Bundles and finite foliations.

D. Cooper * D. D. Long † A. W. Reid ‡

1 Introduction.

By a hyperbolic 3-manifold, we shall always mean a complete orientable hyperbolic 3-manifold of finite volume. We recall that if Γ is a Kleinian group then it is said to be **geometrically finite** if there is a finite-sided convex fundamental domain for the action of Γ on hyperbolic space. Otherwise, Γ is geometrically infinite. If Γ happens to be a surface group, then we say it is **quasi-Fuchsian** if the limit set for the group action is a Jordan curve C and Γ preserves the components of $S^2_\infty \setminus C$. The starting point for this work is the following theorem, which is a combination of theorems due to Marden [10], Thurston [14] and Bonahon [1].

Theorem 1.1 *Suppose that M is a closed orientable hyperbolic 3-manifold. If $g : S \looparrowright M$ is a π_1-injective map of a closed surface into M then exactly one of the two alternatives happens:*
- *The geometrically infinite case: there is a finite cover $\tilde{M}$ of M to which g lifts and can be homotoped to be a homeomorphism onto a fiber of some fibration of $\tilde{M}$ over the circle.*
- *The geometrically finite case: $g_*\pi_1(S)$ is a quasi-Fuchsian group.*

The dichotomy between geometrically finite and geometrically infinite is fundamental and despite the fact that these two cases exhibit widely different behaviour, it seems to be a very difficult problem in general to find a criterion in terms of the image $g(S)$ which distinguishes them.

In this paper we shall give such a criterion which covers a very natural class of π_1-injective immersions in the special case that the hyperbolic 3-manifold is a surface bundle over the circle. To state our main theorem, we need some notation. Suppose that $g : S \looparrowright M$ is a π_1-injective immersion into a closed hyperbolic 3-manifold M which fibers over the circle. The fact that M is a bundle gives some extra structure which comes from the presence of a canonical (up to isotopy) flow $\mathcal{L}$ and two canonical foliations. The flow $\mathcal{L}$ comes from the suspension flow of the product structure on $\{fiber\} \times I$ and the foliations come by suspending the foliations left invariant by the monodromy of the bundle, which is pseudo-Anosov as we assume throughout that the bundle is hyperbolic. To exploit this extra structure, we shall also assume that the image $g(S)$ is transverse to the flow. Then either suspended foliation induces on S a foliation, denoted $\mathcal{F}_S$ and our first theorem is the following:

Theorem 1.2 *Suppose that M is a closed hyperbolic 3-manifold which fibers over the circle and that the immersion $g : S \looparrowright M$ is π_1-injective and transverse to the suspension flow. Then $g_*(\pi_1(S))$ is geometrically infinite if and only if $\mathcal{F}_S$ contains no closed leaf.*

In fact, one can say much more about the structure of the induced foliation in the quasi-Fuchsian case. Define a foliation of a surface to be **finite** if it consists of some finite number of closed leaves and each end of every other leaf spirals towards one of these closed leaves. Then a sharper result is:

*Partially supported by the NSF
†Partially supported by the A. P. Sloan Foundation and the NSF
‡Partially supported by the NSF

Theorem 1.3 *Suppose that M is a closed hyperbolic 3-manifold which fibers over the circle and that the immersion $g : S \looparrowright M$ is π_1-injective and transverse to the suspension flow. Then $g_*(\pi_1(S))$ is a quasi-Fuchsian subgroup if and only if the foliation $\mathcal{F}_S$ is a finite foliation.*

This seems to be the only criterion known for distinguishing the quasi-Fuchsian from the geometrically infinite in terms of the image of the immersion in the bundle M. Moreover, these conditions are checkable in examples - in §4 we give an example of a surface immersed into a bundle which is shown to be quasi-Fuchsian by proving that the resulting foliation has a closed leaf.

Another reason that an understanding of such immersions is interesting is related to the **virtual Betti number** of M, which by definition is

$$sup\{rank(H_1(M_F;\mathbb{R})|\ M_F\ is\ a\ finite\ covering\ of\ M\}$$

The conjecture is that for a closed hyperbolic 3-manifold, the virtual Betti number is infinite; though even going from Betti number zero to Betti number one is an outstanding open problem. One reason that these problems seem to be hard is that in general there is no way known to find a π_1-injective immersion of a closed surface into the manifold. This gives a second reason that bundles are a useful class; it is possible to give constructions which produce many immersions of the type required by the hypothesis of Theorem 1.3. This is done in §2.

One immediate consequence of the truth of the above conjecture would be that every hyperbolic 3-manifold contains an immersion of a surface corresponding to a quasi-Fuchsian surface group; this follows because of the following theorem of Thurston [16]

Theorem 1.4 *If $(M, \partial M)$ is a compact oriented 3-manifold and if $H_2(M, \partial M)$ has rank at least 2 then M possesses at least one incompressible surface which is not the fiber of any fibration.*

In the context of bundles, this leaves only the case of a rank one bundle not covered. We are able to show:

Theorem 1.5 *Every closed hyperbolic surface bundle over the circle contains an immersion of a quasi-Fuchsian surface.*

We conclude with a section which discusses examples. The first, alluded to above, is an immersion shown to be quasi-Fuchsian by these methods, the second example exhibits an embedded quasi-Fuchsian surface which is transverse to the flow coming from a pseudo-Anosov map. We have also included an appendix sketching the proofs of some of the results of [2].

The third author thanks the University of Texas at Austin for its hospitality during the completion of this work.

2 Constructing immersed surfaces.

In this section we give a construction which provides a large number of incompressible immersions inside a bundle which are amenable to the techniques we develop in later sections. Throughout this paper we consider a closed surface F and a pseudo-Anosov homeomorphism θ of F.

Lemma 2.1 *Given any essential simple curve $C \subset F$, there is a finite cover $\tilde{F}$ of F, a map $\tilde{\theta}$ covering θ and a curve $\tilde{C}$ covering C, so that $\tilde{C}$ is disjoint from $\tilde{\theta}(\tilde{C})$. Further, $\tilde{C} \cup \tilde{\theta}(\tilde{C})$ is nonseparating on $\tilde{F}$.*

Proof. Fix an integer $p > 1$ and consider the epimorphism $\pi_1(F) \to H_1(F, \mathbb{Z}_p)$ which determines the finite covering $\pi : \tilde{F}_p \to F$. Observe that any simple closed curve element of the commutator subgroup lifts to a nonseparating curve in $\tilde{F}_p$. This is easily seen since there are only a finite number of such curves up to homeomorphism and for each such curve one can construct by inspection a $\mathbb{Z}_p$-covering for which the curve lifts and becomes nonseparating; now an argument using the transfer map shows that the curve is therefore nonseparating in the covering $\tilde{F}_p$.

Thus, if the given curve C separates, we may replace C by a non-separating curve covering it. Also notice that this covering corresponds to a characteristic subgroup therefore θ is covered by a homeomorphsim of $\tilde{F}_p$. We may now assume that C does not separate.

Given C suppose that $|C \cap \theta(C)| = K$. Fix a curve $\tilde{C}$ covering C (it covers C p-to-1) and a map $\tilde{\theta}$ covering θ. Then $|\tilde{C} \cap \tilde{\theta}(\tilde{C})| \leq pK$, in particular $\tilde{C}$ meets at most pK components of $\pi^{-1}(\theta C)$.

Choose a curve $D \subset F$ which is simple and meets $\theta(C)$ once transversally, and let T be a punctured torus which is a regular neighbourhood of $D \cup \theta(C)$. If $g = genus(F)$ then the index of $H_1(T, \mathbb{Z}_p)$ in $H_1(F, \mathbb{Z}_p)$ is p^{2g-2}. Therefore $\pi^{-1}T$ has p^{2g-2} components, let $\tilde{T}_1$ be one of them. Let $E(\tilde{T}_1)$ be a curve in $\tilde{T}_1$ which covers θC. Now $\tilde{T}_1$ contains a curve $\tilde{D}$ which has intersection number one with E, therefore if $\tilde{T}_2$ is a different component of $p^{-1}T$ then $E(\tilde{T}_1) \neq E(\tilde{T}_2)$ as elements of $H_1(\tilde{F})$. Therefore there are at least p^{2g-2} distinct homology classes of curves which cover $\theta(C)$.

It we choose p so that $p^{2g-2} > pK + 1$ then there is at least one component of $\pi^{-1}\theta(C)$ which is both disjoint from $\tilde{C}$ and not homologous to it. By adjusting our lift of θ we obtain the result. ∎

Lemma 2.2 *If $\tilde{\theta}_1$ and $\tilde{\theta}_2$ are two maps which cover θ on an n-fold covering $\tilde{F}$ of F, then $\tilde{\theta}_1^m = \tilde{\theta}_2^m$ for some $0 < m \leq n$.*

Proof There is a covering transformation τ of $\tilde{F}$ such that $\tilde{\theta}_2 = \tau \circ \tilde{\theta}_1$. Then

$$(\tau \circ \tilde{\theta}_1)^k = \tau_0 \tau_1 \tau_2 \cdots \tau_{k-1} \tilde{\theta}_1^k$$

where $\tau_i = \tilde{\theta}_1^i \tau \tilde{\theta}_1^{-i}$ is a covering transformation. Let G be the group of covering transformations of $\tilde{F}$ thus $|G| \leq n$. Therefore for some k and some $0 < m \leq n$ that

$$\tau_0 \tau_1 \tau_2 \cdots \tau_k = \tau_0 \tau_1 \tau_2 \cdots \tau_{k+m},$$

which implies that

$$1 = \tau_{k+1} \tau_1 \tau_2 \cdots \tau_{k+m},$$

and then conjugating by $\tilde{\theta}^{-(k+1)}$ gives

$$1 = \tau_0 \tau_1 \tau_2 \cdots \tau_{m-1}.$$

Hence

$$\tilde{\theta}_2^m = (\tau \circ \tilde{\theta}_1)^m = \tilde{\theta}_1^m.$$

as required. ∎

Corollary 2.3 *If $\tilde{F}$ is a finite regular cover of F and $\tilde{\theta}$ is map covering θ then the mapping tori M_θ of θ and $M_{\tilde{\theta}}$ of $\tilde{\theta}$ have finite covers which are homeomorphic.*

In what follows we shall 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.

We now introduce a standing definition: M is the mapping torus of θ. We regard M as the space obtained from $F \times [0,1]$ by identifying $(x,1)$ in $F \times 1$ with $(\theta x, 0)$ in $F \times 0$. Projection of

$F \times [0,1]$ onto the second factor induces a map of M to the circle which is a fibration with fiber F. The infinite cyclic covering of M given by this fibration is thus identified with $F \times \mathbb{R}$ and the group of covering transformations is generated by $\tau(x,t) = (\theta x, t+1)$ where $x \in F$ and $t \in \mathbb{R}$. The foliation of $F \times [0,1]$ by intervals has image in M a one-dimensional foliation which we denote by $\mathcal{L}$ called the **suspension flow** on M.

Lemma 2.4 *There is an immersion $g : S \looparrowright M$ which is transverse to $\mathcal{L}$.*

Proof. The inclusion $F \to M$ gives a map $\iota : F_- \to M$ to which it is equal on the interior of F_-. We now isotope ι along the flow $\mathcal{L}$ in annulus neighborhoods of C_+, C_- to obtain a map $g' : F_- \to M$ which is transverse to $\mathcal{L}$ and so that $g'(C_\pm) = \theta(C_\mp)$. A more precise description of this isotopy follows.

Let A_+ be an annulus in F_- with one boundary component C_+ and $h : C_+ \times [0,1] \longrightarrow A_+$ a parameterization so that $h(C_+ \times 1) = C_+$. Let $k : F \times [0,1] \longrightarrow M$ be the quotient map. Then g' is defined on A_+ by $g'(h(x,t)) = k(h(x,t),t)$. We define g' in a similar manner on an annulus neighborhood of C_-, and on the complement of these two annuli $g' = \iota$.

It is clear that g' and ι are isotopic in M by an isotopy having the property that the track of each point during the isotopy lies in a single flow line of $\mathcal{L}$. Then g' factors through a map $g : S \looparrowright M$ which is transverse to the suspension flow. ■

Remark. This is special case of a more general construction. Start with a finite cover $p : \tilde{F} \longrightarrow F$ and two sets of simple closed curves $C_1, \cdots, C_n$ and $D_1, \cdots, D_n$ on $\tilde{F}$ which are mutually disjoint and such that their union does not separate $\tilde{F}$. We also require that $p(D_i) = \theta p(C_i)$ for each $1 \leq i \leq n$. The composition of p with the inclusion of $F \longrightarrow M$ may be homotoped to an immersion $h : \tilde{F} \longrightarrow M$ transverse to $\mathcal{L}$. Let $\tilde{F}_-$ be the surface obtained from $\tilde{F}$ by cutting along the two collections of curves and compactifying as before to get a surface with two boundary curves $(C_i)_\pm$ corresponding to each C_i and two boundary curves $(D_i)_\pm$ corresponding to each D_i. Let $\iota : \tilde{F}_- \longrightarrow M$ be the immersion obtained from h, then as before we isotop ι along the flow in an annulus neighborhood of each $(C_i)_\pm$ to obtain a map g' such that $g'(C_i)_\pm = (D_i)_\mp$. Let S be the closed surface obtained from $\tilde{F}_-$ by identifying $(D_i)_\pm$ with $(C_i)_\mp$ via θ. As before there is an immersion $g : S \looparrowright M$ of S transverse to $\mathcal{L}$.

Lemma 2.5 *The immersion $g : S \looparrowright M$ constructed above is π_1-injective.*

Proof. There is a fibration $M \longrightarrow S^1$ with fiber F and this determines an infinite cyclic covering $\tilde{M} \longrightarrow M$. There is an induced infinite cyclic covering $\tilde{S}$ of S and a map $\tilde{g} : \tilde{S} \longrightarrow \tilde{M}$ covering g. Suppose that γ is an essential loop in S which is null homotopic in M. Then γ lifts to a loop $\tilde{\gamma}$ in $\tilde{S}$.

Using the identification of $\tilde{M}$ with $F \times \mathbb{R}$ we define the subsurface $S_n = \tilde{g}^{-1}(F \times [-n,n])$ of $\tilde{S}$. We may push $g(S)$ slightly in the direction of the flow $\mathcal{L}$ to arrange that $g(S) \cap F \times 0$ consists of two curves parallel to C. Then the boundary of S_n consists of two parallel curves on $F \times n$ and another two parallel curves on $F \times -n$. Now add to $\tilde{g}S_n$ an annulus in $F \times n$ and another in $F \times -n$ to give an immersion $\tilde{g}_n^+ : \tilde{S}_n^+ \looparrowright \tilde{M}$ of a closed surface $\tilde{S}_n^+$ which extends the immersion $\tilde{g}|\tilde{S}_n$. This immersion is also transverse to $\mathcal{L}$.

The composition of the projection $p_1 : F \times \mathbb{R} \longrightarrow F$ and $\tilde{g}_n^+$ is an immersion of a closed surface into another surface and is therefore a covering and so π_1-injective. We may choose n large enough that $\tilde{\gamma}$ is contained in S_n and this implies that $p_1\tilde{g}_n^+(\tilde{\gamma})$ is essential in F and therefore in M, a contradiction. ■

Theorem 2.6 *Every closed hyperbolic surface bundle over the circle contains an immersion of a quasi-Fuchsian surface.*

Proof. By 2.2,2.1,2.3 we may pass to a finite cover of the given bundle in which we may construct an immersion of a surface $g : S \looparrowright M$ as described in Lemma 2.4. From 1.1 we see that either this surface is quasi-Fuchsian (in which case we are done) or else there is a finite covering $p : \tilde{M} \to M$ to which the immersion g lifts to an embedding $\tilde{g}$, and this embedding is a fiber in some fibration of $\tilde{M}$ over the circle.

By Thurston's theorem 1.4 if the rank of $H_2(\tilde{M})$ is at least two there is an embedded surface V which is not a fiber of any fibration of $\tilde{M}$. A theorem of Heil [8] asserts that in a closed irreducible orientable 3-manifold, an embedded nonseparating surface with non-trivial normalizer is a fiber of a fibration. Therefore $\pi_1 V$ must have trivial normaliser and so cannot be geometrically infinite. It follows from Theorem 1.1 $\pi_1 V$ is quasi-Fuchsian.

It therefore suffices to show that the lift of the fiber F of M (which we denote $\tilde{F}$) and the embedding $\tilde{g}(S)$ represent linearly independent classes in $H_2(\tilde{M})$.

To prove this we note that there is a loop α on S such that $g\alpha$ is dual to the fiber F. Thus $\tilde{g}(\alpha)$ is a dual class to the fiber $\tilde{F}$. However since this loop lies in the 2-sided surface $\tilde{g}(S)$ it can be pushed off this surface, and thus cannot be a dual loop for the homology class of $\tilde{g}(S)$. Thus $\tilde{F}$ and $\tilde{g}(S)$ are linearly independent homology classes as required. ∎

Remark. We remark that in general one cannot expect to find a closed *embedded* quasi-Fuchsian surface in a hyperbolic bundle whose homology has rank one. As discussed in the introduction, whenever the homology has rank greater than one, there is always such a surface, in fact it can be chosen to be nonseparating.

For example, consider the 2-bridge knot $K = 6_3$ which has 2-bridge normal form 5/13 (see [7]). It is alternating, and its Alexander polynomial is monic of degree four, from which it follows that this knot is fibered with fiber a once-punctured surface of genus 2. It is easy to check that 0/1-surgery on $S^3 \setminus K$ produces a hyperbolic manifold which fibers over the circle, and necessarily has rank one homology. Let M denote this manifold. Suppose that M did contain a closed embedded quasi-Fuchsian surface (such a surface must separate since $H_2(M)$ is cyclic). Because $S^3 \setminus K$ contains no closed embedded incompressible surfaces, it follows that there is, in addition to the fiber, a separating surface in $S^3 \setminus K$ with slope 0/1. However, this is impossible from the results of [7] which describes all incompressible and boundary incompressible surfaces in 2-bridge knot complements.

3 Finite foliations.

In this section we develop the general picture of π_1-injective immersions which are transverse to the flow. Our point of view is to work in the universal cover of M and flow a pre-image of the immersed surface $\tilde{g}(\tilde{S})$ into a reference copy $\tilde{F}$ of the universal cover of F. We prove that this image is all of $\tilde{F}$ if and only if the immersed surface is geometrically infinite. We show that in the geometrically finite case that the image is a convex set bounded by lifted leaves of the invariant foliations on F. Finally, we examine the foliation on this convex set and show that the induced foliation on $g(S)$ is finite in a certain sense.

Now let $\pi : \tilde{M} \to M$ be the universal cover and $\tilde{\mathcal{L}}$ the lifted flow on $\tilde{M}$. Let $\tilde{F}$ be a component of $\pi^{-1}F$ so that $\tilde{F}$ is a plane which meets every flow line of $\tilde{\mathcal{L}}$ once transversely. Thus we may identify $\tilde{M}/\tilde{\mathcal{L}}$ with $\tilde{F}$ by flowing a point in $\tilde{M}$ until it meets $\tilde{F}$. Let $\phi : \tilde{M} \to \tilde{F}$ be the resulting map which identifies flow lines to points.

Let $g : S \looparrowright M$ be any π_1-injective immersion of a closed orientable surface S into M which is everywhere transverse to the flow $\mathcal{L}$. The results of the previous section show that such immersions are easily constructed. Let $\tilde{S}$ be the universal cover of S and $\tilde{g} : \tilde{S} \to \tilde{M}$ a map covering g.

The measured foliations for θ we denote by $(\mathcal{F}^+, \mu^+)$ and $(\mathcal{F}^-, \mu^-)$ respectively. Our convention for these foliations will be the mnemonic one (as opposed to the convention used by dynamicists)

- the action of θ on the measure of an arc α transverse to $\mathcal{F}^+$ is defined by $\mu^+(\theta\alpha) = \lambda\mu^+(\alpha)$, where $\lambda > 1$ is the dilatation. This pair of measured foliations determines a singular Euclidean structure on F, and makes F into a complete metric space. This metric space is locally isometric to the Euclidean plane except in a neighborhood of the singularities of the foliations. The metric at a k-prong singularity has a cone angle of $(k-2)\pi$. Since $k > 2$ the metric is non-positively curved. It follows that between any two points in F there is a unique shortest arc in any homotopy class. Such an arc is called a geodesic and is a Euclidean geodesic away from the singularities. A short geodesic arc in a neighbourhood of a singularity consists of two Euclidean arcs each having an endpoint on the singularity. The angle between these two arcs is at least π.

The stable and unstable foliations on F lift to $\tilde{F}$ and thus the singular Euclidean structure also lifts to a complete metric on $\tilde{F}$. The action of $\pi_1 M$ on $\tilde{M}$ preserves the flow $\tilde{\mathcal{L}}$ and therefore induces, via ϕ, an action on $\tilde{F}$. In this section we shall examine the action of the surface group $\pi_1(S)$ on the surface $\tilde{F}$. (Here, as in all that follows, we shall suppress reference to the map g_* in the context of fundamental groups.) As a first step we will examine the subset $\phi(\tilde{S})$ of $\tilde{F}$. This set is clearly invariant under $\pi_1(S)$.

Definition 3.1 *A* **leaf box** *is a compact convex subset of $\tilde{F}$ bounded by finitely many segments of stable and unstable leaves. If a leaf box contains no singularity then it has 4 sides. Otherwise it contains precisely one singularity, say k-prong, and has $2k$ sides. An ϵ* **-leaf box neighborhood** *of a point x in $\tilde{F}$ is a leaf box which contains an ϵ-neighborhood of x.*

We will use the notation $L(\pi_1(S), S^2_\infty(\tilde{M}))$ for the (usual) limit set for the action of the subgroup $\pi_1(S)$ on the 2-sphere at infinity.

Lemma 3.2 *Suppose $g : S \looparrowright M$ is incompressibly immersed transverse to the flow $\mathcal{L}$. Then each flow line of $\tilde{\mathcal{L}}$ meets $\tilde{g}\tilde{S}$ at most once.*

Proof. The proof falls into two cases:
Case A. The group $\pi_1(S)$ is quasi-Fuchsian.
Since $\pi_1(S)$ is quasi-Fuchsian the map $\tilde{g}$ extends to a continuous map $\tilde{g} : \tilde{S} \cup S^1_\infty(\tilde{S}) \longrightarrow \tilde{M} \cup S^2_\infty(\tilde{M})$. Since S is transverse to $\mathcal{L}$, an oriented flow line ℓ of $\tilde{\mathcal{L}}$ always crosses $\tilde{g}\tilde{S}$ in the same direction. Again using the fact that $\pi_1(S)$ is quasi-Fuchsian, there is a self homeomorphism ξ of the closed ball $\tilde{M} \cup S^2_\infty(\tilde{M})$ which takes the limit set of $\pi_1(S)$ to a round circle.

We will make use of the following from Cannon Thurston [2]. There is a natural identification of $S^2_\infty(\tilde{M})$ with the space obtained by taking two copies D^+, D^- of $\tilde{F} \cup S^1_\infty(\tilde{F})$ and making the following identifications. The two points, one in each of $D^\pm$, corresponding to a single point in S^1_∞ are identified. Each leaf (regular or singular) of $\tilde{\mathcal{F}}^+$ is identified to a point in D^+. Also each leaf (regular or singular) of $\tilde{\mathcal{F}}^-$ is identified to a point in D^-. Let $q^\pm : D^\pm \longrightarrow S^2_\infty(\tilde{M})$ be the quotient maps.

Cannon and Thurston also show that each flow line ℓ is a quasi-geodesic in the hyperbolic metric. Now $\ell = x \times \mathbb{R}$ for some point x in $\tilde{F}$. The limit points of ℓ on $S^2_\infty(\tilde{M})$ are $q^\pm x$. The map $h^+ : \tilde{F} \to S^2_\infty(\tilde{M})$ given by inclusion into D^+ followed by q^+ is continuous, and equals the map obtained by sending a point x in $\tilde{F}$ to the limit point of ℓ given by flowing along ℓ in the forward direction. h^- is similarly defined by flowing backwards. Observe that the image of $\tilde{F}$ under h^+ is $\pi_1(M)$ invariant. We appeal to the following lemma:

Lemma 3.3 *The image $h^+(\tilde{F})$ is an $\mathbb{R}$-tree.*

Proof. An $\mathbb{R}$-tree is a (nonempty) metric space with the property that (i) Any two points are joined by a unique arc and (ii) With the induced metric this arc is isometric to an interval.

If we take two distinct points, p_1 and p_2 in the quotient space $h^+(\tilde{F})$, then these correspond to a pair of leaves in $\tilde{F}$. To construct an arc in $h^+(\tilde{F})$ which connects them without backtracking, choose any embedded arc α in $\tilde{F}$ running between the relevant leaves. Notice that if two points on α are sufficiently close together, then they cut out a subarc α_i which can be isotoped to an arc α_i^* which lies inside a leaf of $\mathcal{F}^-$. This can be done without changing the leaves which the endpoints of α_i lie on. Notice that the arc α_i^* embeds into the quotient space $h^+(\tilde{F})$ since it is impossible for a leaf of $\mathcal{F}^+$ to meet α_i^* more than once as the leaves are geodesics in the affine metric on $\tilde{F}$.

By subdividing a random arc joining a pair of leaves in $\tilde{F}$, we see that any two points in $h^+(\tilde{F})$ can be connected by an arc. One sees easily that if the points are connected by a very short arc in $\tilde{F}$ then in fact after projection into $h^+(\tilde{F})$ the points are connected by a unique arc, whence a subdivision argument shows that any two points in $h^+(\tilde{F})$ are connected by a unique arc. Further, in an obvious way, we may use the measure of these arcs to produce a metric on $h^+(\tilde{F})$, which is therefore an $\mathbb{R}$ tree. ■

We claim that if B is any leaf box in $\tilde{F}$ then $h^+(B)$ is not contained in the quasi-circle $L(\pi_1(S), S^2_\infty(\tilde{M}))$. The reason is that $h^+(B)$ contains an interval I and if this interval is contained in $L(\pi_1(S), S^2_\infty(\tilde{M}))$ we can use the action of $\pi_1(S)$ to cover the quasi-circle by a finite number of intervals and deduce that the entire quasi-circle embeds inside the image of h^+. However, this is a contradiction since by the Lemma, $h^+(\tilde{F})$ is an $\mathbb{R}$-tree. This proves the claim.

Let B be any compact leaf box, since h^+ is continuous there is an open subset $B' \subset B$ with $h^+(B')$ disjoint from $L(\pi_1(S), S^2_\infty(\tilde{M}))$. By repeating this argument with h^- we see that B' contains a leafbox B'' so that $h^+(B''), h^-(B'')$ are both in the domain of discontinuity of $g_*\pi_1(S)$.

Suppose then that the flowline $x \times \mathbb{R}$ meets $\tilde{g}(\tilde{S})$ more than once, then by transversality the same is true for $y \times \mathbb{R}$ provided that $y \in \tilde{F}$ is sufficiently close to x. The above argument shows that we may choose y so that both endpoint of $\ell \equiv y \times \mathbb{R}$ are in the domain of discontinuity for $g_*\pi_1(S)$. Choose small closed disc neighbourhoods of the endpoints of ℓ which lie inside the domains of discontinuity and use these to attach a 1-handle to $\tilde{M} \cup S^2_\infty(\tilde{M})$to form a solid torus T.

The existence of ξ shows that the closure of $\tilde{g}(\tilde{S})$ either represents the generator of $H_2(T, \partial T)$, or represents the zero element, depending on whether the endpoints of ℓ lie in distinct or the same components of the domain of discontinuity. It is also clear that the cocore of the 1-handle also represents the generator of this group.

However the union of ℓ and the core of the 1-handle is a 1-cycle in T which has algebraic intersection number at least two with $\tilde{g}(\tilde{S})$, but intersection number one with the cocore. This contradiction completes the proof of Case A.

Case B. The subgroup $\pi_1(S)$ is geometrically infinite.
The proof is similar in spirit. As observed in the introduction, $g(S)$ becomes homotopic to a fiber S' in some finite covering M^* of M.

Suppose that if we lift to the covering M_S corresponding to the subgroup $\pi_1(S)$, the lift $\tilde{\gamma}$ of some flowline γ meets the image of the lifted map $g_S : S \longrightarrow M_S$ more than once. To prove the theorem, it suffices to show that this cannot happen. It is well known that the periodic points of a pseudo-Anosov are dense, thus closed flow lines are dense. So by transversality there is a nearby flow line to $\tilde{\gamma}$ which meets $g_S(S)$ more than once and with projection in M^* a closed flowline. Thus we may assume that γ is closed.

We claim that $\tilde{\gamma}$ defines a element of $H^2_c(M_S)$. This falls into two cases: If γ lies in the (normal) subgroup $\pi_1(S) \leq \pi_1(M^*)$, then it lifts to M_S and visibly defines such a cohomology class. If not, then $n = [\gamma] \neq 0$ in $\pi_1(M^*)/\pi_1(S) \cong \mathbb{Z}$ thus the flowline $\tilde{\gamma}$ is invariant under τ^n where τ is the covering transformation of the cyclic covering $M_S \longrightarrow M^*$. Hence $\tilde{\gamma}$ is a properly embedded 1-submanifold in M_S with one end in each end of M_S. Thus $\tilde{\gamma}$ defines such a cohomology class in this case also.

By construction $\tilde{\gamma}$ meets $g_S(S)$ always with the same orientation. Clearly $[g_S(S)] = [S']$ as classes in $H_2(M_S)$. However $\tilde{\gamma}$ defines an element of $H^2_c(M_S)$ and we see that $<\tilde{\gamma}, S'>$ is either zero or 1, since S' separates. On the other hand, by choice of $\tilde{\gamma}, <\tilde{\gamma}, g_S(S)>$ is at least two. This contradiction completes the proof of Case B. ∎

The above result implies that $\tilde{g}$ is an embedding of $\tilde{S}$ into $\tilde{M}$. In the interest of notational simplicity we will use $\tilde{g}$ to identify $\tilde{S}$ with $\tilde{g}\tilde{S}$ except where this is likely to cause confusion.

Corollary 3.4 *The flow map $\phi|\tilde{S}$ is a homeomorphism of $\tilde{S}$ onto its image.*

Corollary 3.5 *The action of $\pi_1(S)$ on $\phi\tilde{S}$ is free and properly discontinuous.*

Here, and elsewhere, the metric we use on $\tilde{F}$ is the singular Euclidean metric given by the measured foliations. We shall use several results concerning geometrically infinite groups due to Cannon and Thurston which are contained in [2]. For the convenience of the reader we include an appendix sketching the proofs of some of these results.

Recall that [2] introduces a certain singular Solv metric. This is a metric on M which is a Riemannian metric on the complement of the suspension of the singularities of the measured foliations on F. It is defined as follows. If $(\mathcal{F}^+, \mu^+), (\mathcal{F}^-, \mu^-)$ are the measured foliations on F and $\lambda > 1$ is the stretch factor of the pseudo-Anosov θ then $\mu^{\pm}(\theta\alpha) = \lambda^{\pm 1}\mu^{\pm}(\alpha)$. Here α is any arc transverse to the measured foliations.

A singular Riemannian metric on $F \times [0,1]$ is defined by

$$ds^2 = \lambda^{-2z}dx^2 + \lambda^{2z}dy^2 + dz^2.$$

Here x, y are local coordinates along the measured foliations $(\mathcal{F}^+, \mu^+), (\mathcal{F}^-, \mu^-)$ respectively, and z is the coordinate in $[0,1]$. The map $(x,1) \mapsto (\theta x, 0)$ of $F \times 1 \longrightarrow F \times 0$ is an isometry of this metric. Thus M inherits a singular Riemannian metric.

The universal cover of M is identified with $\tilde{F} \times \mathbb{R}$ and the metric on M induces a metric on $\tilde{M}$. Furthermore the suspension flow $\tilde{\mathcal{L}}$ agrees with the product structure on $\tilde{F} \times \mathbb{R}$ thus flow lines are geodesics on which x, y are constant. The local z coordinate in M is covered by a global z coordinate on $\tilde{F} \times \mathbb{R}$ given by projection onto the second coordinate. The components of $\pi^{-1}(F)$ are the **horizontal** planes $F \times n$ for integers n. We will use the term the **height** of a point w in $\tilde{M}$ to mean the z coordinate of w, which is therefore also the (signed) distance of w from $\tilde{F} \equiv \tilde{F} \times 0$.

We wish to choose a nice neighborhood of each point in S. Given a point $x \in S$, there is a lift $\tilde{x} \in \tilde{M}$ which lies within a distance 1 of $\tilde{F} \times 0$. Given a (simply connected) neighbourhood $U \subset S$ of x, choose the lift $\tilde{U}$ of U which contains $\tilde{x}$. We will say that U is an **η-leaf box neighbourhood** of x if the flowed image $\phi(\tilde{U}) \subset \tilde{F} \times 0$ is an η-leaf box neighbourhood of $\phi(\tilde{x})$.

We observe that the compactness of S implies that there is an $\epsilon(S) > 0$ so that every point of S has an $\epsilon(S)$-leafbox neighbourhood. The following result implies that points on $\tilde{S}$ with ϕ image in $\tilde{F}$ which are near to the frontier of $\phi\tilde{S}$ are a great height above or below $\tilde{F}$.

Corollary 3.6 *Given a point $\tilde{x}$ in $\tilde{S}$ suppose that $y_{\pm}$ is a point in Frontier$(\phi\tilde{S})$ on the same leaf of $\tilde{\mathcal{F}}^{\pm}$ as $\phi(\tilde{x})$. Define $d(\phi\tilde{x}, y_{\pm})$ to be the distance in $\tilde{F}$ between $\phi(\tilde{x})$ and $y_{\pm}$. Let $z(\tilde{x})$ be the height of $\tilde{x}$. Then $d(\phi\tilde{x}, y_{\pm}) > \epsilon(S)\lambda^{\mp z(\tilde{x})-1}$.*

Proof. Let $\epsilon(S) > 0$ be a leaf box constant provided by the above paragraph. Denote the universal covering projection of S by $\pi_S : \tilde{S} \to S$. Let U be an ϵ-leaf box neighbourhood of $\pi_S(\tilde{x})$ in S and let $\tilde{U}$ be the component of $\pi_S^{-1}(U)$ containing $\tilde{x}$. We may do a homotopy of g to arrange that U is contained in F, this homotopy is covered by a homotopy of $\tilde{g}$ which changes the height of all points by a uniformly bounded amount namely 1. This changes the estimate below by some multiplicative factor L with $\lambda^{-1} < L < \lambda$.

Using the singular Solv metric on $\tilde{M}$ above we see that $\phi\tilde{U}$ is a leaf box centered on $\phi\tilde{x}$ which has a half-width in the $\tilde{\mathcal{F}}^{\pm}$ direction of $\delta = \epsilon\lambda^{\mp z(\tilde{x})}$. Since this leaf box is contained in $\phi\tilde{S}$ it does not contain y in its interior and so $d(\phi\tilde{x}, y_{\pm}) > \delta$. ■

Definition 3.7 *We will call a closed subset P of $\tilde{F}$ a* **polygon** *if:*

- *P is homeomorphic to a closed disc minus part of its boundary*
- *The frontier of P is a union of (possibly infinitely many) bi-infinite geodesics.*

There is a unique geodesic in the singular flat structure between any pair of points in $\tilde{F}$, thus one may talk of convex sets in $\tilde{F}$ and we see that a polygon is convex. It follows as usual that any convex set is a disc minus part of its boundary.

Definition 3.8 *Define $A \equiv A(\tilde{S})$ to be the closure in $\tilde{F}$ of $\phi(\tilde{S})$.*

We further define a **regular leaf polygon** to be a polygon such that each side is a leaf of $\mathcal{F}^{\pm}$ which is **regular** on the inside, that is to say, if there is a singularity on some side ℓ then there are only two prongs of the singularity which lie in the polygon, and these are contained in ℓ.

Proposition 3.9 *$A(\tilde{S})$ is a regular leaf polygon.*

Proof. Since A is visibly path-connected, it suffices to show that the frontier of A consists of geodesics each of which is regular to the inside and completely contained in the frontier of A.

Let y be a point in the frontier of A and let $\tilde{x}$ be a point in $\tilde{S}$ such that $\phi\tilde{x}$ is a very small distance d from y. Let U and $\tilde{U}$ be the leaf box neighborhoods of $\pi_S\tilde{x}$ and $\tilde{x}$ used in Corollary 3.6. Then $\phi\tilde{U}$ has one dimension, say the stable leaf direction, of order d or smaller and therefore the other dimension, the unstable leaf direction, is of order d^{-1}. or larger. This is because vertical displacement in $\tilde{M}$ changes stable and unstable dimensions by reciprocal factors. Thus the segment of the unstable leaf through $\phi(\tilde{x})$ contained in A is very long. As $d \to 0$ we see that the closed set A must contain a bi-infinite unstable leaf ℓ through y. Furthermore, ℓ is the limit of regular leaves in A and is therefore regular on the inside of A.

We claim that ℓ is contained in the frontier of A. Suppose that there is a point z on ℓ with the property that the segment $\overline{y\,z}$ of ℓ is maximal with respect to being in the frontier of A. Let w be a point on ℓ very close to z and in the interior of A. Now $\overline{w\,z}$ is contained in the unstable leaf ℓ so by Corollary 3.6 there is a point $\tilde{w} \in \tilde{S}$ a large distance below w. Let $\tilde{V}$ be a leafbox neighborhood of $\tilde{w}$ then $\phi\tilde{V}$ intersects U and it follows that there is a point v near z such the flow line through v meets $\tilde{S}$ twice, once far above $\tilde{F}$ in $\tilde{U}$ and once far below $\tilde{F}$ in $\tilde{V}$. But this contradicts Lemma 3.2. ■

An element of $\pi_1 F$ gives an isometry of $\tilde{F}$. Since M is fiberd, the fundamental group $\pi_1(M)$ is an HNN extension of the fiber group and the stable letter acts on $\tilde{F}$ by some lift $\tilde{\theta}$ of θ. Thus we see that the action of $\alpha \in \pi_1(M)$ in local coordinates using the lifted measured foliations is multiplication by λ^{-n} in the stable direction and multiplication by λ^{n} in the unstable direction. Here n is the image of α under the homomorphism $\pi_1 M \to \mathbb{Z}$ given by the fibration of M with fiber F. The action of $\pi_1(F)$ extends to the circle at infinity $S^1_\infty(\tilde{F})$ for $\tilde{F}$, as does the map $\tilde{\theta}$ so that the action of $\pi_1(M)$ extends to an action on $\tilde{F} \cup S^1_\infty(\tilde{F})$.

Suppose that G is a subgroup of $\pi_1(M)$. We will temporarily use the notation $L(x, G)$ for the set of accumulation points in $\tilde{F} \cup S^1_\infty(\tilde{F})$ of the orbit of x under G. The subgroup $\pi_1(F)$ of $\pi_1(M)$ is precisely the subgroup which acts by isometries. If G is a subgroup of $\pi_1(F)$ then an easy argument shows that $L(x, G)$ is independent of the choice of a single point x.

However if G is not contained in $\pi_1(F)$ then this need not be the case. For example if G is the cyclic group generated by $\tilde{\theta}$ and if this map has a fixed point p in $\tilde{F}$ then $L(p, G) = p$. Let ℓ^{+} and

ℓ^- be the stable and unstable leaves containing p. Then if $x \neq p$ is a point on one of these leaves then $L(x,G)$ consists of p plus some or all of the limit points of that leaf on $S^1_\infty(\tilde{F})$. Finally if x does not lie on either $\ell^\pm$ then $L(x,G)$ contains at least one limit point on $S^1_\infty(\tilde{F})$ of both ℓ^+ and ℓ^-.

We will use the following two notions of limit set which are independent of the choice of point x.

Definition 3.10 *The* **isometric limit set** $IL(G; S^1_\infty(\tilde{F}))$ *of the subgroup G of $\pi_1(M)$ is $L(x, G \cap \pi_1(F))$. The* **finite limit set** *of G is $L(G;\tilde{F})$ is the set of points in $\tilde{F}$ fixed by some non-trivial element of G.*

Since $\pi_1(S) \cap \pi_1(F) \triangleleft \pi_1(S)$ is a nonelementary group acting on $\tilde{F}$ as a discrete group of isometries the isometric limit set of $\pi_1(S)$ is a subset of the circle at infinity of $\tilde{F}$. It is of interest to us because:

Lemma 3.11 *$IL(\pi_1(S); S^1_\infty(\tilde{F}))$ is $\pi_1(S)$-invariant.*

Proof. Standard arguments show that

$$IL(\pi_1(S); S^1_\infty(\tilde{F})) = closure(\{Fix(\gamma) \mid \gamma \in \pi_1(S) \cap \pi_1(F)\})$$

If $g \in \pi_1(S)$, then $g(Fix(\gamma)) = Fix(g.\gamma.g^{-1})$ and the result follows. ■

We will use the notation $L(A, S^1_\infty(\tilde{F}))$ for the limit points of A in $S^1_\infty(\tilde{F})$.

Theorem 3.12 *$A(\tilde{S})$ is the convex hull of $IL(\pi_1(S); S^1_\infty(\tilde{F}))$.*

Proof. We will write C for the convex hull of $IL(\pi_1(S); S^1_\infty(\tilde{F}))$ and A for $A(\tilde{S})$. Since A is invariant under $\pi_1(S)$ it follows that $L(A, S^1_\infty(\tilde{F})) \supseteq IL(\pi_1(S); S^1_\infty(\tilde{F}))$. Since A is a polygon it is convex and therefore contains C.

From Corollary 3.5 $\pi_1(S)$ acts freely, properly discontinuously on $\phi(\tilde{S})$. Since C is the convex hull of a subset of $S^1_\infty(\tilde{F})$ it is a polygon. Suppose that C is a proper subset of A then there is some point x on some geodesic side ℓ of C contained in the interior of A. It follows that ℓ lies in the interior of A otherwise, since A is a regular leaf polygon, ℓ would exit A.

Thus the orbit of ℓ under $\pi_1(S)$ is in the interior of A. Define C^+ to be $int(C)$ together with the orbit of ℓ, this is convex (hence simply connected) and invariant for the action of $\pi_1(S)$. Moreover, since C^+ is contained in the interior of A, $\pi_1(S)$ acts freely and properly discontinuously on C^+. Thus $C^+/\pi_1(S)$ is a surface with boundary whose fundamental group $\pi_1(S)$ is that of a closed surface, an absurdity. ■

Corollary 3.13

$$L(A(\tilde{S}); S^1_\infty(\tilde{F})) = IL(\pi_1(S); S^1_\infty(\tilde{F})).$$

With these results in hand, we may characterise the geometrically infinite case:

Theorem 3.14 *The following are equivalent:*

1. *The surface group $\pi_1(S)$ is geometrically infinite.*
2. $IL(\pi_1(S); S^1_\infty(\tilde{F})) = S^1_\infty(\tilde{F})$.
3. $A(\tilde{S}) = \tilde{F}$.

Proof. Suppose that $\pi_1(S)$ is geometrically infinite. Then by Theorem 1.1, there is a finite covering of M such that the surface S is homotopic to an embedding and is the fiber of a fibration. This covering corresponds to a subgroup G of finite index in $\pi_1(M)$ therefore

$$IL(G; S^1_\infty(\tilde{F})) = IL(\pi_1(M); S^1_\infty(\tilde{F})) = S^1_\infty(\tilde{F}).$$

Now $\pi_1(S)$ is a nonelementary normal subgroup of G so that

$$IL(\pi_1(S); S^1_\infty(\tilde{F})) = IL(G; S^1_\infty(\tilde{F}))$$

as required. Thus (1) implies (2).

Now suppose that $\pi_1(S)$ is not geometrically infinite so that $\pi_1(S)$ is quasi-Fuchsian. Suppose that $A = \tilde{F}$. The set $\phi(\tilde{S})$ is open and has frontier a union of leaves, thus $\phi(\tilde{S}) = \tilde{F}$, it therefore follows that $\tilde{S}$ meets every flowline.

We now argue in hyperbolic space. Each component of $\pi^{-1}gS$ in the universal cover has limit set a quasi-circle. Fix some component $\tilde{S}$. Fix some closed flowline C in M which defines some free homotopy class containing a unique hyperbolic geodesic γ_C in the hyperbolic metric on M. Notice that the immersed annulus coming from the free homotopy shows that any pre-image $\tilde{\gamma}_C$ of γ_C in $\tilde{M}$ lies within a bounded distance of some pre-image of C.

However $\pi_1(M)$ has dense limit set in the 2-sphere at infinity, so we may choose some element α of $\pi_1(M)$ with a fixed point off the quasi-circle defined by $\tilde{S}$. Translating $\tilde{\gamma}_C$ by a large power of α gives a pre-image of γ_C which is a large distance from $\tilde{S}$. ¿From this it follows that some pre-image of C misses $\tilde{S}$. This contradicts our assumption that $A = \tilde{F}$. Thus (3) implies (1).

The implication (2) implies (3) follows directly from Theorem 3.12 ■

Since the surface S is transverse to $\mathcal{L}$, the suspended foliations $\Omega\mathcal{F}^\pm$ induce foliations on S which we shall denote by $\mathcal{F}_S$, usually suppressing $\pm$.

Lemma 3.15 *Let ℓ be a leaf of either $\tilde{\mathcal{F}}^+$ or $\tilde{\mathcal{F}}^-$. The subgroup of $\pi_1(M)$ which stabilizes ℓ is either trivial or infinite cyclic.*

Proof. Let H be the stabilizer of ℓ, then H contains no non-trivial element of $\pi_1(F)$ otherwise the image of ℓ in F would be a compact leaf contradicting the fact that θ is pseudo-Anosov. Thus the projection of H into $\pi_1(M)/\pi_1(F) \cong \mathbb{Z}$ is injective. ■

Theorem 3.16 *If the surface group $\pi_1(S)$ is geometrically infinite, then $\mathcal{F}_S$ contains no closed leaf.*

Proof. By Theorem 3.14 the polygon $A = \tilde{F} = \phi\tilde{S}$, so by Corollary 3.5 $\pi_1(S)$ acts freely on $\tilde{F}$. However, if $\mathcal{F}_S$ contains a closed leaf ℓ, the covering translation τ corresponding to $[\ell] \in \pi_1(S)$ stabilises some pre-image $\tilde{\ell}$ of ℓ in $\tilde{F}$. Thus τ must be some lift of some power of the pseudo-Anosov by the argument in the proof of Lemma 3.15. But then τ or τ^{-1} is a contraction map of $\tilde{\ell}$, and so has a fixed point on it, contradicting that $\pi_1(S)$ acts freely. ■

Our next goal is to show that in fact the converse to this theorem holds and to gain more information about the foliations in the case that the surface group is quasi-Fuchsian. This requires some preliminary work. As our interest is now the case that $\pi_1(S)$ is quasi-Fuchsian, throughout what follows we shall assume that this is the case, in particular, that the polygon A is a proper subset of $\tilde{F}$ by 3.14. As already observed, the action of $\pi_1(S)$ on $\tilde{S}$ gives rise to a free, properly discontinuous action on $\phi(\tilde{S}) = int(A)$. Our next task is to examine the extension of this action to the closure of $\phi(\tilde{S})$, that is to say, A.

Lemma 3.17 *Let $\{I_n\}_{n\in\mathbb{Z}}$ be a collection of disjoint closed intervals in S^1 then the space obtained by identifying each interval to a point is homeomorphic to a circle.*

Proof. This is a well known fact; we sketch a proof. It suffices to replace S^1 by the unit interval. We do this by cutting open S^1 at a point not contained in any of the given intervals. We will define a continuous function f from the unit interval to itself. Define $f(0) = 0$ and $f(1) = 1$. Then define f on $I_n = [a_n, b_n]$ inductively to be $[f(L) + f(R)]/2$ where L is the closest point to the left of I_n at which f has already been defined and R is the corresponding point on the right. Set $X = \bigcup I_n$, then f is defined and monotonic increasing on X, we extend f over $cl(X)$ by $f(y) = sup\{f(x) \mid x \in X \;\; x \leq y\}$, clearly this extension is continuous and monotonic. The complement of $cl(X)$ is open, so we may extend f continuously by a linear map over each open interval in the complement. It is easy to check that f is continuous, monotonic, and a point pre-image is either one point or some I_n. Therefore f induces a homeomorphism to the unit interval from the unit interval with each I_n identified to a point. ■

According to Cannon and Thurston [2] there is a continuous map

$$CT : \tilde{F} \cup S^1_\infty(\tilde{F}) \longrightarrow \tilde{M} \cup S^2_\infty(\tilde{M})$$

extending the inclusion of $\tilde{F}$ into $\tilde{M}$. The map CT identifies two points on S^1_∞ if and only if they are limit points of the same leaf of the stable or unstable foliation. We will call a point x of $S^1_\infty(\tilde{F})$ **injective** if $CT^{-1}(CTx) = x$. Thus x is injective if and only if it is not the endpoint of any leaf in the stable or unstable foliation.

Let ∂A be the the frontier of A in $\tilde{F} \cup S^1_\infty(\tilde{F})$. Since A is convex, ∂A is a circle. In the case that $\pi_1(S)$ is quasi-Fuchsian, $L(A; S^1_\infty(\tilde{F}))$ is nowhere dense in $S^1_\infty(\tilde{F})$; for example this follows from 3.13 and 3.14. Thus the sides of A in $\tilde{F}$ are dense in ∂A. Let $\partial A/\sim$ be ∂A with each closed side of A identified to a point. A side ℓ of A is contained in leaf of $\tilde{\mathcal{F}}^+$ or $\tilde{\mathcal{F}}^-$ and therefore the two endpoints of ℓ on $S^1_\infty(\tilde{F})$ are identified to a single point in $S^2_\infty(\tilde{M})$.

Lemma 3.18 *Suppose that $\pi_1(S)$ is quasi-Fuchsian then $CT(S^1_\infty \cap \partial A) = L(\pi_1(S), S^2_\infty(\tilde{M}))$ and CT induces a homeomorphism $h : \partial A/\sim \longrightarrow L(\pi_1(S), S^2_\infty(\tilde{M}))$.*

Proof. By Lemma 3.17, we have that $\partial A/\sim$ is a circle. Suppose that h identifies two distinct points $x, y \in \partial A/\sim$. Then x, y separate $\partial A/\sim$ into two arcs I_1, I_2 and we show below that there are two injective points $z_1 \in I_1$ and $z_2 \in I_2$. Thus $CT(\partial A - \{z_1, z_2\})$ is a circle with two distinct points removed hence is not connected. But $(\partial A - \{z_1, z_2\})/\sim$ is the union of two intervals, and the map CT identifies a point x in one interval with a point y in the other interval, thus $CT(\partial A - \{z_1, z_2\})$ is connected yielding a contradiction.

First we establish the existence of one injective point on ∂A. Let α be an essential closed curve representing some element of $g_*\pi_1(S) \cap \pi_1(F)$. Then a line ℓ in $\tilde{F}$ which covers α has two limit points on $S^1_\infty(\tilde{F})$ and it is clear that these points are limit points of A and therefore lie in ∂A. Now it is well known that a leaf of the stable or unstable foliations on $\tilde{F}$ cannot limit on the endpoints of any closed curve; for example this follows from the proof of [3] Lemma 4.5. Let x be one of the limit points of ℓ then x is injective.

We claim that the injective points have dense image in $\partial A/\sim$, and this will complete the proof. Now $\pi_1(S)$ acts on ∂A and it is clear that this action preserves the property of a point being injective. The closure of the orbit of x under $\pi_1(S) \cap \pi_1(F)$ is $IL(\pi_1(S); S^1_\infty(\tilde{F}))$ which by Corollary 3.13, is $L(A, S^1_\infty(\tilde{F}))$. Thus injective points are dense in the complement in $\partial A/\sim$ of the countable set which is the image of the sides of A. Hence injective points are dense in $\partial A/\sim$. ■

Lemma 3.19 *Every regular leaf and every singular leaf in the interior of $A(\tilde{S})$ has at most one limit point on $S^1_\infty(\tilde{F})$, the other limit points are on sides of $A(\tilde{S})$.*

Proof. Let ℓ be a leaf of $\tilde{\mathcal{F}}^\pm$ and suppose that $\ell \cap int(A)$ has two distinct limit points x, y on S^1_∞. A side ℓ' of A is a leaf which is regular on the A side. Suppose that x and y are limit points of ℓ'

also. Then the bi-infinite leaves ℓ and ℓ' in $\tilde{F}$ are parallel. Then the bigon B they bound in $\tilde{F}$ is foliated as a product. The image of B in F is an annulus foliated by circles, but there are no closed leaves in the invariant foliations of a pseudo-Ansosov. Therefore x and y have distinct images in $\partial A/\sim$ and so by the previous lemma $CT(x) \neq CT(y)$. But this implies x and y are not limit points of the same leaf ℓ, a contradiction. ■

Recall that a group action is said to be **wandering** if every point p has a neighbourhood U so that $gU \cap U \neq \phi$ for only finitely many g.

Lemma 3.20 *Each side of A contains at most one point fixed by some non-trivial element of $\pi_1(S)$. The action of $\pi_1(S)$ on $A - L(\pi_1(S), \tilde{F})$ is free and wandering.*

Proof. Suppose that $\alpha \in \pi_1(S)$ has a fixed point x in some leaf ℓ in the boundary of A. Then α preserves A and therefore stabilizes ℓ. It follows that α must be some lift of some power of the pseudo-Anosov and that x is the lift of some periodic point. However a leaf can contain at most one periodic point and this gives the first statement.

We have already seen in Theorem 3.12 that the action of $\pi_1(S)$ is free and properly discontinuous on $int(A)$, so acts freely on $A - L(\pi_1(S), \tilde{F})$. Morevoer to show that the action is wandering we need only consider points x in $\partial A - L(\pi_1(S), \tilde{F})$.

Since A is a regular polygon we may choose a neighbourhood V of x in A to be a quadrilateral, one side of which is an arc of ℓ and the other 3 sides are arcs in $\tilde{\mathcal{F}}^-$ and $\tilde{\mathcal{F}}^+$. For the sake of definiteness we shall suppose that ℓ is contained in a leaf in $\tilde{\mathcal{F}}^+$ thus the two arcs β_1, β_2 of ∂V adjacent to ℓ are in $\tilde{\mathcal{F}}^-$ and the remaining arc of ∂V is in $\tilde{\mathcal{F}}^+$. We may choose V so small that there is no singularity in the interior of V and using Lemma 3.15 we may arrange that the orbit of V under the stabilizer of ℓ is disjoint from V. Refer to Figure 1.

Insert Figure 1.

Finally we choose V so that $length(\beta_i)$ is very small compared to the distance of x from every side of A other than ℓ.

Suppose now that $\tau(V)$ meets V for some $\tau \in \pi_1(S)$. By choice of V it follows that τ does not stabilize ℓ. Thus $\tau(V)$ meets a side $\tau(\ell)$ of V other than ℓ. It follows that $\tau(\beta_i)$ is much longer than β_i. If we now lengthen β_i slightly the images $\tau(\beta_i)$ move towards ℓ very rapidly and so will eventually leave A; a contradiction. ■

Lemma 3.21 *Let ℓ be any side of A then $stab(\ell) \cong \mathbb{Z}$.*

Proof. By Lemma 3.15, the stabilizer of a leaf is $\mathbb{Z}$ or trivial, so suppose that ℓ is a side of A and that $stab(\ell)$ is trivial.

Without loss, assume that ℓ is an unstable leaf and let A' be the union of $interior(A)$ and the orbit of ℓ under $\pi_1(S)$. Set $Y = A'/\pi_1(S)$. It follows from Lemma 3.20 that the map $A' \to Y$ is covering map and so Y is a nonclosed surface which however may fail to be Hausdorff. We shall show that we may combine Lemma 3.19 together with the assumption that the stabiliser of ℓ is trivial to deduce that Y is Hausdorff. But this implies that $\pi_1(S) \cong \pi_1(Y)$ is a free group, contradicting the fact that S is a closed surface, not S^2.

To show that Y is Hausdorff, it suffices to show that given two distinct points x, y in A' there are neighborhoods U of x and V of y, such that the $\pi_1(S)$-orbit of U is disjoint from V. By Corollary 3.5 the action of $\pi_1(S)$ on $int(A)$ is free and properly discontinuous. Therefore it suffices to consider the case that x lies in some side of A which we may assume to be ℓ. Let ℓ_x, ℓ_y be the intersection of the stable leaves through x, y with A'. By Lemma 3.19 ℓ_y has at most one limit point z on S^1_∞.

There are a finite number of elements in $\pi_1(S)$ which move any endpoint of ℓ_x onto a specified side of A' because $stab(\ell)$ is trivial. Therefore we may focus attention on those elements α of $\pi_1(S)$

with the property that $\alpha(\ell_x)$ does not have any endpoint on a side which also contains an endpoint of ℓ_y.

Choose V so every stable leaf through V meets a side of A containing an endpoint of ℓ_y. This is possible because ℓ_y has at most one limit point which is not an endpoint on some side of A. We may choose V so that there is a positive lower bound δ to the distance between every point in V and every side of A not containing y. We may choose U so that the diameter of U is small compared to δ. Suppose that $\alpha \in \Sigma$ and αU meets V, let ℓ' be the side of A' which contains αx then the distance of αx from V is much larger then the diameter of U therefore α expands the stable sides of U. Refer to Figure 2.

Insert Figure 2.

Then αx is very close to z and αU contains a segment $[x, w]$ of a stable leaf running from αx to a point $w \in V$. By enlarging U slightly, we see that w hits a side of A' containing an endpoint of ℓ_y. But this is absurd. ■

We may now prove our promised converse to Theorem 3.16. We shall define a foliation to be **finite** if it consists of a finite number of closed leaves and every other leaf spirals towards one of these closed leaves at every end. The reason for the terminology is that a foliation is finite if and only if the lamination obtained by straightening it has only a finite number of leaves. Finite foliations also arose in Fenley [5] in the context of depth 1 foliations of 3-manifolds. Nonetheless, the connection between Theorem 1.3 and Fenley's remains unclear. In particular, even if the surface in Theorem 1.3 is embedded and is a compact leaf of a depth 1 foliation, it is not proven that the two theorems give the same finite foliation, though presumably this is the case. At any rate, the methods are somewhat different.

Theorem 3.22 *Let $\mathcal{F}_S^+, \mathcal{F}_S^-$ be the foliations on S obtained by intersecting with S the suspension of the foliations $\mathcal{F}^+, \mathcal{F}^-$ on F. If $\pi_1(S)$ is quasi-Fuchsian then these are finite foliations.*

Proof. We recall that it is a consequence of our work thus far that the fact that the surface group is quasi-Fuchsian shows that the polygon A is not the whole disc and that each component of the frontier is a leaf of either the stable or unstable foliation, regular to the inside and has cyclic stabiliser.

We first observe that there is a closed leaf in $\mathcal{F}_S^+$ corresponding to each side of A which is (contained in) an unstable leaf of $\tilde{F}$. This is because if ℓ_u is an unstable side of A then by lemma 3.21 there is a non-trivial element α of $\pi_1(S)$ which stabilizes ℓ_u. Now α must have a fixed point x on ℓ_u, let ℓ_s be the intersection with A of the stable leaf through x. Then ℓ_s must be regular in A and have a limit point on $S^1_\infty(\tilde{F})$ because ℓ_s is stabilized by α. The pre-image under the flow map ϕ of ℓ_s in $\tilde{S}$ projects to a closed leaf in $\mathcal{F}_S^+$.

Let ℓ_s' be the intersection with A of any stable leaf in $\tilde{F}$, and suppose that ℓ_s' contains no singularity. By Lemma 3.19, ℓ_s' has at least one limit point on an unstable side ℓ_u of A. Let ℓ_s be the leaf constructed above for ℓ_u. Then the leaf in S corresponding to ℓ_s' is asymptotic in one direction to the closed leaf corresponding to ℓ_s.

We have just shown that the foliation $\mathcal{F}_S^+$ has the property that every regular leaf is asymptotic in at least one direction to a closed leaf. Therefore the corresponding lamination has this property for every nonclosed leaf.

Now it follows from [3] Lemma 4.6 that a leaf of a geodesic lamination which accumulates on a closed leaf is isolated, and we deduce that every nonclosed leaf of this lamination is isolated. Every isolated leaf is the side of some principal region of the complement of the lamination. Therefore if $g = genus(S)$ then there are at most $12g - 12$ isolated leaves and at most $3g - 3$ closed leaves. Moreover, an isolated leaf must accumulate in either direction on some non-isolated leaf which must therefore be a closed leaf. It follows that in the lamination either end of every nonclosed leaf accumulates to a closed leaf, so that in the universal cover, a lift of any nonclosed leaf must share both endpoints with the lift of a closed leaf. Since endpoints are not changed in the passage from foliation to lamination, we deduce that the foliation is finite. ■

4 Two Examples.

Example 1.

We construct an example of an immersion into a hyperbolic bundle of the type that we have been discussing and use our results to show that the immersion is quasi-Fuchsian. In order to make this as simple as possible, we construct a pseudo-Anosov map of a genus two surface which moves a curve disjoint from its image. The method is that of [9], but for the convenience of the reader, we develop it from first principles.

Let $\theta' : T' \longrightarrow T'$ be an Anosov diffeomorphism of the torus $T' \equiv \mathbb{R}^2/\mathbb{Z} \times 2\mathbb{Z}$. Let γ be an arc embedded in T' with endpoints x_0 and x_1 which are fixed points of θ' and let $p : F \longrightarrow T'$ be the 2-fold branched cover of T' branched over x_0, x_1 using the branch cut γ. Then θ' is covered by a homeomorphism $\overline{\theta} : F \longrightarrow F$, and $\overline{\theta}$ is pseudo-Anosov. To see this, let $\mathcal{F}^+(T'), \mathcal{F}^-(T')$ be the invariant measured foliations on T', then the pull-back of these via p to F gives two measured foliations $\mathcal{F}^+(F), \mathcal{F}^-(F)$ on F. These foliations each have two 4-prong singularities, one lying above each of x_0, x_1. The transverse measures on F are also obtained from those on T' by pull-back, and it is clear that in local coordinates the action of $\overline{\theta}$ is isometrically conjugate to the action of θ', where the metrics involved are those determined by the measured foliations on F and T'. Therefore the eigenvalues of θ' and $\overline{\theta}$ are the same. In particular this implies that no power of $\overline{\theta}$ can stabilise any finite curve system and $\overline{\theta}$ is therefore pseudo-Anosov.

Now suppose that there is a free orientation preserving involution τ on T' which commutes with θ' and swaps x_0 with x_1. Then $T = T'/\tau$ is also a torus and θ' covers a map θ on T. Furthermore F is a 4-fold irregular branched cover of T branched over the single point $x_0 \equiv x_1$ in T.

We will now apply this construction in a particular case. We will regard $T = \mathbb{R}^2/\mathbb{Z}^2$ and let θ be the linear homeomorhism on T induced by the linear map on $\mathbb{R}^2$ with matrix

$$\begin{pmatrix} 1 & 2 \\ 2 & 5 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$$

Define two simple closed curves m, ℓ on T to be the images of the x-axis and y-axis respectively. Let $\tau(m), \tau(\ell)$ be the Dehn twist around these curves, then

$$\tau(\ell) = \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix} \qquad \tau(m) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

thus θ is isotopic to the product of Dehn twists $\tau(\ell)^{-2}\tau(m)^2$.

Observe that a Dehn twist $\tau(\alpha)^{np}$ about the core curve α of an annulus A lifts to an n-fold covering $\tilde{A}$ of A to give a Dehn twist $\tau(\tilde{\alpha})^p$ about the core curve $\tilde{\alpha}$ of $\tilde{A}$. We can use this observation to calculate the lift of any mapping class of a compact surface to a branched or unbranched cover of that surface by expressing the mapping class as a product of Dehn twists.

Now we form the 2-fold cover $p_1 : T' \longrightarrow T$ by unwrapping in the ℓ direction, more formally this is the cover corresponding to the subgroup $C =< \ell^2, m >$ of $\pi_1(T)$. Let $\tilde{\ell} = p_1^{-1}\ell$ and let $\tilde{m}_1, \tilde{m}_2$ be the 2 components of $p_1^{-1}m$. Then $\tau(\ell)^{-2}$ is covered by $\tau(\tilde{\ell})^{-1}$ and $\tau(m)^2$ is covered by $\tau(\tilde{m}_1)^2\tau(\tilde{m}_2)^2$. There are 2 possible choices for each of these covering maps, we will choose all maps to fix a specified pre-image of the origin. Then $\theta' = \tau(\tilde{\ell})^{-1}\tau(\tilde{m}_1)^2\tau(\tilde{m}_2)^2$.

Finally we form the branched cover $p_2 : F \longrightarrow T'$. A fundamental domain for T' is the rectangle $R = [0,1] \times [0,2]$ in $\mathbb{R}^2$, and we choose γ to be the vertical arc in this rectangle $\{(0,t) : 0 \leq t \leq 1\}$

thus $x_0 = (0,0)$ and $x_1 = (0,1)$. We label $\tilde{m}_1, \tilde{m}_2$ so that $\tilde{m}_1$ is $\frac{3}{2} \times [0,1]$ of R and then $\tilde{m}_1$ is $\frac{1}{2} \times [0,1]$. Thus γ is disjoint from $\tilde{m}_1$ and meets $\tilde{m}_2$ once.

Thus $p_2^{-1}\tilde{m}_1$ has two components which we label M_1, M_1' and $p_2^{-1}\tilde{m}_2$ has one component which we label M_2. Also $p_2^{-1}\tilde{\ell}$ has two components L, L'. It now follows that

$$\overline{\theta} = \tau(L)^{-1}\tau(L')^{-1}\tau(M_1)^2\tau(M_1')^2\tau(M_2).$$

The curves M_1, M_1', M_2, L, L' on F are shown in Figure 3.

Insert Figure 3.

Let $\psi : F \longrightarrow F$ be the branched covering involution associated to the 2-fold branched cover $p_2 : F \longrightarrow T'$. Since ψ is an isometry of the metric on F determined by the invariant measures for $\overline{\theta}$ it follows that $\Theta = \psi \circ \overline{\theta}$ is also pseudo-Anosov.

Lemma 4.1 *The curve M_1 is disjoint from $\Theta(M_1) \simeq L' \cdot M_1'$.*

Proof. We will use the above factorisation of $\overline{\theta}$ as a product of Dehn twists. Since M_1, M_1', M_2 are disjoint, Dehn twist about these curves do not move M_1. Furthermore L and L' are disjoint and therefore Dehn twist about these curves commute. Since M_1 is disjoint from L' we see that

$$\overline{\theta}C = \tau(L)^{-1}M_1 = L \cdot M_1.$$

Referring to Figure 3, it is easily seen that $\psi(L \cdot M_1) = L' \cdot M_1'$ is disjoint from M_1. Thus ΘM_1 is disjoint from M_1 as claimed. We remark that M_1 may be replaced by M_1' in this argument. ∎

We now wish to calculate the invariant foliations on F. This is done by exhibiting a fundamental domain $\mathcal{D}$ for the universal cover $\tilde{F}$ of F. Now F is a 4-fold branched cover of T, and we can choose $\mathcal{D}$ so that it is mapped injectively onto its image in the universal cover $\mathbb{R}^2$ of T. A choice for this image of $\mathcal{D}$ is shown in Figure 4A.

Insert Figures 4A & 4B.

It consists of 4 fundamental domains for T, each of which is a unit square with corners in $\mathbb{Z}^2$. The branch points x_0, x_1 are lattice points and therefore the singularities of the foliations on F correspond to some of the corners of $\mathcal{D}$.

We now use the construction in Lemma 2.4 to define a surface $S(F, C, \Theta)$. Let $C = M_1$ this curve in F has image $\{(0,t) : \ 0 \le t \le 1\ \}$ and ΘC has image $\{(1+t, 2t-2) : \ 0 \le t \le 1\ \}$. Take a copy of F and cut F along C and ΘC to obtain F_-. Then we cross-join $C_\pm$ and $\Theta C_\mp$ to get $S(F, C, \Theta)$ and an immersion $g : S \longrightarrow M$ where M is the mapping torus for Θ.

Notice that the foliation on S is obtained by taking the foliation on F, restricting to F_- and then glueing using Θ to get S. The identification of $C_\pm$ with $\Theta C_\mp$ corresponds in $\mathcal{D}$ to identifying $(0,t)$ on C to $(1+t, 2t-2)$.

Theorem 4.2 *The induced foliations on the surface $S(F, C, \Theta)$ have closed leaves, therefore S is quasi-Fuchsian.*

Proof. The eigenvalues of θ are $\lambda_\pm = (3 \pm 2\sqrt{2})$ and the corresponding eigenvectors are $v_\pm = (1, (1 \pm \sqrt{2}))$. The foliation with eigenvector v_+ is shown in Figure 4B. Inspection of this figure reveals that every flow line starting on $\Theta C_\pm$ in $\mathcal{D}$ ends in some proper subset of $C_\mp$. In particular, this implies that there is a closed leaf on S which meets C_+ exactly once. It follows from Theorem 3.22 that immersion S is quasi-Fuchsian.

This theorem also implies that the other foliation also has a closed leaf; finding this closed leaf is left as an (easy) exercise for the reader. ■

Example 2.

One question raised by our results is the question of when an *embedded* quasi-Fuchsian surface can be isotoped so as to be transverse to the flow $\mathcal{L}$. This seems to be unresolved in general, although some partial results are presented by Mosher in [11]. In particular his *Weak transverse surface theorem* gives a sufficient condition for when a 2 dimensional homology class contains such a representative. Here we give a simple construction for embedded quasi-Fuchsian surfaces which are tranverse to the flow associated to the suspension of a pseudo-Anosov map. An example of a quasi-Fuchsian surface which is transverse to a flow on a compact manifold is given in [12], but the flow in this case is not the suspension flow of a pseudo-Anosov map.

Of course such an example cannot exist without some restrictions, for we have the following well known observation from the theory of dynamical systems:

Lemma 4.3 *Suppose that S is an embedded surface transverse to the foliation $\mathcal{L}$ and that S meets every flowline at least once. Then S is a fiber of a fibration of M.*

Proof. For the purposes of this proof it is convenient to think of the foliation $\mathcal{L}$ as the orbits of a flow $\{\phi_t\}$ on the manifold M. The main claim is that every point flows forward onto S; once this has been established it will follow that M can be reconstructed as the mapping cylinder of the first return map $S \rightarrow S$ and thus that S is a fiber in a fibration of M.

To this end, given any point $p \in M$ define $W(p)$ to be the set of all possible accumulation points of the set $\{ \phi_{t_n}(p) \mid t_n \rightarrow \infty \}$. Since M is compact, this set is nonempty. Then the set $W(p)$ is actually invariant under the flow, since if $q \in W(p)$ and any K is given, $\phi_K(q) = lim_{n\rightarrow\infty}\phi_K(\phi_{t_n}(p)) = lim_{n\rightarrow\infty}\phi_{K+t_n}(p)$ and the righthandside lies in $W(p)$ by definition. Thus $W(p)$ is a union of flowlines.

Suppose that we were able to choose some p so that the forward orbit did not meet S. Since S meets the flow transversely, it follows that the set $W(p)$ does not meet some neighbourhood of S, and we deduce that there is a flowline which does not meet S, a contradiction. ■

Thus an embedded surface which is transverse to $\mathcal{L}$ meets every flowline if and only if it is geometrically infinite. In fact it follows from results of Fried [6] that any embedded surface transverse to the suspension flow of a pseudo-Anosov map associated to fibration $F \rightarrow M \rightarrow S^1$ must lie in the closure of the face of the Thurston norm containing $[F]$.

The reason is that Fried proves that one can characterise the representatives of such a face as those embedded surfaces with positive intersection numbers with a finite number of (carefully chosen) closed flowlines. Therefore if S is any embedded surface transverse to $\mathcal{L}$, then for any positive integer n, the class $[nS + F]$ meets every flowline and in particular all the closed orbits, so must lie in the interior of the face containing $[F]$. From this it follows that $[S]$ lies in the closure of the face.

We now show how to construct examples of embedded surfaces which are transverse to the suspension flow of a pseudo-Anosov, but are quasi-Fuchsian. The above remarks show that such a surface is in the boundary of the face of the Thurston norm containing the fibration.

We need the following construction: Suppose that $\theta : S \rightarrow S$ is a pseudo-Anosov map which has the property that one can find a nonseparating oriented simple closed curve C with (a) C is disjoint from $\theta(C)$ and (b) With the given orientations C and $\theta(C)$ represent the same nonzero class in $H_1(S)$.

Under these circumstances, we may form an embedded surface S which is transverse to the flow in the following way: Let S_0 be one of the subsurfaces of F whose boundary components are C

and $\theta(C)$. Regard this as a subsurface $S_0 \times \{0\} \subset F \times \{0\} \subset F \times I$ and adjoin an annulus $C \times I$. This surface then closes up to be an orientable surface in M, which one sees easily can be made transverse to the flow.

We claim that this surface is actually incompressible; in fact we show the stronger fact that it is a Thurston norm minimising representative of the homology class. The reason is that if we do double curve sums with F, then we obtain a new surface transverse to the flow - so that although this new surface may be disconnected, it cannot involve any $2-$sphere components. Therefore the Euler characteristic of $S + nF$ is $\chi(S) + n\chi(F)$. If we sum with enough copies of F we obtain a class lying in the fiber face, which is thus an integral multiple of a class represented by a fiber. It follows that for very large n, $S + nF$ is a Thurston norm minimising surface. However, if the surface S were not norm minimising, we could replace it by a representative with smaller Euler characteristic and double curve sum with the same number of copies of F to lower the Euler characteristic of $S + nF$, a contradiction.

We shall now give a construction for examples of the type promised above. Suppose that F is a hyperbolic orbifold with underlying topological space a torus, and that $\theta : F \rightarrow F$ is a pseudo-Anosov map with the properties described above as well as the property that it fixes at least one of the cone points. (We shall show below that such examples may be constructed rather easily.) The surface S_0 of the paragraph above must be a topological annulus with some number of cone points. This number cannot be zero or all, else (looking to the other homology cobordism in F if necessary) we find a free homotopy between C and $\theta(C)$ which is impossible for a pseudo-Anosov map.

Form the orbifold M containing the surface S as above, choosing as S_0 the side of the homology cobordism which does not contain a fixed cone point. Note that this guarantees that the closed flowline through this cone point misses the surface S. By passing to a torsion free subgroup of finite index, we find a surface with all the promised properties, which misses the full preimage of this flowline and therefore is quasi-Fuchsian.

It remains to verify the existence of such maps. One example which one can compute explicitly comes from using the monodromy of the Borromean rings. One finds in this case that there is a curve which is moved disjoint from its image, so that by doing appropriately large $(0, k)$ surgeries, all the above hypotheses can be arranged.

More generally, such examples may be constructed as follows. We refer to Figure 5A: This depicts a surface F which is a torus with a single boundary component glued to a sphere with at least four boundary components.

Insert Figures 5A & 5B.

Choose a pseudo-Anosov map of the braid group of the disc and denote it by ξ. We also have a map of F to itself which we shall denote by p_α; it is an element of the braid group of F which is constructed by moving the first puncture around F and back to its starting position along a path such as that indicated by α in Figure 5A. With the choice of path shown it is clear that the image of the curve C under the composition $p_\alpha \xi$ is the curve denoted in Figure 5B by $\theta(C)$. In order to arrange that the other curve on the torus is moved into the disc part we should choose some other path β using the second puncture. It is clear our choices can be made so that $\theta = p_\beta p_\alpha \xi$ is pseudo-Anosov and carries the curve C to the curve $\theta(C)$ as shown in Figure 5B. Filling in the punctures on F will now give the orbifold example which was promised.

5 Appendix : The Cannon-Thurston map.

For the convenience of the reader, we will review part of the work of Cannon and Thurston on geometrically infinite surfaces [2]. Let M be a closed hyperbolic 3-manifold which fibers over the circle with fiber F and pseudo-Anosov monodromy $\theta : F \longrightarrow F$. Let $\lambda > 1$ be the stretch factor for

θ and $(\mathcal{F}^+, \mu^+)$ and $(\mathcal{F}^-, \mu^-)$ the stable and unstable projectively invariant measured foliations for θ respectively. If α is an arc transverse to these foliations then

$$\mu^{\pm}(\theta\alpha) = \lambda^{\pm 1}\mu^{\pm}(\alpha).$$

This means that θ contracts segments of stable leaves and expands segments of unstable leaves. Both these foliations have the same singular points $\mathcal{S}$ which are finite in number, and we denote the punctured surface obtained by removing the singularities by $F^* = F - \mathcal{S}$. The measured foliations define a Euclidean metric dt on F^* as follows. We adopt the convention that (x, y) are local coordinates on F^* where:

- x is constant along a leaf of $\mathcal{F}^-$
- y is constant along a leaf of $\mathcal{F}^+$.

Thus μ^+ is locally dx and μ^- is locally dy. Using these local coordinates θ expands the x direction and contracts the y direction and is given locally by the matrix

$$\begin{pmatrix} \lambda^{-1} & 0 \\ 0 & \lambda \end{pmatrix}$$

relative to local xy coordinates around $w \in F^*$ and θw. The metric completion of F^* is a singular Euclidean metric on F. A neighborhood of a singular point in F corresponding to a k-prong singularity of both $\mathcal{F}^+$ and $\mathcal{F}^-$ is isometric to the metric space obtained by taking a k-fold branched cover of the Euclidean plane branched around the origin and then quotienting by the isometry $x \mapsto -x$.

Now put a Riemannian metric ds on $F^* \times [0, 1]$ given by the formula

$$ds^2 = \lambda^{-2z}dx^2 + \lambda^{2z}dy^2 + dz^2$$

where z is the coordinate on $[0, 1]$. This is a Solv metric, see [13] and is chosen so that the map

$$F^* \times 1 \longrightarrow F^* \times 0 \qquad (w, 1) \mapsto (\theta w, 0)$$

is an isometry. This induces a Solv metric on $F^* \times I/(w, 1) \equiv (\theta w, 0)$ and the metric completion of this is a Solv metric with singularities on M.

Since M is compact, the metric ds on M and the hyperbolic metric dh on M are Lipschitz related thus there is $K > 1$ with

$$K^{-1}dh \leq ds \leq Kdh.$$

Lifting these metrics to the universal cover

$$\pi : \tilde{M} \longrightarrow M$$

the same comparison applies to the lifted metrics.

Let ℓ be a regular leaf of $\mathcal{F}^{\pm}$ in the universal cover $\tilde{F}$ of F. Given a point w in $\tilde{F}$ there is a unique dt geodesic γ in $\tilde{F}$ which contains w and is orthogonal to ℓ. This is because dt has non-positive curvature. Let $\pi(w)$ be the point of intersection of γ and ℓ then the map

$$p : \tilde{F} \longrightarrow \ell$$

is called **orthogonal projection** onto ℓ and does not increase dt distance. We identify $\tilde{F} \times \mathbb{R}$ with $\tilde{M}$ and extend the above map to

$$p : \tilde{F} \times \mathbb{R} \longrightarrow \ell \times \mathbb{R} \qquad p(w, t) = (pw, t).$$

The formula for the Solv metric on $\tilde{M}$ shows that this map also does not increase ds distance. It follows that the shortest path in $\tilde{F} \times \mathbb{R}$ between two points in $\ell \times \mathbb{R}$ lies in $\ell \times \mathbb{R}$. Thus $\ell \times \mathbb{R}$ separates

$\tilde{F} \times \mathbb{R}$ into two components with closure A and B with the property that they are ds **convex** in the following sense. Given two points in A a shortest path between them lies entirely in A. We call A and B the **half-spaces** associated to ℓ.

Fix a hyperbolic metric on F then use this to identify $\tilde{F}$ with $\mathbb{H}^2$. Now identify $\mathbb{H}^2$ with the interior of the closed unit disc B^2 and identify $\mathbb{H}^3$ with the interior of the closed unit ball B^3 in the usual way. We will use de for the Euclidean metric on both B^2 and B^3. A continuous map

$$f : \tilde{F} \longrightarrow \mathbb{H}^3$$

has a continuous extension to

$$\overline{f} : B^2 \longrightarrow B^3$$

if and only if f is uniformly continuous with respect to the Euclidean metric de. We will apply this to the map f which covers the inclusion of F into M.

Let c be the center of B^2 and let ℓ be a leaf of $\mathcal{F}^{\pm}$ in $\tilde{F}$ which does not contain c. Let A be the half-space associated to ℓ that does not contain c. Suppose that the ds distance between c and A is D. If x and y are two points in A and γ is a ds-shortest path between them, then γ is contained in A and so $f\gamma$ is contained in $f(A)$. Now $f\gamma$ lies within some distance L of the unique hyperbolic geodesic δ between fx and fy. Since γ is a ds shortest path and f is a K quasi-isometry it follows that $f\gamma$ is a K-quasi-geodesic and so L depends only on K. Thus the distance of fc from δ is at least $D/K - L$. For D sufficiently large it follows that fA is contained in a set of small de-size. This implies the f is de-uniformly continuous.

Thus the inclusion $\tilde{F} \longrightarrow \tilde{M}$ has a continuous extension to the circle at infinity. Since $\pi_1(F)$ is normal in $\pi_1(M)$ they have the same limit set, thus $S^1_\infty(\tilde{F})$ maps onto $S^2_\infty(\tilde{M})$. We claim that two points in $S^1_\infty(\tilde{F})$ have the same image if and only if they are limit points of the same $\mathcal{F}^{\pm}$-leaf. Let ℓ be a leaf of $\mathcal{F}^+$, the formula for the singular Solv metric restricted to $\ell \times \mathbb{R}$ is

$$ds^2 = \lambda^{-2z} dx^2 + dz^2.$$

Suppose that $|x_2 - x_1|$ is large then the ds distance between two points $(x_1, 0)$ and $(x_2, 0)$ on $\ell \times 0$ is approximately $2 \log_\lambda |x_2 - x_1|$. An approximation to a ds-shortest path γ between them consists of two vertical intervals $x_i \times [0, log_\lambda |x_2 - x_1|]$ together with the horizontal segment of $\ell \times \log_\lambda |x_2 - x_1|$ between $x_1 \times \log_\lambda |x_2 - x_1|$ and $x_2 \times \log_\lambda |x_2 - x_1|$. If x_1 and x_2 are a large ds distance from c then γ is also a large distance from c. It follows that $f\gamma$ is a large dh-distance from fc and since fc is almost a K-quasi-goedesic it has small de diameter in B^3. This proves the if part of the claim. Next we prove the only if part of the claim.

Let λ^+, λ^- be regular leaves of $\mathcal{F}^+, \mathcal{F}^-$ respectively which intersect in a point x. The intersection of a half space associated with λ^+ with one for λ^- is called a **quarter space** and the line $\ell = x \times \mathbb{R}$ is called the **axis** of the quarter space. There are 4 quarter spaces associated to λ^+, λ^- and two of these which intersect only in ℓ are called **opposite** quarter spaces.

Lemma 5.1 *Suppose that A', B' are two opposite quarter spaces which contain quarter spaces A, B respectively in their interiors. Then A and B have disjoint closures in the sphere at infinity.*

Proof. The proof involves showing there is a compact subset K of $\tilde{M}$ such that every dh-geodesic segment with endpoints in A and B meets K. ■

Assuming this consider two points a, b on $S^1_\infty(\tilde{F})$ which are not limit points of the same leaf. There are quarter spaces A, B as in the lemma such that $a \in cl(A)$ and $b \in cl(B)$ and this proves the claim.

Let λ^+ be a regular leaf of $\mathcal{F}^+$ then the ds-metric restricted to $\lambda^+ \times \mathbb{R}$ is the hyperbolic metric and the lines $\lambda^+ \times t$ are horocycles which curve upwards. A vertical line $x \times \mathbb{R}$ is a ds geodesic which is orthogonal to these horocycles and therefore converges in the upwards direction to the point at infinity where these horocycles limit.

This shows that the map which flows a point on $\tilde{F} \times 0$ upwards to the sphere at infinity is the map which sends a point $x \in \tilde{F}$ to the point on the sphere at infinity to which the Cannon-Thurston map sends the leaf λ^+ containing x.

References

[1] F. Bonahon. *Bouts des variétés hyperboliques de dimension 3.* Ann. of Math. 124 (1986) pp. 71-158.

[2] J. Cannon and W.P.Thurston. *Group Invariant Peano Curves.* Preprint.

[3] A.J. Casson and S. A. Bleiler. *Automorphisms of Surfaces after Nielsen and Thurston.* C.U.P. (1988).

[4] M. Culler and P. B. Shalen. *Varieties of group representations and splittings of 3-manifolds.* Ann. of Math. 117 (1983) pp. 109-146.

[5] S. Fenley. *Depth one foliations in hyperbolic 3-manifolds.* Ph.D. Dissertation, Princeton, 1989

[6] D. Fried. *Fibrations over S^1 with pseudo-Anosov monodromy.* Exposé 14 in Travaux de Thurston sur les surfaces, Astérisque, 66 - 67 (1979)

[7] A. Hatcher and W. Thurston. *Incompressible surfaces in 2-bridge knot complements.* Invent. Math. 79 (1985) 225-246.

[8] W. Heil. *Normalizers of incompressible surfaces in 3-manifolds.* Glasnik Matematički. 16 (1981) pp. 145 - 150.

[9] D.D. Long. *Constructing pseudo-Anosov maps.* Knots and manifolds. Ed. D. Rolfsen. Springer-Verlag Lecture Notes No. 1144, pp. 108 - 114.

[10] A. Marden *The geometry of finitely generated Kleinian groups.* Ann. of Math. 99 (1974) pp. 383 - 461

[11] L. Mosher *Surfaces and branched surfaces transverse to pseudo-Anosov flows on 3-manifolds.* J. Diff. Geom. 34 (1991) pp. 1 - 36.

[12] L. Mosher *Examples of quasi-geodesic flows on hyperbolic 3-manifolds.* Topology 90, Ohio State university Math. Research Inst. Pub. 1, pp. 227 - 241.

[13] G.P. Scott *The geometries of 3-manifolds.* Bull. L.M.S. 15 (1983) pp. 401 - 487.

[14] W.P. Thurston. *The Geometry and Topology of 3-manifolds.* Princeton University mimeographed notes. (1979)

[15] W.P. Thurston. *Hyperbolic structures on 3-manifolds II: Surface groups and 3-manifolds which fiber over the circle.* Preprint.

[16] W.P. Thurston. *A norm for the homology of 3-manifolds.* Memoirs of the A.M.S. vol 339 pp. 100-130

[17] W.P. Thurston. *Three dimensional manifolds, Kleinian groups and hyperbolic geometry.* Bull. A.M.S. vol 6 (1982) pp. 357 - 381.

[18] F. Waldhausen. *On irreducible three-manifolds which are sufficiently large.* Ann. of Math. 87 (1968) pp. 56-68.

Cooper and Long
Department of Mathematics,
University of California,
Santa Barbara, Ca. 93106 USA

Reid
Department of Pure Mathematics,
University of Cambridge,
Cambridge CB2 1SB
UK

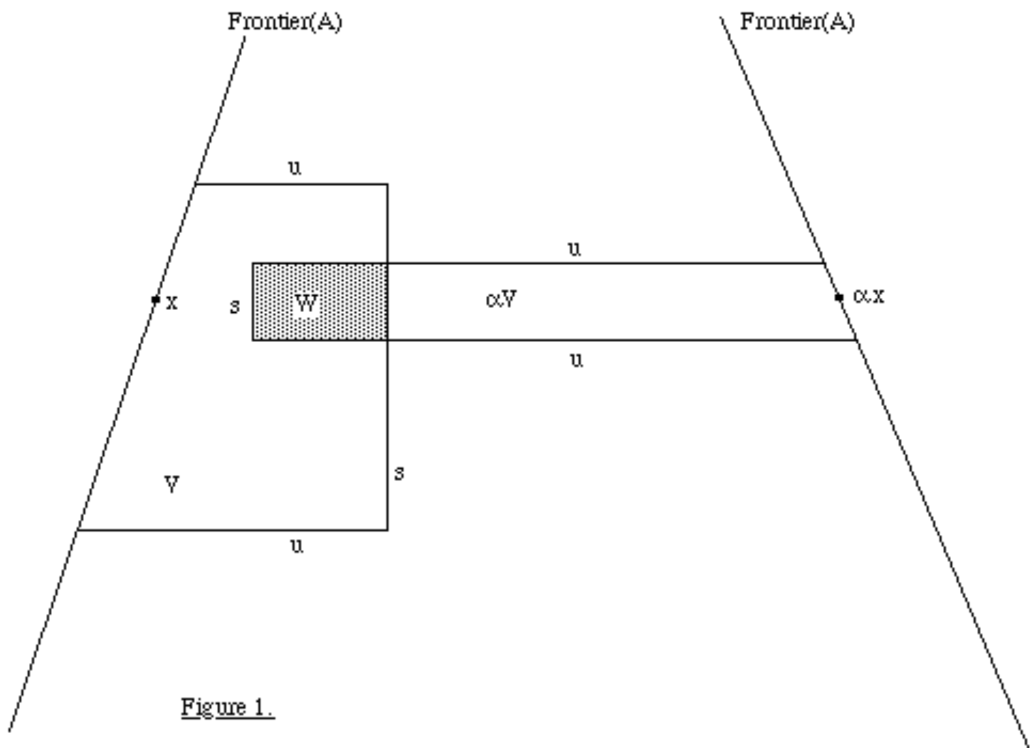

Figure 1.

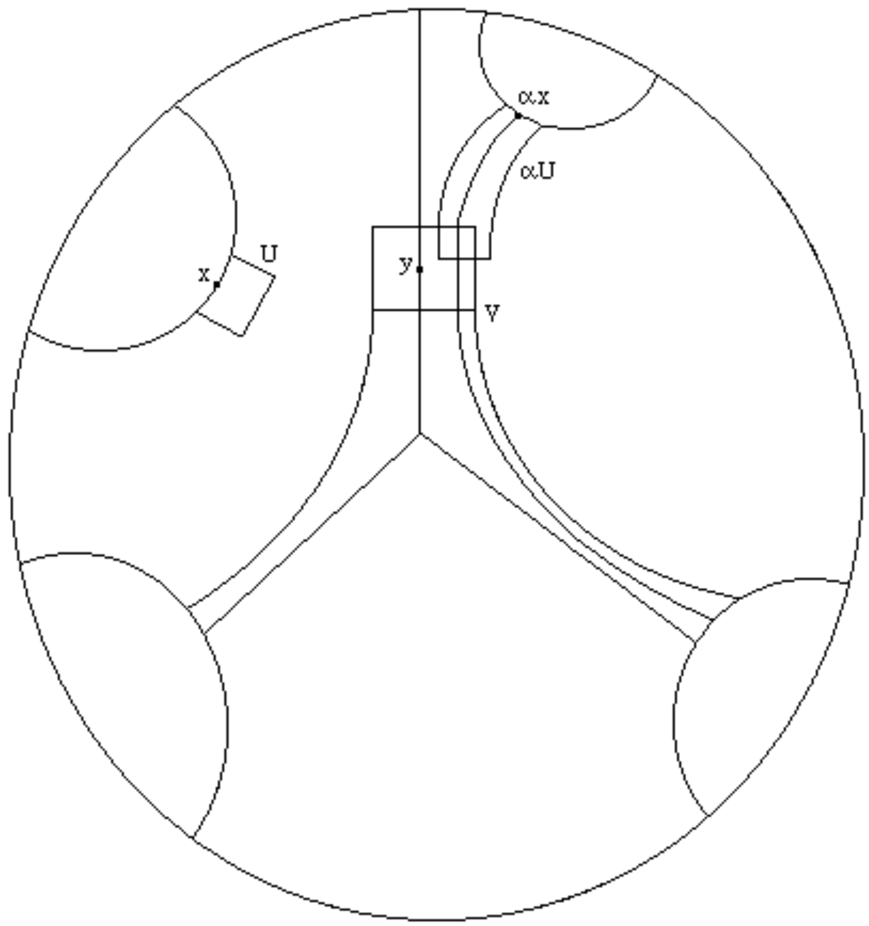

Figure 2.

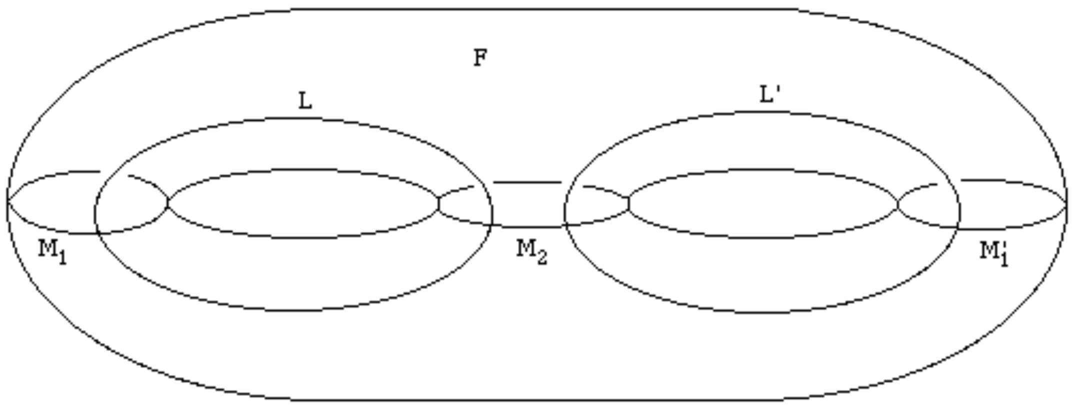
F
L
L'
M_1
M_2
M_1'

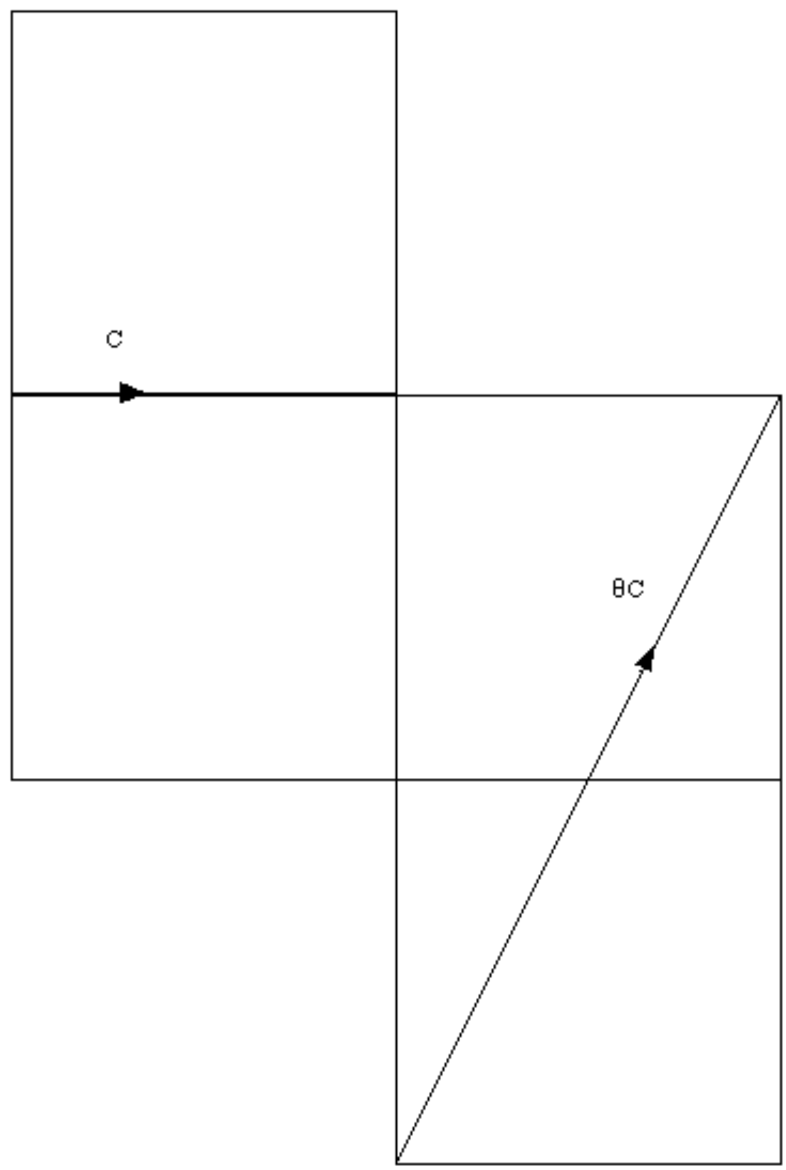
C
θC

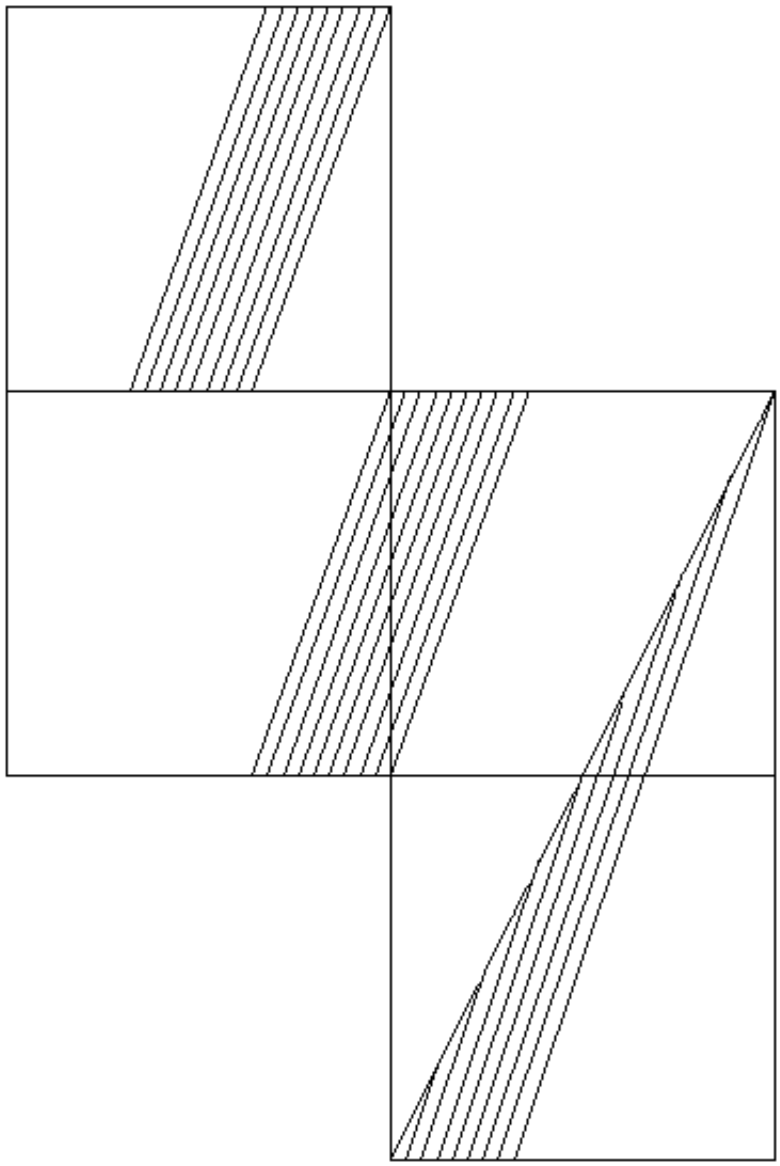

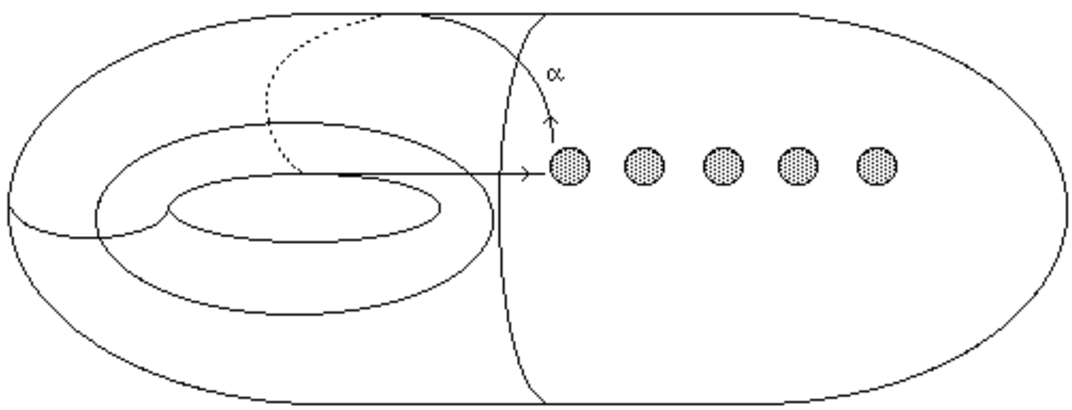

Figure 5A.

Figure 5B.