with some remarks of a general nature. If we base all fundamental groups at L and let the closed flowline through L be the element t, this shows that $\langle \pi_1(S(-2)), t^{10}\rangle$ is a subgroup of finite index (in fact, the index is 10) inside $\pi_1(M(\theta))$ and if we pass to the covering corresponding to this subgroup, the surface $S(-2)$ lifts to an embedding. An easy computation shows that $H_1(M(-2); \mathbb{Q}) \cong \mathbb{Q}$; in particular, the surface $S(-2)$ shows that rank of the rational homology of $M(-2)$ can be increased by passing to a finite sheeted covering. It is also easily shown that the subgroup $\langle \pi_1(S(-2)), t^{10}\rangle$ is not normal, so that the associated covering is an irregular 10-fold covering of $M(-2)$.

Some of the power of this method comes from the fact that once the similarity exchange map has been computed, this can be used to compute data associated with the monodromy of the fibration $S(-2)$. Such computations seem to be the first such of their type for closed

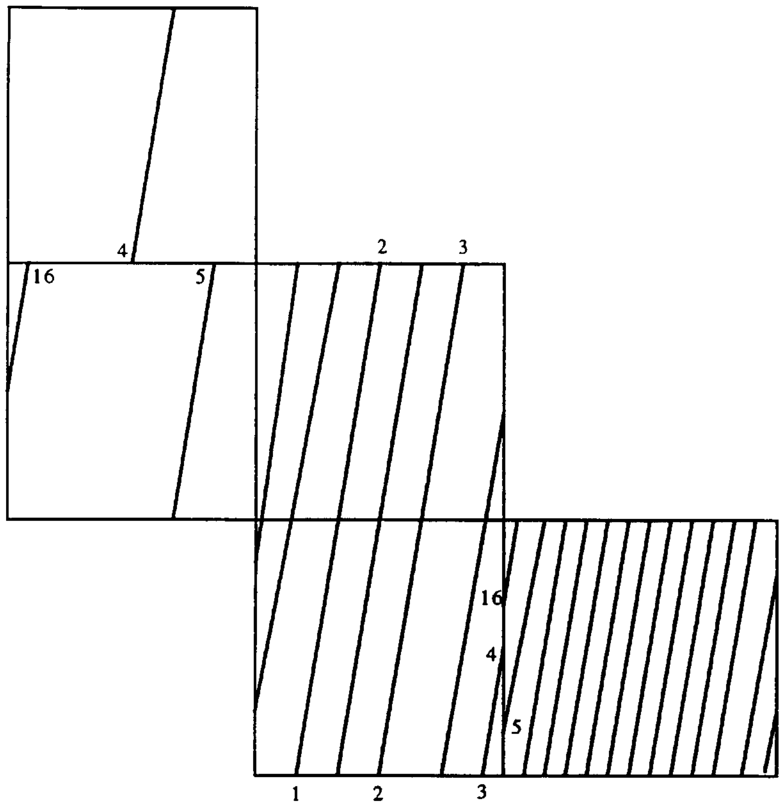

Fig. 9.

hyperbolic manifolds. Indeed, although there is an example of a non-Haken hyperbolic 3-manifold which is known to contain an immersion of a geometrically infinite surface [14], the immersion constructed there is far from explicit, relying on certain arithmetic coincidences. One pleasing aspect of the above construction is that it does give the first explicit immersion of a geometrically infinite surface into a hyperbolic 3-manifold distinct from a fibration.

Moreover, for the example above, further calculation shows that the dilatation of the monodromy of the fibration $S(-2)$ is the largest root of the equation

$$1 - 1110416t - 1285122t^2 - 1110416t^3 + t^4 = 0.$$

A similar calculation gives an explicit description of the measure. Notice that the fields $\mathbb{Q}(\sqrt{3})$ and $\mathbb{Q}(t)$ do not seem to be related in any obvious way.

One of the reasons such an example is interesting is that it seems to be a rare phenomenon that a surface obtained from a fibre by a single cut and cross join is geometrically infinite. Indeed, we may generalise as follows. Let F_1 be any covering of F to which we can find a lift $\tilde{C}$ of C and a lift $\widetilde{\theta(C)}$ of $\theta(C)$. (Note that we do not require that the map θ lifts.) Then we can immerse the covering space F_1 inside $M(\theta)$ transverse to the flow and do a single cut and cross join on F_1 using the curves $\tilde{C}$ and $\widetilde{\theta(C)}$. This surface which we denote by $S = S(\tilde{C}), \widetilde{\theta(C)}, F_1)$, is transverse to the flow. We ask:

Question. Is there always a covering F_1 as above for which $S = S(\tilde{C}, \widetilde{\theta(C)}, F_1)$ is geometrically infinite?

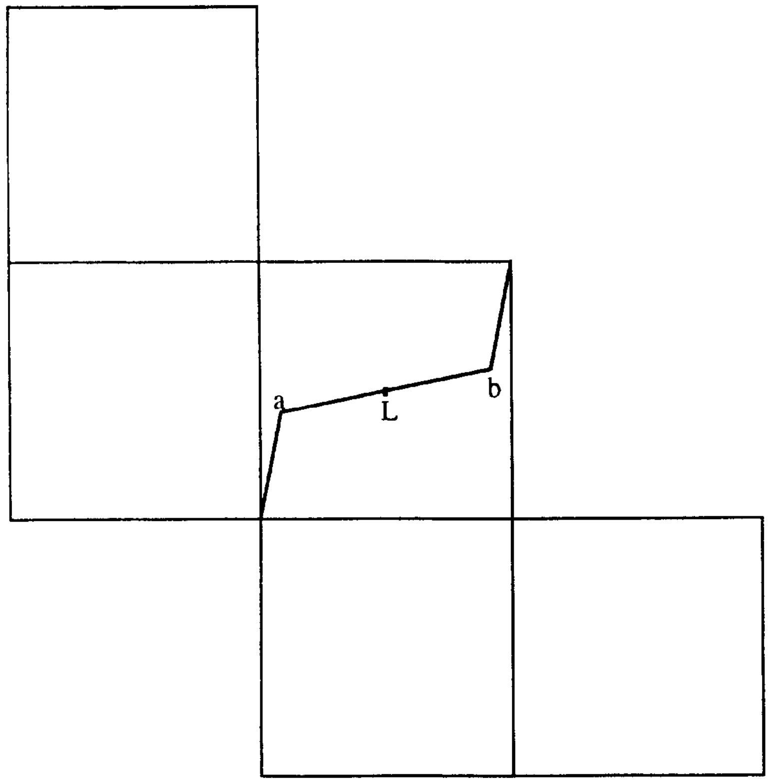

Fig. 10.

We conjecture the answer is yes; however, to emphasise further our interest in constructing geometrically infinite surfaces by a single cut and cross join, we have shown:

THEOREM 3.3. *Suppose that $\theta: F \to F$ is a pseudo-Anosov map for which we can find a simple closed curve C with* (a) *$C \cup \theta(C)$ is non separating, and* (b) *C, $\theta(C)$ and $\theta^2(C)$ are all disjoint. Then either*

(1) *the immersion $g: S(C, \theta(C), F) \looparrowright M(\theta)$ cannot be lifted to become an embedding in any finite covering of $M(\theta)$,*

or

(2) *there is a finite sheeted covering $p: F_1 \to F$, a map θ_1 covering θ and a lift C_1 of C so that the immersion $S = S(C_1, \theta_1(C_1), F_1) \subset M(\theta)$ is geometrically infinite.*

Proof. Without loss of generality, there is a point $p \in F$ fixed by the monodromy θ, and fundamental groups referred to will be assumed to be based at p.

Let F_- be the complement in F of open regular neighbourhoods of C and θ(C). As usual, we regard $g(S)$ as built out of F_- together with a pair of annuli embedded in $M(\theta)$ transverse to the flow.

Suppose that the immersion $g(S)$ can be lifted to an embedding in a finite cover which we denote by M_0, which we regard as the mapping torus of $\tilde{\theta}: F_0 \to F_0$. Since $\pi_1(F_-)$ is a subgroup of $\pi_1(S)$ it lifts to M_0, and gives an embedded subsurface X of the lifted fibre F_0. F_- has two boundary components $\tilde{C}_1$ and $\tilde{C}_2$ which cover two components of ∂F_-, and two boundary components $\widetilde{\theta(C)}_1$ and $\widetilde{\theta(C)}_2$ which cover the remaining two components of ∂F_-. Since $g(S)$ lifts to M_0 this lift is recovered by adjoining two annuli to X. It follows that

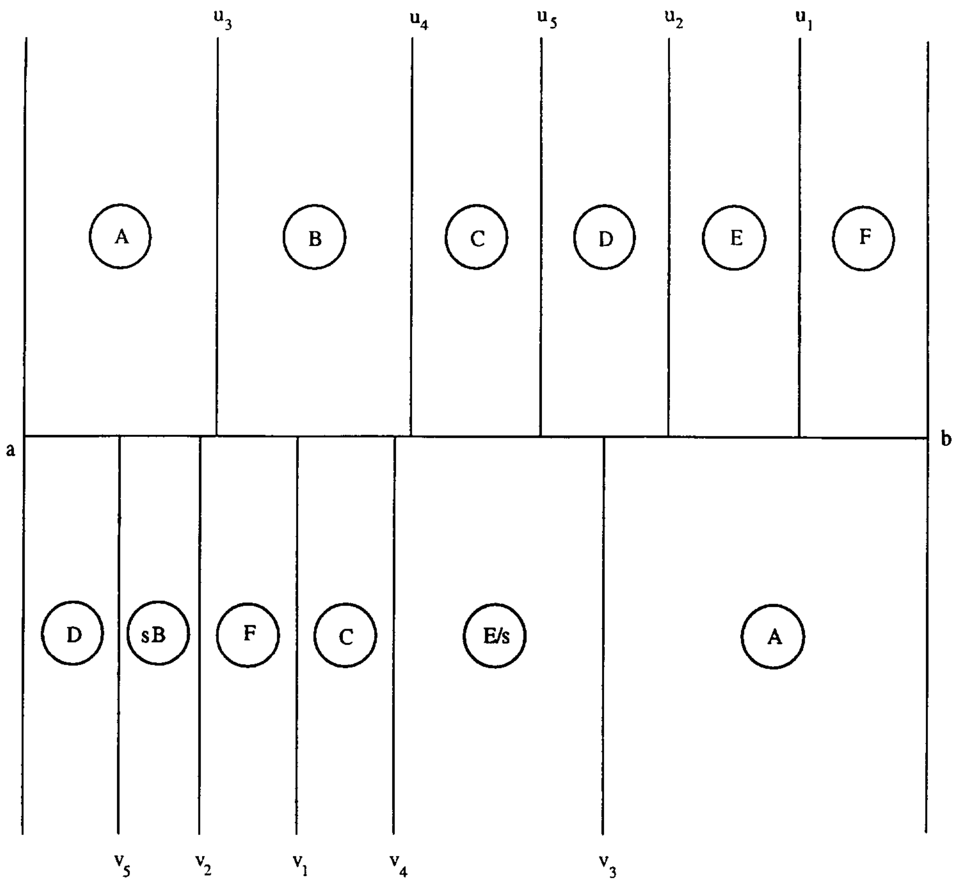

$u_1 = 2 - 5s$	$v_1 = -3 + 12s$	$u_1 = 620105 - 2314262s$	$v_1 = -1749542 + 6529381s$
$u_2 = 3 - 9s$	$v_2 = -2 + 8s$	$u_2 = 1184824 - 4421822s$	$v_2 = -1184823 + 4421821s$
$u_3 = 1 - 3s$	$v_3 = 2s$	$u_3 = 358018 - 1336140s$	$v_3 = -358017 + 1336139s$
$u_4 = 1 - 2s$	$v_4 = s$	$u_4 = 206702 - 771421s$	$v_4 = -206701 + 771420s$
$u_5 = 4 - 13s$	$v_5 = -1 + 4s$	$u_5 = 1749543 - 6529382s$	$v_5 = -620104 + 2314261s$

Fig. 11.

$\tilde{\theta}\{\tilde{C}_1, \tilde{C}_2\} = \{\widetilde{\theta(C)}_1, \widetilde{\theta(C)}_2\}$. After renumbering if necessary, we can assume that $\tilde{\theta}(\tilde{C}_i) = \widetilde{\theta(C)}_i$ for $i = 1, 2$.

Let Y be the closure of the complement of X in F_0; this may be disconnected. We form a new surface $S(Y)$ by attaching a pair of annuli to Y in the obvious way. Notice that $S(Y)$ is transverse to the flow and is therefore incompressible. We claim $S(Y)$ is geometrically infinite. It is well-known that it suffices to show that $S(Y)$ meets every flowline (see [1, Theorem 3.14]). Now $S(Y)$ meets every flowline except possibly those that meet X. On the other hand, it follows from hypothesis (b) that preimages of C, $\theta(C)$ and $\theta^2(C)$, are all disjoint, which in turn implies $\theta(X) \subset Y$, so that $S(Y)$ meets every flowline, justifying the claim.

If p_0 is the covering projection of M_0 over $M(\theta)$, then the restriction of p_0 to Y can be perturbed slightly so that the image is immersed transverse to the flow which two of the boundary components are close to, and either side of $C \subset F \times \{\frac{1}{2}\}$, and similarly the other two boundary components are copies of $\theta(C)$. Working in $F \times I$ as opposed to $M(\theta)$ we now define a surface F_1 in $F \times I$ by attaching two annuli, one between $p(\tilde{C}_1)$ and $p(\tilde{C}_2)$ and the other between $p(\widetilde{\theta(C)}_1)$ and $p(\widetilde{\theta(C)}_2)$. Since F_1 is immersed transverse to the I-direction it is a covering of F. Moreover by construction, the surface $S = S(C_1, \theta_1(C_1), F_1)$ lifts to the surface $S(Y)$ and is therefore geometrically infinite. Q.E.D.

Remark. (i) Condition (b) is not a serious restriction, since it can always be arranged in a finite sheeted covering of F to which the monodromy θ lifts.

(ii) It seems possible that for sufficiently complicated monodromies, condition (2) might fail. This would give hyperbolic manifolds whose fundamental groups are not subgroup separable, even on a geometrically finite surface group.

REFERENCES

1. D. Cooper, D. D. Long and A. W. Reid: Bundles and finite foliations, *Invent. Math.* **118** (1994), 255–283.
2. B. Mangum: Incompressible surfaces and pseudo-Anosov flows, preprint.
3. A. Marden: The geometry of finitely generated Kleinian groups, *Ann. Math.* **99** (1974), 383–461.
4. F. Bonahon: Bouts des variétés hyperboliques de dimension 3, *Ann. Math.* **124** (1986), 71–158.
5. T. Soma: Virtual fibers in hyperbolic 3-manifolds, *Topology Appl.* **41** (1991), 179–192.
6. A. Hatcher and U. Oertel: Affine lamination spaces for surfaces, preprint
7. H. Masur: Interval exchange maps and measured foliations, *Ann. Math.* **115** (1982), 169–200.
8. W. Veech: Gauss measures for transformations on the space of interval exchange maps, *Ann. Math.* **115** (1982), 201–242.
9. A. Fathi, F. Laudenbach and V. Poenaru: Travaux de Thurston sur les surfaces, *Astérisque* **66–67** (1979).
10. D. D. Long and U. Oertel: Hyperbolic surface bundles, preprint.
11. S. Lang: *Algebra*, Addison-Wesley, Reading, MA.
12. D. D. Long: Constructing pseudo-Anosov maps., in *Knots and manifolds*, D. Rolfsen, Ed., SLNM **1144**, Springer, Berlin (1983), pp. 108–114.
13. D. Fried: Fibrations over S^1 with pseudo-Anosov monodromy, Exposé 14 in Travaux de Thurston sur les surfaces, Astérisque, 66–67 (1979).
14. A. W. Reid: A non-Haken hyperbolic 3-manifold covered by a surface bundle, *Pacific J. Math.* **167** (1995), 163–182.

Department of Mathematics
University of California
Santa Barbara, CA 93106, U.S.A.

Department of Pure Mathematics
University of Cambridge
Cambridge CB2 1SB, U.K.

Current address of D. D. Long:
Department of Mathematics
University of Texas
Austin, TX 78712, U.S.A.