

MINIMAL LEAVES IN FOLIATIONS

BY
DANIEL M. CASS

ABSTRACT. The paper defines a property of open Riemannian manifolds, called *quasi-homogeneity*. This property is quasi-isometry invariant and is shown to hold for any manifold which appears as a minimal leaf in a foliation. Examples are given of surfaces which are not quasi-homogeneous. One such is the well-known noncompact leaf of Reeb's foliation of S^3 . These surfaces have bounded geometry.

Let \mathcal{F} be a foliation of a compact manifold M^n by leaves of dimension p . When M is given a metric, a leaf L of \mathcal{F} receives an induced metric which depends both on the metric on M and on the way L is imbedded in M . To get a geometric property of L which does *not* depend on the metric on M we define:

DEFINITION 0.1. The quasi-isometry class $[L]$ of a p -dimensional Riemannian manifold L is the set of images of L under diffeomorphisms ϕ of "bounded dilation": i.e., under diffeomorphisms ϕ satisfying, for some positive (k, K) , $k < \|\phi_*v\|/\|v\| < K$ for all $v \in \tau(L)$.

Any change in the metric on M is equivalent to a diffeomorphism $\Phi: M \rightarrow M'$ carrying \mathcal{F} to a foliation \mathcal{F}' of M' , and the leaf L to a leaf L' . Because M is compact, Φ has bounded dilation, therefore so does $\phi: L \rightarrow L'$ induced by Φ . We see from this that if $L \in \mathcal{F}$ and is given the induced metric from some metric on M , then $[L]$ is independent of the metric chosen on M . This allows us to pose the following question.

QUESTION 0.2. *Appearance question for leaves.* Which quasi-isometry classes $[L]$ contain a representative L which appears as a leaf in some foliation \mathcal{F} of some Riemannian manifold M^m ?

Notice that in this question we do *not* prescribe the codimension $q = m - p$. This question was raised by Sullivan (Proc. Internat. Congress Math., 1974). M. Gromov has shown that if L is Riemannian and has "bounded geometry" then some member of $[L]$ admits an immersion into a "branched manifold"; however, it is still an open question whether "bounded geometry" implies appearance as a leaf of a foliation.

This paper gives no answer for Question 0.2; however, we give a partial answer to a more restricted question, depending on the following notions.

DEFINITION 0.3. A *minimal set* in a foliation is a collection of leaves which is closed in the M -topology and has no proper subsets which are closed unions of leaves. A *minimal leaf* is a leaf in a minimal set.

Received by the editors July 28, 1983 and, in revised form, January 4, 1984.
1980 *Mathematics Subject Classification.* Primary 57R30.

Key words and phrases. Dilation, quasi-isometry class, minimal set of leaves, recurrent.

©1985 American Mathematical Society
0002-9947/85 \$1.00 + \$.25 per page

An application of Zorn's Lemma assures the existence of minimal leaves in any foliation, so any results about minimal leaves will apply to any foliation whatsoever. Further, any leaf comes close to some minimal leaf, as we see from:

DEFINITION 0.4. If $L \in \mathcal{F}$ and K_i is a sequence of compact subsets of L such that $K_i \subset K_{i+1}$ and $\bigcup K_i = L$, then the *asymptote of L* , denoted $A(L)$, consists of the intersection of the closures in M of the sets $L - K_i$.

$A(L)$ is a closed union of leaves, each point of which is approached by a sequence of points of L . $A(L)$ may or may not contain L , according to which L is called "improper" or "proper". In any case, Zorn's Lemma shows that $A(L)$, if nonempty, must contain a minimal set, and so at least one minimal leaf. ($A(L)$ is empty exactly when L is a compact leaf. A compact leaf is a minimal set.)

We see that minimal leaves exist in any foliation, and also in any asymptote of any noncompact leaf. So similar to 0.2 we should pose

QUESTION 0.5. *Appearance question for minimal leaves.* Which quasi-isometry classes $[L]$ contain a representative L which appears as a minimal leaf in some foliation \mathcal{F} of some Riemannian manifold M^m ?

Theorem 2 of this paper will provide a restriction on the quasi-isometry classes which have representatives appearing as minimal leaves. The theorem implies for example that the noncompact leaves in Reeb's foliation of S^3 are of a quasi-isometry type which cannot appear in a foliation (of whatever codimension) as minimal leaves. They are in this sense "essentially nonminimal".

A leaf of linear growth, which in codimension one must be an "infinite repetition" of a compact segment (Cantwell and Conlon), need not be a repetition in higher codimension, as we will show in Example 4.1. The noncompact leaves of this example are "essentially nonminimal" and appear there as nonproper leaves. The noncompact leaves of Reeb's foliation, on the other hand, cannot appear in a foliation as nonproper leaves—they are "essentially proper" as well as "essentially nonminimal".

As is standard, $\tau(\mathcal{F})$, the tangent bundle to \mathcal{F} , is a subbundle of $\tau(M)$ consisting of p -planes tangent to leaves. We identify M with the zero-section of $\tau(M)$. We have the exponential map $\exp: \tau(M) \rightarrow M$ defined in a neighborhood of M .


A key element in what follows will be the existence of "maps between nearby leaves". To this end we need to fix in advance an $\varepsilon^* > 0$ together with a subbundle \tilde{n} of $\tau(M)$ which is complementary to $\tau(\mathcal{F})$ and has the following property:

1.1 The disc bundle $\tilde{n}_{\varepsilon^*}$ (i.e. the ε^* -neighborhood of the zero-section M of \tilde{n}), when restricted to a leaf, is locally embedded by \exp . More precisely, for any leaf L of \mathcal{F} and any point $x \in L$, x has a d -neighborhood N such that $\exp: \tilde{n}_{\varepsilon^*}(N) \rightarrow M$ is an embedding. (The d -metric is the metric along leaves.)

If the leaves of \mathcal{F} are C^2 smooth, then the normal bundle to the foliation will serve as \tilde{n} . Even if the leaves are only C^1 smooth, there still exists such a complement \tilde{n} .


If K is a subset of a leaf of \mathcal{F} , we will write $\tilde{n}_{\varepsilon}(K)$ for $\exp(\tilde{n}_{\varepsilon}(K))$, thinking of it as a locally imbedded disc bundle.

We would like to use the disc bundle $\tilde{n}_{\varepsilon^*}$ in order to construct maps via path-lifting from balls of radius A in leaves to balls in nearby leaves. So we need a lemma which guarantees the existence of lifts of paths of length less than A :



σ_1 and σ_2 represent different points of $\tilde{B}(x, A)$ although the same point in the universal cover.

FIGURE 1.3



The leaves near x can pull away from L_x , but not too fast, if they start within δ of x .

FIGURE 1.4

LEMMA 1.2. *Given ε less than the ε^* of 1.1 and $A > 0$ however large, we can find $\delta > 0$ such that, for all $x \in M$ and all $y \in \tilde{n}_\delta(x)$, any path σ_x in L_x = leaf through x of length less than A will lift via the disc bundle $\tilde{n}_\varepsilon(B(x, A))$ to a path $\bar{\sigma}_y$ in L_y .*

This can be shown using a plaque decomposition of \mathcal{F} . The idea is that there is a universal bound on the ratio between greatest and least distances along discs $\tilde{n}_\varepsilon(x)$ for two plaques in a flow box.

The universal cover of a leaf may contain two paths of length less than A which are homotopic only through paths that become longer than A . So in order to make our proposed path-lifting maps well defined, we need the notion of “ A -cover balls” which follows:

DEFINITION 1.3. An A -cover ball in the leaf L_x centered at x , which we denote $\tilde{B}(x, A)$, consists of equivalence classes of paths from x in L_x of length less than A , where the equivalence is homotopy rel endpoints through paths of length less than A (see Figure 1.3).

With this definition a straightforward argument involving a plaque decomposition of \mathcal{F} and Lemma 1.2 for small ε produce:

LEMMA 1.4. *Given $\varepsilon < \varepsilon^*$ of 1.1 and $A > 0$ however large, there is a δ such that for all $y \in \tilde{n}_\delta(x)$ the path lifting map $f: \tilde{B}(x, A) \rightarrow L_y$ taking $\sigma_x \mapsto \bar{\sigma}_y(1) \in L_y$ is a local embedding with dilation $(2^{-\varepsilon}, 2^\varepsilon)$ (see Figure 1.4).*

To determine whether L is a minimal leaf of \mathcal{F} , Definition 0.3 seems to leave us only one procedure: We must examine the whole set of leaves in the closure of L and see if this set contains any proper closed subsets which are unions of leaves. Theorem 1 gives us an alternative, since it states that the minimal leaves

are precisely the *recurrent* leaves, and whether a leaf L is or is not recurrent depends only on the embedding $L \hookrightarrow M$ and not on the other leaves in the closure of L . The definition of recurrent follows.

DEFINITION 2.1. Suppose M is given a metric inducing a distance function \hat{d} on M , which we denote by d when it is restricted to a leaf. Then L is *recurrent* provided that for each $\varepsilon > 0$ there is a sufficiently large T such that

$$L \subset B_{\hat{d}}(B_d(x, T), \varepsilon) \quad \text{for all } x \in L.$$

The force of the definition is that T depends only on ε and not on the point $x \in L$. Note that a leaf remains “recurrent” under any change of metric on M . Theorem 1 then states simply, for any leaf L of any foliation \mathcal{F} :

THEOREM 1. *L is minimal if and only if L is recurrent.*

This theorem was stated and proved by Birkhoff for flows, i.e. foliations with one-dimensional leaves. Plante has shown the forward implication for the codimension one case (Topology **12**, p. 179). Plante’s proof resembles the following one; however, in general codimension one needs to use an immersed microbundle complementary to the tangent bundle of the foliation, so we include the details.

The theorem allows us to replace “minimal” with “recurrent”, a notion which is one step closer to being intrinsic to the geometry of L . There still remains, however, the reference to \hat{d} , the metric on M ; this reference will be eliminated in Theorem 2 and its accompanying Definition 3.3.

FORWARD IMPLICATION. A minimal leaf must be recurrent.

Assume L is minimal and has closure Σ . This means that Σ is a minimal set of leaves, or equivalently that if L' is any leaf in Σ then $\text{closure}(L') = \Sigma$. Now assume that L is not recurrent. Then it is possible to find $\varepsilon > 0$, $T_n \rightarrow \infty$, and points $p_n, q_n \in L$, such that

$$(*) \quad q_n \text{ is further than } \varepsilon \text{ (in the } \hat{d}\text{-distance) from the leaf-ball } B_d(p_n, T_n).$$

Since Σ is a compact subset of M we can pass to subsequences and assume $p_n \rightarrow p$, $q_n \rightarrow q$, where $p, q \in \Sigma$. Consider the leaf L_p passing through p .

Claim. Every point of L_p is at least at distance $\varepsilon/3$ from the point q in the \hat{d} -distance.

This claim implies a contradiction to minimality of Σ_1 since the closure of L_p is then a closed union of leaves, which is a proper subset of Σ because it fails to contain q . Hence a proof of the Claim will finish the proof that a minimal leaf must be recurrent.

PROOF OF CLAIM. Pick a point on L_p and connect it to p via a path σ with $\sigma(0) = p$. If n is large enough, Lemma 1.2 implies σ will lift via $\tilde{n}_{\varepsilon^*}$ to a path $\bar{\sigma}$ in L = leaf through p_n . The endpoint $\bar{\sigma}(1)$ will lie in $B_d(p_n, T_n)$ as long as T_n is larger than say $2\|\sigma\|$. Further increasing n we can ensure (putting $z = \bar{\sigma}(1)$) that $\hat{d}(z, x) < \varepsilon/3$, $\hat{d}(q_n, q) < \varepsilon/3$, and $(*)$ gives that $\hat{d}(z, q)_n > \varepsilon$ (see Figure 2.1).

To finish the proof of the Claim, we apply the triangle inequality to the sequence of points (z, x, q, q_n) of M using the metric \hat{d} . (We suppress the “ \hat{d} ” in the distances.)

$$(z, q_n) < (z, x) + (x, q) + (q, q_n).$$




FIGURE 2.1

When we reverse this inequality we obtain

$$(x, q) > (z, q_n) - (z, x) - (q, q_n) > \varepsilon - \varepsilon/3 - \varepsilon/3 = \varepsilon/3.$$

Thus every point x of L_p is at least at distance $\varepsilon/3$ from q as stated in the Claim. (It is clear that if a leaf is at a positive distance from a point, then so is its closure.)

REVERSE IMPLICATION. A recurrent leaf must be minimal.

Assume L is a recurrent leaf with closure Σ , and that L is not minimal. Then some leaf K of Σ has closure $\bar{K} \neq \Sigma$. L is not contained in \bar{K} , otherwise we would have $\bar{K} = \Sigma$. So pick a point x on L at a positive distance ε_1 (\hat{d} -distance) from the compact set \bar{K} . Pick also a point y on the leaf K . Since K is a leaf of $\Sigma = \text{closure of } L$, we can pick a sequence y_n of points of $\hat{n}_{\varepsilon_1}(y) \cap L$, with $y_n \rightarrow y$.

The recurrence assumption produces $T(\varepsilon_1/3)$ large enough that any ball of radius T in L approximates the whole leaf to within $\varepsilon_1/3$. Now apply Lemma 1.2 with its ε taken to be the present $\varepsilon_1/3$, and its A taken to be the present T . It is then clear that for large enough n , paths σ in $B_d(y_n, T)$ of length less than T will lift via \tilde{n}_{ε} to paths $\bar{\sigma}$ in $B_d(y, T)$ satisfying

$$\hat{d}(\sigma(t), \bar{\sigma}(t)) < \varepsilon_1/3, \quad \text{for each } t \in [0, 1].$$

The ball $B_d(y_n, T)$ must approximate all of L to within $\varepsilon_1/3$. In particular, x is within $\varepsilon_1/3$ of some point z of $B_d(y_n, T)$, in the sense of the distance \hat{d} . But if σ is a path in L starting from y_n and going to z of length $< T$, then the endpoint of its lift $z' = \bar{\sigma}(1)$ will be within $\varepsilon_1/3$ of the point z . This gives (see Figure 2.2)

$$(x, z') < (x, z) + (z, z') < \varepsilon_1/3 + \varepsilon_1/3 = 2\varepsilon_1/3,$$

contradicting the fact that x is at distance ε_1 from K and $z' \in K$. This contradiction completes the proof of the reverse implication, ending the proof of Theorem 1.

The recurrence property, which involves both distance functions d and \hat{d} , and therefore the geometry of the embedding $L \rightarrow M$, nonetheless *implies* a property which is intrinsic to the geometry of L . This property we will call “quasi-homogeneity”. To define it we need the notion of an “ A -cover ball” (Definition 1.3) and a related notion:

DEFINITION 3.1. Suppose $\tilde{B}(x, A)$ is an A -cover ball in a Riemannian manifold L . Then an immersion f of $\tilde{B}(x, A)$ into another manifold is called *endpoint-monic* provided whenever $f(\sigma_1) = f(\sigma_2)$, the paths σ_1 and σ_2 must end at the same point (in L).




FIGURE 2.2

That such an immersion f be “endpoint-monik” is weaker than that it be monik, for it allows the possibility of two paths from x to y in L , which represent different equivalence classes in $\tilde{B}(x, A)$, to be carried by f to the same point.

A good example of an f which is endpoint-monik but not monik is the covering map $\tilde{B}(x, A) \rightarrow B(x, A)$. Since the covering map assigns to an equivalence class of paths their common endpoint, it is clearly endpoint monik. That it can fail to be monik is clear from Figure 1.3. Here the two paths σ_1 and σ_2 from x_0 are distinct (inequivalent) in $\tilde{B}(x, A)$ but the covering map sends both to their endpoint y .

Now we can define “quasi-homogeneous”.

DEFINITION 3.2. Let L be a p -dimensional Riemannian manifold. L is called *quasi-homogeneous* provided there exist dilation bounds (k, K) such that:

(1) For any positive r there is positive R satisfying: Any r -cover ball $\tilde{B}(x, r) \subset L$ admits an immersion $\tilde{B}(x, r) \xrightarrow{f} B(y, R)$ into any ball $B(y, R) \subset L$, where all such f share dilation bounds (k, K) .

(2) For any fixed r -cover ball $\tilde{B}(x, r) \subset L$ there is R' so large that there is an endpoint-monik immersion $\tilde{B}(x, r) \rightarrow B(y, R')$ into any ball $B(y, R') \subset L$.

Notice that the endpoint-monik immersion asserted to exist in (2) might need an R' which is much larger than the R of (1).

It is easily shown that “quasi-homogeneous” is quasi-isometry invariant. With the above definition, we now have:

THEOREM 2. *A minimal leaf must be quasi-homogeneous in the metric it inherits from M .*

PROOF OF THEOREM 2. By Theorem 1 we can assume L is recurrent. We will show it is quasi-homogeneous for dilation bounds $(\frac{1}{2}, 2)$. So let r be a given positive number (which can be assumed to be greater than 1); we must produce R so as to satisfy the conditions of “quasi-homogeneity”.

Set $r_1 = r+1$. We apply Lemma 1.4 to r_1 and any ϵ small enough that $[2^{-\epsilon}, 2^\epsilon] \subset [\frac{1}{2}, 2]$. This ensures that the maps f which come from path-lifting will all have dilation bounds $(\frac{1}{2}, 2)$. Specifically, Lemma 1.4 gives us $\delta_1 > 0$ (we also assume




FIGURE 3.1

$\delta < 1/2$) small enough so that whenever x, y are points of M with $y \in \tilde{n}_{\delta_1}(x)$, we can always define a path-lifting map $f: \tilde{B}_d(x, r_1) \rightarrow L_y$ as stated in the lemma.

Now we choose δ_2 small enough so that for $x \in M$ we have

$$B_d(x, \delta_2) \subseteq \tilde{n}_{\delta_1}(B_d(x, \delta_1)).$$

Such a δ_2 exists because M is compact and is the zero-section of $\tilde{n}_{\varepsilon^*}$. We now apply recurrence with its $\varepsilon = \delta_2$. This produces $T(\delta_2)$ with the property that any d -ball of radius $T(\delta_2)$ in L is δ_2 -close (in the \hat{d} -distance) to the entire leaf L . Finally we set $R = T(\delta_2) + 2r_1$. We claim that this R satisfies the condition of quasi-homogeneity (part 1).

To check this, assume we are given $\tilde{B}_d(x, r)$ and $B_d(y, R)$, both situated at random in the leaf L . By recurrence, some point y_1 of $B_d(y, T(\delta_2))$ is δ_2 -close to the point x . This implies there is $x_1 \in L_x$ with $y_1 \in \tilde{n}_{\delta_1}(x_1)$, and $d(x, x_1) < \delta_1$ (see Figure 3.1). Our choice of δ_1 then gives the existence of a path-lifting map $f: \tilde{B}_d(x_1, r_1) \rightarrow L_{y_1}$. Since f must have dilation bounds $(\frac{1}{2}, 2)$ it follows that the image $f(\tilde{B}_d(x_1, r_1))$ is contained in $B_d(y_1, 2r_1)$. Finally, our choice of r_1 implies that $\tilde{B}_d(x, r) \subset \tilde{B}_d(x_1, r_1)$, since $d(x, x_1) < \delta_1 < 1/2$, and our choice of R implies $B_d(y_1, 2r_1) \subset B_d(y, R)$. Thus the map required is the restriction of the above f to the r -cover ball $\tilde{B}_d(x, r)$.

For the proof of part (2) we consider a *fixed* ball $B(x, r)$. The idea is to follow through the part (1) proof, adjusting $\varepsilon, \delta_1, \delta_2, T$ along the way. To begin with, when we apply Lemma 1.4 to $r_1 = r + 1$ and ε , we make sure ε is not only small enough that $[2^{-\varepsilon}, 2^\varepsilon] \subset [1/2, 2]$ and $\varepsilon < 1/2$, but, in addition, ε is small enough that $\tilde{n}_\varepsilon(B(x, r_1))$ is *embedded* by the exponential map for the fixed ball $B(x, r_1)$. This can be done since $B(x, r_1)$ is contained in a compact subset of L and $\tilde{n}_{\varepsilon^*}$ is a local embedding when restricted to leaves.

Thus any map $f: \tilde{B}(x, r_1) \rightarrow L_y$ defined via path-lifting within this restricted disc bundle $\tilde{n}_\varepsilon(B(x, r_1))$ will be endpoint-monic, since the fibers $\tilde{n}_\varepsilon(z)$ for $z \in B(x, r_1)$ are pairwise disjoint.

Following through the proof of part (1) and adjusting, we see that since this part (2) ε is smaller than its part (1) counterpart, therefore δ_1 becomes smaller, and δ_2 in turn becomes smaller, and so $T(\delta_2)$ becomes larger, and so $R' = T(\delta_2) + 2r_1$ is now larger than the R of part (1). \square




FIGURE 4.1




FIGURE 4.2

EXAMPLE 4.1. There exists a codimension two foliation \mathcal{F} of a compact manifold M^4 by surfaces such that:

- (1) \mathcal{F} has one compact leaf T , a torus.
- (2) T is in the asymptote of every leaf of \mathcal{F} (in particular, T is the only minimal leaf of \mathcal{F}).
- (3) Each leaf of \mathcal{F} other than T has two ends and is of linear growth.
- (4) At least one end of each leaf other than T has the quasi-isometry type of a cylinder $R^+ \times S^1$, with handles attached, spaced at intervals d_k , where the d_k are unbounded as k increases.

See Figure 4.1.

Before describing this example, we indicate why the noncompact leaves of this foliation have quasi-isometry types which are not quasi-homogeneous, or in other words, why these leaves, considered as surfaces, are “essentially nonminimal”.

Recall that a noncompact leaf in our example will have at least one end with the quasi-isometry type of a cylinder with handles attached at intervals d_k , where the d_k are unbounded as k increases. The other end of our leaf L , if it is not of this type, is an infinite cylinder. These facts will be seen from the construction (see Figure 4.2). If L is such a leaf, the following holds:

Short homotopy fact. There is a limit on the size A of balls $B(x, A)$ contained in simply connected regions of L . It follows that there is c_1 such that any two homotopic paths γ_1, γ_2 of length less than A in L admit a homotopy through paths γ_t of length less than $c_1 A$. In particular, a null-homotopic closed loop of length less than A can be shrunk to the trivial loop through paths of length less than $c_1 A$. The cylinders which occur between the handles on L are arbitrarily long and of constant width. Let C_i be a sequence of these cylinders whose lengths steadily approach infinity. Now we fix a basepoint x_0 on L . For any $r > 0$ (we also take $r > 4W$, where W = circumference of cylinders) there exists a value $i(r)$ such that for $i \geq i(r)$ there are maps $f_i: \tilde{B}(x_0, r) \rightarrow C_i$ (guaranteed by homogeneity, part (2)) which carry the basepoint x_0 into the “middle third” of the cylinder C_i (see Figure 4.3). Let now γ_i be, for each i , a nontrivial closed loop in C_i based at $f_i(x_0)$. (We can take all these γ_i to have the same length.) Each path γ_i for i large enough will then pull back via f_i (since $r > 4W$) to a path $\bar{\gamma}_i$ which begins at x_0 ; we contend $\bar{\gamma}_i$ must be closed. Otherwise one violates that f_i is endpoint-monic: Consider the point $x_1 = \bar{\gamma}_i(1/2)$, and set $x_2 = \bar{\gamma}_i(1)$. The two halves of $\bar{\gamma}_i$ going from x_1 to




FIGURE 4.3

x_0 and x_1 to x_2 are then paths σ_1 and σ_2 which end at different points in L and yet are carried by f_i to the same point $f_i(x_0)$. (Since the lengths of γ_i are equal it is easy to justify this “change of basepoint” in the definition of endpoint-monic, provided i is large enough.)

So far we have for any i large enough, a closed loop γ_i on C_i which pulls back via f_i to a closed loop $\bar{\gamma}_i$ based at x_0 . What we need is a *single* nontrivial closed loop $\bar{\gamma}$ based on x_0 which, for infinitely many i , is carried via f_i to a nontrivial closed loop in C_i . We argue as follows. There is a bound on the lengths $\|\bar{\gamma}_i\|$. We claim that infinitely many of them must fall in the same nontrivial class $[\bar{\gamma}]$ of a closed loop $\bar{\gamma}$ based at x_0 . (If almost all $\bar{\gamma}_i$ were trivial we could use the “short homotopy fact” to show that a homotopy $\bar{\gamma}_i \sim 0$ would be carried by f_i to a homotopy $\gamma_i \sim 0$ which does not go off the end of C_i , for large enough i .) But infinitely many nontrivial loops $\bar{\gamma}_i$ with a bound on their lengths $\|\bar{\gamma}_i\|$ can represent only finitely many classes in $\pi_1(L, x_0)$. (This “pigeon hole principle for bounded sets of loops” holds for any leaf in a foliation.) The Claim is shown. We pass to the subsequence of C_i carrying those loops in what follows.

$\pi_1(L, x_0)$ is a countably generated free group. We fix a set of generators and express $[\bar{\gamma}]$ in terms of them. Then we pick two generators α and β not used in that expression. This means that the loop $\delta = (\alpha\bar{\gamma}\alpha^{-1})(\beta\bar{\gamma}\beta^{-1})(\alpha\bar{\gamma}^{-1}\alpha^{-1})(\beta\bar{\gamma}^{-1}\beta^{-1})$ is nontrivial. Since $\bar{\gamma}$ goes by f_i to a closed loop in C_i , so does $\alpha\bar{\gamma}\alpha^{-1}$ and $\beta\bar{\gamma}\beta^{-1}$. But δ is a commutator in C_i , and $\pi_1(C_i)$ is abelian. For large enough i , a homotopy of this commutator to the trivial loop at a point of C_i will pull back via f_i to a homotopy of δ to the trivial loop at x_0 , contradicting the fact that δ is nontrivial.

The existence of these pullbacks of homotopies for long enough cylinders C_i is guaranteed by the fact that the commutator δ has a length bound l independent of i and so the “short homotopy fact” bounds the homotopies by $c_1 l$. We must also choose our initial r large enough that $B(x_0, r)$ contains the loop δ and its homotopy, but this still leaves infinitely many C_i to work with, to get the contradiction. This completes our remarks about why the noncompact leaves in our example will fail to be quasi-homogeneous; it is high time we constructed the example.

CONSTRUCTION OF EXAMPLE 4.1. Start with the torus T^2 viewed as the unit square in the x - y plane with the usual side identifications. Pick an angle θ which is an irrational multiple of π . Let \vec{a} denote the unit vector field on T^2 all of whose vectors make the angle θ with the x -axis. Each orbit of \vec{a} is dense in T^2 .

Choose a point x_0 in the interior of the square and a small ε -disc around x_0 ; call it D . Let f be a C^∞ function defined on D which is zero only at x_0 and is 1 off a subdisc containing x_0 . We now slow down the vector field \vec{a} near x_0 by multiplying lengths by f :

$$\vec{g}(x, y) = \begin{cases} f(x, y) \cdot \vec{a} & \text{if } (x, y) \in D, \\ \vec{a} & \text{otherwise.} \end{cases}$$




FIGURE 4.4

As it stands, \vec{g} is a flow on the unit square in R^2 . Use the same letter \vec{g} to denote the flow on the unit cube in R^3 where now $\vec{g}(x, y, z) = \vec{g}(x, y)$. This is the old flow acting trivially in the z -direction.

We also use the unit flow in the z -direction, called $\partial/\partial z$. This is a vector field on the unit cube which points straight up, having no x - or y -components.

We now add these two flows to obtain $\vec{w} = \vec{g} + \partial/\partial z$. This is a flow on the unit cube in R^3 which is invariant under unit translations of the three axes and so defines a flow on the torus T^3 viewed as the cube with identifications.

The flow \vec{w} now has no stationary points. Adjust the lengths of the vectors so that \vec{w} is a unit flow. There will be one closed orbit in \vec{w} , namely $\{x_0\} \times (z\text{-axis})$. There are two special types of orbit in w which at $+\infty$ or at $-\infty$ limit on the closed orbit. All other orbits have the property that with time $\rightarrow \pm\infty$ they spend arbitrarily long times near the closed orbit, only to escape eventually, cross the x - z plane some number of times, and then return again near the closed orbit. Note that the closer an orbit passes to the closed orbit, the more vertical it becomes, making more and more passes in the vertical direction before escaping finally from the cylinder $D \times (z\text{-axis})$, only to return in time to this cylinder. Figure 4.4 illustrates this behavior.

For definiteness, assume the cylinder $D \times (z\text{-axis})$ lies in the left half $x < 1/2$ of the cube. Then on the right half of the cube the flow \vec{w} goes in a constant direction. Pick now four numbers $a < b < c < d$ between $1/2$ and 1 , where $d - a$ is small. We want to adjust \vec{w} to a flow \vec{w}_2 by straightening it out to go in the $\partial/\partial x$ direction between b and c . That is:

- (1) Outside $(a, d) \times (y\text{-axis}) \times (z\text{-axis})$, $\vec{w}_2 = \vec{w}$.
- (2) Each orbit of \vec{w}_2 enters and leaves the "slab" of (1) at the same points as the corresponding orbit of \vec{w} .
- (3) In the interval $[b, c]$ crossed with the y - z plane, the perturbed flow \vec{w}_2 is just $\partial/\partial x$.




FIGURE 4.5

It is clear we can change \vec{w} to \vec{w}_2 in this way, so that \vec{w}_2 is C^∞ . The purpose of this change is to make the flow especially simple in the region $[b, c] \times (y\text{-}z \text{ plane})$, because the next step is to perform a surgery operation involving the tori $\{b_1\} \times (y\text{-}z \text{ plane})$ and $\{b_2\} \times (y\text{-}z \text{ plane})$, where $b < b_1 < b_2 < c$, so that these tori are now at right angles to the flow \vec{w}_2 (see Figure 4.5).

The tori just mentioned will be denoted by $T^2(b_1)$ and $T^2(b_2)$. They are situated in the unit cube of R^3 perpendicular to the x -axis, and so also to the flow \vec{w}_2 , which is $\partial/\partial x$ in the slab $[b, c] \times (y-z \text{ plane})$; this slab contains both tori. The orbits of \vec{w}_2 have the property that they cross the pair of tori infinitely often in either direction (except for the special orbits limiting on the closed orbit, which do so in one direction and cross the tori infinitely often in the other). They spend arbitrarily long times between crossings for time going to $\pm\infty$, in the sense that given a long time T we can find a time interval $(t, t+T)$ in which the orbit remains in the cylinder $D \times (z\text{-axis})$, and so stays away from the pair of tori. After we perform the surgery, these long periods, in which an orbit stays away from the pair of tori, will produce long cylinders in the surface which “comes from” the orbit.

We are constructing a foliation of a manifold M^4 . We will obtain M^4 by removing tubes around $T^2(b_1)$ and $T^2(b_2)$, which are two-dimensional compact submanifolds of $T^4 = T^3 \times S^1$. Then we will identify the two holes in T^4 along the circles bounding the two tubes, after a flip of orientation.

So first consider $T^4 = T^3 \times S^1$ foliated by two-dimensional surfaces (orbit of \vec{w}_2) $\times S^1$. In this foliation there is one toral leaf, coming from the closed orbit of \vec{w}_2 . All other leaves are infinite cylinders, which course around T^4 similarly to the way the orbits of \vec{w}_2 course around T^3 .

The tori $T^2(b_1)$ and $T^2(b_2)$ are thought of as at the level $t = 1/2$, where t is the parameter of S^1 in the product $T^4 = T^3 \times S^1$. Thus each torus is orthogonal to the $\partial/\partial x$ and $\partial/\partial t$ directions. Each cylindrical leaf crosses the two tori, intersecting each in one point: these points have coordinates $(b_i, y, z, \frac{1}{2})$, $i = 1, 2$. Here y and z depend on the leaf, or rather on the particular passage of the leaf through the two tori.

Now we choose a very small radius r , so that the intervals of radius r around




FIGURE 4.6




FIGURE 4.7

b_1 and b_2 are disjoint and contained in the interval $[b, c]$ where the flow \vec{w}_2 goes in the $\partial/\partial\gamma$ direction). This r will be the radius of the tubular neighborhoods removed around the tori. The tube around $T^2(b_1)$, for instance, consists of a union of two-dimensional discs, with centers at the points of $T^2(b_1)$ and radius r , and situated in the x - t plane. That is, for each point $(b_1, y, z, \frac{1}{2})$ of $T^2(b_1)$, the two-disc in question is the following subset of T^4 :

$$D_{y,z}^1 = \{(x, y, z, t) \in T^4: (x - b_1)^2 + (t - 1/2)^2 \leq r^2\}.$$

The union of these discs for y, z ranging over the y - z plane constitutes the tubular neighborhood of $T^2(b_1)$. The tube around $T^2(b_2)$ is defined analogously, using discs $D_{y,z}^2$.

To complete the construction, we remove the interiors of the tubes around the two tori, then identify the two boundary components of T^4 —(two tubes). This identification should reverse orientation and sew together the corresponding circles, $\partial D_{y,z}^1$ and $\partial D_{y,z}^2$.

The reason for the flip in orientation is so we end up with a cylinder-with-handles as typical leaf (see Figure 4.6). We have already indicated why the distance between handles is unbounded. This completes the construction.

Finally we would like to show, as we have claimed, that the noncompact leaves L of Reeb's foliation of S^3 , which have the quasi-isometry type of infinite cylinders capped on one end are not quasi-homogeneous. Let x_0 be the point in the center of the cap, as in Figure 4.7. Define $B_n = B(x_0, n)$. Note that the n -cover ball $\tilde{B}(x_0, n)$ is identical with $B(x_0, n)$. Let C_k be cylinders of length k which occur in L far enough out (towards the right in Figure 4.7). For each n , quasi-homogeneity guarantees the existence of immersions $B_n \xrightarrow{f} C_k$ for all $k \geq$ some k_0 depending on n .

There are closed loops γ , as in Figure 4.7 all along any C_k , and all such γ have the same length, which we may take to be 1. Let (k, K) be the dilation bounds shared by the immersions $B_n \xrightarrow{f} C_k$.

Since γ has length 1, $f^{-1}(\gamma)$ if it exists can have length at most K . Consider B_n for some $n \geq K$. Using part (2) of quasi-homogeneity for this fixed ball B_n , we see that for large enough k the closed loop $\gamma \subset C_k$ will pull back via f to some closed loop $\bar{\gamma}$ in B_n . But such a loop $\bar{\gamma}$ in B_n can be shrunk to a constant loop in B_n , through paths of length less than some constant depending on (k, K) , n , but independent of the index “ k ” of C_k . This latter homotopy will then be carried

by f , for all sufficiently large k , to a homotopy of γ in C_k to a trivial loop in C_k through paths of length less than some other constant independent of k . This is impossible for k sufficiently large.

The argument that the noncompact leaves of Reeb's foliation are "essentially proper" is essentially the same as that just given; one need only note that a non-proper leaf must pass closer and closer to any part of itself, in particular to the initial segments B_n . Since the entire remote region of the surface is cylindrical, the parts that come close can be taken to be an increasing sequence of cylinders C_k , just as in the above proof. The same contradiction results.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF SAN DIEGO, ALCALA PARK, SAN DIEGO, CALIFORNIA

Current address: 6775 Glidden Street, San Diego, California 92111